
Introduction to Concurrent Programming Using Processes and Pthreads

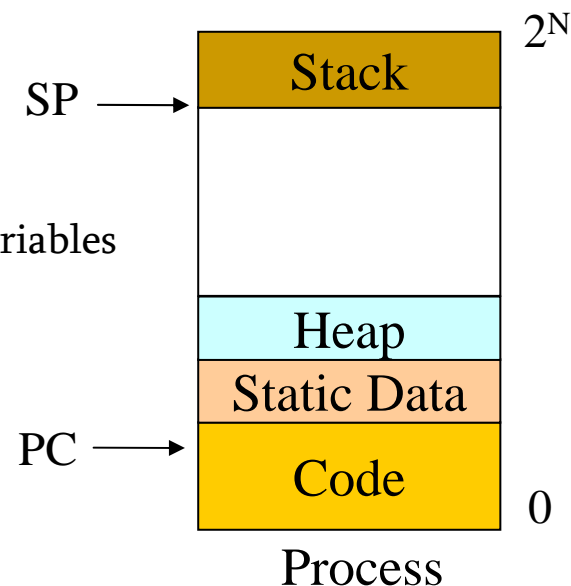
Professor: Panagiota Fatourou

TA: Eleftherios Kosmas

CSD - March 2012

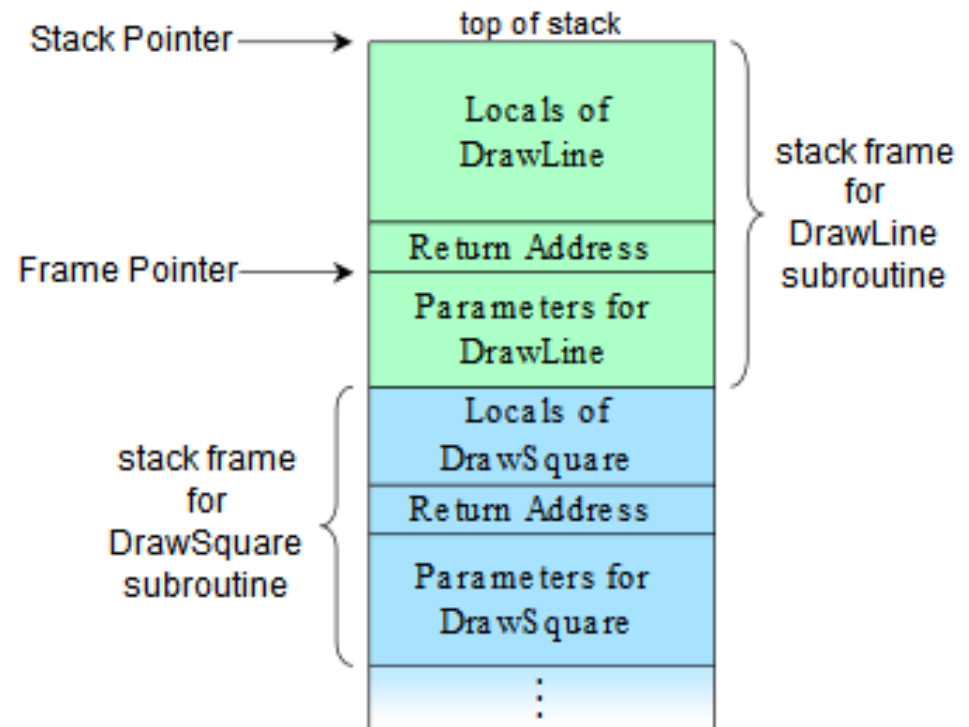
Process

- A **process** is an **instance** of a program that is being executed
 - It consists of, or it owns: (PCB - Process Control Block)
 - an executable **machine code**
 - **memory** (some private address space)
 - input and output data
 - **call stack**: keeps track of active subroutines
 - function parameters, return addresses, and local variables
 - **heap**: dynamically allocated memory
 - **static data**: global variables
 - file descriptors, socket descriptors, etc.
 - processor state (or **context**),
 - e.g., content of CPU registers
- (Program Counter - PC, Stack Pointer - SP, numeric)
- Allows program to act as if it owns the machine
 - **Multitasking** using **timesharing**
 - context switch, scheduling



Process - Call Stack

- It is composed of **stack frames** or **activation records**
- Each stack frame
 - corresponds to an active subroutine
 - is machine dependent



Multitasking

- Parent processes create children processes by calling `fork()`
- **`pid_t fork()`**
 - spawns a new child process
 - with separate address space
 - it has an exact copy of all the memory segments of the parent process
 - it has a new pid
 - returns
 - 0 to child
 - the pid of the newly created child to parent
- **`pid_t wait(int *status)`**
 - suspends execution of the calling process until one of its children terminates
 - status indicates reason for termination

Multitasking

- `pid_t waitpid (pid_t pid, int *status, int options)`
 - `pid=-1`: any child process
 - `pid>0`: specific child process
 - `pid=0`: any child process with some process group id
 - options can be used to wait or to check and proceed
- `int exit (int *status)`
 - executed by a child process when it wants to terminate
 - makes status available to parent

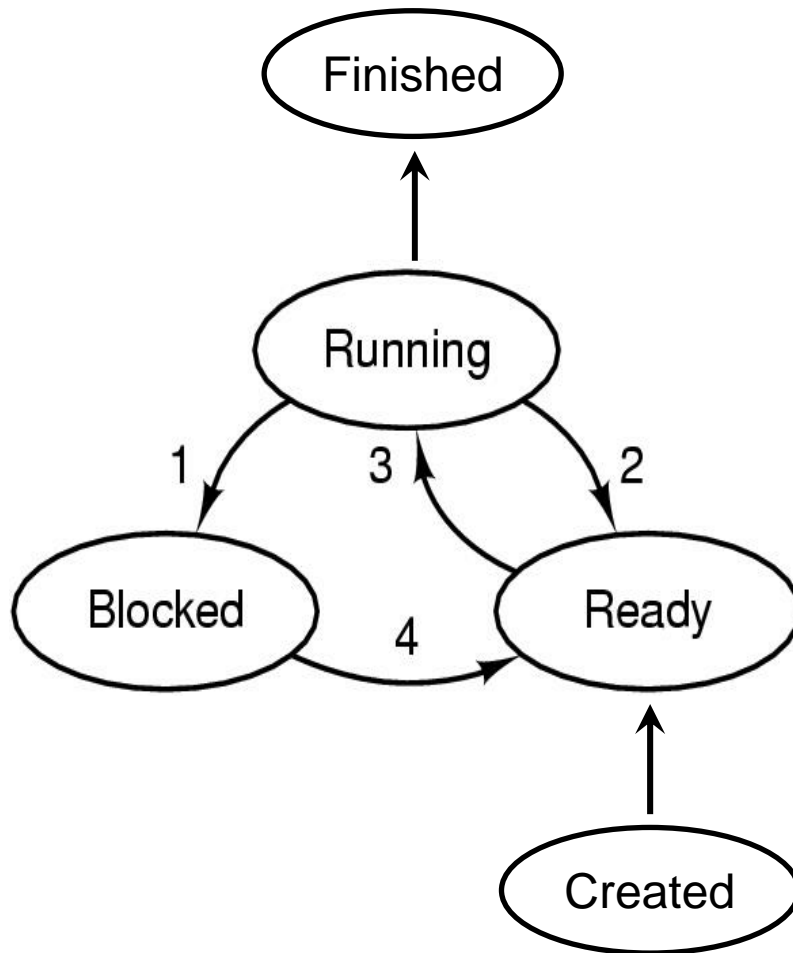
Multitasking - Example

```
1. #include <stdio.h> /* printf, stderr, fprintf */
2. #include <sys/types.h> /* pid_t */
3. #include <unistd.h> /* _exit, fork */
4. #include <stdlib.h> /* exit */
5. #include <errno.h> /* errno */
6.
7. int main(void)
8. {
9.     pid_t pid;
10.    int i;
11.
12.    for (i=0; i<10; i++) {
13.        /* Output from both the child and the parent process will be written to the standard output, as they both run at the same
14.         * time. */
15.        pid = fork();
16.        if (pid == -1) {
17.            /* Error: When fork() returns -1, an error happened (for example, number of processes reached the limit). */
18.            fprintf(stderr, "can't fork, error %d\n", errno);
19.            exit (EXIT_FAILURE);
20.        }
21.        else if (pid = 0) break; /* When fork() returns 0, we are in the child process. */
22.    }
23.    if (pid == 0) { /* When fork() returns 0, we are in the child process. */
24.        ...
25.        exit(0);
26.    }
27.    else { ... /* When fork() returns a positive number, we are in the parent process */
```

Multitasking - Example

```
1. #include <stdio.h> /* printf, stderr, fprintf */
2. #include <sys/types.h> /* pid_t */
3. #include <unistd.h> /* _exit, fork */
4. #include <stdlib.h> /* exit */
5. #include <errno.h> /* errno */
6.
7. int main(void)
8. {
9.     pid_t pid;
10.    for (i=0; i<10; i++) {
11.        pid = fork();
12.        ...
13.        else if (pid = 0) break; /* When fork() returns 0, we are in the child process. */
14.    }
15.
16.    if (pid == 0) { /* Child process: When fork() returns 0, we are in the child process. */
17.        ...
18.        exit(0);
19.    }
20.    else { /* When fork() returns a positive number, we are in the parent process */
21.        for(i=0; i<10; i++) {
22.            pid_t child = wait(0);
23.            printf ("Child with pid [%d] terminated\n", child);
24.        }
25.    }
26.    return 0;
27. }
```

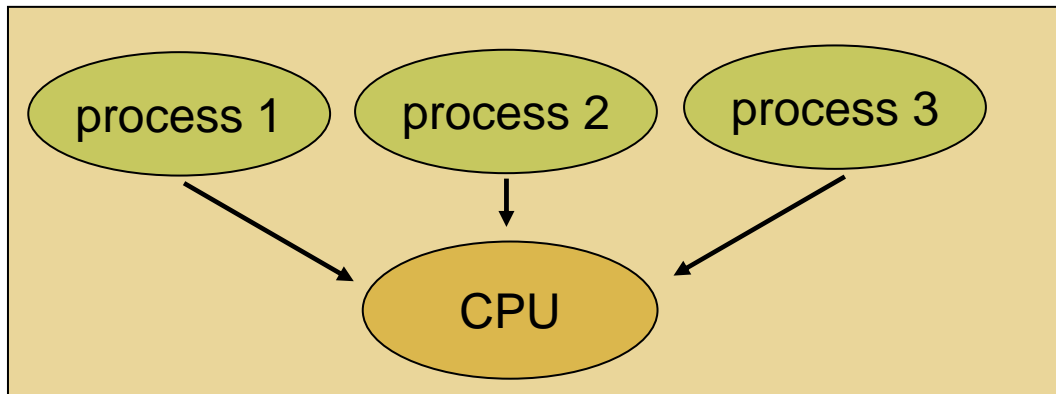
Multitasking - Process States



1. Process blocks for input
2. Scheduler picks another process
3. Scheduler picks this process
4. Input becomes available

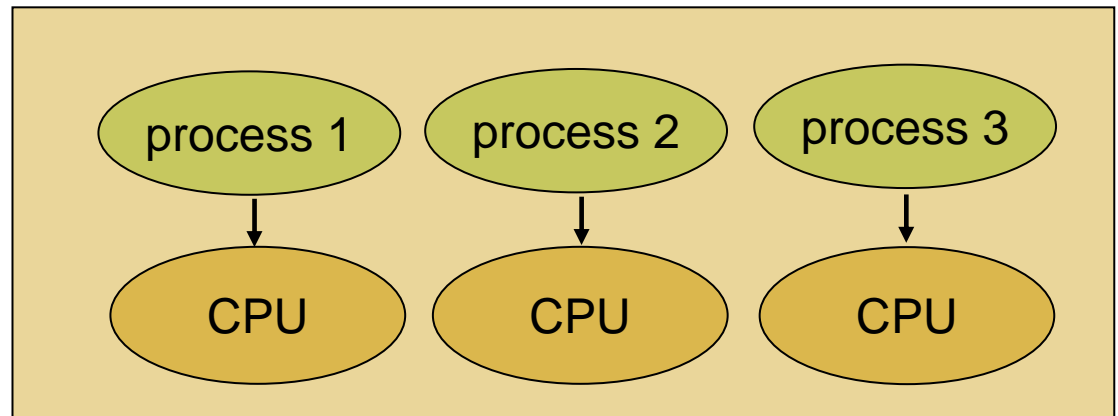
Multiprocessing

- When applying multitasking to a multicore (multiprocessor) machine we get **multiprocessing**



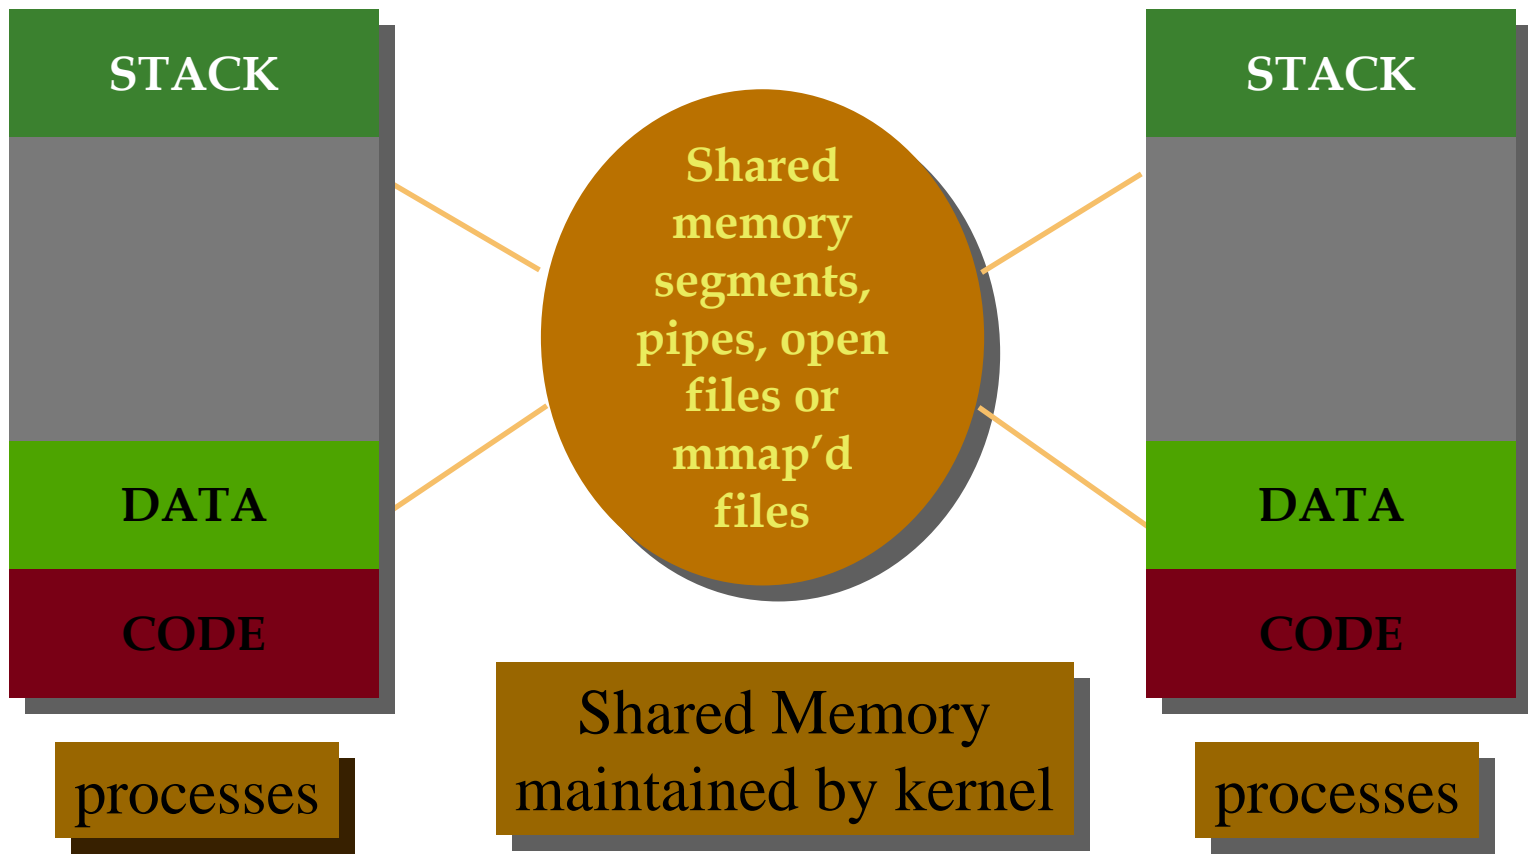
Multitasking

Multiprocessing



Inter Process Communication (IPC)

- processes do not share anything implicitly
- explicit actions are required to achieve IPC



Shared Memory IPC

- One process will create a memory portion which other processes (if permitted) can access
- `shmget ()`: is used to create a shared memory segment
 - a shared memory segment is described by a control structure (in `<sys/shm.h>`) with a unique ID that points to an area of physical memory
- `shmctl ()`:
 - used by the original owner of a shared memory segment can assign (or revoke) ownership to another user
 - used by processes with permission to perform various control functions
- `shmat ()`: attach a shared segment to a process address space
 - Once attached, the process can read or write to the segment
- `shmdt ()`: detach

Shared Memory IPC

- `int shmget(key_t key, size_t size, int shmflg):`
 - returns the **identifier** of shared memory segment associated with the value of the argument `key`
 - it is also used to get the **identifier** of an **existing** shared segment
 - `size`: the size in bytes of the requested shared memory
 - `shmflg`: specifies the initial access permissions and creation control flags

```
1. #include <sys/types.h>
2. #include <sys/ipc.h>
3. #include <sys/shm.h>
4. ...
5. key_t key; int shmflg; int shmid; int size;
6. ...
7. key = ...; size = ...; shmflg = ...;
8. if ((shmid = shmget (key, size, shmflg)) == -1) {
9.     perror("shmget: shmget failed"); exit(1);
10. }
11. else {
12.     fprintf(stderr, "shmget: shmget returned %d\n", shmid); exit(0);
13. }
14. ...
```

Shared Memory IPC

- `int shmctl(int shmid, int cmd, struct shmids *buf):`
 - is used to alter the permissions and other characteristics of a shared memory segment
 - **cmd**: SHM_LOCK, SHM_UNLOCK, IPC_STAT, IPC_SET, IPC_RMID
 - **buf**: shared memory data structure to hold results

```
1.  #include <sys/types.h>
2.  #include <sys/ipc.h>
3.  #include <sys/shm.h>
4.  ...
5.  int cmd; int shmid; struct shmids shmids;
6.  ...
7.  shmid = ...; cmd = ...;
8.  if ((rtrn = shmctl(shmid, cmd, shmids)) == -1) {
9.      perror("shmctl: shmctl failed");
10.     exit(1);
11. }
12. ...
```

Shared Memory IPC

- `void *shmat(int shmid, const void *shmaddr, int shmflg):`
 - returns a pointer, **shmaddr**, to the head of the shared segment associated with a valid **shmid**
 - **shmflag**: flags used on attach
- `int shmdt(const void *shmaddr):`
 - detaches the shared memory segment located at the address indicated by **shmaddr**

```
1.  #include <sys/types.h>
2.  #include <sys/ipc.h>
3.  #include <sys/shm.h>

4.  static struct state {
5.      int shmid; char *shmaddr; int shmflg;
6.  } ap[MAXnap]; /* State of current attached segments. */
7.  int nap; /* Number of currently attached segments. */
8.  ...
9.  char *addr; /* address work variable */
10. register int i; /* work area */
11. register struct state *p; /* ptr to current state entry */
12. ...
13. p = &ap[nap++];
14. p->shmid = ...; p->shmaddr = ...; p->shmflg = ...
```

Shared Memory IPC

- `void *shmat(int shmids, const void *shmaddr, int shmflg):`
- `int shmdt(const void *shmaddr):`

```
1.  p->shmaddr = shmat(p->shmids, p->shmaddr, p->shmflg);
2.  if(p->shmaddr == (char *)-1) {
3.      perror("shmop: shmat failed");
4.      nap--;
5.  }
6.  else fprintf(stderr, "shmop: shmat returned %#8.8x\n", p->shmaddr);
7.  ...
8.  i = shmdt(addr);
9.  if(i == -1) {
10.     perror("shmop: shmdt failed");
11. }
12. else {
13.     fprintf(stderr, "shmop: shmdt returned %d\n", i);
14.     for (p = ap, i = nap; i--; p++)
15.         if (p->shmaddr == addr) *p = ap[--nap];
16. }
17. ...
```

Shared Memory IPC - Example - Server

```
1. #include <sys/types.h>
2. #include <sys/ipc.h>
3. #include <sys/shm.h>
4. #include <stdio.h>

5. #define SHMSZ      27

6. main()
7. {
8.     char c, *shm, *s; int shmid; key_t key;
9.
10.    key = 1234;
11.    shmid = shmget(key, SHMSZ, IPC_CREAT | 0666);
12.    shm = shmat(shmid, NULL, 0);

13.    s = shm;
14.    for (c = 'a'; c <= 'z'; c++)
15.        *s++ = c;
16.    *s = NULL;

17.    while (*shm != '*')
18.        sleep(1);

19.    shmctl(shmid, IPC_RMID, 0)
20.    exit(0);
21. }
```


Shared Memory IPC - Example - Client

```
1. #include <sys/types.h>
2. #include <sys/ipc.h>
3. #include <sys/shm.h>
4. #include <stdio.h>

5. #define SHMSZ      27

6. main()
7. {
8.     int shmid; key_t key; char *shm, *s;

9.     key = 1234;
10.    shmid = shmget(key, SHMSZ, 0666);
11.    shm = shmat(shmid, NULL, 0);

12.    for (s = shm; *s != NULL; s++)
13.        putchar(*s);
14.    putchar('\n');

15.    *shm = '*';

16.    exit(0);
17. }
```

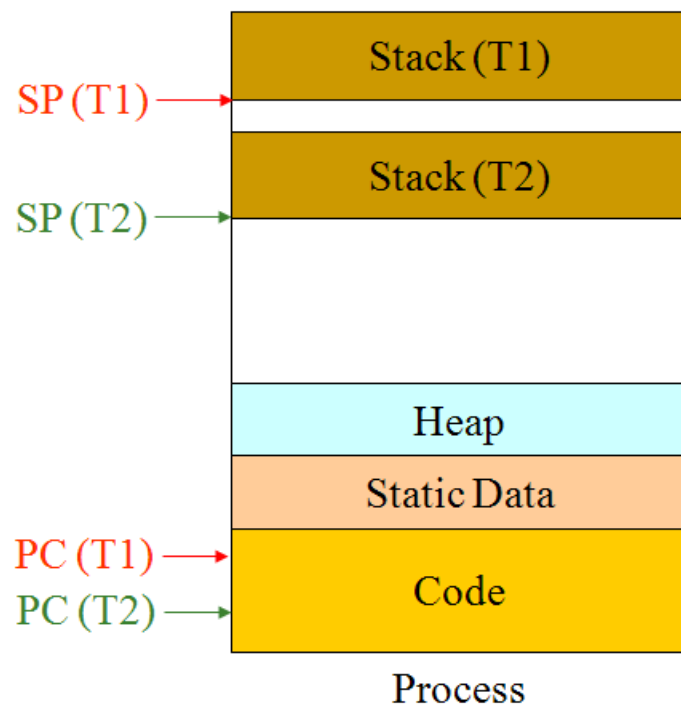

Threads

- A process is the **heaviest** unit of kernel scheduling
 - **creating** a new process is costly
 - data structures needed to be allocated and initialized
 - expensive **context switch**
 - **communication** among processes is costly, since it goes through the OS
 - IPC
 - overhead of system calls and copying data
- A process consists of:
 - i. a collection of resources
 - the code & address space, open files, etc.
 - ii. a **thread** of execution
 - the current state that operates on these resources

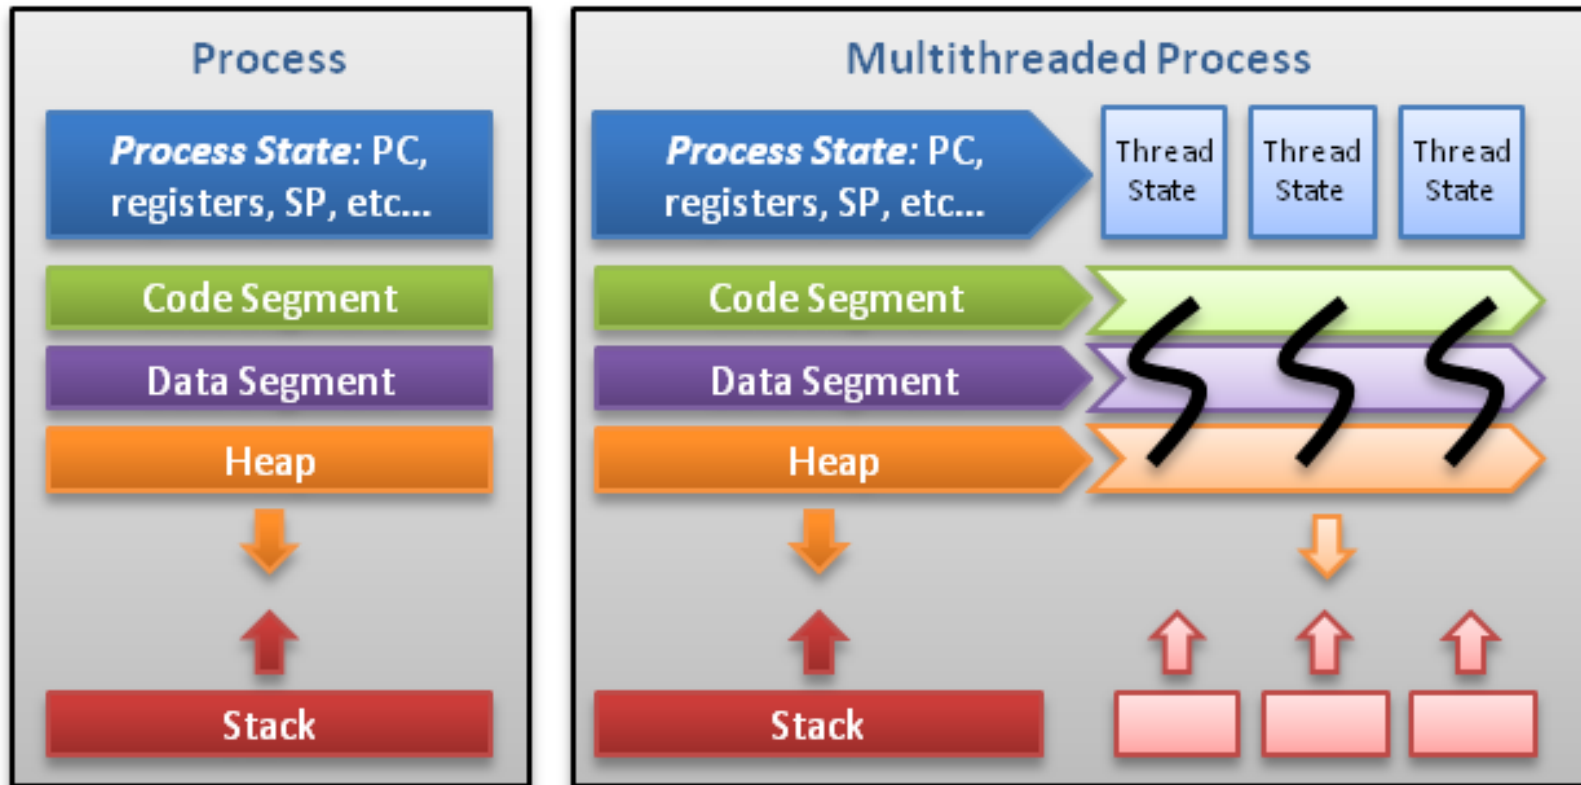
 The idea is to let multiple threads share a common address space

Threads

- Threads share the same memory (global variables, heap, file descriptors, etc.)
- Threads own a stack (including thread- local storage) and a copy of the registers (including PC and SP)
- Threads are executed in parallel
 - using time slices, in a single core machine
 - or really in parallel, in a multicore machine



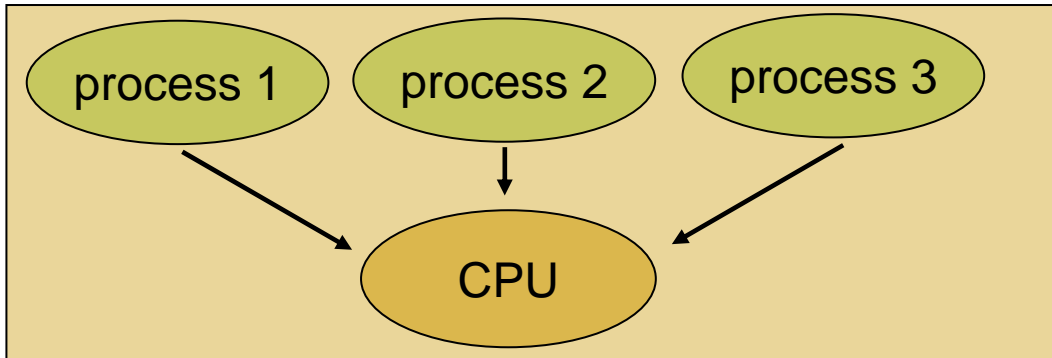
Single vs. Multi threaded



Threads contain only necessary information, such as a stack (for local variables, function arguments, return values), a copy of the registers, program counter and any thread-specific data to allow them to be scheduled individually. Other data is shared within the process between all threads.


Multithreading

- A way for program to split itself into multiple running tasks




Multitasking

Multiprocessing



Multithreading



Threads vs. Processes

- Threads
 - ✓ easier to create and destroy
 - ✓ inter-thread communication is cheaper
 - can use process memory and may not need (for user-level threads) to context switch
 - ✓ provide faster context switch
 - ✗ not secure: a thread can write the memory used by another thread
- Processes
 - ✓ secure: one process cannot corrupt another process
 - ✗ inter-process communication is expensive: need to context switch

Threads

- A **kernel thread** is the lightest unit of kernel scheduling
- Each process contains at least one kernel thread
- The kernel
 - may (on may not) assign one thread to each logical core, resulting to different models:
 - 1:1, Kernel-level threading
 - N:1, User-level threading
 - M:N, Hybrid threading
 - can swap out threads that get blocked
 - ✗ kernel threads take much longer than user threads to be swapped

Threads

- **User threads** are managed and scheduled in userspace
 - may run on top of several kernel threads to benefit from multi-processors
 - ✓ fast to create and manage
 - ✗ can not take full advantage of multithreading
 - ✗ they get blocked when all of their associated kernel threads are blocked, even if there are some user threads that are ready to run

POSIX Threads or Pthreads API

- **Thread management:** The first class of functions work directly on threads - creating, terminating, joining, etc.
- **Mutexes:** provide for creating, destroying, locking and unlocking mutexes.
- **Condition variables:** include functions to create, destroy, wait and signal based upon specified variable values.
- **Barriers**

Pthreads - Thread Management

- `int pthread_create(pthread_t *thread, pthread_attr_t *attr, void *(*start_routine)(void *), void *arg):`
 - `thread`: the actual thread object that contains pthread id
 - `attr`: attributes to apply to this thread
 - `start_routine`: the function this thread executes
 - `arg`: arguments to pass to thread function above
- `void pthread_exit(void *value_ptr):`
 - terminates the thread and provides `value_ptr` available to any `pthread_join()` call
- `int pthread_join(pthread_t thread, void **value_ptr):`
 - suspends the calling thread to wait for successful termination of `thread`
 - `value_ptr`: data passed from the terminating thread's call to `pthread_exit()`

Pthreads - Thread Management - Example

```
1.  #include <stdio.h>
2.  #include <stdlib.h>
3.  #include <pthread.h>
4.
5.  #define NUM_THREADS 2

                                     /*create thread argument struct for thr_func() */
6.  typedef struct _thread_data_t {
7.      int tid;
8.      double stuff;
9.  } thread_data_t;

10. void *thr_func(void *arg) { /* thread function */
11.     thread_data_t *data = (thread_data_t *)arg;

12.     printf("hello from thr_func, thread id: %d\n", data->tid);

13.     pthread_exit(NULL);
14. }
```

Pthreads - Thread Management - Example

```
1.  int main(int argc, char **argv) {
2.      pthread_t thr[NUM_THREADS];
3.      int i, rc;
4.
5.          /* create a thread_data_t argument array */
6.      thread_data_t thr_data[NUM_THREADS];
7.      for (i = 0; i < NUM_THREADS; ++i) { /* create threads */
8.          thr_data[i].tid = i;
9.          if ((rc = pthread_create(&thr[i], NULL, thr_func, &thr_data[i]))) {
10.             fprintf(stderr, "error: pthread_create, rc: %d\n", rc);
11.             return EXIT_FAILURE;
12.         }
13.     }
14.
15.     for (i = 0; i < NUM_THREADS; ++i) /* block until all threads complete */
16.         pthread_join(thr[i], NULL);
17.
18.     return EXIT_SUCCESS;
19. }
```

Pthreads - Thread Management

- `int pthread_create(pthread_t *thread, pthread_attr_t *attr, void *(*start_routine)(void *), void *arg):`
 - `attr`: attributes to apply to this thread

☞ Attributes can be specified using the following functions:

- `int pthread_attr_init(pthread_attr_t *attr)`
- `int pthread_attr_setdetachstate(pthread_attr_t *attr, int detachstate)`
- `int pthread_attr_setguardsize(pthread_attr_t *attr, size_t guardsize)`
- `int pthread_attr_setinheritsched(pthread_attr_t *attr, int inheritsched)`
- `int pthread_attr_setschedparam(pthread_attr_t *attr, const struct sched_param *param)`
- `int pthread_attr_setschedpolicy(pthread_attr_t *attr, int policy)`
- `int pthread_attr_setscope(pthread_attr_t *attr, int contentionscope)`
- `int pthread_attr_setstackaddr(pthread_attr_t *attr, void *stackaddr)`
- `int pthread_attr_setstacksize(pthread_attr_t *attr, size_t stacksize)`

☞ Attributes can be retrieved via corresponding get functions

Pthreads - Mutexes

- `int pthread_mutex_init (pthread_mutex_t *mutex, const pthread_mutexattr_t *mutexattr):`
 - initializes **mutex**
 - attributes for the mutex can be given through **mutexattr**
 - use NULL, to specify default attributes
- `int pthread_mutex_lock(pthread_mutex_t *mutex)`
 - blocks until mutex lock is acquired
- `int pthread_mutex_trylock(pthread_mutex_t *mutex)`
 - non-blocking, may return without acquiring the mutex lock
- `int pthread_mutex_unlock(pthread_mutex_t *mutex)`

Pthreads - Mutexes - Example

```
1.     ...
2.     typedef struct _thread_data_t {
6.         int tid;
7.         double stuff;
8.     } thread_data_t;

9.     double shared_x;                /* shared data between threads */
10.    pthread_mutex_t lock_x;         /* shared data between threads */

11.    void *thr_func(void *arg) {     /* thread function */
12.        thread_data_t *data = (thread_data_t *)arg;

13.        printf("hello from thr_func, thread id: %d\n", data->tid);

14.                /* get mutex before modifying and printing shared_x */
15.        pthread_mutex_lock(&lock_x);
16.        shared_x += data->stuff;
17.        printf("x = %f\n", shared_x);
18.        pthread_mutex_unlock(&lock_x);

19.        pthread_exit(NULL);
20.    }
```

Pthreads - Mutexes - Example

```
1.  int main(int argc, char **argv) {
2.      pthread_t thr[NUM_THREADS]; int i, rc;
3.      thread_data_t thr_data[NUM_THREADS];
4.
5.      shared_x = 0; /* initialize shared data */
6.          /* initialize pthread mutex protecting "shared_x" */
7.      pthread_mutex_init(&lock_x, NULL);
8.
9.      for (i = 0; i < NUM_THREADS; ++i) {          /* create threads */
10.         thr_data[i].tid = i;
11.         thr_data[i].stuff = (i + 1) * NUM_THREADS;
12.         if ((rc = pthread_create(&thr[i], NULL, thr_func, &thr_data[i]))) {
13.             fprintf(stderr, "error: pthread_create, rc: %d\n", rc);
14.             return EXIT_FAILURE;
15.         }
16.     }
17.
18.     for (i = 0; i < NUM_THREADS; ++i) /* block until all threads complete */
19.         pthread_join(thr[i], NULL);
20.
21.     return EXIT_SUCCESS;
22. }
```


Pthreads - Condition Variables

- `int pthread_cond_init(pthread_cond_t *cond, pthread_condattr_t *cond_attr):`
 - initialized condition variable `cond`
 - attributes for `cond` can be given through `cond_attr`
 - use NULL, to specify default attributes
- `int pthread_cond_wait(pthread_cond_t *cond, pthread_mutex_t *mutex)`
 - puts the current thread to sleep, waiting on `cond` for `mutex` to be released
- `int pthread_cond_signal(pthread_cond_t *cond)`
 - signals one thread out of the possibly many sleeping threads waiting on `cond` to wakeup
- `int pthread_cond_broadcast(pthread_cond_t *cond)`
 - signals all threads waiting on `cond` to wakeup

Pthreads - Condition Variables - Example

```
1. void *thr_func1(void *arg) {
2.     pthread_mutex_lock(&count_lock); /*thread code blocks here until MAX_COUNT is reached*/
3.     while (count < MAX_COUNT)
4.         pthread_cond_wait(&count_cond, &count_lock);
5.     pthread_mutex_unlock(&count_lock);
6.     ...
7.     pthread_exit(NULL);
8. }
9.
10. /*some other thread code that signals a waiting thread that MAX_COUNT has been reached*/
11. void *thr_func2(void *arg) {
12.     pthread_mutex_lock(&count_lock);
13.
14.     /* some code here that does interesting stuff and modifies count */
15.
16.     if (count == MAX_COUNT) {
17.         pthread_mutex_unlock(&count_lock);
18.         pthread_cond_signal(&count_cond);
19.     }
20.     else pthread_mutex_unlock(&count_lock);
21.
22.     pthread_exit(NULL);
23. }
```

Pthreads - Barrier

- `int pthread_barrier_init(pthread_barrier_t *barrier, pthread_barrierattr_t *barrier_attr, unsigned int count)`
 - initialized **barrier**
 - attributes for **barrier** can be given through **barrier_attr**
 - use NULL, to specify default attributes
 - **count**: defines the number threads that must join the barrier for the barrier to reach completion and unblock all threads waiting at the barrier
- `int pthread_barrier_wait(pthread_barrier_t *barrier)`

Pthreads - Additional Useful Functions

- `pthread_kill()`, can be used to deliver signals to specific threads.
- `pthread_self()`, returns a handle on the calling thread.
- `pthread_equal()`, compares for equality between two pthread ids
- `pthread_once()`, can be used to ensure that an initializing function within a thread is only run once.

Pthreads - How to compile

- C : `gcc -lpthread multithreaded_app.c`
- C++ : `g++ -lpthread multithreaded_app.cxx`
- Other platforms and compilers may require `-pthread` flag instead