

No. 3

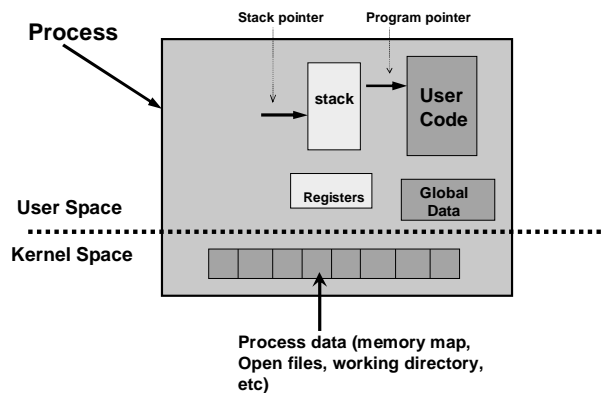
Thread

Prof. Hui Jiang
Dept of Computer Science and Engineering
York University


Process vs. Thread

- Traditional process contains a single stream of control.
(one process can do one thing at a time)
- Multithreaded process: contains several different streams of control.
Each stream is called a thread of this process
(multithreaded process can do multiple jobs simultaneously)
- A multi-threaded process contains several threads.
- Each thread includes:
 - A thread ID
 - A program counter
 - A register set
 - A stack & stack pointer
- All threads in a process share:
 - Code section & data section
 - OS resources (memory map, open devices, accounting, etc.)


One single-threaded Process



Multiple single-threaded Process




One multi-threaded Process



Comparison

- One single-threaded process:
 - can do one thing at a time
- Multiple single-threaded processes:
 - can do many things at the same time
- One multi-threaded process
 - Also can do many things at the same time
- Why multiple thread??
 - Multi-threaded process requires less OS resources (memory)
 - More efficient for OS to handle threads than process

Multithreading



Multithreading Example on a Uniprocessor

Benefits to use threads

- Threads occupy less memory than processes.
- Takes less time to create a new thread than a process.
- Less time to terminate a thread than a process.
- Less time to switch between two threads within the same process.
- Since threads within the same process share memory and files, they can communicate with each other without invoking the kernel.

Reentrant and thread-safe code

- To be thread safe, the program must be reentrant:
 - Program never modifies itself.
 - No use of static/global data.
 - Each Function calling keeps track of its own progress.

Non-reentrant C code

```
int delta;

int diff (int x, int y)
{
    delta = y - x;

    if (delta < 0) delta = -delta;

    return delta;
}
```

Reentrant C code

```
int diff (int x, int y)
{
    int delta;

    delta = y - x;


    if (delta < 0) delta = -delta;

    return delta;
}
```

User Thread

- User thread: supported above the kernel and implemented by a thread library in user space.
 - The library supports thread creation, scheduling, management with no support from the kernel.
 - User threads are fast to create and manage (no need to make a system call to trap to the kernel).
 - The kernel is not aware of the existence of threads.
 - User thread must be mapped to the kernel to execute.
 - Examples:
 - POSIX Pthread
 - Mach C-threads
 - Solaris UI-threads


Pure User Thread: many-to-one mapping




Kernel Threads

- Kernel threads are supported directly by OS.
- The kernel performs thread creation, scheduling, and management in the kernel space.
- Slow to maintain (need system call to kernel space).
- Each kernel thread can run totally independently:
 - One thread blocks, the kernel will schedule another thread to run.
 - Several kernel threads can run in parallel if many CPU's are available.
- OS to support kernel thread:
 - Windows NT/2000/XP
 - Solaris 2
 - Linux


Pure Kernel Thread: one-to-one mapping



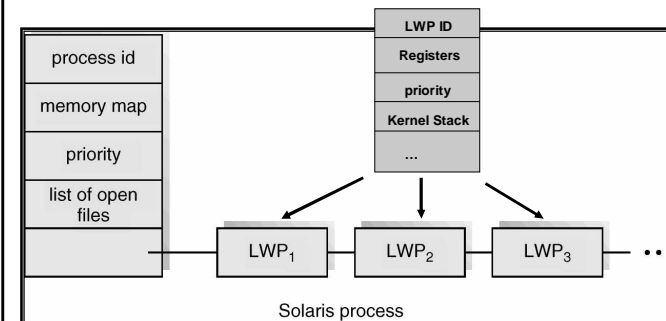
Combined Model: many-to-many mapping



Solaris Threads



Thread data structure in Solaris



Threading Issues

- ***fork()* and *exec()* implementation**
 - One thread in a process call *fork()*, it duplicates all threads in the process or just one calling thread.
 - One thread calls *exec()*, it will replace the entire process
- **Thread cancellation: terminating a thread before it finishes.**
 - Asynchronous cancellation
 - Deferred cancellation
- **Signal Handling**
 - Deliver the signal to the thread to which the signal applies.
 - Deliver the signal to every thread in the process
 - Deliver the signal to certain threads in the process
 - Assign a specific thread to receive all signals for the process

Thread Pools

- Create a number of threads at process start-up, place them into a pool, where they sit and wait for work.
- When the process receives a request, it awakens a thread from the pool, and serves the request immediately.
- Once the thread completes, it returns to the pool.
- If the pool contains no available thread, the server waits until one becomes free.
- **Benefits of thread pools:**
 - Faster to service a request.
 - Thread pool limits the total number of threads in system (no overload).

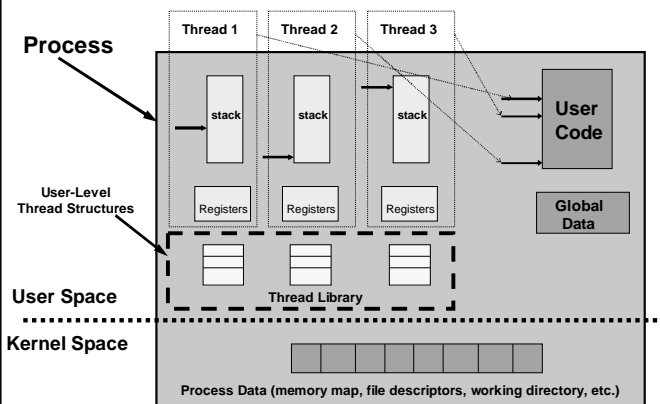
Linux Thread

- Linux uses pure kernel thread method with the one-to-one mapping.
- *fork()* creates a new process
 - Create a new memory space for new process
 - Copy from the address space of the calling process
- *clone()* simulates *fork()*, but
 - It does not create new memory space
 - The new process shares the same address space of the original process
 - → two processes sharing the same memory space (something like thread)

User Threads: Pthreads


- A POSIX standard (IEEE 1003.1c) API for thread creation and synchronization
- API specifies behavior of the thread library, implementation is up to development of the library
- Common in UNIX operating systems (Solaris, Linux, Mac OS X)

Multi-threaded Process



Multiple-thread programming

- User thread vs. kernel thread
- Multi-threaded programming with POSIX thread (Pthread)



POSIX Thread (1)

- Creation and termination

```
#include <pthread.h>
```

```
pthread_create(pthread_t *thread, const pthread_attr_t  
*attr, void *(*start) (void *), void *argv) ;
```

```
pthread_exit(void *value_ptr) ;
```

POSIX thread(2)

- Wait for another thread to terminate

```
pthread_join(pthread_t thread, void **value_ptr) ;
```

- Cancellation

```
pthread_cancel(pthread_t thread) ;
```

- Others

```
pthread_self(void) ;
```

```
pthread_detach(pthread_t thread) ;
```

```
pthread_attr_init(pthread_attr_t *attr) ;
```

Example 1: thread.c

- Example: thread.c (How to use Pthread)
- Two threads:
 - *main()* thread
 - *runner()* thread


Example 2: alarm.c

- Example 1: alarm.c (No thread)
- Example 2: alarm_fork.c (multiple process)
- Example 3: alarm_thread.c (multiple thread)


Three Models to use Threads

- Pipeline
 - Assembly line: each thread repeatedly performs the same operation on a sequence of data sets, passing each result to another thread for next step.
- Work Crew
 - Each thread performs an operation on its own data independently, then combine all results to get the final.
- Client/Server
 - A client contacts with an independent server for each job.

Pipeline



Work Crew



Client/Server

