

CS 134:
Operating Systems
Processes

2013-05-17 CS34

CS 134:
Operating Systems
Processes

Overview

Processes

- Processes in Unix
- Implementation
- States

Threads

- Concepts
- Uses
- Models
- Design

2013-05-17 CS34

└ Overview

Overview

Processes
Processes in Unix
Implementation
States

Threads
Concepts
Uses
Models
Design

Processes & Concurrency

2013-05-17
CS34
└ Processes
└ Processes & Concurrency

What is a process?

What is concurrency?

What *is* a process?

What *is* concurrency?

User's View of Processes

A fundamental OS abstraction

- ▶ But details vary from OS to OS:
 - ▶ Batch system—Jobs
 - ▶ Time-shared systems—User programs or tasks
- ▶ Common idea: Process = “A program in execution”
- ▶ Processes have a degree of independence from each other
 - ▶ Possibly only allowed communicate through designated mechanisms
 - ▶ One errant processes should not affect other unrelated ones

2013-05-17

 CS34
 └ Processes

└ User's View of Processes

User's View of Processes

- A fundamental OS abstraction
- ▶ But details vary from OS to OS:
 - Batch system—Jobs
 - Time-shared systems—User programs or tasks
 - ▶ Common idea: Process = “A program in execution”
 - ▶ Processes have a degree of independence from each other
 - Possibly only allowed communicate through designated mechanisms
 - One errant processes should not affect other unrelated ones

Class Exercise

2013-05-17

CS34

└ Processes

└ Class Exercise

Class Exercise

What makes up a process? ("A process has...")

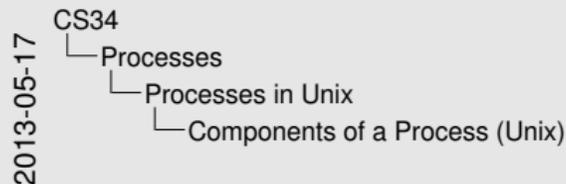
- In general
- On a typical POSIX system

What makes up a process? ("A process has...")

- ▶ In general
- ▶ On a typical POSIX system

Components of a Process (Unix)

- ▶ Execution state
 - ▶ Registers
 - ▶ Program counter
 - ▶ Program status word
 - ▶ Stack pointer
- ▶ Scheduling information
 - ▶ Process state
 - ▶ Priority
 - ▶ Class, etc.
- ▶ Memory
 - ▶ Text area
 - ▶ Data area
 - ▶ Stack area
- ▶ Security/Authentication Info
 - ▶ User ID
 - ▶ Group ID
- ▶ I/O State
 - ▶ File descriptors
 - ▶ Working directory
 - ▶ Root directory
- ▶ Event Notifications
 - ▶ Signals waiting
 - ▶ Signal mask
 - ▶ Time of next alarm
- ▶ Other
 - ▶ Process ID
 - ▶ Parent process
 - ▶ Process group
 - ▶ Controlling terminal
 - ▶ Start time
 - ▶ CPU time
 - ▶ Children's CPU time



Components of a Process (Unix)

- Execution state
 - Registers
 - Program counter
 - Program status word
 - Stack pointer
- Scheduling information
 - Process state
 - Priority
 - Class, etc.
- Memory
 - Text area
 - Data area
 - Stack area
- Security/Authentication Info
 - User ID
 - Group ID
- I/O State
 - File descriptors
 - Working directory
 - Root directory
- Event Notifications
 - Signals waiting
 - Signal mask
 - Time of next alarm
- Other
 - Process ID
 - Parent process
 - Process group
 - Controlling terminal
 - Start time
 - CPU time
 - Children's CPU time

Processes under UNIX

2013-05-17

- CS34
 - Processes
 - Processes in Unix
 - Processes under UNIX

Processes under UNIX

Processes:

- Create with `fork`
 - Exit with `exit`
 - Replace "process image" with `execve`
- Multiple processes may be active at any one time (compare w/ unprogrammed system)

Processes:

- ▶ Create with `fork`
- ▶ Exit with `exit`
- ▶ Replace "process image" with `execve`

Multiple processes may be active at any one time (compare w/ unprogrammed system)

Class Question

2013-05-17
CS34
└─ Processes
 └─ Processes in Unix
 └─ Class Question

Class Question

If there's `fork`, should we have `join`?

If there's `fork`, should we have `join`?

Process Implementation

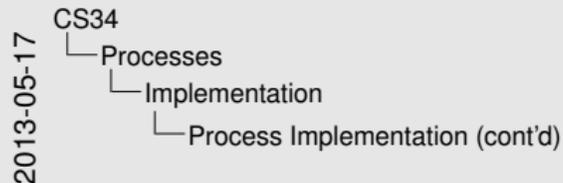
2013-05-17

- CS34
 - Processes
 - Implementation
 - Process Implementation

How does the OS implement the process abstraction?

How does the OS implement the process abstraction?

Process Implementation (cont'd)



Process Implementation (cont'd)

The OS needs to maintain a process image for each process:

- Process's address space, containing:
 - Program code
 - Program data
 - Processor stack
- Housekeeping information (PCB)
 - One of most important is process state

The OS needs to maintain a *process image* for each process:

- ▶ Process's address space, containing:
 - ▶ Program code
 - ▶ Program data
 - ▶ Processor stack
- ▶ Housekeeping information (*PCB*)
 - ▶ One of most important is *process state*

A Two-State Process Model

2013-05-17

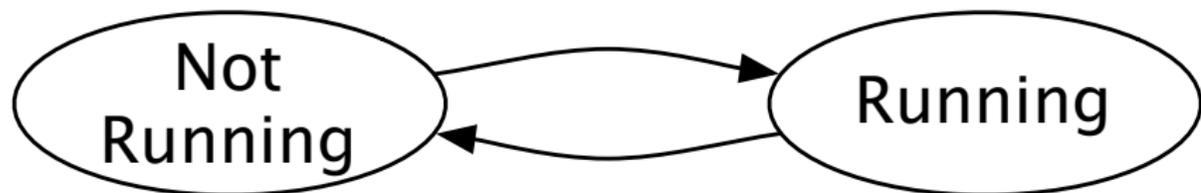
- CS34
 - Processes
 - States
 - A Two-State Process Model

A Two-State Process Model

Simplest model for processes:

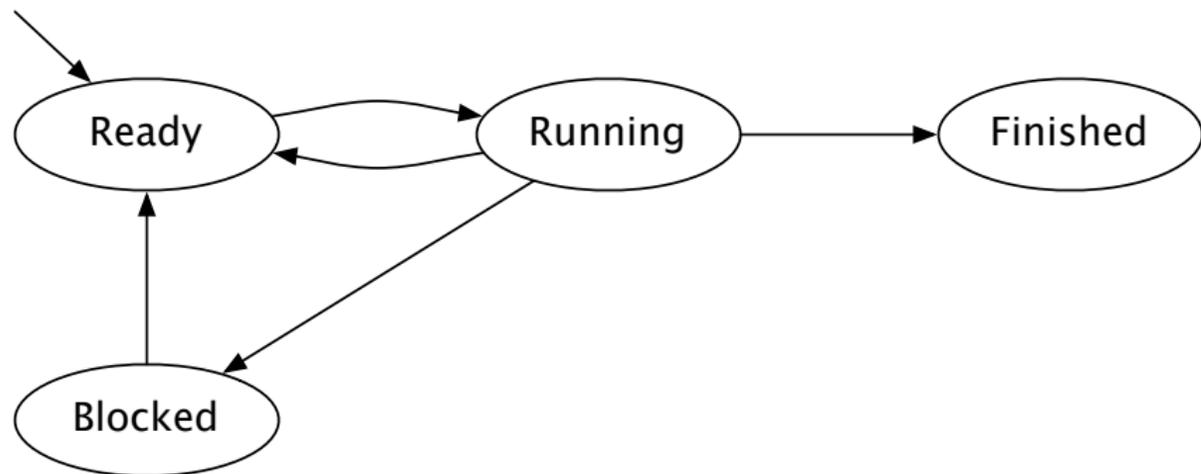


Simplest model for processes:



A Four-State Process Model

More useful model for processes:



2013-05-17

CS34
├── Processes
│ └── States
│ └── A Four-State Process Model

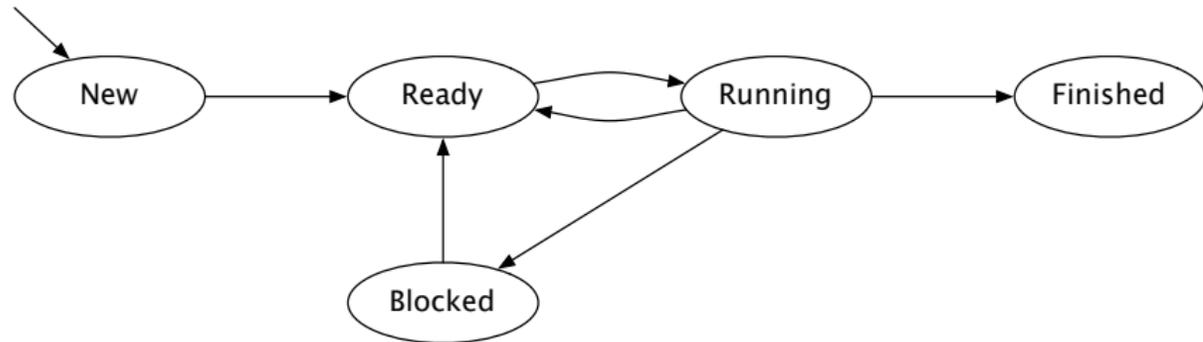
A Four-State Process Model

More useful model for processes:



A Five-State Process Model

Five states can model additional needs of batch systems:



Scheduler queues:

- ▶ **Ready queue:** Processes ready and waiting to execute.
- ▶ **New queue:** Processes waiting to be created

2013-05-17

CS34
├── Processes
│ ├── States
│ └── A Five-State Process Model

A Five-State Process Model

Five states can model additional needs of batch systems:



Scheduler queues:

- ▶ **Ready queue:** Processes ready and waiting to execute.
- ▶ **New queue:** Processes waiting to be created

Generalizing Processes

2013-05-17

- CS34
 - Threads
 - Concepts
 - Generalizing Processes

Generalizing Processes

Simple view of process is

Address space

+ Thread of execution

Does the mapping need to be one-to-one?

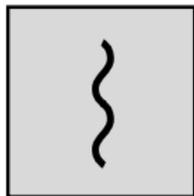
Simple view of process is

Address space

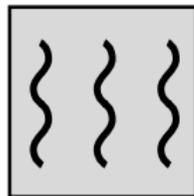
+ Thread of execution

Does the mapping need to be one-to-one?

Possible Mappings



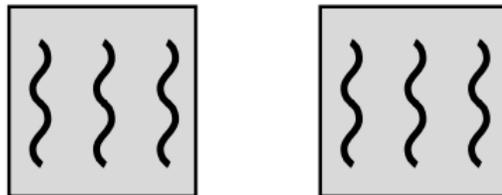
one process
one thread



one process
multiple threads



multiple processes
one thread per process



multiple processes
multiple threads per process

2013-05-17
CS34
├── Threads
│ ├── Concepts
│ └── Possible Mappings

Possible Mappings



one process

one thread



one process

multiple threads



multiple processes

one thread per process



multiple processes

one thread per process



multiple processes

multiple threads per process

Threads

2013-05-17
CS34
└─ Threads
 └─ Concepts
 └─ Threads

Threads

Motivation:

- ▶ Traditional processes: Virtual uniprocessor machine
- ▶ Multithreaded processes: Virtual multiprocessor machine

Uses of Threads

2013-05-17
CS34
└─ Threads
 └─ Uses
 └─ Uses of Threads

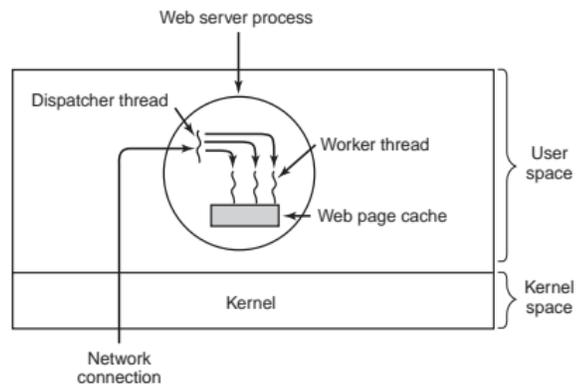
Uses of Threads

- Various reasons why people use threads
- ▶ Performing foreground and background work
 - ▶ Supporting asynchronous processing
 - ▶ Speeding execution
 - ▶ Organizing programs

Various reasons why people use threads

- ▶ Performing foreground and background work
- ▶ Supporting asynchronous processing
- ▶ Speeding execution
- ▶ Organizing programs

Uses of Threads—Example



```

/* Dispatcher Thread */
for ( ; ; ) {
    url = get_next_request();
    handoff_work(url);
}

```

```

/* Worker Thread */ \
for ( ; ; ) {
    url = wait_for_work();
    page = look_in_cache(url);
    if (page == NULL)
        page = generate_page(url);
    send_page(page);
}

```

2013-05-17

CS34
 Threads
 Uses
 Uses of Threads—Example

Uses of Threads—Example



```

/* Dispatcher Thread */
for ( ; ; ) {
    url = get_next_request();
    handoff_work(url);
}

/* Worker Thread */ \
for ( ; ; ) {
    url = wait_for_work();
    page = look_in_cache(url);
    if (page == NULL)
        page = generate_page(url);
    send_page(page);
}

```

Class Exercise

2013-05-17
CS34
├── Threads
│ ├── Uses
│ └── Class Exercise

Class Exercise

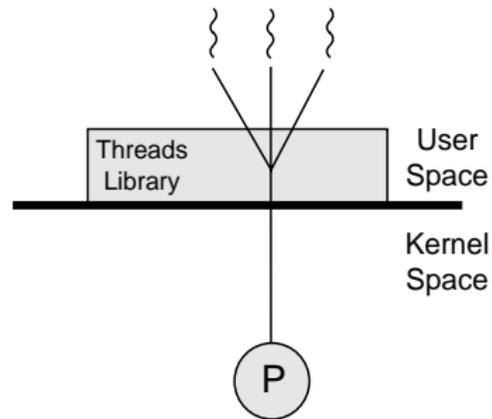
Can an application implement threads without built-in thread support in the OS?

If so, what does it need from the OS to support threads?

Can an application implement threads without built-in thread support in the OS?

If so, what *does* it need from the OS to support threads?

Model for User Threads



Pure user-level

Key:

-  User-level thread
-  Kernel-level thread
-  Process

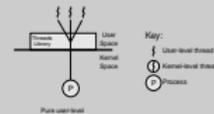
Class Exercise

What are the pros and cons of this approach?

2013-05-17

CS34
 └─ Threads
 └─ Models
 └─ Model for User Threads

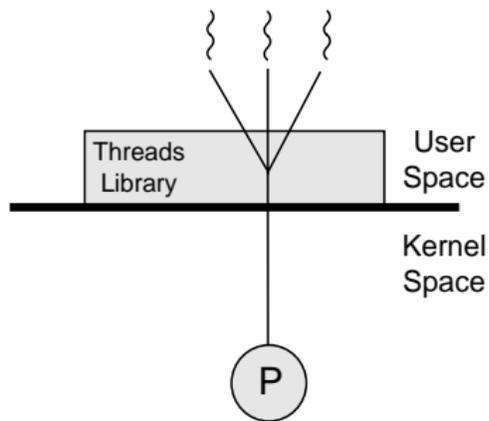
Model for User Threads



Class Exercise
 What are the pros and cons of this approach?

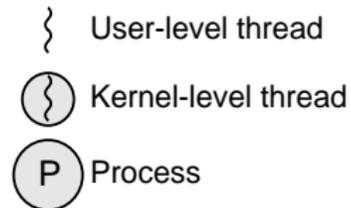
So, maybe we should put the threads in the kernel?

Model for User Threads



Pure user-level

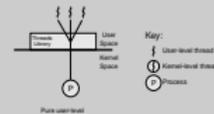
Key:



- + No kernel overhead for thread library calls
- + Own scheduler = Application-specific scheduling policy?
- I/O issues
- Can't (easily) take advantage of multiprocessing

2013-05-17
 CS34
 Threads
 Models
 Model for User Threads

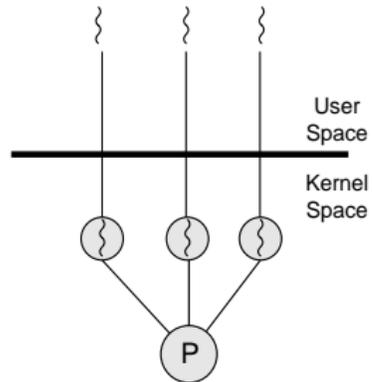
Model for User Threads



- No kernel overhead for thread library calls
- Own scheduler = Application-specific scheduling policy?
- I/O issues
- Can't (easily) take advantage of multiprocessing

So, maybe we should put the threads in the kernel?

Model for Kernel-Level Threads



Pure kernel-level

Key:

-  User-level thread
-  Kernel-level thread
-  Process

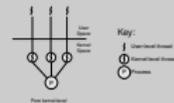
Class Exercise

What are the pros and cons of this approach?

2013-05-17

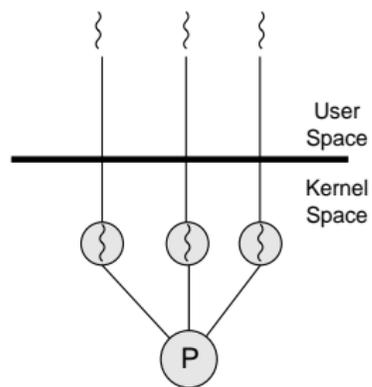
CS34
 Threads
 Models
 Model for Kernel-Level Threads

Model for Kernel-Level Threads



Class Exercise
 What are the pros and cons of this approach?

Model for Kernel-Level Threads



Pure kernel-level

Key:

-  User-level thread
-  Kernel-level thread
-  Process

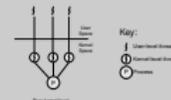
Now we have kernel overheads:

- ▶ Kernel data structures
- ▶ Mode switch to kernel

2013-05-17

CS34
 Threads
 Models
 Model for Kernel-Level Threads

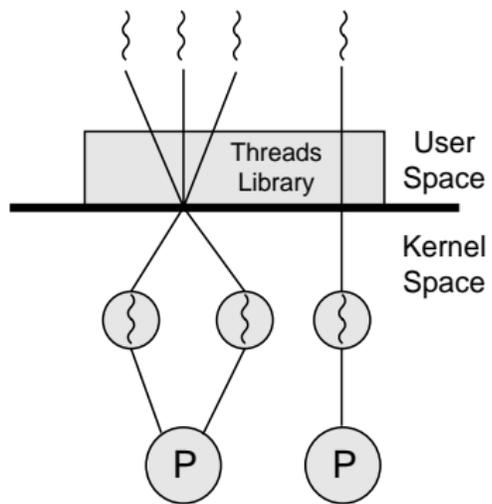
Model for Kernel-Level Threads



Now we have kernel overheads:

- ▶ Kernel data structures
- ▶ Mode switch to kernel

Hybrid Thread Schemes



Combined

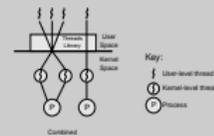
Key:

-  User-level thread
-  Kernel-level thread
-  Process

2013-05-17

CS34
 Threads
 Models
 Hybrid Thread Schemes

Hybrid Thread Schemes



Class Exercise

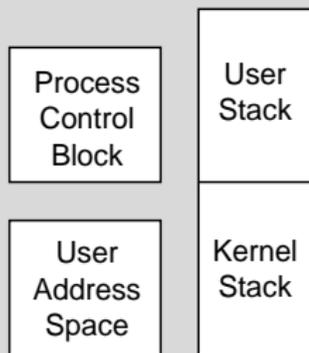
What are the pros and cons of this approach?

Class Exercise

What are the pros and cons of this approach?

Traditional vs. Multithreaded Processes

Single-Threaded Process Model



2013-05-17

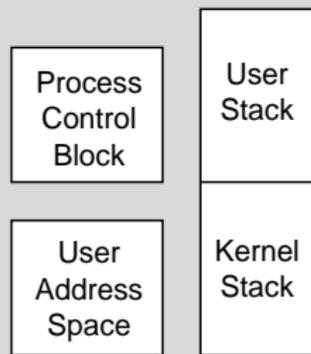
- CS34
 - Threads
 - Models
 - Traditional vs. Multithreaded Processes

Traditional vs. Multithreaded Processes

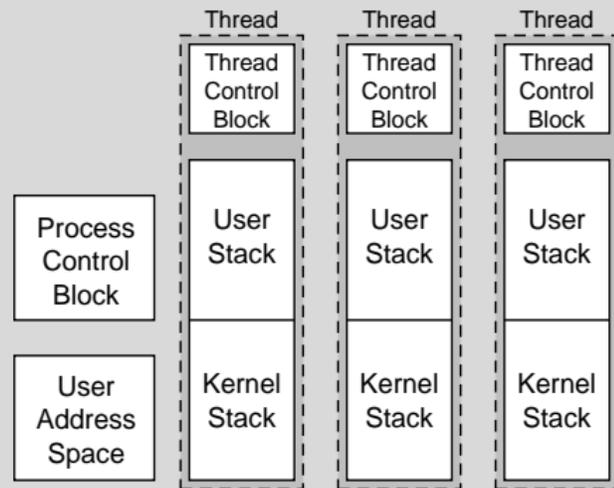


Traditional vs. Multithreaded Processes

Single-Threaded Process Model



Multithreaded Process Model

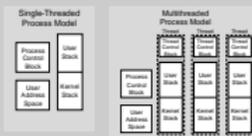


Class Question

But what's per-process and what's per-thread?

2013-05-17
 CS34
 Threads
 Models
 Traditional vs. Multithreaded Processes

Traditional vs. Multithreaded Processes



Class Question
 But what's per-process and what's per-thread?

Per-Process vs. Per-Thread—*You Decide* . . .

- ▶ Execution state
 - ▶ Registers
 - ▶ Program counter
 - ▶ Program status word
 - ▶ Stack pointer
- ▶ Scheduling information
 - ▶ Process state
 - ▶ Priority
 - ▶ Class, etc.
- ▶ Memory
 - ▶ Text area
 - ▶ Data area
 - ▶ Stack area
- ▶ Security/Authentication Info
 - ▶ User ID
 - ▶ Group ID
- ▶ I/O State
 - ▶ File descriptors
 - ▶ Working directory
 - ▶ Root directory
- ▶ Event Notifications
 - ▶ Signals waiting
 - ▶ Signal mask
 - ▶ Time of next alarm
- ▶ Other
 - ▶ Process ID
 - ▶ Parent process
 - ▶ Process group
 - ▶ Controlling terminal
 - ▶ Start time
 - ▶ CPU time
 - ▶ Children's CPU time

2013-05-17

CS34
 └─ Threads
 └─ Design
 └─ Per-Process vs. Per-Thread—*You Decide* . . .

Per-Process vs. Per-Thread—*You Decide* . . .

- Execution state
 - Registers
 - Program counter
 - Program status word
 - Stack pointer
- Scheduling information
 - Process state
 - Priority
 - Class, etc.
- Memory
 - Text area
 - Data area
 - Stack area
- Security/Authentication Info
 - User ID
 - Group ID
- I/O State
 - File descriptors
 - Working directory
 - Root directory
- Event Notifications
 - Signals waiting
 - Signal mask
 - Time of next alarm
- Other
 - Process ID
 - Parent process
 - Process group
 - Controlling terminal
 - Start time
 - CPU time
 - Children's CPU time