

---

## Threads and Processes — Part 1

COMP3231/COMP9201 Operating Systems

2003/S2

---

Slide 1

### MAJOR REQUIREMENTS OF AN OS

- Interleave the execution of several programs
    - to maximize utilization of CPU and other resources while providing reasonable response time
    - to support multiple user working interactively
    - for convenience (e.g., compile program while editing other file)
  - Allocate resources required for execution of programs
  - Support communication between executing programs
- 

Slide 2

Previously, we listed several definitions of the term **Process**:

- \* A program in execution
- \* An instance of a program running on a computer
- \* A unit of execution characterised by
  - a single, sequential thread of execution
  - a current state
  - an associated set of system resources (memory, devices, files)
- \* Unit of resource ownership

Slide 3

Many applications consist of more than one thread of execution which share resources

⇒ distinction between **thread** and **process**

---

### PROCESSES AND THREADS

**Process:**

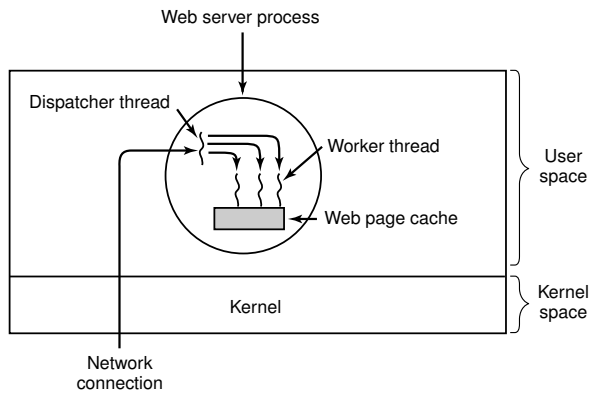
- "Owner" of resources allocated for individual program execution
- Can encompass more than one **thread** of execution
  - Outlook, Evolution: different threads for calendar, mail components etc

Slide 4

**Thread:**

- Unit of execution
  - Belongs to a process
  - Can be traced
    - list the sequence of instructions that execute
-

### EXAMPLE: WEB SERVER



Slide 5

### SINGLE-THREADED WEB SERVER IMPLEMENTATIONS

#### → Sequential processing of requests:

- web server gets request, processes it, accepts next request
- CPU idle while data retrieved from disk
- Poor performance

#### → Finite-State Machine:

- use non-blocking read
- program records state of current request
- gets next event
- on reply (signal) from disk, fetches and processes data
- good performance, complicated to implement and debug

#### → Processes instead of Threads

- Communicate by sharing data, messages


Slide 6

### ADVANTAGES OF THREADS

- ① Program does not stall when one of its operations blocks
  - save contents of a page to disk while downloading other page
- ② Overhead for thread creation and destruction is less than for processes (depending on implementation, can be about a factor of 100 faster)
- ③ Simplification of programming model
- ④ Performance gains on machines with multiple CPU's

Slide 7

### THREADS AND PROCESSES



Slide 8

## THREADS AND PROCESSES

→ Single process, single thread

- MS-DOS, old MacOS

→ Single process, multiple threads

- OS/161 as distributed

Slide 9

→ Multiple processes, single thread

- traditional Unix

→ Multiple processes, multiple threads

- modern Unices (Solaris, Linux), Windows-2000

Note: Literature (incl. textbooks) often do not cleanly distinguish those concepts (for historical reasons)!

Logical traces of threads:

1	5000		18	100
2	5001		19	101
3	5002		20	102
4	5003		21	103
5	5004		22	104
6	5005		23	105
<hr/>				
7	100	Time out	24	12000
8	101		25	12001
9	102		26	12002
10	103		27	12003
11	104		28	12004
12	105		29	12005
<hr/>				
13	8000		30	100
14	8001		31	101
15	8002		32	102
16	8003		33	103
17	8004	I/O request	34	104
			35	105

Logical traces of threads:

5000	8000	12000	
5001	8001	12001	5000: Starting address of code for Thread A
5002	8002	12002	
5003	8003	12003	8000: Starting address of code for Thread B
5004		12004	
5005		12005	12000: Starting address of code for Thread C
5006		12006	
5007		12007	
5008		12008	
5009		12009	
5010		12010	
5011		12011	


Thread A      Thread B      Thread C

## THREAD STATES

Three states (may be more, depending on implementation):

- ① **Running**: currently active, using CPU
- ② **Ready**: runnable, waiting to be scheduled
- ③ **Blocked**: waiting for an event to occur (I/O, alarm)

Slide 12



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

## REASONS FOR LEAVING THE RUNNING STATE

Slide 13

- Thread terminates
  - `exit()` system call (voluntary termination)
  - killed by another thread
  - killed by OS (due to `exception`)
- Thread cannot continue execution
  - blocked waiting for event (I/O)
- OS decides to give someone else a chance
  - requires the OS to be invoked
    - via system call or exception
    - via interrupt
- Thread voluntarily gives another thread a chance
  - `yield()` system call


## NON-RUNNING THREADS

Slide 14

- Many separate reasons for a thread not running
  - another thread is running on the CPU
  - thread is blocked (waiting for an event)
  - thread is in initialisation phase (during creation)
  - thread is being cleaned up (during `exit`, `kill`)
- Dispatching ought to be fast
  - Shouldn't search through all threads to find runnable one
  - Achieved by distinguishing more thread states

## SEPARATE QUEUES

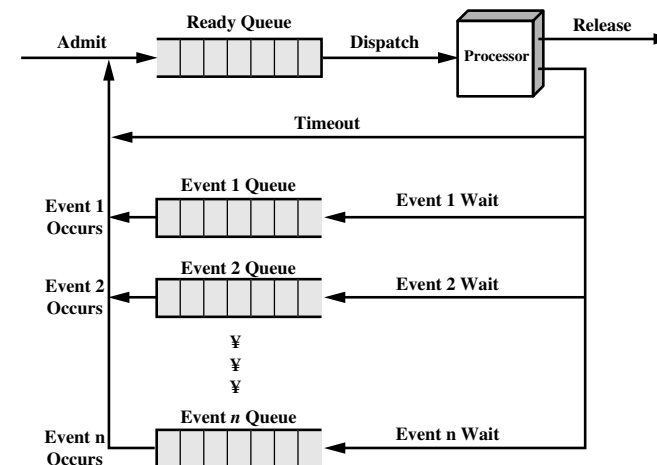
Slide 15



- Simplifies scheduler's job
- How about `wakeup` of blocked thread when event occurs?

## Multiple wait queues:

Slide 16



## COOPERATIVE VS. PREEMPTIVE MULTITHREADING

### Cooperative multithreading:

- Threads determine exact order of execution
- Use `yield()` to switch between threads
- Problems if thread doesn't yield (e.g., buggy)

Slide 17

### Preemptive multitasking:

- OS **preempts** thread's execution after some time
- Only guaranteed to work if H/W provides **timer interrupt**
- Implies **unpredictable execution sequence!**
  - thread switch can happen between **any** two instructions
  - threads may require **concurrency control**

## PROCESSES AND THREADS

The OS stores information about Threads and Processes in Thread Control Block (TCB) and Process Control Block (PCB)

- PCBs stored in process table
- TCBs stored in thread table

Slide 19

	Process	Thread
Address Space	✓	
Registers		✓
Program Counter		✓
Stack		✓
Open Files	✓	
State		✓
Signals and Handlers	✓	
Accounting Info	✓	✓
Global Variables		

## USER-LEVEL OPERATIONS ON THREADS IN OS/161

→ Start a new thread in OS/161

```
thread_fork(const char *   name,
            void          *   data1,
            unsigned long  data2,
            void          (* func)(void *, unsigned long),
            struct thread **ret);
```

Slide 18

→ Terminate thread

- `thread_exit()`

→ Yield CPU

- `thread_yield()`

→ Synchronisation:

- `thread_sleep(const void *addr)`
- `thread_wakeup(const void *addr)`