



# 5. Threads

*Sungyoung Lee*

*College of Engineering  
KyungHee University*



# Contents

- n *Overview*
- n *Multithreading Models*
- n *Threading Issues*
- n *Pthreads*
- n *Solaris 2 Threads*
- n *Windows 2000 Threads*
- n *Linux Threads*
- n *Java Threads*

## n Heavy-weight

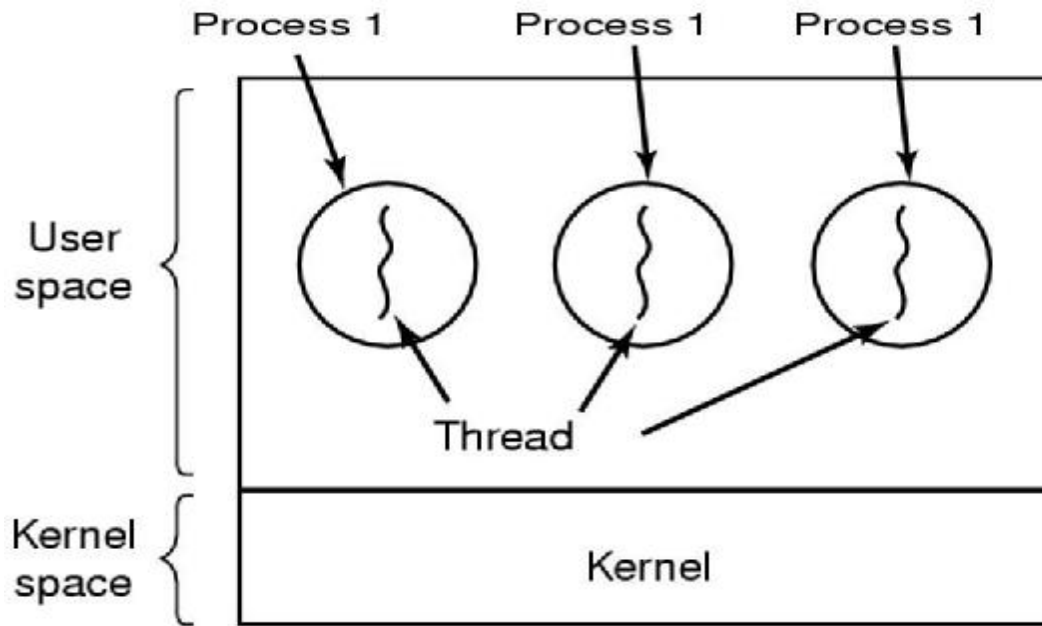
- ü A process includes many things:
  - § An address space (all the code and data pages)
  - § OS resources (e.g., open files) and accounting info.
  - § Hardware execution state (PC, SP, registers, etc.)
- ü Creating a new process is costly because all of the data structures must be allocated and initialized
  - § Linux: over 100 fields in `task_struct`  
(excluding page tables, etc.)
- ü Inter-process communication is costly, since it must usually go through the OS
  - § Overhead of system calls and copying data

# Thread Concept: Key Idea


## n Separate the concept of a process from its execution state

- ü Process: address space, resources, other general process attributes (e.g., privileges)
- ü Execution state: PC, SP, registers, etc.
  
- ü This execution state is usually called
  - § a thread of control,
  - § a thread, or
  - § a lightweight process (LWP)

# Thread Concept: Key Idea (Cont'd)

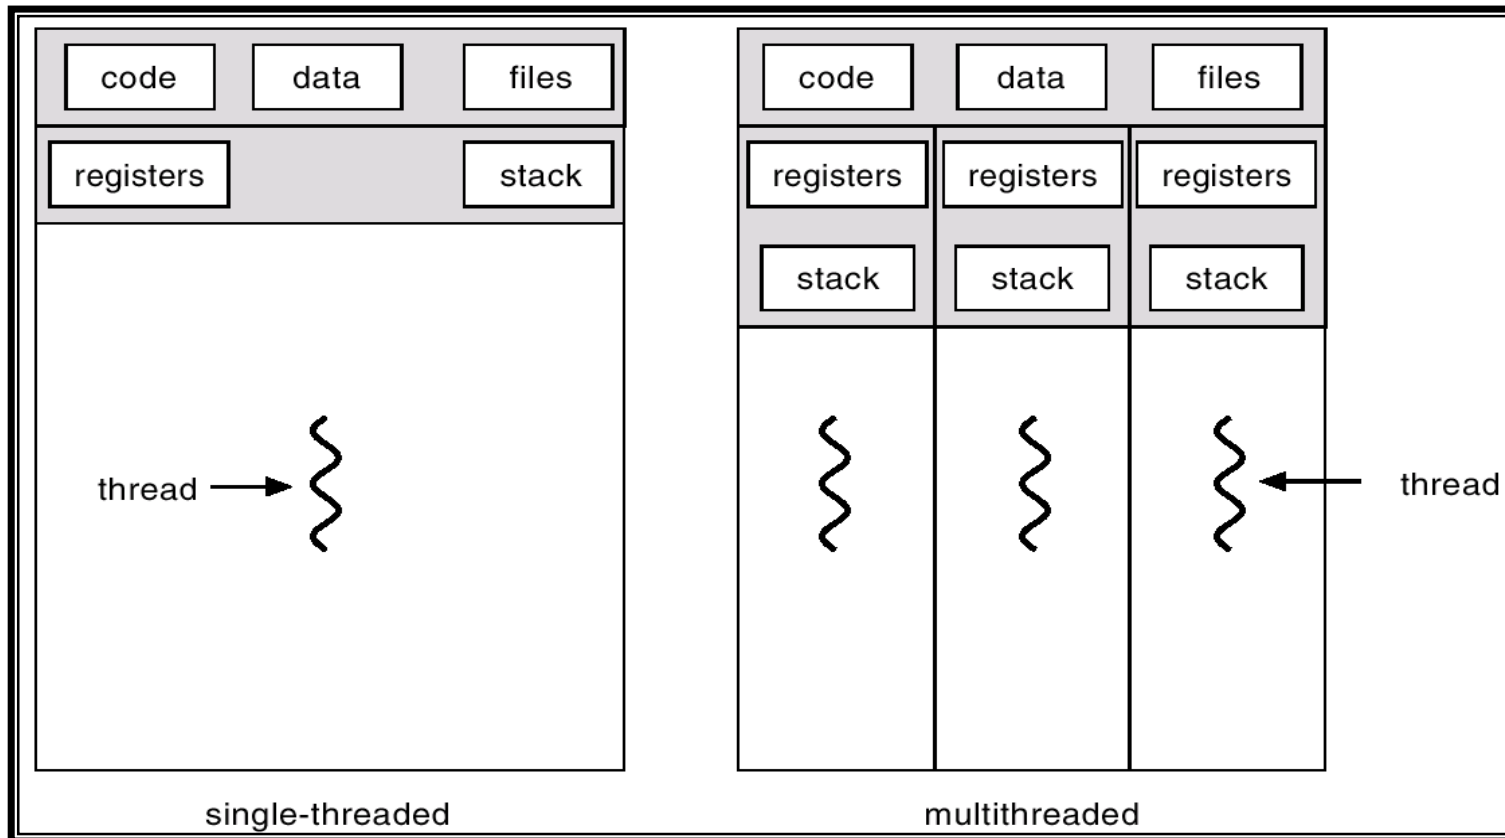


(a)



(b)

# Single and Multithreaded Processes



# What is a Thread?

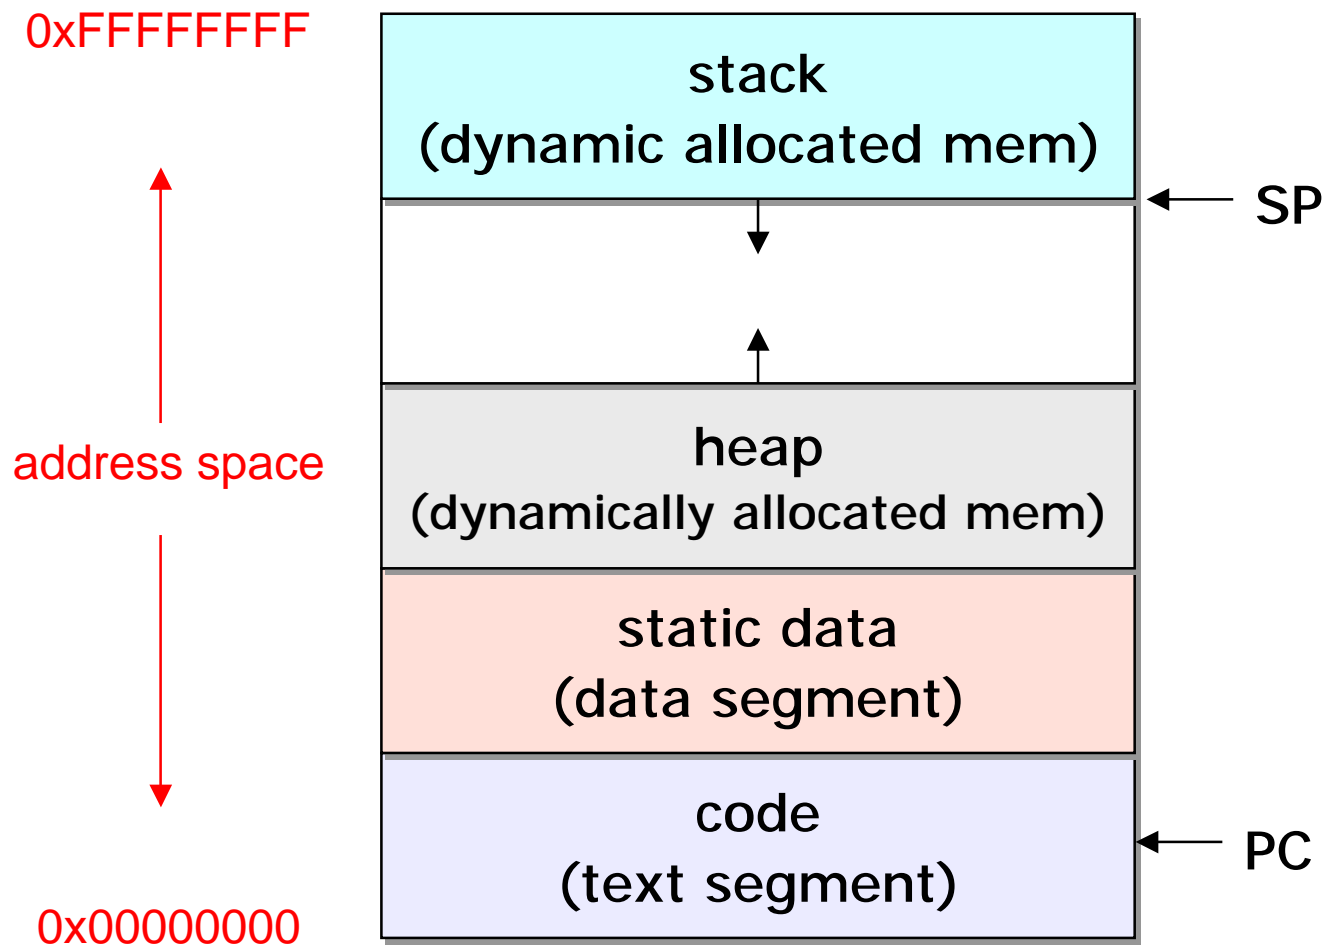
- n A *thread* (or *lightweight process*) is a basic unit of CPU utilization; it consists of:
  - ü program counter
  - ü register set
  - ü stack space
  
- n A thread shares with its peer threads its:
  - ü code section
  - ü data section
  - ü operating-system resources
  - ü collectively known as a *task* or process
  
- n A traditional or *heavyweight* process is equal to a task with one thread

## n Processes vs. Threads

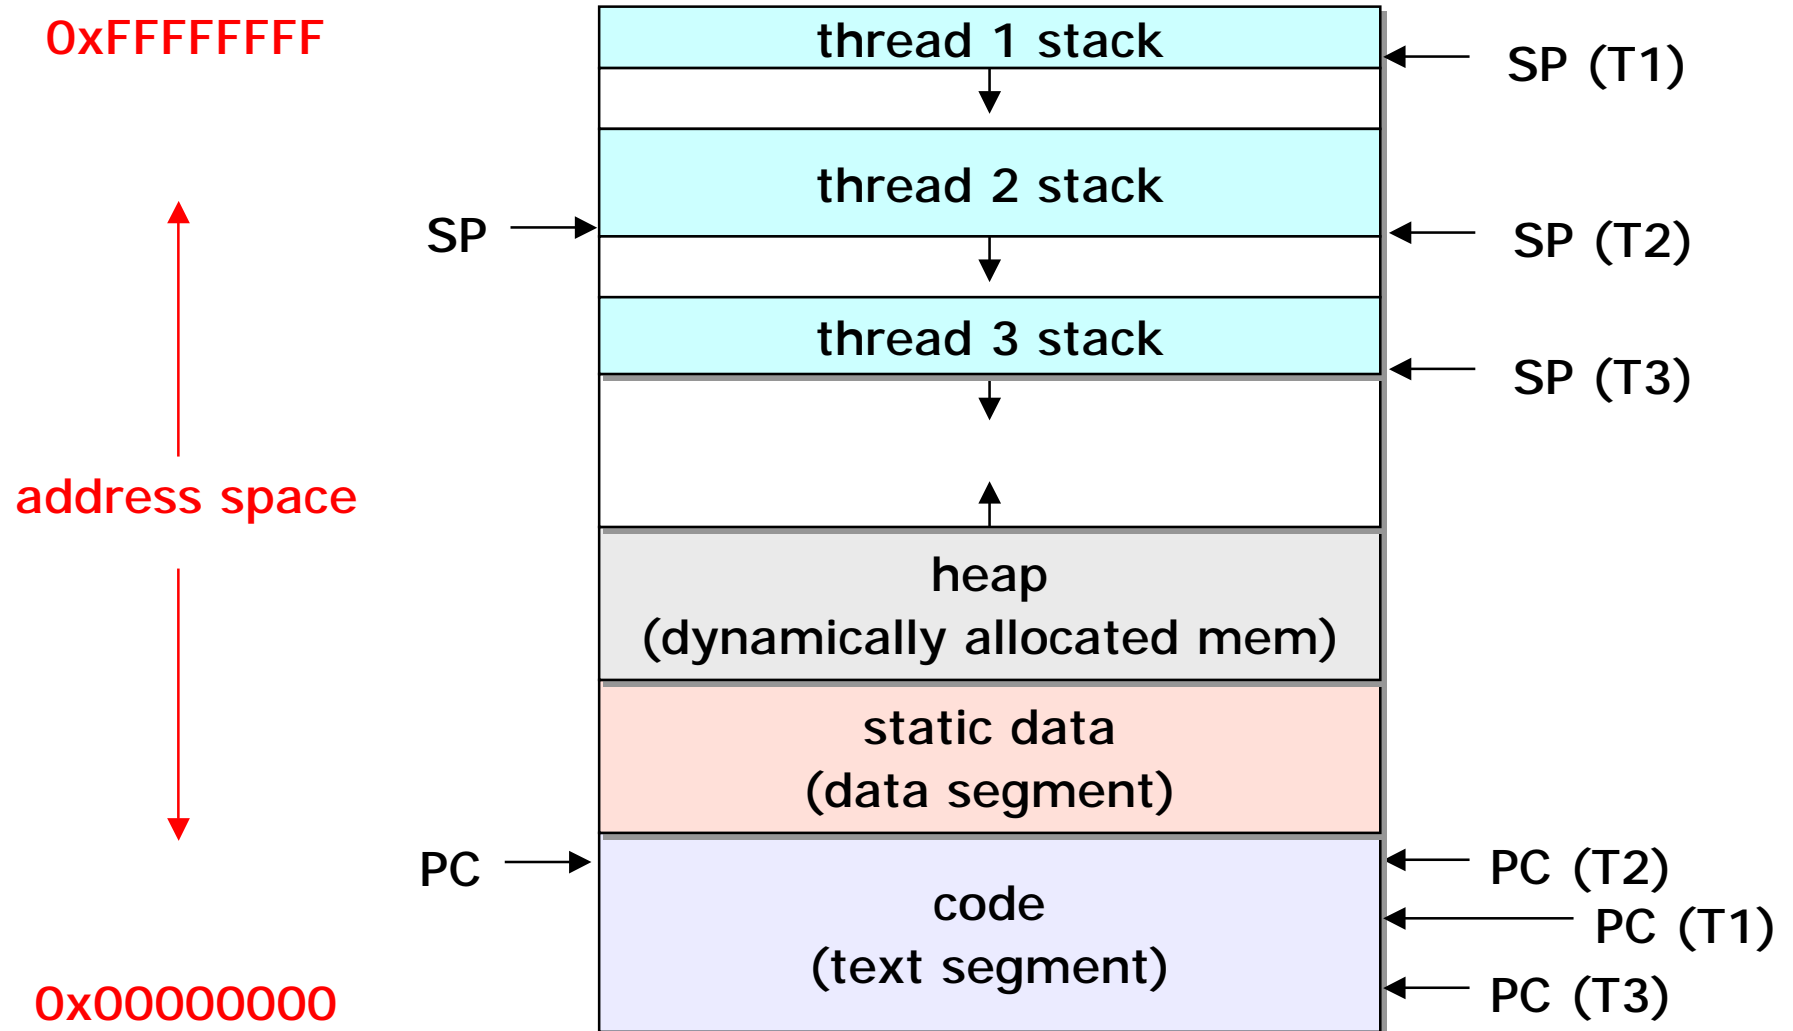
- ü A thread is bound to a single process
- ü A process, however, can have multiple threads
- ü Sharing data between threads is cheap: all see the same address space
- ü Threads become the unit of scheduling
- ü Processes are now containers in which threads execute
- ü Processes become static, threads are the dynamic entities



# Process Address Space



# Address Space with Threads



# Concurrent Servers: Processes

## n Web server example

- ü Using fork() to create new processes to handle requests in parallel is overkill for such a simple task.

```
While (1) {  
    int sock = accept();  
    if ((pid = fork()) == 0) {  
        /* Handle client request */  
    } else {  
        /* Close socket */  
    }  
}
```

# Concurrent Servers: Threads

## n Using threads

- ü We can create a new thread for each request

```
webserver ()
{
    While (1) {
        int sock = accept();
        thread_fork (handle_request, sock);
    }
}
handle_request (int sock)
{
    /* Process request */
    close (sock);
}
```

# Benefits


- n Responsiveness
- n Resource Sharing
- n Economy
- n Utilization of MP Architectures

# User Threads

n Thread management done by user-level threads library

n Examples

- ü POSIX *Pthreads*
- ü Mach *C-threads*
- ü Solaris *threads*




# Kernel Threads

## n Supported by the Kernel

- ü thread creation and management requires system calls

## n Examples

- ü Windows 95/98/NT/2000
- ü Solaris
- ü Tru64 UNIX
- ü BeOS
- ü Linux



# User-level Threads vs. Kernel-level Threads

## n User-level threads

- ü The user-level threads library implements thread operations
- ü They are small and fast
- ü User-level threads are invisible to the OS
- ü OS may make poor decisions
  - § e.g. blocking I/O
- ü Thread scheduling
  - § Non-preemptive scheduling: `yield()`
  - § Preemptive scheduling: timer through signal

## n Kernel-level threads

- ü All thread operations are implemented in the kernel
- ü The OS schedules all of the threads in a system
- ü Kernel threads are cheaper than processes
- ü They can still be too expensive




# Multithreading Models

- n Many-to-One
- n One-to-One
- n Many-to-Many

# Many-to-One

- n Many user-level threads mapped to single kernel thread
- n Used on systems that do not support kernel threads




# One-to-One

n Each user-level thread maps to kernel thread

n Examples


ü Windows 95/98/NT/2000

ü OS/2



# Many-to-Many Model

- n Allows many user level threads to be mapped to many kernel threads
- n Allows the operating system to create a sufficient number of kernel threads
- n Solaris 2
- n Windows NT/2000 with the *ThreadFiber* package



# Threading Issues

- n Semantics of fork() and exec() system calls
  - ü Two versions of fork()
- n Thread cancellation
  - ü Asynchronous cancellation
  - ü Deferred cancellation
- n Signal handling
  - ü To the thread to which the signal applies
  - ü To every thread in the process
  - ü To certain threads in the process
  - ü Assign a specific thread to receive all signals for the process
- n Thread pools
  - ü Create a number of threads at process startup
- n Thread specific data

# Pthreads

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

## n POSIX-style threads

- ü Pthreads
- ü DCE threads (early version of Pthreads)
- ü Unix International (UI) threads (Solaris threads)
  - § Sun Solaris 2, SCO Unixware 2

## n Microsoft-style threads

- ü Win32 threads
  - § Microsoft Windows 98/NT/2000/XP
- ü OS/2 threads
  - § IBM OS/2

## n Thread creation/termination

```
int pthread_create (pthread_t *tid,  
                  pthread_attr_t *attr,  
                  void *(start_routine)(void *),  
                  void *arg);
```

```
void pthread_exit (void *retval);
```

```
int pthread_join (pthread_t tid,  
                 void **thread_return);
```



## n Mutexes

```
int pthread_mutex_init  
    (pthread_mutex_t *mutex,  
     const pthread_mutexattr_t *mattr);
```

```
void pthread_mutex_destroy  
    (pthread_mutex_t *mutex);
```

```
void pthread_mutex_lock  
    (pthread_mutex_t *mutex);
```

```
void pthread_mutex_unlock  
    (pthread_mutex_t *mutex);
```

# Pthreads (Cont'd)

## n Condition variables

```
int pthread_cond_init  
    (pthread_cond_t *cond,  
     const pthread_condattr_t *cattr);
```

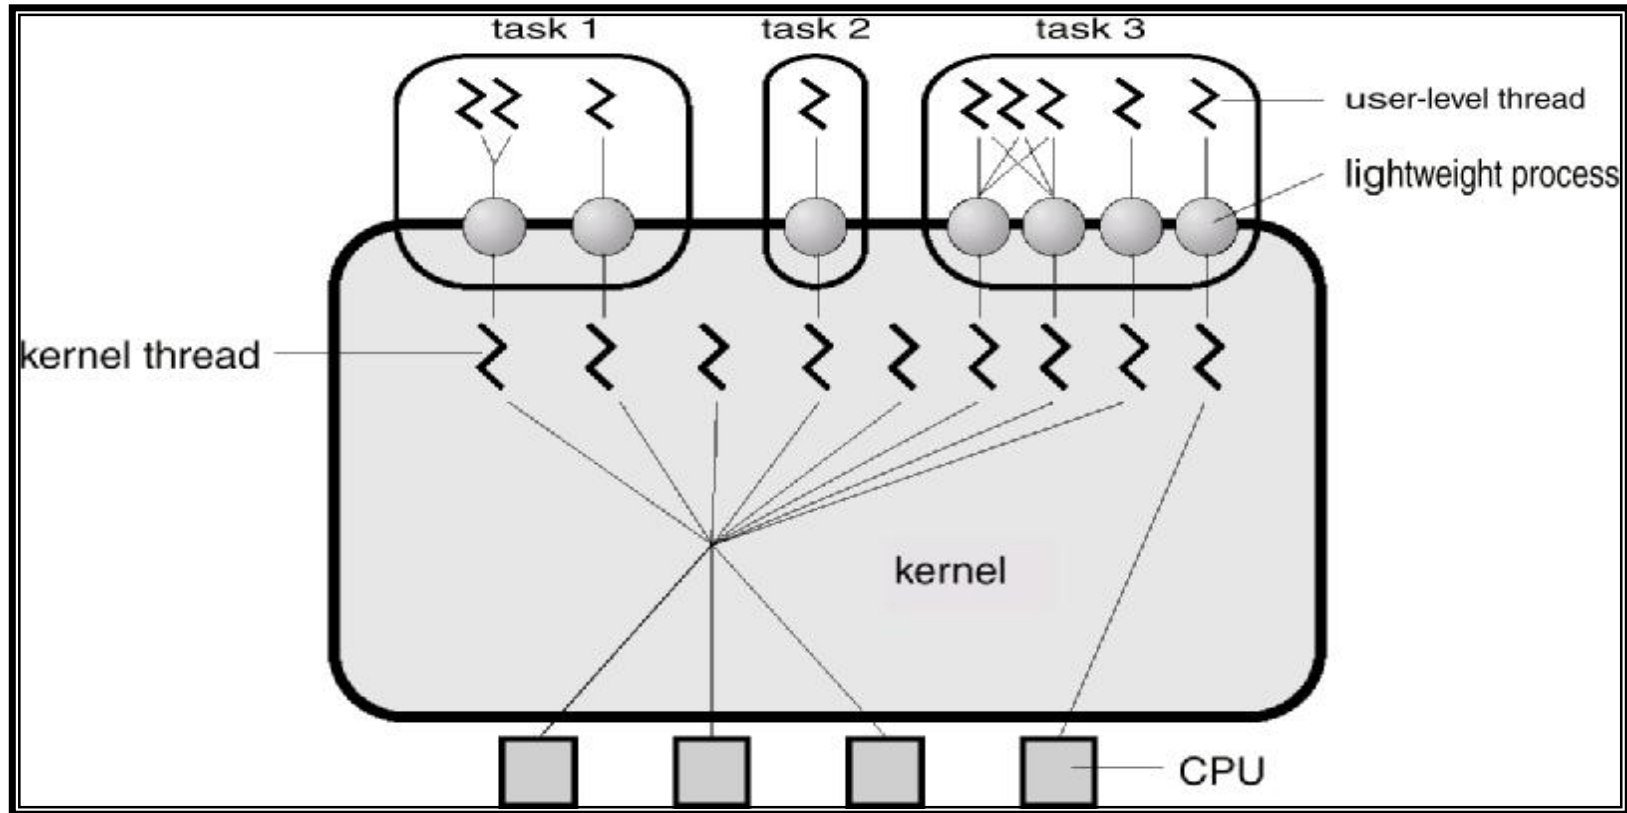
```
void pthread_cond_destroy  
    (pthread_cond_t *cond);
```

```
void pthread_cond_wait  
    (pthread_cond_t *cond,  
     pthread_mutex_t *mutex);
```

```
void pthread_cond_signal  
    (pthread_cond_t *cond);
```

```
void pthread_cond_broadcast  
    (pthread_cond_t *cond);
```

# Solaris 2 Threads




## n LWP (Lightweight Process)

- ü A virtual CPU for executing code or system calls
- ü Each process contains at least one LWP
- ü Each LWP is connected to exactly one kernel-level thread
- ü Each LWP is separately dispatched by the kernel, may
  - § perform independent system calls
  - § incur independent page faults
  - § run in parallel on a multiprocessor, etc.
- ü The thread library dynamically adjusts the number of LWPs in the pool to ensure the best performance for the application
- ü It also “ages” LWPs and deletes them when they are unused for a long time.
- ü An LWP is a kernel data structure

## n *For implementing many-to-many model*

# Solaris Process



# Windows 2000 Threads

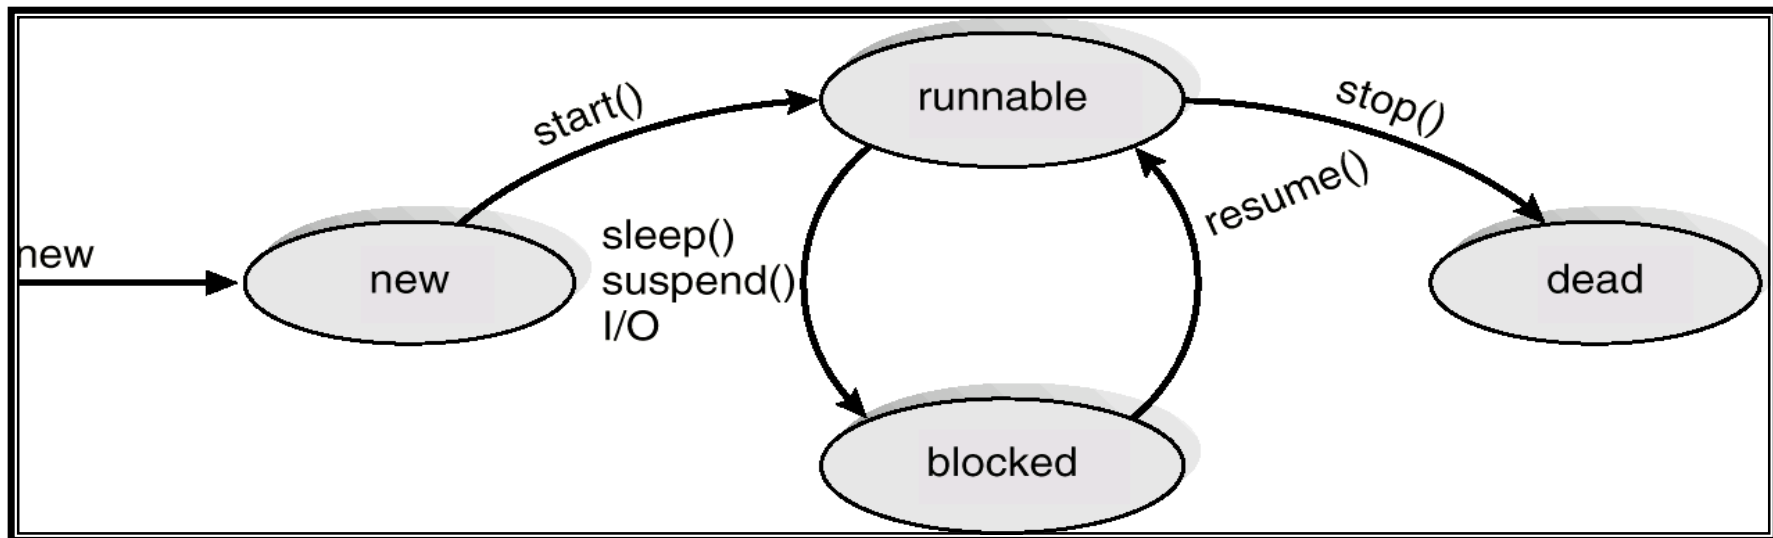
- n Implements the one-to-one mapping
  
- n Each thread contains
  - ü a thread id
  - ü register set
  - ü separate user and kernel stacks
  - ü private data storage area
  
- n Cf) Fibers
  - ü Fibers are often called “lightweight” threads
  - ü Fibers are invisible to the kernel
  - ü Fibers provide a functionality of the many-to-many model

# Linux Threads

- n Linux refers to them as *tasks* rather than *threads*
- n Thread creation is done through `clone()` system call
- n `Clone()` allows a child task to share the address space of the parent task (process)
- n *So, there exist POSIX compatibility problems*
- n *Approaches for POSIX compliance*
  - ü Linux 2.4 introduces a concept of “thread groups”
  - ü NPTL (Native POSIX Threading Library) – by RedHat
    - § 1:1 model
  - ü NGPT (Next Generation POSIX Threading) – by IBM
    - § M:N model

# Java Threads

- n Java threads may be created by:
  - ü Extending Thread class
  - ü Implementing the Runnable interface
- n Java threads are managed by the JVM
- n Java thread states





# Threads Design Space

