

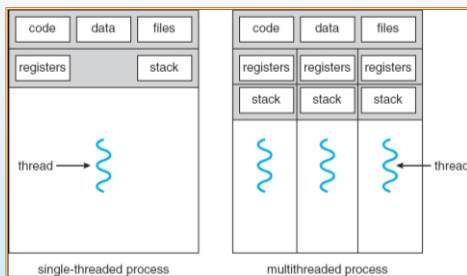
## Threads

### Concepts and Implementation

## Threads

- Multithreading Models
- Threading Issues
- Why Threads
- Thread libraries:
  - POSIX Pthreads
  - Win32 threads
  - Java threads

## Single and Multithreaded Processes



## Benefits

- Responsiveness
- Resource (memory, file, etc.) Sharing
- Economy (less synchronization overhead)
- Utilization of MP Architectures

## Benefits-1

- A traditional OS process is one thread of execution only.
- Multiple threads allow multiple thread of computation in the same address space.
- Decomposing applications into concurrent threads, may reduce process or application blocking
- It may even simplify the programming
- Easier to create and destroy threads
- They may offer performance gain
- It is possible to utilize multiple CPUs

## Benefits-2

- Performance gains from multiprocessing hardware (parallelism)
  - Different threads can run on different processors simultaneously with no special input from the user and no effort on the part of the programmer

### Benefits -3

- An application that uses multiple processes can be replaced by an application that uses multiple threads to accomplish the same tasks, if possible.


### User/Kernel Threads

- User level thread management done using user-level threads library
  - Primary user level thread libraries:
    - ▶ UNIX and Linux: POSIX Pthreads
    - ▶ Windows: Win32 threads
    - ▶ JVM: Java threads
- Kernel level threads are managed by the operating system itself


### User and Kernel level Multithreading Models: Done by the kernel

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

### Many-to-One Model




### One-to-one Model




### Many-to-Many Model

- Allows many user level threads to be mapped to many kernel threads
- Allows the operating system to create a sufficient number of kernel threads
- Solaris prior to version 9
- Windows NT/2000 with the *ThreadFiber* package

## Many-to-Many Model




## Implementing Threads in User Space



## Implementing Threads in User Space

- A user-level threads package:
- Kernel knows nothing about threads.
- The user level library can manage a thread table to allow orderly execution. If a thread blocks its state is saved in the table and a ready thread is allowed to execute...
- The application has its own thread scheduler...
- **Result: Efficient implementation is very hard to achieve, virtually impossible...**


## Kernel level threads



## Kernel level threads

- A threads package managed by the kernel:
- Kernel has a thread table to keep record of all the threads in the system
- A thread library is part of the system call library.
- Kernel level thread management is much more expensive, as blocking is handled similar to multiple process case concept...

## Hybrid Implementations



Multiplexing user-level threads onto kernel-level threads: using advantages of user/kernel level threads

## Thread Cancellation

- Terminating a thread before it has finished
- Two general approaches:
  - **Asynchronous cancellation** terminates the target thread immediately
  - **Deferred cancellation** allows the target thread to periodically check if it should be cancelled

Operating System Concepts

## Thread Pools

- Create a number of threads in a pool where they await to be activated.
- Advantages:
  - Usually slightly faster to service a request with an existing thread than create a new thread
  - Allows the number of threads in the application(s) to be bound to the size of the pool

Operating System Concepts

## Thread Specific Data

- Each thread can have its own copy of data
- This is useful when there is no control over the thread creation process (i.e., when using a thread pool)

Operating System Concepts

## Scheduler Activations

- Scheduler maintains an appropriate number of kernel threads allocated to an application
- Scheduler activations provide **upcalls**: a communication mechanism from the kernel to the thread library
  - This communication allows an application to maintain the correct number kernel threads

Operating System Concepts

## What is Pthreads

- In order to take full advantage of the capabilities provided by threads, a standardized programming interface was required.
  - For UNIX systems, this interface has been specified by the IEEE POSIX 1003.1c standard (1995).
  - Implementations adhering to this standard are referred to as POSIX threads, or Pthreads.
  - Most hardware vendors now offer Pthreads in addition to their proprietary API's.
- The POSIX standard has continued to evolve and undergo revisions, including the Pthreads specification.

Operating System Concepts

## Pthreads

- Some useful links:
  - [standards.ieee.org/findstds/standard/1003.1-2008.html](http://standards.ieee.org/findstds/standard/1003.1-2008.html)
  - [www.opengroup.org/austin/papers/posix\\_faq.html](http://www.opengroup.org/austin/papers/posix_faq.html)
  - [www.unix.org/version3/ieee\\_std.html](http://www.unix.org/version3/ieee_std.html)
- 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)
- Pthreads are defined as a set of C language programming types and procedure calls, implemented with a pthread.h header/include file and a thread library

Operating System Concepts

## The Thread Model (2)

Per process items	Per thread items
Address space	Program counter
Global variables	Registers
Open files	Stack
Child processes	State
Pending alarms	
Signals and signal handlers	
Accounting information	

- Per process: Items shared by all threads in a process
- Per thread: Items private to each thread


## Common Reason for Threads (1)

- When compared to the cost of creating and managing a process, a thread can be created with much less operating system overhead. The difference could be up to 20:1
- Threaded applications offer potential performance gains and practical advantages over non-threaded applications :
  - Overlapping CPU work with I/O:
    - ▶ A program may have sections where it is performing a long I/O operation. While one thread is waiting for an I/O system call to complete, CPU intensive work can be performed by other threads.

## Common Reason for Threads (2)

- Priority/real-time scheduling:
  - tasks which are more important can be scheduled to supersede or interrupt lower priority tasks.
- Asynchronous event handling:
  - interleaving. For example, a web server can both transfer data from previous requests and manage the arrival of new requests.
  - A perfect example is the typical web browser, where many interleaved tasks can be happening at the same time, and where tasks can vary in priority.


## Example: WINDOWS: Start button/ All apps/Windows Administrative Tools/Resource Monitor.




## Suitability for thread implementation

- Work that can be executed, or data that can be operated on, by multiple tasks simultaneously:
- Block for potentially long I/O waits
- Use many CPU cycles in some places but not others
- Must respond to asynchronous events
- Some work is more important than other work (priority interrupts)

## Thread Usage (1)



## Thread Usage: web server



- A multithreaded Web server: maintains a cache to hold the frequently requested pages

## Thread Usage (3)

### Web Dispatcher Thread

```
while (TRUE) {
    get_next_request(&buf);
    handoff_work(&buf);
}
```

(a)

### Worker thread: client

```
while (TRUE) {
    wait_for_work(&buf)
    look_for_page_in_cache(&buf, &page);
    if (page_not_in_cache(&page))
        read_page_from_disk(&buf, &page);
    return_page(&page);
}
```

(b)

- Rough outline of code for the web server
  - (a) Dispatcher thread: handles the incoming web page requests
  - (b) Worker thread, when woken up checks to see if the requested page in the cache, if not it reads from the disk before delivering

One worker would have been inefficient as it has to block until IO is complete. Thus multiple worker implementation should be preferred.


## common models for threading

- **Manager/worker:** a single thread, the *manager* assigns work to other threads, the *workers*. Typically, the manager handles all input and parcels out work to the other tasks.
- **Pipeline:** a task is broken into a series of suboperations, each of which is handled in series, but concurrently, by a different thread. An automobile assembly line best describes this model.
- **Peer:** similar to the manager/worker model, but after the main thread creates other threads, it participates in the work just like others.


## Shared Memory Model

- All threads have access to the same global, shared memory
- Threads also have their own private data
- Programmers are responsible for synchronizing access (protecting) globally shared data.

## Multi-thread appearance



## User Level Thread Scheduling Example




Possible: A1, A2, A3, A1, A2, A3  
Not possible: A1, B1, A2, B2, A3, B3

Possible scheduling of user-level threads

- 50-msec process quantum
- threads run 5 msec/CPU burst

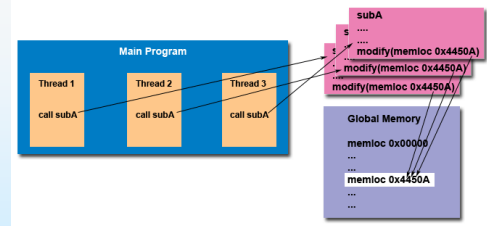
## Kernel Level Thread Scheduling Example



Possible scheduling of kernel-level threads

- 50-msec process quantum
- threads run 5 msec per CPU burst

## Thread-Safeness



- Suppose that your application creates several threads, each of which makes a call to the same library routine
- There is problem: This not thread safe

## Thread Limits

- Although the Pthreads API is an ANSI/IEEE standard, implementations can, and usually do, vary in ways not specified by the standard.
- Because of this, a program that runs fine on one platform, may fail or produce wrong results on another platform.
- For example, the maximum number of threads permitted, and the default thread stack size are two important limits to consider when designing your program.

## pthread Example

```

.....
* FILE: pthread.c
* DESCRIPTION:
*   A "hello world" Pthreads program. Demonstrates thread creation and
*   termination.
* AUTHOR: Brian Borner
* LAST REVISION: 08/04/11
.....
#include <pthread.h>
#include <stdio.h>
#include <unistd.h>
#define NUM_THREADS 5

void *PrintHello(void *threadid)
{
    unsigned long tid;
    tid = (unsigned long)threadid;
    printf("Hello World %d ms. Thread %lu!\n", tid,
        pthread_self());
}

int main(int argc, char *argv[])
{
    pthread_t threads[NUM_THREADS];
    int rc;
    long t;
    for(t=0;t<NUM_THREADS;t++)
        printf("main: creating thread %d\n", t);
        if (pthread_create(&threads[t], NULL, PrintHello, (void *)t)
            != 0)
            printf("ERROR: return code from pthread_create() is %d\n", rc);
        else
            t++;
}

/* Last thing that main() should do */
pthread_exit(NULL);
}
    
```

## Passing Arguments to Threads

- The pthread\_create() routine permits the programmer to pass one argument to the thread start routine.
- For cases where multiple arguments must be passed, this limitation is easily overcome by creating a structure which contains all of the arguments, and then passing a pointer to that structure in the pthread\_create() routine.
- All arguments must be passed by reference and cast to (void \*).

## Thread Multiple Argument Passing Example

```

.....
* FILE: pthread_arg.c
* DESCRIPTION:
*   A "hello world" Pthreads program. Demonstrates passing multiple
*   arguments to a thread start routine.
* AUTHOR: Brian Borner
* LAST REVISION: 08/04/11
.....
#include <pthread.h>
#include <stdio.h>
#include <unistd.h>
#define NUM_THREADS 5

struct thread_arg {
    int thread_id;
    int arg1;
    int arg2;
};

void *PrintHello(void *threadid)
{
    struct thread_arg *arg;
    arg = (struct thread_arg *)threadid;
    printf("Hello World %d ms. Thread %d, arg1=%d, arg2=%d\n",
        pthread_self(), arg->thread_id, arg->arg1, arg->arg2);
}

int main(int argc, char *argv[])
{
    pthread_t threads[NUM_THREADS];
    int rc;
    long t;
    struct thread_arg arg;
    for(t=0;t<NUM_THREADS;t++)
        printf("main: creating thread %d\n", t);
        if (pthread_create(&threads[t], NULL, PrintHello, (void *)
            &arg) != 0)
            printf("ERROR: return code from pthread_create() is %d\n", rc);
        else
            t++;
}

/* Last thing that main() should do */
pthread_exit(NULL);
}
    
```