


CS 261 Fall 2017

Mike Lam, Professor



Threads

Parallel computing

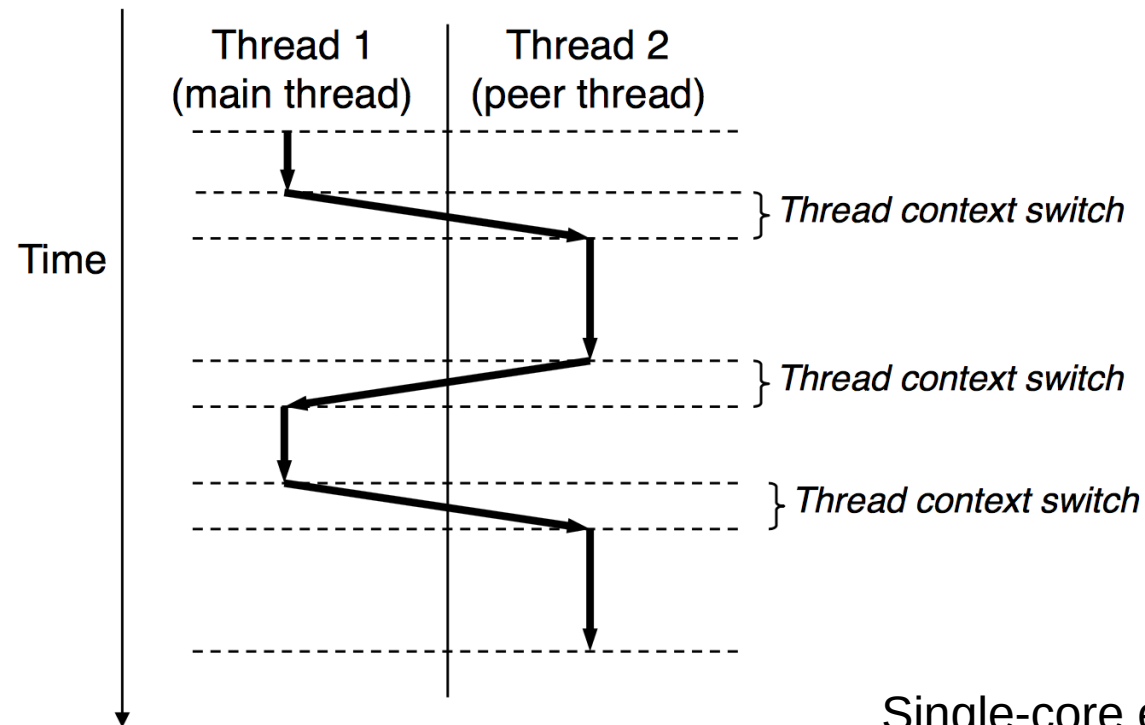
- Goal: **concurrent** or **parallel** computing
 - Take advantage of multiple hardware units to solve multiple problems simultaneously
- Motivations:
 - Maintain high utilization during slow I/O downtime
 - Maintain UI responsiveness during computation
 - Respond simultaneously to multiple realtime events
 - Split up a large problem and solve sub-pieces concurrently to achieve faster time-to-solution (**strong scaling**)
 - Solve larger problems by adding more hardware (**weak scaling**)

Parallel computing

- **Process**: currently-executing program
 - Code and state (PC, stack, data, heap)
 - Private address space
- **Thread**: unit of execution or logical flow
 - Exists within the context of a single process
 - Shares code/data/heap/files w/ other threads
 - Keeps private PC, stack, and registers
 - Stacks are technically shared, but harder to access

Threads

- One **main** thread for each process
 - Can create multiple **peer** threads



Single-core example

POSIX threads

- **Pthreads** – POSIX standard interface for threads in C
 - Not part of the standard library
 - Requires “-lpthread” flag during linking
 - `pthread_create`: spawn a new child thread
 - `pthread_t` struct for storing thread info
 - attributes (or NULL)
 - thread work routine (function pointer)
 - thread routine parameter (void*, can be NULL)
 - `pthread_self`: get current thread ID
 - `pthread_exit`: terminate current thread
 - can also terminate implicitly by returning from the thread routine
 - `pthread_join`: wait for another thread to terminate

Thread creation example

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


void* work (void* arg)
{
    printf("Hello from work routine!\n");
    return NULL;
}

int main ()
{
    printf("Spawning single child ...\n");

    pthread_t child;
    pthread_create(&child, NULL, work, NULL);
    pthread_join(child, NULL);

    printf("Done!\n");

    return 0;
}
```



Shared memory


- Variables in threaded programs
 - Global variables (shared, single static copy)
 - Local variables (multiple copies, one on each stack)
 - Technically still shared if in memory, but harder to access
 - Not shared if cached in register
 - Safer to assume they're private; this is conventional
 - Local static variables (shared, single static copy)
 - Heap-allocated variables (shared, dynamic)

Processes vs. threads

- **Process**: currently-executing program
 - Code and state (PC, stack, data, heap)
 - Created via system call (fork); parent and child continue equally
 - **Private address space** not shared w/ other processes
 - **Advantages: isolation, safety, and mutual exclusion**
- **Thread**: unit of execution or logical flow
 - Exists within a process context
 - Created via library call (pthread_create); child runs separate routine
 - **Shared address space** w/ other threads
 - Private PC, registers, condition codes, and stack
 - **Advantages: faster context switching, more shared resources**

Parallel patterns

- Common pattern: **master/worker** threads
 - One original (**main thread**) creates multiple “child” threads
 - Each worker thread does a chunk of the work
 - Coordinate via shared global data structures
 - Keep as much data as possible local to the thread
 - Main thread waits for workers, then aggregates results



Issues with shared memory

- Nondeterminism
- Data races and deadlock

foo:

```
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

x:

```
    .quad 0
```

thread1



thread2



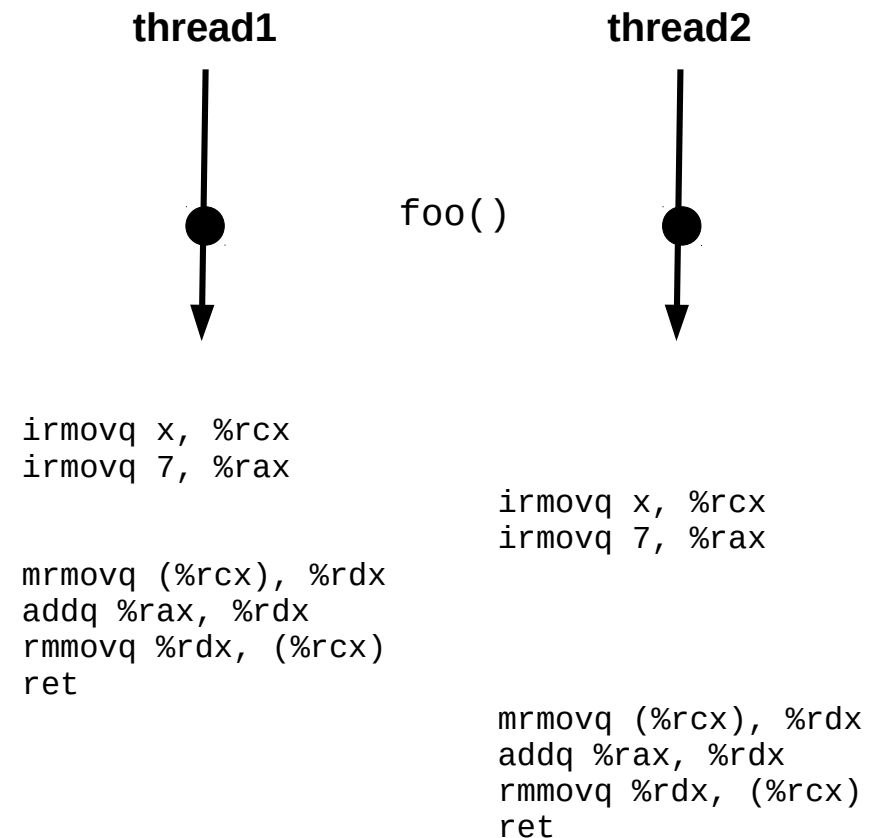
foo()

Issues with shared memory

- Nondeterminism
- Data races and deadlock

```
foo:
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

```
x:
    .quad 0
```




This interleaving is ok.

Issues with shared memory

- Nondeterminism
- Data races and deadlock

```
foo:
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

```
x:
    .quad 0
```




Issues with shared memory

- Nondeterminism
- Data races and deadlock

```
foo:
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

```
x:
    .quad 0
```



PROBLEM!

Issues with shared memory

- Nondeterminism
- Data races and deadlock

foo:

```
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

x:

```
    .quad 0
```

thread1



thread2



foo()

```
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
```

```
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

```
    irmovq x, %rcx
    irmovq 7, %rax
    mrmovq (%rcx), %rdx
```

```
    addq %rax, %rdx
    rmmovq %rdx, (%rcx)
    ret
```

This will be a major topic in CS 361.

Mutual exclusion

- Fixing a data race requires some form of **mutual exclusion**
 - Only one thread at a time should update a particular memory region
 - In Pthreads, this can be accomplished using either a **mutex** or a **semaphore** (more details in CS 361)
 - However, these mechanism introduce overhead!
 - Threads must perform additional checks before updating memory
 - Some threads may have to pause and wait before they may continue
 - If not implemented carefully, the additional overhead may defeat the purpose of using multiple threads
 - Efficient **parallel and distributed computing** can be hard!

Automatic parallelism

- Wouldn't it be great if the compiler could automatically parallelize our programs?
 - This is a HARD problem
 - In some cases, it is (kind of) possible
 - Approach #1: code annotations in existing language
 - Example: OpenMP (CS 450, **CS 470**)
 - Approach #2: new language designed for parallelism
 - Example: HPF and Chapel (CS 430, **CS 470**)

```
int a[100];  
#pragma omp parallel for  
int i;  
for (i=0; i < 100; i++)  
    a[i] = i*i;
```

OpenMP example

```
var a: [100] int;  
forall i in 0..100 do  
    a[i] = i*i;
```

Chapel example

Parallel systems

- **Uniprogramming** / batch (1950s) - **CS 261**
 - One process at a time w/ complete control of CPU
 - Minimal OS (mostly for launching programs)
- **Multiprogramming** / multitasking / time sharing (1960s) - CS 261, **CS 450**
 - Multiple processes taking turns on a single CPU
 - Increased utilization, lower response time
 - OS handles scheduling and context switching
- (Symmetric) **multiprocessing** (1970s) - **CS 361**, CS 450, CS 470
 - Multiple processes share multiple CPUs or cores
 - Increased throughput, increased parallelism
 - OS handles scheduling, context switching, and communication
- **Distributed** processing (1980s and onward) - CS 361, **CS 470**
 - Multiple processes share multiple computers
 - Massive scaling; OS no longer sufficient (other middleware required)