

## CSE 451 11au Module 2 - Multithreaded Programming

CSE 451 11au 10/1/11

II-1

Derived from materials copyright © 2010 Thomas W. Døeppner.

## Part 1

### Brief introduction to threads

CSE 451 11au 10/1/11

II-2

Derived from materials copyright © 2010 Thomas W. Døeppner.

## What Is a Thread?


- A thread is an execution context
  - All the state required to execute according to the rules you're accustomed to
- "Process" vs. "Thread"
  - Saying "process" indicates you want to highlight protection boundaries
    - E.g., a virtual address space
  - Saying "thread" indicates you want to highlight concurrency / cooperation
    - The threads of a computation share an address space

CSE 451 11au 10/1/11

II-3

Derived from materials copyright © 2010 Thomas W. Døeppner.

## The Virtual Address Space



CSE 451 11au 10/1/11

II-4

Derived from materials copyright © 2010 Thomas W. Døeppner.

## Execution Context


- A program counter
  - Where is the thread currently executing?
- A stack
  - Procedure invocation history and
  - Local variables
- Additionally, some metadata is specific to each thread
  - e.g., the tid (thread id)

CSE 451 11au 10/1/11

II-5

Derived from materials copyright © 2010 Thomas W. Døeppner.

## The Virtual Address Space



CSE 451 11au 10/1/11

II-6

Derived from materials copyright © 2010 Thomas W. Døeppner.

## Implications

- A process can use no more cores at a time than it has threads
- (Procedure) local variables are thread local
  - As always, local variables created by entry into a procedure are accessible only to that execution of that procedure

## Uses of Threads

- Parallelism
  - Goal: Higher performance via execution on multiple cores
- Concurrency
  - Goal: Make it easier to write correct code
    - Threads offer a kind of modularity, much as procedures and packages do
    - The modularity in this case is "execution state"

## Rest of this module


- We'll be most interested in concurrency
  - It would be impossible to write a functional OS without it
  - It would be much harder to write many user-level applications as well
- We'll be looking at a particular thread package
  - `pthread` (POSIX threads)

## Part 2

### Concurrency


### Basic pthreads

## Why Threads?



Many things are easier to do with threads  
Many things run faster with threads

## A Simple Example: rlogind



## Life Without Threads

```
void rlogin(int r_in, int r_out, int l_in, int l_out) {
    eof = 0;
    int want_read = 0, want_write = 0;
    int want_read = 1, want_write = 1;
    int eof = 0, size, size_wrt;
    char buf[BSIZE], buf2[BSIZE];

    read(r_in, buf, BSIZE);
    read(l_in, buf2, BSIZE);
    while (!eof) {
        FD_ZERO(&fd_set);
        FD_SET(r_in, &fd_set);
        FD_SET(l_in, &fd_set);
        if (want_read)
            FD_SET(r_in, &fd_set);
        if (want_write)
            FD_SET(l_out, &fd_set);
        select(MAXFD, &fd_set, &fd_set, &fd_set, &fd_set);
        if (FD_ISSET(r_in, &fd_set)) {
            if (size = read(r_in, buf, BSIZE)) > 0 {
                want_read = 1;
                want_write = 1;
            } else
                eof = 1;
        }
        if (FD_ISSET(l_in, &fd_set)) {
            if (size = read(l_in, buf2, BSIZE)) > 0 {
                want_read = 1;
                want_write = 1;
            } else if (eof)
                eof = 1;
        }
        if (FD_ISSET(l_out, &fd_set)) {
            if (write(l_out, buf, size) > 0)
                want_read = 1;
            else if (eof)
                eof = 1;
        }
    }
}
```

CSE 451 11au 10/1/11

II-13

Copyright © 2010 Thomas W. Doepfner. All rights reserved.

## Life With Threads

```
incoming(int r_in, int l_out) {
    int eof = 0;
    char buf[BSIZE];
    int size;

    while (!eof) {
        size = read(r_in, buf, BSIZE);
        if (size <= 0)
            eof = 1;
        if (write(l_out, buf, size) <= 0)
            eof = 1;
    }
}

outgoing(int l_in, int r_out) {
    int eof = 0;
    char buf[BSIZE];
    int size;


    while (!eof) {
        size = read(l_in, buf, BSIZE);
        if (size <= 0)
            eof = 1;
        if (write(r_out, buf, size) <= 0)
            eof = 1;
    }
}
```

CSE 451 11au 10/1/11

II-14

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Single-Threaded Database Server




CSE 451 11au 10/1/11

II-15

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Multithreaded Database Server



CSE 451 11au 10/1/11

II-16

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Creating a Thread

```
start_servers() {
    pthread_t thread;
    int i;
    for (i=0; i<nr_of_server_threads; i++)
        pthread_create(&thread, // thread ID
                      0, // attributes (0 means defaults)
                      server, // start routine
                      argument); // arg to start routine
}

void *server(void *arg) {
    while(1) {
        /* get and handle request */
    }
}
```

CSE 451 11au 10/1/11

II-17

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Complications

```
rlogin(int r_in, int r_out, int l_in, int l_out) {
    pthread_t in_thread, out_thread;
    pthread_create(&in_thread,
                  0,
                  incoming,
                  r_in, l_out); // Can't do this ...
    pthread_create(&out_thread,
                  0,
                  outgoing,
                  l_in, r_out); // Can't do this ...

    /* How do we wait till they're done? */
}
```

CSE 451 11au 10/1/11

II-18

Derived from materials copyright © 2010 Thomas W. Doepfner.

## #1: Multiple Arguments

```
typedef struct {
    int first, second;
} two_ints_t;

rlogind(int r_in, int r_out, int l_in, int l_out) {
    pthread_t in_thread, out_thread;
    two_ints_t in={r_in, l_in}, out={l_in, r_out};
    pthread_create(&in_thread,
        0,
        incoming,
        &in);
    ...
}
```

CSE 451 11au 10/1/11

II-19

Derived from materials copyright © 2010 Thomas W. Doepfner.

## #2: When Are Threads Done?

```
rlogind(int r_in, int r_out, int l_in, int l_out) {
    pthread_t in_thread, out_thread;
    two_ints_t in={r_in, l_in}, out={l_in, r_out};

    pthread_create(&in_thread, 0, incoming, &in);
    pthread_create(&out_thread, 0, outgoing, &out);

    pthread_join(in_thread, NULL);
    pthread_join(out_thread, NULL);
}
```

CSE 451 11au 10/1/11

II-20

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Thread Termination

- Threads terminates by:
  - `pthread_exit((void *) value);`  
or
  - `return((void *) value);`  
from its entry point procedure
- Creating thread gets return value:
  - `pthread_join(thread, (void **) &value);`

CSE 451 11au 10/1/11

II-21

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Sometimes You Don't Want To Wait: Detached Threads

```
start_servers() {
    pthread_t thread;
    int i;
    for (i=0; i<nr_of_server_threads; i++) {
        pthread_create(&thread, 0, server, 0);
        pthread_detach(thread);
    }
    ...
}

server() {
    ...
}
```

CSE 451 11au 10/1/11

II-22

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Thread Attributes

```
pthread_t thread;
pthread_attr_t thr_attr;

pthread_attr_init(&thr_attr);

...
/* establish some attributes */
...

pthread_create(&thread, &thr_attr, startroutine, arg);

...
```

CSE 451 11au 10/1/11

II-23

Derived from materials copyright © 2010 Thomas W. Doepfner.

## Stack Size

```
pthread_t thread;
pthread_attr_t thr_attr;

pthread_attr_init(&thr_attr);
pthread_attr_setstacksize(&thr_attr, 20*1024*1024);

...

pthread_create(&thread, &thr_attr, startroutine,
arg);

...
```

CSE 451 11au 10/1/11

II-24

Derived from materials copyright © 2010 Thomas W. Doepfner.

## pthread Program: C = AB

```
#include <stdio.h>          main() {
#include <pthread.h>        int i;
#include <string.h>         pthread_t thr[M];
                             int error;
#define M      3
#define N      4           /* initialize the matrices ... */
#define P      5
                             ...
int A[M][N];
int B[N][P];
int C[M][P];

void *matmult(void *);
```

CSE 451 11au 10/1/11

II-25

Derived from materials copyright © 2010 Thomas W. Droeppner.

## matrix multiply (cont.)

```
// create the worker threads - one per row
for (i=0; i<M; i++) {
    if (error = pthread_create( &thr[i], 0, matmult, (void *)i)) {
        fprintf(stderr, "pthread_create: %s", strerror(error));
        exit(1);
    }
}

// wait for workers to finish their jobs
for (i=0; i<M; i++)
    pthread_join(thr[i], 0)
    /* print the results ... */
}
```

CSE 451 11au 10/1/11

II-26

Derived from materials copyright © 2010 Thomas W. Droeppner.

## matrix multiply (cont.)

```
// each thread executes this routine
void *matmult(void *arg) {
    int row = (int)arg; // gens a warning on 64-bit
    systems
    int col;
    int i;
    int t;

    for (col=0; col < P; col++) {
        t = 0;
        for (i=0; i<N; i++) {
            t += A[row][i] * B[i][col];
        }
        C[row][col] = t;
    }
    return(0);
}
```

CSE 451 11au 10/1/11

II-27

Derived from materials copyright © 2010 Thomas W. Droeppner.

## Compiling It

```
Linux% gcc -o mat mat.c -pthread
```

*Failure to specify the pthreads library results  
in a linker error.*

CSE 451 11au 10/1/11

II-28

Derived from materials copyright © 2010 Thomas W. Droeppner.