

## **Tutorial: Introduction to POSIX Threads**

#### Stéphane Zuckerman Haitao Wei Guang R. Gao

Computer Architecture & Parallel Systems Laboratory Electrical & Computer Engineering Dept. University of Delaware 140 Evans Hall Newark,DE 19716, USA {szuckerm, hwei, ggao}@udel.edu

Friday, September19, 2014

## Outline



## Introduction

- An Introduction to Multithreading
- Processes and Threads Layouts in Memory

## **PTHREADS Basics**

- A Short Introduction to POSIX Threads
- PTHREADS Program Examples

## 3 Where to Learn More



#### **Processes: a Definition**

A process is a set of instructions with its own memory space which is accessed privately. A process is composed of a sequence of instructions (its code), as well as input and output sets (its data). Accessing the memory allocated to a process is in general forbidden unless specific mechanisms are being used, such as inter-process communication functions (IPCs).

#### **Threads: a Definition**

A thread is a sequence of code that is part of a process. Consequently, processes contain at least one thread. All threads belonging to the same process share the same address space, and thus can access the same memory locations.



#### **Process**

- A list of instructions
- Some memory to access with the guarantee it is exclusive to the process
  - A stack to store current values with which to compute
  - A heap to store bigger objects that don't fit in the stack

#### Thread

- A list of instructions
- A memory space
  - A stack to store current values with which to compute (private to the thread)
  - Some heap space, shared between threads belonging to the same process

Various Kinds of Multithreading



- User threads
- Kernel threads
- Hybrid ( $M \times N$ ) threads













Slides inspired by M. Pérache's multithreading course





Zuckerman et al.

Slides inspired by M. Pérache's multithreading course





Zuckerman et al.













Slides inspired by M. Pérache's multithreading course





Zuckerman et al.

#### **Characteristics of User Threads**



- 1 thread per kernel process
- Simple to implement
- Threads libraries were initially implemented this way
- Very fast: fully running in user space
- Not really suited to SMP and CMP architectures
- Usually handle system calls badly
- Example of "popular" user thread library: GNU Pth

Slides inspired by M. Pérache's multithreading course





Zuckerman et al.





























## **Characteristics of Kernel Threads**



- ► *N* kernel threads
- Well suited to SMP and CMP architectures
- Handles system calls nicely
- Completely managed at the system level
- Complex to implement
- Slower than user threads (overheads due to entering kernel space)
- Example of "popular" user thread libraries: Windows Threads, LinuxThreads, NPTL

Slides inspired by M. Pérache's multithreading course





Zuckerman et al.

















Slides inspired by M. Pérache's multithreading course





Zuckerman et al.





Slides inspired by M. Pérache's multithreading course





Zuckerman et al.









#### **Characteristics of Hybrid Threads**



- M kernel threads and N user threads: hybrid threads are also called M × N threads (or sometimes M : N threads)
- Well suited to SMP and CMP architectures
- Most Complex to implement
- Two schedulers:
  - Kernel Space Scheduler
  - User Space Scheduler
- Efficient
- Handles system calls "well enough" (better than user threads, less than kernel threads)
- Examples of M × N thread libraries: Solaris' default thread library (until Solaris v10), MPC, most efficient implementations of OpenMP's runtime system.

## Process Layout in Memory An Example Implementation in the Linux OS





Memory

Structure

## **Thread Layout in Memory** An Example Implementation in the Linux OS



2

1

0



#### Structure

luckerman et al.	PThreads

## A Thread's Characteristics An Example Implementation in the Linux OS



- All threads share the same address space
- A thread's stack never grows (except for Thread 0)
- A thread's stack is located in the heap (except for Thread 0)
- Global variables are shared by all threads
- Threads communicate directly through memory

## A Short Introduction to POSIX Threads



- Based on the IEEE POSIX 1003.1 standard
- Any POSIX-compliant system (*i.e.*, UNIX and Linux at the very least) implement the PTHREAD standard:
  - Linux implements PTHREADS using kernel threads
  - Solaris used to implement PTHREADS as an  $M \times N$  library, but now it is implemented as a kernel thread library
  - OpenBSD used to have a user-level PTHREAD library, but now uses kernel-level one
  - There are a few third-party libraries to provide a source compatibility with PTHREADS on MS-Windows systems
- ► Are PTHREADS lightweight processes?
  - Well, a lightweight process, in essence, is a kernel thread. So if your PTHREAD library is implemented as kernel threads, then yes.
  - In general, the answer is "it depends"



- How to create and destroy threads
- How to make threads synchronize with each other



pthread_t	A PTHREAD descriptor and ID
pthread_mutex_t	A lock for PTHREADS
pthread_cond_t	A conditional variable. It is necessarily associated
	with a mutex
pthread_attr_t	Descriptor for a PTHREAD's properties
	( <i>e.g.</i> , scheduling hints)
pthread_mutexattr_t	Descriptor for mutex' properties ( <i>e.g.</i> ,
	private to the process or shared between processes;
	recursive or not; etc.)
pthread_condattr_t	Descriptor for a condition variable ( <i>e.g.</i> , private
	to the process, or shared between processes)

#### PTHREADS: Basic Functions Creation and Destruction



#### Creation

Creates a new PTHREAD, using its descriptor reference, the required attributes (or NULL for default attributes), a function pointer, and an argument pointer. The function returns 0 if it succeeded, and -1 otherwise. The descriptor is filled and becomes "active" if the call succeeded.

### **Destruction**

int pthread\_join( pthread\_t tid, void\*\* retval )
Waits for the PTHREAD with ID tid to return, and stores its return
value retval. If retval is NULL, the return value is discarded.
pthread\_join returns 0 on success, and -1 otherwise.

# Note: Calling exit (3) from *any* thread will terminate the whole process, and thus all threads will also terminate!

Zuckerman et al.

PThreads



#### void pthread\_exit( void\* retval )

Exits from the thread calling the function. If retval is not NULL, it contains the return value of the thread to pthread\_join (see below).

#### pthread\_t pthread\_self( void )

Retrieves a thread's own ID. Note: pthread\_t, while often implemented as an integer, does not have to be!

#### A First PTHREAD Example



#### Hello, World! ... Headers and worker function

```
#include <stdio.h> // for snprintf(), fprintf(), printf(), puts()
#include <stdlib.h> // for exit()
#include <errno.h> // for errno (duh!)
#include <pthread.h> // for pthread_*
#define MAX NUM WORKERS 4UL
typedef struct worker_id_s { unsigned long id } worker_id_t;
void* worker(void* arg)
   // Remember, pthread_t objects are descriptors, not just IDs!
   worker_id_t* self = (worker_id_t*) arg; // Retrieving my ID
    char hello[100]; // To print the message
    int err = snprintf(hello, sizeof(hello),
                       "[%lu]\t_Hello,_World!\n", self->id);
    if (err < 0) { perror("snprintf"); exit(errno); }</pre>
   puts(hello);
    return arg; // so that the "master" thread
                // knows which thread has returned
```

#### A First PTHREAD Example

ELAWARE.

Hello, World! ...main

```
#define ERR_MSG(prefix,...) \
    fprintf(stderr,prefix ",%lu_out_of,%lu_threads",___VA_ARGS___)
int main(void) {
 pthread t workers [ MAX NUM WORKERS ];
 worker_id_t worker_ids [ MAX_NUM_WORKERS ];
 puts("[main]\tCreating_workers...\n");
  for (unsigned long i = 0; i < MAX_NUM_WORKERS; ++i) {</pre>
   worker_ids[i].id = i;
    if (0 != pthread create(&workers[i], NULL, worker, &worker ids[i]))
      { ERR MSG("Could not create thread", i, MAX NUM WORKERS);
        exit(errno); }
  puts("[main] \tJoining.the.workers...\n");
  for (unsigned long i = 0; i < MAX NUM WORKERS; ++i) {
    worker id t* wid = (worker id t*) retval;
    if (0 != pthread_join(workers[i], (void**) &retval))
     ERR_MSG("Could_not_join_thread", i, MAX_NUM_WORKERS);
    else
     printf("[main]\tWorker N.%lu has returned!\n", wid->id);
  return 0;}
```

## A First PTHREAD Example

**U**ELAWAR

#### Hello, World! ... Output

#### **Compilation Process**

```
gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c hello.c
gcc -o hello hello.o -lpthread
```

... Don't forget to link with the PTHREAD library!

... And the output:

#### Output of ./hello

```
[main] Creating workers...
[0] Hello, World!
[main] Joining the workers...
[2] Hello, World!
[main] Worker N.0 has returned!
[1] Hello, World!
[3] Hello, World!
[main] Worker N.1 has returned!
[main] Worker N.2 has returned!
[main] Worker N.3 has returned!
```



```
#ifndef BAD_GLOBAL_SUM_H
#define BAD_GLOBAL_SUM_H
#include <stdio.h>
#include <stdlib.h>
#include "utils.h"
typedef struct bad global sum s {
   unsigned long *value;
 bad_global_sum_t;
#endif // BAD_GLOBAL_SUM_H
```

Figure : bad\_global\_sum.h

kerman et a
-------------



```
#include "bad_global_sum.h"
#define MAX NUM WORKERS 20UL
typedef unsigned long ulong t;
void* bad sum(void* frame) {
    bad global sum t* pgs = (bad global sum t*) frame;
   ++*pgs->value;
   return NULL:
int main(void) {
   pthread t threads [ MAX NUM WORKERS ];
   bad global sum t frames [ MAX NUM WORKERS ];
   ulong_t counter = 0;
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i) {</pre>
        frames[i].value = &counter;
        spthread create(&threads[i],NULL,bad sum,&frames[i]);
    for (ulong t i = 0; i < MAX_NUM_WORKERS; ++i)</pre>
        spthread join(threads[i],NULL);
    printf("%lu threads were running. Sum final value: %lu\n", MAX_NUM_WORKERS, counter);
    return 0;
```

#### Figure : bad\_sum\_pthreads.c

Zuckerman et al.	PThreads	24 / 41



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c bad\_sum\_pthreads.c
gcc -o badsum bad\_sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c bad\_sum\_pthreads.c
gcc -o badsum bad\_sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!

#### Output of ./badsum

szuckerm@evans201g:bad\$ ./badsum
20 threads were running. Sum final value: 20

Hey, it's working!



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c bad\_sum\_pthreads.c
gcc -o badsum bad\_sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!

#### Output of ./badsum

szuckerm@evans201g:bad\$ ./badsum
20 threads were running. Sum final value: 20

#### Hey, it's working!

#### Multiple executions of . /badsum

szuckerm@evans201g:bad\$ (for i in `seq 100';do ./badsum ;done)|uniq 20 threads were running. Sum final value: 20 20 threads were running. Sum final value: 19 20 threads were running. Sum final value: 20 20 threads were running. Sum final value: 19 20 threads were running. Sum final value: 20



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c bad\_sum\_pthreads.c
gcc -o badsum bad\_sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!

#### Output of ./badsum

szuckerm@evans201g:bad\$ ./badsum
20 threads were running. Sum final value: 20

#### Hey, it's working!

#### Multiple executions of . /badsum

szuckerm@evans201g:bad\$ (for i in `seq 100`;do ./badsum ;done)|uniq 20 threads were running. Sum final value: 20 20 threads were running. Sum final value: 19 20 threads were running. Sum final value: 20 20 threads were running. Sum final value: 20 20 threads were running. Sum final value: 20

#### Waiiiiiit a minute...

Zuckerman et al.

#### **Incrementing a Global Counter**

**Fixing the Implementation** 

#### **Mutexes**

A MUTual EXclusive object (or mutex) is a synchronization object which is either owned by a single thread, or by no-one. It is the basic block to create critical sections.

```
#ifndef GLOBAL_SUM_H
#define GLOBAL_SUM_H
#include <stdio.h>
#include <stdlib.h>
#include "utils.h"
typedef struct global_sum_s {
    unsigned long *value;
    pthread_mutex_t *lock;
} global_sum_t;
#endif // GLOBAL SUM H
```

#### Figure: global\_sum.h

Zuc	kerman	et al.	



## **Incrementing a Global Counter**



#### Fixing the Implementation (2)

```
#include "global_sum.h"
#define MAX NUM WORKERS 20UL
typedef unsigned long ulong t;
void* sum(void* frame) {
   global sum t* gs = (global sum t*) frame;
   spthread_mutex_lock ( gs->lock ); /* Critical section starts here */
   ++*gs->value;
   spthread mutex unlock (gs->lock); /* Critical section ends here */
   return NULL:
int main(void) {
   pthread_t threads [ MAX_NUM WORKERS ];
   global sum t frames [ MAX NUM WORKERS ];
   ulong t counter = 0;
   pthread mutex t m = PTHREAD MUTEX INITIALIZER;
   for (ulong t i = 0; i < MAX NUM WORKERS; ++i) {</pre>
        frames[i] = (global_sum_t) { .value = &counter, .lock = &m };
       spthread create(&threads[i],NULL,sum,&frames[i]);
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i)</pre>
       spthread join(threads[i],NULL);
   printf("%lu threads were running. Sum final value: %lu\n", MAX_NUM_WORKERS, counter);
   return 0;
                              Figure : sum_pthreads.c
```

_		
/110	(ormon	- ot o
Z I II.:	Nei Itali	H A

#### **Incrementing a Global Counter** Fixing the Implementation (3)



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c sum\_pthreads.c
gcc -o sum sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!

#### Multiple executions of ./sum

szuckerm@evans201g:good\$ (for i in `seq 100`;do ./sum ;done)|uniq 20 threads were running. Sum final value: 20

#### **Incrementing a Global Counter** Fixing the Implementation (3)



#### **Compilation Process**

gcc -Wall -Wextra -pedantic -Werror -O3 -std=c99 -c sum\_pthreads.c
gcc -o sum sum\_pthreads.o -lpthread

#### ... Don't forget to link with the PTHREAD library!

#### Multiple executions of ./sum

szuckerm@evans201g:good\$ (for i in `seq 100`;do ./sum ;done)|uniq 20 threads were running. Sum final value: 20

#### Fixed!



## **Condition variables**

Condition variables are used when threads are waiting on a specific event. When the event occurs, the code where it the event was realized *signals* a condition variable, either to wake up one of the threads waiting on the event, or all of them.

#### **Examples of Events to Be Worth Signaling**

- Availability of a resource, e.g.:
  - A file descriptor for a network connection,
  - A file descriptor for accessing (reading or writing) a regular file,
  - Any device handle, really
- A specific input provided by the user (string provided by the user, etc.)
- ▶ etc.

#### Reacting on Specific Events II Condition Variables



## High-Level Explanation: Waiting on a Condition

- A condition variable is always associated with a mutex
- Ito wait on an event, a thread must first acquire the mutex, then
- Call int pthread\_cond\_wait ( pthread\_cond\_t\* cond, pthread\_mutex\_t\* mutex )
- If the call succeeds, then the thread releases the mutex
- When the condition variable is signaled, *if* the thread which was "asleep" is re-awakened, the system first returns ownership of the mutex back to it



## **High-Level Explanation: Signaling an Event Has Occurred** There are two function calls to perform this function:

> int pthread\_cond\_signal( pthread\_cond\_t\* cond )

- To signal a single thread that the event has occurred. Note: there is no guarantee as to *which* thread will wake
- int pthread\_cond\_broadcast( pthread\_cond\_t\* cond )
  - To signal all threads that the event has occurred.

### **Reacting on Specific Events**

#### **Condition Variables**

```
#ifndef BARRIER H
#define BARRIER H
#define SET BARRIER MSG(...) ∖
    snprintf(buffer, sizeof(buffer), ___VA_ARGS__)
#define NOT_LAST_TO_REACH \
    "[%lu]\tI'm_NOT_the_last_one_to_reach_the_barrier!"
#define LAST TO REACH
    "[%lu]\tI.am.the.last.to.reach.the.barrier!.Waking.up.the.others."
typedef struct barrier s {
   pthread mutex t *lock;
   pthread_cond_t *cond;
   ulong t *count;
} barrier t;
typedef struct context s {
   barrier t* barrier;
   ulong t id;
context t;
#endif // BARRIER H
```

Figure · hannian h

PThreads



## **Reacting on Specific Events**

#### **Condition Variables (2)**



41

```
#include "barrier.h"
void* worker(void* frame);
int main(void) {
   pthread_t threads [ MAX_NUM WORKERS ];
   context_t contexts [ MAX_NUM WORKERS ];
   pthread mutex t m = PTHREAD MUTEX INITIALIZER;
   pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
   ulong t count = MAX NUM WORKERS;
   barrier t barrier = {.lock = \&m, .cond = \&cond, .count = \&count};
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i)</pre>
       contexts[i] = (context_t) { .barrier = &barrier, .id = i };
       spthread create(&threads[i],NULL,worker,&contexts[i]);
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i)</pre>
       spthread join(threads[i],NULL);
   return 0;
```

#### Figure : barrier\_main.c

Zuckerman et al.	PThreads	33 /

## Reacting on Specific Events

Condition Variables (3)

```
#include "barrier.h"
void* worker(void* frame) {
   char buffer[81];
    context_t* c = (context_t*) frame;
   printf("[%lu]\tReaching.the.barrier...\n",c->id);
    spthread_mutex_lock ( c->barrier->lock );
    --*c->barrier->count;
    if (*c->barrier->count > 0) {
        SET BARRIER MSG(NOT LAST TO REACH, c->id);
        spthread cond wait ( c->barrier->cond, c->barrier->lock );
    } else {
        SET_BARRIER_MSG(LAST_TO_REACH, c->id);
   puts(buffer);
    spthread mutex unlock ( c->barrier->lock );
   pthread cond broadcast ( c->barrier->cond );
   printf("[%lu]\tAfter_the_barrier\n", c->id);
    return NULL;
                          Figure : barrier.c
```

_			
2110	korman		<u>_</u>
<b>2</b> uu	Nerman	61	- CU I



#### Reacting on Specific Events Condition Variables (4)



szuckerm@evans201g:condvar\$ gcc -Wall -Wextra -pedantic -Werror -03 -std=c99 -c barrier.c szuckerm@evans201g:condvar\$ gcc -o barrier barrier.o -lpthread szuckerm@evans201g:condvar\$ ./barrier [0] Reaching the barrier ... [2] Reaching the barrier ... [1] Reaching the barrier ... [3] Reaching the barrier ... [4] Reaching the barrier ... [5] Reaching the barrier ... [7] Reaching the barrier... [6] Reaching the barrier ... [6] I am the last to reach the barrier! Waking up the others. [6] After the barrier [0] I'm NOT the last one to reach the barrier! [0] After the barrier [1] I'm NOT the last one to reach the barrier! [1] After the barrier [2] I'm NOT the last one to reach the barrier! [2] After the barrier [3] I'm NOT the last one to reach the barrier! [3] After the barrier [4] I'm NOT the last one to reach the barrier! [4] After the barrier [5] I'm NOT the last one to reach the barrier! [5] After the barrier [7] I'm NOT the last one to reach the barrier! [7] After the barrier

**Creating Barriers More Easily** 



"Hey, barriers are nice! I wish I could have a more practical construct, though."



- "Hey, barriers are nice! I wish I could have a more practical construct, though."
- ▶ ... Well actually, did I tell you about PTHREAD barriers?

#### pthread barrier t and its associated functions

- int pthread\_barrier\_init( pthread\_barrier\_t restrict\* barrier, const pthread\_barrierattr\_t \*restrict attr, unsigned count )
- int pthread\_barrier\_destroy( pthread\_barrier\_t restrict\* barrier )
- int pthread\_barrier\_wait( pthread\_barrier\_t restrict\* barrier )

#### Updated Barrier Program Using PTHREAD Barriers



#endif // BARRIER\_H

Figure : pth\_barrier.h

```
#include "barrier.h"
void* worker(void* frame) {
    context_t* c = (context_t*) frame;
    printf("[%lu]\tReaching_the_barrier...\n",c->id);
    spthread_barrier_wait( c->barrier );
    printf("[%lu]\tAfter_the_barrier\n", c->id);
    return NULL;
}
Figure : pth_barrier.c (1)
```



_		
7110	kormai	n of a
Zuu	Nermai	1616

#### Updated Barrier Program Using PTHREAD Barriers (2)



```
#include "barrier.h"
int main(void) {
   pthread t threads [ MAX NUM WORKERS ];
   context_t contexts [ MAX_NUM_WORKERS ];
   ulong t count = MAX_NUM_WORKERS;
   pthread barrier t barrier;
    spthread barrier init(&barrier,NULL,count);
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i)</pre>
        contexts[i] = (context t) \{ .barrier = \& barrier, .id = i \};
        spthread_create(&threads[i],NULL,worker,&contexts[i]);
    for (ulong t i = 0; i < MAX NUM WORKERS; ++i)</pre>
        spthread join(threads[i],NULL);
    spthread barrier destroy(&barrier);
    return 0;
                      Figure : pth_barrier.c (2)
```

## Learning More About Multi-Threading and PTHREADS



#### Books (from most theoretical to most practical)

- ▶ Tanenbaum, Modern Operating Systems
- Herlihy and Shavit, The Art of Multiprocessor Programming
- Bovet and Cesati, Understanding the Linux Kernel, Second Edition
- Stevens and Rago, Advanced Programming in the UNIX Environment, 3rd Edition

#### **Internet Resources**

- "POSIX Threads Programmings" at https://computing.llnl.gov/tutorials/pthreads/
- "Multithreaded Programming (POSIX pthreads Tutorial)" at http://randu.org/tutorials/threads/

#### **Food for Thoughts**

- Sutter, "The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software" (available at http://www.gotw.ca/publications/concurrency-ddj.htm)
- Lee, "The Problem with Threads" (available at http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-1.pdf)
- Boehm, "Threads Cannot Be Implemented As a Library" (available at www.hpl.hp.com/techreports/2004/HPL-2004-209.pdf)

#### **References I**



Boehm, Hans-J. "Threads Cannot Be Implemented As a Library". In: SIGPLAN Not. 40.6 (June 2005), pp. 261-268. ISSN: 0362-1340. DOI: 10.1145/1064978.1065042.URL: http://doi.acm.org/10.1145/1064978.1065042. Bovet, Daniel and Marco Cesati. Understanding the Linux Kernel, Second Edition. Ed. by Andy Oram. 2nd ed. Sebastopol, CA, USA: O'Reilly & Associates, Inc., 2002. ISBN: 0596002130. Herlihy, Maurice and Nir Shavit. The Art of Multiprocessor Programming. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2008. ISBN: 0123705916, 9780123705914. Lee, Edward A. "The Problem with Threads". In: Computer 39.5 (May 2006), pp. 33-42. ISSN: 0018-9162. DOI: 10.1109/MC.2006.180. URL: http://dx.doi.org/10.1109/MC.2006.180.

#### **References II**



- Stevens, Richard W. and Steven A. Rago. Advanced Programming in the UNIX Environment, 3rd Edition. Indianapolis, IN, USA: Addison-Wesley Professional, 2013. ISBN: 0321637739, 9780321637734.
- Sutter, Herb. "The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software". In: *Dr. Dobb's Journal* 30.3 (2005).
- Tanenbaum, Andrew S. *Modern Operating Systems*. 3rd. Upper Saddle River, NJ, USA: Prentice Hall Press, 2007. ISBN: 9780136006633.