

PA193 - Semminar on concurrency

Miroslav Jaroš

1st March 2022

Quick quiz

- What is concurrency?
- Thread x Process difference
- Who provides threads?

Quick quiz

- What is concurrency?
 - Concurrency is the decomposability property of a program, algorithm, or problem into order-independent or partially-ordered components or units.
- Thread x Process difference
 - Thread is minimal runnable unit for OS that can be deployed on processor, unlike process it does not own virtual memory, but uses virtual memory of parent process.
 - Every process has at least one thread, which is main for execution.
- Who provides threads?
 - Threads are provided by OS, although many languages have their own implementation due to optimization.

Concurrency and security

- Modern processors are built as pipelined and superscalar and multicore
 - Pipelined processors have instruction execution split into different stages
 - e.g. Instruction Fetch (IF), Instruction Decode (ID), Execution (EX), Memory access (MEM), Write Back (WB)
 - Superscalar processors scales individual stages to execute multiple instructions at a same time
 - This leads to effects like speculative execution, when processor enables instructions to be executed out of order.
 - And then, Specter and Meltdown happened:
<https://meltdownattack.com/>

Concurrency and security 2.

- But not only hardware is problematic. The idea of operating system stands on concurrency:
 - Operating system is responsible for process creation and planning even on single core processors.
 - But operating system is responsible for access control.
- Then there are high level applications.
 - Programmers typically rely on some properties of other dependencies - like transactions in databases.
 - But in large scale applications the execution of actions more resembles some process.
 - Are we sure that the side effects of the process are secured enough.

Race condition

- Imagine we want to increase variable up to a certain number and we decide to use threads.
- The process is simple, unless the number is x we do $x++$ and we do this in several threads.
- How would this code look in the assembly language?

```
volatile int shared_var = 0;
void increase()
{
    while (shared_var < 10000) {
        shared_var++;
    }
}
```

- The dining philosophers problem
 - Five philosophers are sitting around table each with one plate.
 - Between each plate is one chopstick
 - Each philosopher needs two chopsticks to eat.
 - The philosophers don't communicate between each other and they can either dine or think.
 - What if all of them will be extremely hungry and since all are right handed will take the right chopstick as a first every time.
 - Will they eat or die of hunger?

Threads

UNIX

- Pthread library part of POSIX
- `#include <pthread.h>`

WINDOWS

- Defined in WIN32 API
- `#include <windows.h>`

MULTI PLATFORM

- Since C11 and C++11 standards are threads part of standard library
- Qt Framework
- Boost library

LANGUAGES

- **GO**: goroutines
- **ERLANG**: processes
- **ADA**: tasks
- **Java**: `java.lang.Thread`
- **Python**: thread and threading module

Pthread library

- Part of POSIX library
- Provides basic interface for Thread management and mutual exclusion techniques
- All types and functions are prefixed with pthread string
- Needs to be compiled with `-pthread` argument, to link pthread library into binary
- All types and functions are described in `pthread.h` header file

Pthread library

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

- Creates new thread, which will start execution in `start_routine` function
- **thread** in/out attribute, after `pthread_create` is called, it's set to threads identifier
- **attr** thread attributes, typically passed `NULL`
- **start_routine** entry point of newly created thread
- **arg** arguments passed to `start_routine`
- `man 3 pthread_create`

```
int pthread_join(pthread_t thread, void **retval);
```

- Waits for thread to end execution and collect return value
- **thread** thread identifier, set by `pthread_create`
- **retval** if not set to `NULL` `pthread_join` will store `start_routine` return value in it.
- `man 3 pthread_join`

Helgrind

- Part of valgrind tool for dynamic analysis
- Designed to find bugs in threaded code
- Executed similarly to memcheck
- `valgrind --tool=helgrind ./your_code`
- Your code should be compiled with debugging symbols “`gcc -g ...`”
- <http://valgrind.org/docs/manual/hg-manual.html>

CRITICAL SECTION

- Point of code where shared resource is manipulated.
- Must be executed exclusively - only one thread at time
- **Even read operations must be exclusive**
 - Context switch can happen in the middle of read operation
 - Then data can be inconsistent
- Goal is to make critical section as small as possible
- Use mutual exclusion to achieve exclusivity

Mutual exclusions

- Posix defines several methods of mutual exclusion
- Mutex - **M**utual **E**xclusion
- Condition variable
- Semaphore

Mutex

- Object which can be in two states, locked and unlocked
- When thread wants to enter critical section, it locks mutex
- When other thread tries to lock mutex, the execution will be stopped and will wait until mutex is unlocked by blocking thread
- When thread is leaving critical section, it unlocks the mutex
- `man 3 pthread_mutex_lock`

```
#include <pthread.h>
int pthread_mutex_destroy(pthread_mutex_t *mutex);
int pthread_mutex_init(pthread_mutex_t *restrict mutex,
    const pthread_mutexattr_t *restrict attr);
pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
int pthread_mutex_lock(pthread_mutex_t *mutex);
int pthread_mutex_trylock(pthread_mutex_t *mutex);
int pthread_mutex_unlock(pthread_mutex_t *mutex);
```

Condition Variable

- Critical section can be entered after condition is met
- Typically in producer-consumer applications, where consumer needs to wait for producer
- Consumer locks mutex, but finds, that it cannot enter critical section
- It calls `pthread_cond_wait` and sleeps, mutex is unlocked
- When producer creates new resource, it calls `pthread_cond_signal`
- All threads waiting for condition are waked and tries to obtain lock, check condition and if its met, they enter critical section with locked mutex
- `man 3 pthread_cond_init`

```
#include <pthread.h>
int pthread_cond_destroy(pthread_cond_t *cond);
int pthread_cond_init(pthread_cond_t *restrict cond,
    const pthread_condattr_t *restrict attr);
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
int pthread_cond_signal(pthread_cond_t *cond);
int pthread_cond_wait(pthread_cond_t *restrict cond,
    pthread_mutex_t *restrict mutex);
```

Semaphores

- Integer value that identifies how many resources are consumable
- Every time the new resource is created, or released the counter is increased
- Every time resource is consumed the counter is decreased.
- Thread that tries to use resource sleeps until resource is allocated
- This allows multiple threads enter critical section when there are enough resources
- Sometimes it needs to be used with mutex, due to possible inconsistencies.
- `man 3 sem_init`

```
#include <semaphore.h>
```

```
int sem_init(sem_t *sem, int pshared, unsigned int value);  
int sem_destroy(sem_t *sem);  
int sem_post(sem_t *sem);  
int sem_wait(sem_t *sem);
```


Tasks

Work on any UNIX computer

- 1 Create program which will increase one variable to 10000 from 3 different threads
- 2 Increase number of threads to 100 and wait for problems to appear
- 3 Try to find problems with helgrind
- 4 Add locking to your program
- 5 Try helgrind to find possible race conditions
- 6 Modify your code to create deadlock
- 7 Test it with helgrind
- 8 Fix your code, so deadlock won't happen.

Tasks II.

- 1 Create a program, where one thread is putting random numbers into the array and several others are verifying if those numbers are primes.
- 2 The time needed to produce the numbers should be unpredictable (use random numbers again) so the worker threads will need to wait for new elements to appear
- 3 Use an appropriate technique for mutual exclusion to avoid busy waiting.
- 4 Try helgrind to navigate you through the “hell” of the problems.

Conclusion

- **Don't be afraid of threads**
- **Use threads in your applications**
- You should keep in mind dangers that concurrency can create in your code
- Always try to make critical sections as minimal as possible
- Use mutexes, semaphores and other tools to avoid race conditions, deadlocks and other possible issues
- Check your code with helgrind, it can save you many hours of debugging

Conclusion

- Write your code with concurrency in mind, you might not want to write concurrent library, but someone will eventually try to use it with threads
- Many frameworks and libraries uses threads, even though you don't know it
- Last but not least: **Test your code!**
 - Multi threaded applications are hard to debug, you need to be sure, that particular function/method is doing what it should do!