## **Introduction to Lab 3**

Mahdad Davari <mahdad.davari@it.uu.se>

Division of Computer Systems Dept. of Information Technology Uppsala University

2013-10-28

# In this lab

- Parallel programming using pthreads
- A real-world example of parallelising sequential code for multicores (Gauss-Seidel numerical method)
- Putting all together parallel programming techniques, data access pattern, and cache usage when writing parallel code for multicores
- A lot of mathematics :~(

## Parallel programming paradigms

- Process-level
  - MPI (Message Passing Interface) library
    - Process creation
    - Explicit message "send" and "receive", synchronizations, barriers
- Thread-level
  - POSIX Threads (pthreads) library
    - Explicit thread creation and control by programmer
  - Open Multi-Processing (OpenMP) Compiler Extension
    - Implicit thread creation and control by compiler

## Some pthread types and routines ...

#include <pthread.h>

pthread\_t thread;

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

Parameters:

thread Where to store the thread ID. attr Attributes for the thread, NULL defaults. start\_routine Procedure to call in the new thread. arg Argument passed to start\_routine

#### **Return Value:**

0 if successful, error number otherwise.

#include <pthread.h>

Waiting for threads to terminate

Parameters:

thread Thread to wait for.

value\_ptr Pointer to variable to store return value in, NULL to discard return value.

#### Return Value:

0 if successful, error number otherwise.

```
Thread creation
An example
   #include <pthread.h>
   #include < stdio.h>
   static void *my_thread(void *arg) {
            printf("Hello_Threads!\n");
           return NULL;
   }
   int main(int argc, char * argv[]) {
            pthread t thread;
            pthread_create(&thread, NULL,
                           my_thread, NULL);
            pthread_join(thread, NULL);
           return 0;
   }
```

#include <pthread.h>

pthread\_mutex\_t mutex;

Mutex initialization the easy way, uses default attributes. No need for explicit cleanup.

#### Parameters:

mutex Pointer to mutex to initialize.

attr Pointer to mutex attributes, NULL for default attributes.

#### Return Value:

0 if successful, error number otherwise.

#### Parameters:

mutex Pointer to mutex to destroy.

#### **Return Value:**

0 if successful, error number otherwise.

#include <pthread.h>

Parameters:

mutex Pointer to mutex to lock or unlock.

#### **Return Value:**

0 if successful, error number otherwise.

```
Mutexes
Example
   static int balance = 512;
   static pthread mutex t balance mutex =
           PTHREAD MUTEX INITIALIZER;
   static int withdraw (int amount) {
           int ret = 0;
           pthread_mutex_lock(&balance_mutex);
           if (balance > amount) {
                   balance -= amount;
                   ret = amount;
           pthread mutex unlock(&balance mutex);
           return ret;
```

}

#include <pthread.h>

pthread\_barrier\_t barrier;

Note:

Barriers are *optional* in the Posix specification. **Parameters:** 

barrier Pointer to barrier to initialize.

attr Pointer to barrier attributes, NULL for defaults.

count Number of threads to wait for.

#### Return Value:

0 if successful, error number otherwise.

#### Parameters:

barrier Pointer to barrier to destroy. Return Value: 0 if successful, error number otherwise.

#include <pthread.h>

pthread\_barrier\_t barrier;

Parameters:

barrier Pointer to barrier to wait for.

#### **Return Value:**

PTHREAD\_BARRIER\_SERIAL\_THREAD or 0 on success, error number otherwise.

#include <pthread.h>

pthread\_barrier\_t barrier;

void \* thread() {

// some computation

pthread\_barrier\_wait(&barrier);

// rest of computation

return;

}

#include <pthread.h>

pthread\_barrier\_t barrier;

void \* thread() {

// some computation

int ret = pthread\_barrier\_wait(&barrier);

```
if (ret != 0 && ret != PTHREAD_BARRIER_SERIAL_THREAD) {
    // error handling routine
} else if (ret == PTHREAD_BARRIER_SERIAL_THREAD) {
    // serial task
}
```

// rest of computation

return;

}

## More information on pthread ...

- Lawrence Livermore National Laboratory <sup>1</sup>
- Single Unix Specification<sup>2</sup>
- Your local system's manual pages:
  - host\$ man man
  - host\$ man pthreads
- Variety of material available online

1. https://computing.llnl.gov/tutorials/pthreads/

2. http://www.unix.org/single\_unix\_specification/

- Gauss-Seidel method:
  - An iterative method for solving linear systems of equations<sup>1</sup>

 $A\mathbf{x} = \mathbf{b}$ 

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix}.$$

1. http://en.wikipedia.org/

- Gauss-Seidel method:<sup>1</sup>
  - matrix A is decomposed into a lower triangular component , and a strictly upper triangular component U:

$$A = L_{*} + U \quad \text{where} \quad L_{*} = \begin{bmatrix} a_{11} & 0 & \cdots & 0 \\ a_{21} & a_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix}, \quad U = \begin{bmatrix} 0 & a_{12} & \cdots & a_{1n} \\ 0 & 0 & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix}$$
$$A\mathbf{x} = \mathbf{b} \cdot$$
$$L_{*}\mathbf{x} = \mathbf{b} - U\mathbf{x}$$
$$\mathbf{x}^{(k+1)} = L_{*}^{-1}(\mathbf{b} - U\mathbf{x}^{(k)}).$$

1. http://en.wikipedia.org/

• Gauss-Seidel method: <sup>1</sup> solved sequentially using forward substitution

$$x_i^{(k+1)} = \frac{1}{a_{ii}} \left( b_i - \sum_{j < i} a_{ij} x_j^{(k+1)} - \sum_{j > i} a_{ij} x_j^{(k)} \right), \quad i, j = 1, 2, \dots, n.$$

$$\begin{split} x_1 &= \frac{b_1}{l_{1,1}}, \\ x_2 &= \frac{b_2 - l_{2,1} x_1}{l_{2,2}}, \\ &\vdots \\ x_m &= \frac{b_m - \sum_{i=1}^{m-1} l_{m,i} x_i}{l_{m,m}}. \end{split}$$

1. http://en.wikipedia.org/

Gauss-Seidel method in lab3:

$$u_{i,j}^{k+1} = \frac{u_{i-1,j}^{k+1} + u_{i,j-1}^{k+1} + u_{i+1,j}^{k} + u_{i,j+1}^{k}}{4}$$

Algorithm 1 Gauss-Seidel solver for the Laplace equation on an  $n \times m$  matrix, with the tolerance t.

```
Require: n, m \ge 2

repeat

e \leftarrow 0

for i = 1 to n - 1 do

for j = 1 to m - 1 do

v \leftarrow \frac{u_{i-1,j} + u_{i,j-1} + u_{i+1,j} + u_{i,j+1}}{4}

e \leftarrow e + |u_{i,j} - v|

u_{i,j} \leftarrow v

end for

end for

until e \le t
```

• Access pattern for serial version:



• Access pattern for parallel version:



## In this lab

- We will complete the synchronization between threads that implement the parallel version of Gauss-Seidel solver.
- lab3 skeleton files available on Piazza (a revised version uploaded) and extract the package:
  - Makefile: automates compilation and run
    - make gs\_seq
    - make gs\_pth
    - make test
  - gs\_common.c: glue code for running lab3 it is recommended to have a quick look into this file (the beginning) to see how default values are defined.
  - gs\_interface.h
  - gsi\_seq.c: sequential implementation of GS sweep
  - gsi\_pth.c: your parallel version of GS sweep
  - solution.c: don't peek!

## In this lab

- To test your solution:
  - make test



#### Please ignore "test did NOT converge!"

gsi\_calculate () {

}

pthread\_create (&threads[t].thread, NULL, thread\_compute, &threads[t]);

pthread\_join (...);

typedef struct {
 int thread\_id;
 pthread\_t thread;

volatile double error;

/\* TASK: Do you need any thread local state for synchronization? \*/
 /\* !END\_SOLUTION! \*/
} thread\_info\_t;

gsi\_calculate () {

}

pthread\_create (&threads[t].thread, NULL, thread\_compute, &threads[t]);

pthread\_join (...);

thread\_compute (...) {

}

}

// TASK compute bounds

for (iteration = 0; < gs\_iterations && global\_error > gs\_tolerance; ++) {

thread\_sweep(tid, iter, lbound, rbound);

// TASK update global error  $\rightarrow$  which thread should update this ?

// TASK barrier: all threads should finish the current iteration before next one

thread\_sweep (...) {

}

}

```
threads[tid].error = 0.0;
```

```
for (int row = 1; row < gs_size - 1; row++) {
```

// TASK when should we start sweeping a row ?

```
for (int col = lbound; col < rbound; col++) {
    // sweep the row
}</pre>
```

// TASK tell the thread to the right to continue

}

We use four threads to parallelize GS sweep •



```
for (int row = 1; row < gs size - 1; row++) {
    for (int col = lbound; col < rbound; col++) {
         gs_matrix[GS_INDEX(row + 1, col)] +
         gs_matrix[GS_INDEX(row - 1, col)] +
         gs matrix[GS INDEX(row, col + 1)] +
         gs matrix[GS INDEX(row, col - 1)]);
```

matrix size n (n columns)

- Pay attention to gsi\_init() and gsi\_finish()
- Define your barrier before initialization and use
- Notice that first thread does not need to synchronize with its left thread (use tid)
- To debug, set DEBUG to 1 in gsi\_pth.c (#define DEBUG 1)
- Lab tasks review
  - Implement synchronization between threads working on the same row in matrix
  - Barrier implementation at the end of each iteration
    - extra: a solution without using barrier
  - Speed-up your design by modifying how error variable is stored (modify thread\_info\_t)

# Lab sign-up and groups

Sign up using doodle: piazza  $\rightarrow$  Course Page  $\rightarrow$  Resources  $\rightarrow$  General Resources

Lab Preparation: 2013-10-28, 13:00 ~ 17:00, room 1412

Group A: 2013-10-29, 08:00 ~ 12:00, room 1412 Group B: 2013-10-31, 08:00 ~ 12:00, room 1412 Group C: 2013-11-01, 13:00 ~ 17:00, room 1412

Good luck!