

Excercise Sheet 1

(To be processed until 11.11.)

Lecture Parallel Processing Winter Term 2025/26

Exercise 1: Run-Time Analysis of a Sequential Program

Before a program (which typically is not known in all its details) is parallelized, it is common practice to use performance analysis tools in order to identify the most compute intensive parts of the program code.

A simple, sampling-based tool for this purpose is the program `gprof`, which measures for each procedure (or method) P in the program how often it is called and how much time is required to process P . In doing so, the tool differentiates between

- the *inclusive time* of P , which also includes the processing time of the procedures called by P (*children*), and
- the *exclusive time* (or *self time*) of P , which does not include the processing time of these calls.

When, e.g., the procedure

```
void funcA() {  
    for (int i=0; i<N; i++) { ... }  
    funcB();  
}
```

is invoked, the *exclusive time* of `funcA` is only the runtime of the `for` loop, while the *inclusive time* also includes the runtime of `funcB`.

In this exercise, you should familiarize yourself with `gprof`. First, compile the file `example.cpp` (in the archive [u01eFiles.zip](#)¹ on the lecture's web page) with instrumentation for `gprof`:²

```
g++ -pg -g -o example example.cpp
```

Start the program with `./example`. During the execution, the gathered profiling data is written to a file `gmon.out`. A readable form of this data is then obtained using the command `gprof ./example`. Thoroughly look at the output, in particular the included explanations, and answer the following questions:

- Which function requires the most computing time (*exclusive*) and would therefore be the first candidate for parallelization?
 - What is the maximum speedup you could achieve, if only the body of this function were parallelized?
 - How often is the function called?
 - What other functions does the function call?
- Which function is most frequently called and from which other functions?
- How is the computation time of the function `func5` composed?

Exercise 2 (Compulsory): Data Dependences, Programming with C++ Threads

Submit until Tuesday, November 11th, 10:00 via moodle

Look at the file `compute.cpp` in the archive [u01eFiles.zip](#)³ on the lecture's web page. First, analyse the data dependences between the different function calls. The functions `f1` to `f7` do not have any side effects, thus they, e.g., do

¹<http://www.bs.informatik.uni-siegen.de/web/wismueller/v1/pv/u01eFiles.zip>

²By the way, the code in this file is a nice example of *obfuscation*, to prevent you from inspecting the (rather short) source code.

³<http://www.bs.informatik.uni-siegen.de/web/wismueller/v1/pv/u01eFiles.zip>

not access global variables.

- a) Based on the data dependences, determine which calls of the functions f_1 to f_7 can be executed in parallel. Try to remove the occurring anti dependences and output dependences (if any) by renaming one or several variables.
- b) Draw a task graph (see Section 2.7.7 of the lecture) representing the dependences between the different tasks, where in this case a task corresponds to a function call. In the example, all functions have an execution time of 1s. What is the execution time of the parallel program?

Document to be submitted: PDF file with task graph and answer.

- c) Parallelize the program using C++ Threads. If necessary, implement a proper synchronization using condition variables (see the comment in the code for a hint). Pay attention to ensure that the printed result is *exactly* the same as with the sequential program.

Document to be submitted: modified C++ file 'compute.cpp'.

Exercise 3 (Compulsory): Parallelization of Quicksort using C++ Threads

Submit until Tuesday, November 11th, 10:00 via moodle

In this exercise, the Quicksort algorithm is to be parallelized (see Section 2.7.4 of the lecture). The basic idea of Quicksort is as follows:

- From a given array, an element is selected as a pivot or a reference element, e.g., the element from the center.
- A larger element is sought from below and a smaller one from above.
- If such items are found, they are placed incorrectly and are swapped.
- This search and swap will continue until 'below' and 'above' meet. After that, there are only smaller (or equal) elements below, and only larger (or equal) elements above.
- These two areas must now be sorted in turn, which translates to two recursive calls to quicksort.

Have a look at the code in the archive [u01eFiles.zip](#)⁴ on the lecture's web page. The procedure

```
void quicksort(int *a, int lo, int hi)
```

in `qsort.cpp` sorts the portion of the array starting at index value `lo` up to and including index value `hi` in ascending order. In addition, some auxiliary functions are implemented:

- `void initialize(int *a, int n)` initializes an array `a` of length `n` with random numbers.
- `int checkSorted(int *a, int n)` checks if an array `a` of length `n` is sorted.
- `void printArray(int *a, int n)` can be used as necessary to print a (small!) array.

- a) Compile the program and run it. Test the code for different array lengths: 10'000'000, 20'000'000, 40'000'000 and 80'000'000 elements). Record the time that the program requires for sorting the array in each case.
- b) Rewrite the program that one of the two recursive invocations of `quicksort` will execute in a new thread, but only if the array length for this invocation is greater than a constant `MINSIZE`. Try different values for `MINSIZE` (recommended starting value: 1,000,000).

Document to be submitted: modified C++ file 'qsort.cpp'.

- c) Compare the performance of your implementation with the measurements from a), using a table. Calculate the speedup and interpret your results.

Document to be submitted: PDF file with measurement results.

For support, you can use the following tutorial on C++ threads:

<https://solarianprogrammer.com/2011/12/16/cpp-11-thread-tutorial>

For motivated students: The C library offers a function `qsort` (documentation: `man qsort`):

```
void qsort(void *base, size_t nmemb, size_t size,  
          int(*compar)(const void *, const void *))
```

Compare the performance of your implementation against the function `qsort` in the C library.

⁴<http://www.bs.informatik.uni-siegen.de/web/wismueller/v1/pv/u01eFiles.zip>