

## Exercise Sheet 2

(Deadline: 12.12.2023)

## Parallel Processing Winter Term 2023/24

**Preparation:** Download the program codes for this exercise sheet from the Internet:

<https://www.bs.informatik.uni-siegen.de/web/wismueller/vl/pv/u2Files.zip>

### Exercise 1: Dependence analysis

Consider the following program fragment:

```
double b[10], a[10], c[10];
for (int i=1; i<10; i++) {
    a[i] = (b[i-1] + b[i]) / 2; // S1
    b[i] = a[7] * c[i];        // S2
    c[i-1] = 2 * b[i] ;        // S3
}
```

Which of the following statements are correct and which are incorrect?

- a) There is a **true** dependence between **different** loop iterations from S1 to S2.
- b) There is an **anti** dependence within a **single** loop iteration from S1 to S2.
- c) There is a **true** dependence within a **single** loop iteration from S2 to S1.
- d) There is a **true** dependence between **different** loop iterations from S2 to S1.
- e) There is an **anti** dependence between **different** loop iterations from S2 to S1.
- f) There is a dependence from S1 to S3.
- g) There is an **anti** dependence between **different** loop iterations from S2 to S3.
- h) There is an **output** dependence within a **single** loop iteration from S2 to S3.
- i) There is a **true** dependence between **different** loop iterations from S3 to S2.
- j) The loop does not contain any **output** dependencies.

*Note: this is a question similar to what you can expect in the exam.*

*Hint: six statements are correct, four are incorrect.*

### Exercise 2: Loop parallelization

In `loops.cpp`, you'll find some loops that process arrays. Analyze these loops and decide whether they can be parallelized (without synchronization). If necessary, you can also reorganize the loops to eliminate dependences.

Test your analysis by parallelizing the appropriate loops using the OpenMP directive `parallel for` (do not use any other directives). If you compile the program (using `make`) and run it, the code in `main.cpp` checks if your parallel code still calculates exactly the same result, and prints an error message if not. The code in `main.cpp` must not be changed!

### Exercise 3 (Compulsory Exercise for 5 LP and 6 LP): Parallelization of a simple optimization code with OpenMP

In this exercise, you have to parallelize a simple optimization code. We assume that there is a task (e.g., assembling a car) that can be solved in several different ways (configurations). For each configuration, we compute the cost (e.g., by simulating the assembly process), and then calculate the minimum cost of all examined configurations. For simplicity, the code does not determine which configuration actually leads to the minimal cost.

Obviously, the computation of the cost can be carried out independently (and thus, concurrently) for each configuration, so the problem is easy to parallelize.

a) In `optimize1.cpp`, you find a simple program that just computes the minimum cost with four possible configurations. Since there is no loop in this program, you could use either

- the `single` directive (see Sect. 2.7.5 of the lecture),
- the `sections` directive (see Sect. 2.8.1), or
- the `task` directive (see Sect. 2.8.2)

to execute the four calls to `computeCostForConfig...()` in parallel.

Parallelize the program using one of the approaches given above and determine the speed-up. In addition, think about the difference between these approaches, especially, how the computational work is distributed to the existing threads. Shortly describe these differences in a comment at the beginning of your code.

b) In `optimize2.cpp`, a loop is used to find the minimum cost for 100 different configurations. Parallelize this loop and determine the speed-up.

Note that the execution time of the function `computeCost()` varies for different configurations (i.e., it depends on the function's parameter), so use a suitable strategy for scheduling the loop iterations to threads (see Sect. 2.5.1).

Do not modify any code in the file `functions.cpp`!

#### Exercise 4 (Compulsory Exercise for 6 LP): Parallelization of a numerical integration using OpenMP

The code in `integrate.cpp` calculates the integral of a given function  $f(x)$  between 0 and 1 using the center rule. The function  $f(x)$  computes the expression  $4/(1+x^2)$ , so the integral has exactly the value  $\pi$ . The file contains the integration code four times: one version (implemented in the function `Serial_Integration()`) should remain sequential for comparison, the other three (in the functions `Parallel_Integration1()` to `Parallel_Integration3()`) you should parallelize as follows:

1. using an OpenMP reduction,
2. using the OpenMP `ordered` directive,
3. by splitting the loop into a parallel and a sequential part.

The `main()` routine calls all four functions and prints their result, the error, the runtime, and the speedup, if applicable. The `makefile` contains the necessary commands to compile and link the program (you only have to invoke the command `make`).

Measure how much time each of your functions requires. Try different values for the number of threads and intervals (recommended starting value: 100,000) and compare the times for the sequential and parallel implementations. Interpret your results.

In the second variant (with `ordered` directive), consider how you can improve the runtime by a suitable scheduling of the loop!

The first parallel version (with the OpenMP reduction) does not compute exactly the same result as the sequential code. Why not? For versions 2 and 3, make sure that the result matches exactly with the sequential version!

#### Exercise 5 (Compulsory Exercise for 5 LP and 6 LP): Parallelization of the Jacobi method using OpenMP

In this exercise we consider the Jacobi method for the iterative solution of a boundary value problem. In Jacobi iteration, a new matrix of values is computed from a given one by assigning to each inner element of the new matrix the mean value of the four neighboring elements of the old matrix. The iteration is repeated until the maximum change of an element is below a predefined limit.

Parallelize the sequential code for the Jacobi iteration in the file `solver-jacobi.cpp` using OpenMP directives. Use an OpenMP reduction to compute the maximum change in the array (variable `diff`).<sup>1</sup>

Test the program with different matrix sizes and thread counts. At the end, the program prints the values of several matrix elements. Note that in your parallel version these results must be **exactly** the same (that is, in **all** decimal places) as in the sequential program!

For further verification, you can also view the output (in the file `Matrix.txt`) graphically, using the Java program `ViewMatrix`. This output is, however, only created with a matrix size up to 1000.

Conduct a performance analysis (if possible, using Scalasca) and try to optimize your program as far as possible, e.g., by eliminating barriers. Determine the achievable speedup for different numbers of threads.

## Exercise 6: Parallelization of the Gauss/Seidel method using OpenMP

In this exercise, the Gauss/Seidel method is to be parallelized. In contrast to the Jacobi method, the Gauss/Seidel method uses only one matrix, in which each element is replaced by the mean value of its neighbor elements. The value of the left and the upper neighbor already originates from the current iteration, so that data dependences result.

In this exercise you should parallelize the method by restructuring the loops, such that the matrix is traversed diagonally (see Section 2.6 of the lecture slides).

Parallelize the sequential code of the Gauss/Seidel relaxation in the file `solver-gauss.cpp` using OpenMP directives. The code given for this exercise differs from Exercise 5 only by the different implementation of the function `solver()`. For simplicity, this solver does not iterate until the maximum change is below a threshold, but instead estimates the necessary number of iterations in advance.

For the parallelization, also pay attention to the notes in Exercise 5.

## Exercise 7 (For Motivated Students): Pipelined parallelization of the Gauss/Seidel method using OpenMP

If the number of iterations is known in advance, the Gauss/Seidel method can also be parallelized in a pipeline-like manner (see Section 2.6 of the lecture slides). For the parallel execution of the  $i$  loop, a synchronization must be provided which ensures the following conditions (induced by the data dependences): before the  $i$ -th row is computed in iteration  $k$ ,

1. the  $(i-1)$ -th row in iteration  $k$  and
2. the  $(i+1)$ -th row in iteration  $k-1$

must have been computed. Since OpenMP version 4.5, the `ordered` directive supports a `depend` clause which enables to specify a synchronisation safeguarding the specified dependencies (see the note at the end of section 2.7.6. “Execution in program order” of the lecture slides).

Parallelize the Gauss/Seidel method using the above synchronization. Compare the achievable speedup with the parallelization from Exercise 6. Conduct a performance analysis (if possible, using Scalasca) and try to optimize your program further.

---

<sup>1</sup>Reductions of this form are directly supported by OpenMP only as of version 4.0. Therefore, if necessary, use the method shown in the lecture slides at the end of Section 2.7.3 (“Reduction operations”).