

# Parallel Computing

## Exercise 5 (June 20)

Wolfgang Schreiner  
Wolfgang.Schreiner@risc.jku.at

The result is to be submitted by the deadline stated above via the Moodle interface. If the assignment has been elaborated in a collaboration of two students, only one of them shall upload the assignment (indicating of course on the cover page the collaboration partner).

The submitted result is as a .zip or .tgz file which contains

- a PDF file with
  - a cover page with the title of the course, your name(s), Matrikelnummer(s), and email-address(es),
  - the source code of the sequential program,
  - the demonstration of a sample solution of the program,
  - the source code of the parallel program,
  - the demonstration of a sample solution of the program,
  - a benchmark of the sequential and of the parallel program.
- the source (.c/.java) files of the sequential and of the parallel program.

## MPI Programming

The goal of this exercise is to develop in MPI a distributed memory solution to one of the problems specified in Exercise 1, Gaussian Elimination or Traveling Salesman. As the base of your parallel programming effort, you may use the sequential program you have developed in Exercise 1; you may also write a new sequential program or ask one of your colleagues for one. Our default assumption is that the programs for this assignment are written in C, using the official MPI binding for the parallel solution.

Having said this, you may also write your sequential program in Java and use for the parallelization one of the MPI bindings for Java provided by

- OpenMPI: <https://www.open-mpi.org/faq/?category=java>
- FastMPJ: <http://fastmpj.com>
- MPJ Express: <http://mpj-express.org>

OpenMPI is available at the course machine (see `module avail`), without guarantee of a functional Java interface; newer versions of this package respectively the other packages are to be installed on your own.

However, neither do we recommend to solve this assignment in Java nor will we be able or willing to give any support for the use of Java with MPI.

### Alternative A: Gaussian Elimination

For this alternative, you only need to implement a solution with floating point numbers as the element type (type `long double`). Furthermore, you may assume that the number  $P$  of processes divides the matrix dimension  $N$  exactly.

- The program starts by distributing the system  $A, b$  row-wise among the  $P$  processes in a round-robin fashion (i.e. process 0 receives rows  $0, P, 2P, \dots$ , process 1, receives rows  $1, P+1, 2P+1, \dots$ , and so on). By this distribution, we ensure that the workload is evenly shared in the later phases of the triangulation (when the non-zero part of  $A$  becomes small).

To distribute  $A$ , process 0 constructs a correspondingly permuted version  $A', b'$  of the system to scatter the values among all processes (by a single call of `MPI_Scatter`).

- For performing the triangulation, the program runs in  $N$  iterations, where in iteration  $i$  process  $p = i\%P$  broadcasts row  $i$  to all other processes (`MPI_Bcast`). Each process then uses this row to update all the rows of the system for which it is responsible.

To simplify the program, you may assume that  $A(i, i)$  is different from 0 (if this should not be the case, you may abort the computation).

- For performing the back-substitution, the program runs in  $N$  iterations where in each iteration the process  $p = N\%i$  that holds the newly computed result  $x[i]$  broadcasts this value to all other processes (`MPI_Bcast`). Each process then uses this value to remove one unknown from all the rows of the system for which it is responsible.

- Finally, since process 0 has received all values that were broadcast during back-substitution, it can determine the result  $x$ .

### **Alternative B: Traveling Salesman**

For this alternative, implement a version where the root serves as a manager of the tasks to be distributed among additional  $P$  worker processes:

- The root manages the global pool of paths from which each worker may query new work, i.e., the worker receives a partial path which it has to extend in all possible ways for a potential optimum solution using a local pool of paths.
- Whenever a worker queries the manager for a partial path, the manager provides as part of its answer the length of the shortest cycle found so far which the worker may use to prune its search.
- Whenever a worker finds a potentially shorter cycle, it informs the manager about this cycle and receives as an answer the length of the current shortest cycle which it may subsequently use for pruning.
- If the manager has only one partial path left in its pool, it extends this path in all possible ways such that subsequent queries by multiple workers may be addressed.
- If the manager runs out of work; it informs each client upon its next request about this fact, such that the clients may subsequently terminate; when all clients have been informed, the manager may terminate.

### **Both Alternatives: Benchmarking**

Benchmark the programs (both the sequential and the parallel one) as in Exercise 1; you may use the MPI function `double MPI_Wtime()` which returns the current wall clock time in seconds. Report the results as in Exercise 1.