

Your Special Challenge

Develop a parallel sorting application with C and MPI

Sorting is a fundamental problem in computer science. Unfortunately, for most serial sorting algorithms, it is challenging to implement versions that benefit from the potential speedup of parallel computers with distributed memory.

The odd-even transposition sort is one example of a sorting algorithm that is useful for parallel systems. The serial version of the odd-even transposition sort algorithm works this way:

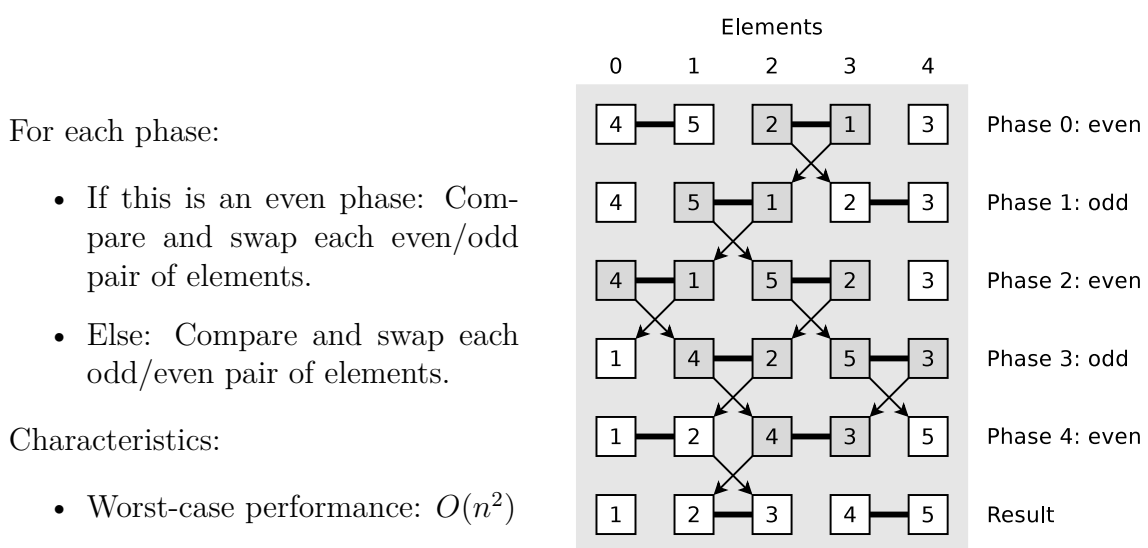


Figure 1: Odd-even sort example [1]

The individual steps of your task

1. **Investigate** how to implement with the programming language C and with MPI an application, that does the Odd-even transposition sort in parallel. Find out which parts of the sorting process can be parallelized. Literature provides many helpful sources, covering the Odd-even transposition sort algorithm [2][3][4][5] and how to implement it in C with MPI [6][7][8].
2. **Deploy** a MPI environment for testing and development on physical hardware, inside virtual machines [9][10] or inside a public cloud infrastructure service like EC2 [11][12][13].
3. **Develop** your application with C and MPI.

4. **Test** your application on your own MPI cluster and test it on our 128 node cluster with the 512 CPU cores. Test your application...
 - with sufficient large problem sizes (this means you prior need to generate some data sets of several MB in size) and
 - with different numbers of cores (1, 2, 4, 8, 16... 512) andcalculate the speedup. Your outcomes may prove some of the laws and limitations we already discussed during class in slide set 1.
5. **Create** a presentation (max. 30 Minutes) with maximum 15 slides and additionally a live demonstration. Give your presentation during class or during an exercise session.

Some final words to motivate you

In real life, one or two weeks is a typical amount of time to do such a task in parallel to your daily tasks. In this course you have several weeks and you are not alone.

Last but not least, questions of the exam will cover this special task!

References

- [1] T. Bräunl, *Parallele Programmierung: Eine Einführung*. Vieweg, 1993.
- [2] D. Knuth, *The Art of Computer Programming – Volume 3 – Sorting and Searching*. Addison-Wesley, 2nd edition ed., 1998.
- [3] “Odd–even sort.” http://en.wikipedia.org/wiki/Odd-even_sort.
- [4] “Odd-even transposition sort.” <http://www.iti.fh-flensburg.de/lang/algorithmen/sortieren/networks/oetsen.htm>.
- [5] “Parallel Algorithms – sorting.” <http://www.dcc.fc.up.pt/~fds/aulas/PPD/1112/sorting.pdf>.
- [6] “Parallel Sorting.” <http://cs.umw.edu/~finlayson/class/fall14/cpsc425/notes/18-sorting.html>.
- [7] “How does MPI Odd-Even sort work?.” <http://stackoverflow.com/questions/23633916/how-does-mpi-odd-even-sort-work>.
- [8] “MPICH – Odd-Even Transposition Sort.” <http://larryvdh.com/blog/mpich-odd-even-transposition-sort/>.

- [9] “MPI Tutorials 01.” <https://www.youtube.com/watch?v=2rpWEZY0aPo>.
- [10] “Create test MPI cluster using VirtualBox images.” <http://www.zamaudio.com/?p=837>.
- [11] “Launching an Amazon EC2 MPI Cluster.” <http://mpitutorial.com/tutorials/launching-an-amazon-ec2-mpi-cluster/>.
- [12] “Using MPICH in Amazon EC2.” https://wiki.mpich.org/mpich/index.php/Using_MPICH_in_Amazon_EC2.
- [13] “MPI on EC2.” <https://www.youtube.com/watch?v=049CVTGH01k>.