

Parallel Programming 2016, Assignment 2: Parallel Minimum Spanning Tree

Deadline: Friday, December 2 before 23:59 hours.

In this second programming assignment, you will implement a parallel Minimum Spanning Tree (MST) algorithm. Your algorithm should be based on Borůvka's algorithm that was discussed in class. After implementation you will perform a small benchmark using three test graphs.

The assignment has to be completed individually. You are expected to hand in a tarball containing your source code and a report briefly describing your implementation of the algorithm and the benchmark results in PDF format. The assignments can be handed by e-mail to *ppi2016(at)handin(dot)liacs(dot)nl*.

1 Implementation

Your algorithm needs to be implemented in the C/C++ language, using MPI for interprocess communication. The target platform is the DAS-4 cluster. Parallelization is *only* to be done by distributing the work over different nodes. So, per node only one CPU core and one CPU thread will be used. The maximum degree of parallel is limited to 16.

Note that your implementation should also handle disconnected graphs, resulting in multiple spanning trees. Note also that as output you should report the total weight of each minimum spanning tree as well as the edges which belong to the spanning tree.

2 Benchmarking

The benchmark needs to be performed on the DAS-4 system at Leiden University. You should benchmark with 2, 4, 8 and 16 nodes to see how well the performance of your algorithm scales as more compute nodes are added. As test data, we will again use sparse matrices from the University of Florida Sparse Matrix Collection (<http://www.cise.ufl.edu/research/sparse/matrices/>). The following matrices, that represent undirected weighted graphs, were selected:

Belcastro/mouse_gene (29M nnz)
GHS_psdef/ldoor (42M nnz)
Schenk/nlpkkt240 (760M nnz)

During code development you can use the smaller *Gaertner/nopoly* matrix for testing.

Include the obtained benchmark results in your report, together with a brief discussion.

3 Links

- DAS-4 website with information on job submission and starting MPI programs: <http://www.cs.vu.nl/das4/jobs.shtml>.
- Online MPI API documentation: <http://www.open-mpi.org/doc/v1.8/>. The API documentation is also available on DAS-4 in the form of man pages after executing the `module load openmpi/gcc` command.
- MPI tutorial from Lawrence Livermore National Laboratory (LANL): <https://computing.llnl.gov/tutorials/mpi/>.