
Exercise Sheet 5

Submit until Friday, July 8 at 2:00pm

Exercise 1 (6 points)

Implement contraction hierarchies as explained in the lecture. To be able to achieve a reasonable performance, you should, as a minimum, make use of the following heuristics in the precomputation:

- (1) Limit the search space of the Dijkstra computations used to determine whether a shortcut is needed by implementing *Improvement 1* (upper bound on the cost) and *Improvement 3* (upper bound on the search space size) from the lecture;
- (2) Order the nodes by edge difference, and maintain these edge differences using the heuristics explained in the lecture (*lazy*, *neighbour*, *periodic*). You may first try *periodic* only (at reasonable intervals, so that the precomputation takes reasonable time), but if too many shortcuts get inserted, be prepared to also implement *lazy* and *neighbour*.

For query processing, it suffices if you compute the cost of the shortest path, you need not produce the sequence of arcs. (Although you may, of course, if you wish.)

Exercise 2 (3 points)

Compute average query times for the same dataset and number of queries as for the previous exercise sheets. Report both precomputation time and average query times on the Wiki. (There will be a new table, dedicated to this exercise sheet, for this.)

Exercise 3 (1 point)

Don't forget your *feedback-exercise-sheet-5.txt*, we really want to know how you are doing!