

Distributed Systems, Summer Term 2020

Exercise Sheet 12

1 Borůvka's MST Algorithm in the Strongly Sublinear MPC Model

We consider the strongly sublinear memory regime of the MPC model, i.e., given an input graph with n nodes and m edges, the set of edges is uniformly distributed among $O(m/S)$ machines with a memory size $S = n^\alpha$ for an $0 < \alpha < 1$. To simplify the calculations, we will assume that initially, the edges are partitioned among at most $4m/S$ machines.

In class, we sketched how to implement one phase of Borůvka's MST algorithm in $O(1)$ time in this setting. We will now look at this more closely. We first recap what a single phase of the algorithm has to achieve:

1. The nodes are partitioned into the already computed fragments of the MST. Assume that each MST fragment has a unique fragment ID and that at the beginning of the phase, for each edge $\{u, v\} \in E$, the machine that stores the edge $\{u, v\} \in E$ also knows the fragment IDs of the nodes u and v .
2. For each fragment F , we need to find the minimum weight edge that connects a node in F with a node in another fragment. (Recall that we assume that the edge weights are distinct and therefore the minimum weight outgoing edge of a fragment F is unique.)
3. We randomly label each fragment either red or blue (independently with probability $1/2$ each). If fragment F is red and its minimum weight outgoing edge e connects to a blue fragment F' , we add e to the MST. The new fragments after the phase are given by the forest that is induced by the old fragments and the newly added edges.
4. The algorithm needs to compute a new unique fragment ID for each of the new fragments and it needs to make sure that for every edge $\{u, v\} \in E$, the machine that stores the edge $\{u, v\} \in E$ also knows the new fragment IDs of the nodes u and v .

For implementing a phase, we intend to use a structure consisting of $O(\log_S m)$ levels of $O(m/S)$ machines each (for simplicity, assume that we have disjoint sets of machines for the different levels). We number the levels from bottom to top from 1 to $\ell = O(\log_S m)$. Assume that level 1 is exactly given by the $O(m/S)$ machines that initially store the graph. We say that a machine μ on a level i is responsible for a fragment ID x if it participates in the computation for the fragment with ID x . On level 1, a machine μ is responsible for fragment ID x if and only if it initially stores an edge $\{u, v\}$ such that either u or v are in fragment with ID x .

For each fragment ID x and each level i for $1 \leq i < \ell$, each machine μ on level i that is responsible for x has a parent machine w.r.t. fragment ID x on level $i + 1$. On level $i \geq 2$, a machine μ is called responsible for fragment ID x if there exists a machine μ' on level $i - 1$ such that μ is the parent w.r.t. x of μ' . The indegree of a machine μ on level $i \geq 2$ is the number of pairs (μ', x) such that μ is the parent w.r.t. x of machine μ' on level $i - 1$. We make the following assumptions:

1. The indegree of every machine μ on some level $i \geq 2$ is at most S

2. For each fragment ID x , the number of machines responsible for fragment ID x on level $i \geq 1$ is $m/(S/4)^i$ (for simplicity, we ignore rounding non-integer values to integer values). Further, on the top level ℓ , there is exactly one machine that is responsible for each fragment ID x .
- (a) Argue in detail how a single phase of the MST algorithm can be carried out in $O(1)$ rounds in the MPC model if above structure is given.
- (b) As the second part of the question, we discuss how to build the structure on machines described above. Assume that the machines have access to shared randomness. Assume that on each level we use $4m/S$ machines. For each fragment ID x , the set of machines responsible for fragment ID x on level $i \geq 2$ is a uniformly random subset of size $m/(S/4)^i$ of the $4m/S$ machines of level i (on level ℓ , we choose one random machine for each fragment ID x). Such a set can be chosen by using the shared randomness (you do not have to worry about the details of this). Each machine on level $1 \leq i < \ell$ chooses its parents on level $i + 1$ as follows. For each fragment ID x for which a machine μ on level i is responsible, the machine μ chooses the parent μ' w.r.t. x on level $i + 1$ uniformly at random among the machines responsible for x on level $i + 1$. Show that with this randomized construction, for every level i and every machine μ on level i , the expected indegree of μ is at most $S/2$.