Locality Sensitive Distributed Computing

Final Exam, Jan. 2008

Answer five of the following eight questions.

Question 1

Consider an n-vertex network $G = (V, E)$, $V = \{v_1, \ldots, v_n\}$. The *Individual messages (IM)* task requires vertex v_1 to deliver a (distinct) $\log n$-bit message to every other vertex in the network, along some pre-specified shortest route. Prove or disprove each of the following claims regarding the message complexity of the problem.

1. $Comm(IM) = O(nD \log n)$ (or, there exists a constant $c > 0$ such that for every network G as above, $Comm(IM, G) \leq cnD \log n$).
2. $Comm(IM) = \Omega(nD \log n)$ (or, there exists a constant $c > 0$ such that for every $n \geq 1$ there exists an n-vertex network G as above for which $Comm(IM, G) \geq cnD \log n$).
3. There exists a constant $c > 0$ such that for every network G as above, $Comm(IM, G) \geq cnD \log n$.

Question 2

1. Give a distributed algorithm for counting the number of nodes in a rooted tree T, initiated at the root.
2. Extend your algorithm to an arbitrary graph G.
3. Give a distributed algorithm for counting the number of nodes in each layer of the rooted tree T separately.

Analyze the time and message complexities of your algorithms.

Question 3

1. Describe a distributed synchronous algorithm based on DFS for counting the number of vertices in the network $G = (V, E)$ and informing the outcome to all the vertices of G. The algorithm should function correctly even when invoked by a number of initiators.
2. What are the time and message complexities of your algorithm assuming it was invoked by K initiators?
3. Does your algorithm work in the asynchronous model?

Question 4

Describe an optimal routing scheme RS (with $Stretch(RS) = 1$ and $Memory(RS) = O(\log n)$ per node) for the unweighted $\sqrt{n}$ by $\sqrt{n}$ 2-dimensional n-vertex grid and for the unweighted d-dimensional n-vertex hypercube (for $n = 2^d$).

Question 5

Modify Algorithm *BasicPart* so that in addition to the constructed partition, it also selects an edge set $\check{E}$ as follows. Initially set $\check{E}$ to $\emptyset$. Whenever completing the construction of a cluster S, for every neighboring node $v \in \Gamma^v(S)$, select one edge connecting v to some neighbor w in S, and place it in $\check{E}$.

1. Prove that the resulting set $\check{E}$ is of cardinality at most $n^{1+1/\kappa}$.
2. Suppose that we took into the set $\check{E}$ *every* edge $e = (v, w) \in E$ connecting nodes $v \in S$ and $w \notin S$. Explain why the bound on the cardinality of $\check{E}$ no longer holds.

Question 6

Consider an execution of Procedure *Part*. For each iteration i of the main loop, let Z_i be the final cluster generated by the procedure, and let Y_i be its kernel, taken into the output cover $\mathcal{DT} = \{Y_1, \ldots, Y_q\}$. Prove the following two properties.

1. There are no two iterations i, j such that both $Y_i \cap Z_j \neq \emptyset$ and $Y_j \cap Z_i \neq \emptyset$.
2. There is no "cycle of nonempty intersections," namely, there are no t iterations $i_0, i_1, \ldots, i_{t-1}$ such that $Y_{i_j} \cap Z_{i_{j+1 \bmod t}} \neq \emptyset$ for every $0 \leq j \leq t-1$.

Question 7

Prove or disprove the following "weighted analog" of the Basic Spanner Lemma:

> For a weighted graph $G = (V, E, \omega)$,
> the subgraph $G' = (V, E', \omega)$ is a κ-*spanner* of G iff
> for every $e = (u, v) \in E$, $dist_{G'}(u, v) \leq \kappa \cdot \omega(e)$.

Question 8

In the distributed implementation of the basic partition algorithm *BasicPart*,

1. Prove that $Time(ClusterCons) = O(n)$,
2. Improve the bounds on $Time(NextCtr)$ and $Comm(NextCtr)$ or provide an example proving their tightness,
3. Provide an example proving the tightness of the bounds on $Time(RepEdge)$ and $Comm(RepEdge)$.