Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

FS 2011

Prof. R. Wattenhofer
R. Eidenbenz

Principles of Distributed Computing
Unanswered Questions from Question Session

Q: 3rd remark p17: How would Leader Election work if “synchronous start and the non-uniformity assumptions are dropped”?

A: This was shown in the lecture. A node would be waken up at latest, when receiving a message from a neighbor. One possible, very slow algorithm would be that nodes send messages around the ring with their id as long as they didn't see any lower id. Nodes forward messages if the contained id is a leader candidate, i.e., it is smaller than any id seen so far. However, they do not forward them immediately, but messages from a node with id i wait at every hop for 2^i rounds. Thus, the message from node with lowest id will be at least twice as fast as the message from the node with the second smallest id and so forth. This yields a message complexity in $O(n)$.

Q: Proof lemma 9.8: Why is the probability that the long-range contact of w points to x in $\Omega(1/2^{2j+4}\log n)$?

A: From Lemma 9.6 we know that for $\alpha = 2$, the probability that w points to x is in $\Theta(1/(d(w,x)^2 \log n))$. As $d(w,x) < 2^{j+2}$ we immediately get a probability in $\Omega(1/2^{2j+4}\log n)$.

Q: Proof Lemma 9.8: Why is the number of nodes in B_j at least 2^{2j-1} ?

A: B_j is the set of all nodes with distance at most 2^j to the target t. In the worst case, t is in a corner. Ignoring long range contacts, B_j then looks like a square with side length 2^j cut in half along the diagonal. The entire square would contain $(2^j)^2$ nodes. We have only half the square. Hence $(2^j)^2/2 = 2^{2j-1}$.

Q: Ex 9.1b): For the sentence " The linearity of the expectation value gives us $E[X]$.......". Here, is the meaning of X the same with that of a)?

A: Not exactly. X is the random variable indicating the number of “good” links out of S. A link is *good* if it exclusively contributes 5 new nodes to the set we are looking at, i.e., if the link ends in x neither x nor its 4 grid neighbors should be in S, nor should x and its neighbors overlap with any of the 5 nodes contributed by any other link out of S.

Q: Ex 91.b): The sample solution proves that $Pr[X <= (1-\delta)E[X]] \leq 1/n^c$ and concludes that the number of good choice will be in $(1-o(1))|S|$ w.h.p. But w.h.p. means some event occurs with probability $p >= 1 - 1/n^c$. Thus I think only the format such as $Pr[....] >= 1 - 1/n^c$ can get the conclusion. In addition, what does “good” choice mean?

A: If we show that $Pr[X \leq k] \leq p$ then we automatically show that $Pr[X > k] > 1 - p$.
For the second part, see answer above.

Q: Reading Assignment, p21, Lemma 7: “...since the nodes of cluster C_0 are not necessary to obtain a feasible vertex cover...”. Why are these nodes not necessary?

A: The construction of the lower bound graph is such that the optimal solution would add all nodes in the neighboring clusters of C_0 to the vertex set. Thus all edges into C_0 are already covered.

Q: Reading Assignment, p21 Lemma 7: Is this δ_i the same δ_i as on the labels of the arcs of the Cluster tree? What do the δ_i mean? What is the relationship between δ_i and δ_{i+1}?

A: Yes, the δ_i correspond to the arc labels. If you read the caption of Fig. 2 on page 15 you'll see the function of δ_i. They count the number of links. On page 21, equation (7), the δ_i are explicitly defined based on a single parameter δ. The relation of δ_{i+1} to δ_i is stated immediately after: $\delta_{i+1} = 2^i\delta \cdot \delta_i$.