



Principles of Distributed Computing

Exercise 9: Sample Solution

1 Diameter of the Augmented Grid

- a) As our target node set is of size $\Omega(\log^2 n)$ and link targets are distributed uniformly at random over all n nodes, each link connects to the target set with probability $p \in \Omega((\log^2 n)/n)$. Thus, for sufficiently large¹ n , the probability that $n/\log n$ many links miss the set is bounded by

$$(1 - p)^{n/\log n} \leq e^{-pn/\log n} \in e^{-\Omega(\log n)} = \frac{1}{n^{\Omega(1)}}.$$

Now we exploit the power of the Big- O notation. Choosing a sufficiently large multiplicative constant in front of the $(n/\log n)$ -term, this becomes a bound of $1/n^c$, and choosing a large additive constant, we make sure that the bound holds also for the values of n that are not “sufficiently large”. Thus, the probability that at least one link enters the set of $\Omega(\log^2 n)$ nodes is at least $1 - 1/n^c$, i.e., this event occurs w.h.p..

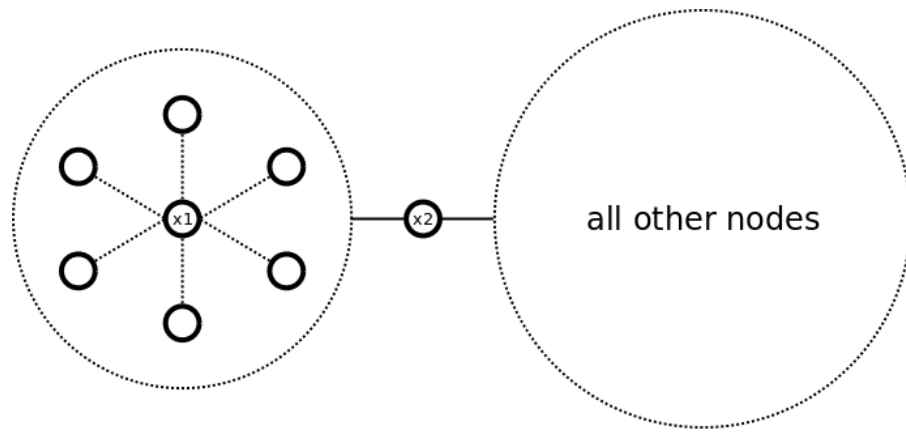
In order to obtain the same result using a Chernoff bound, let X_i , $i \in \{1, \dots, l\}$, where $l \in O(n/\log n)$ is the number of considered links, be random variables that are 1 if the i^{th} link ends in the set (i.e., with the probability p from above) and 0 otherwise. Defining $X := \sum_{i=1}^l X_i$, we get that $\mathbb{E}[X] = pl$. Picking a constant $C > 0$ and properly adapting the constants in the $O(n/\log n)$ -term, we get that $\mathbb{E}[X] \geq C \log n$. Thus, the Chernoff bound states that for sufficiently large values of C and n (we need to cope with the fact that we do not know the constants in the O -term in the Chernoff bound), we have that $|X - \mathbb{E}[X]| < C \log n$ w.h.p. Because $P[X > 0] \geq P[|X - \mathbb{E}[X]| < \mathbb{E}[X]]$, we know that $X > 0$ also holds w.h.p., which means that at least one link points to our target set with high probability.²

- b) Because $|S| \in o(n)$, also $O(|S|) \subset o(n)$, i.e., the union of the set S ($|S|$ nodes) with the destinations of the $|S|$ random links and all grid neighbors of such nodes (at most $5|S|$ many nodes) has $o(n)$ nodes (because $O(|S|) \subset o(n)$). Thus, always $n - o(n) = (1 - o(1))n$ nodes can be found which neither have been visited themselves nor have any neighbors that have been visited so far. Hence, regardless of the choice of the set S and any random links leaving S we have (sequentially) examined up to now, any uniformly independent random choice will contribute 5 new nodes with some probability $p \in 1 - o(1)$. The linearity of the expectation value gives us $\mathbb{E}[X] \in (1 - o(1))|S|$. Now we use a Chernoff bound (Bound 2) on the number of such “good” choices and set δ to $\frac{1}{\sqrt{\log n}}$:

$$\Pr[X \leq (1 - \delta)\mathbb{E}[X]] \leq e^{-\mathbb{E}[X]\delta^2/2} \leq e^{-\Omega(\log n)} \leq \frac{1}{n^c}$$

¹This phrase means for some constant n_0 , the statement will hold for all $n \geq n_0$.

²Small values of n are again dealt with by the additive constant in the O -notation. In general, it is always feasible to assume that n is “sufficiently large” when proving asymptotic statements.



In a graph with n nodes and an approximate power law node-degree distribution, the highest degree of a node can be less than $\frac{n}{2}$. In the depicted network node x_1 is the node with the highest degree and all nodes adjacent to x_1 are arranged in a star. This star is connected through only one node to the rest of the graph. This means that node x_2 is closer to more than $\frac{n}{2}$ of the nodes.