ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

FS 2010 Prof. R. Wattenhofer / C. Lenzen

Principles of Distributed Computing
Exercise 10

For this exercise, you may use the following *Chernoff bound*.[1]

Theorem 1 (Chernoff's Bound)
Let $X := \sum_{i=1}^{N} X_i$ *be the sum of* N *independent* $0-1$ *random variables* X_i*. Then* with high probability[2] *(w.h.p.)*

$$|X - \mathbb{E}[X]| \in O\left(\log n + \sqrt{\mathbb{E}[X]\log n}\right).$$

1 Diameter of the Augmented Grid

Recall the network from the lecture where nodes were arranged in a grid and each node had an additional directed link to an uniformly and independently at random drawn node in the network (i.e., $\alpha = 0$). In the lecture, a proof of the fact that such a network has diameter $O(\log n)$ w.h.p. was sketched. We will now fill in the details.

a) Show that $O(cn/\log n)$ many nodes are enough to guarantee with high probability that at least one of their random links connects to a given set of $\Omega(\log^2 n)$ nodes. Prove this (i) by direct calculation and (ii) using Chernoff's bound.

Hint: Use that $1 - p \leq e^{-p}$ for any p.

b) Suppose for some node set S we have that $|S| \in \Omega((\log n)^2) \cap o(n)$ and denote by H the set of nodes hit by their random links. Prove that H and together with its grid neighbors contains w.h.p. $(5 - o(1))|S|$ nodes!

Hint: Observe that *independently* of all previous random choices, each new link has at least a certain probability p of connecting to a node whose complete neighborhood has not been reached yet. Then use Chernoff's bound on the sum of $|S|$ many variables.

c) Infer from **b)** that starting from $\Omega(\log^2 n)$ nodes, with each hop the number of reached nodes w.h.p. more than doubles, as long as we have still $O(n/\log n)$ nodes (regardless of the constants in the O-notation).

Hint: Play with the constant c in the definition of w.h.p. and use the union bound.

d) Conclude that the diameter of the network is w.h.p. in $O(\log n)$.

[1] Chernoff-type and similar probability bounds are very powerful tools that allowed to design a plethora of randomized algorithms that *almost* guarantee success. Frequently this "almost" makes a huge difference in e.g. running time and/or approximation quality.

[2] I.e., with probability at least $1 - 1/n^c$ for any fixed constant $c > 0$. The "fixed" here indicates that we typically see c as a constant with regard to O-notation. Thus, if someone names a constant c, then for instance the statement "the algorithm requires $O(\log n)$ rounds with high probability" translates to "the algorithm requires $O(\log n)$ rounds with probability at least $1 - 1/n^c$", where the $O(\log n)$ could mean e.g. $5c^2 \log n + e^{c^2} + 1001$.

2 Greedy Routing in the Augmented Grid

Consider now $\alpha = 2$, i.e., the random link of node u connects it to node w with probability $d(u,w)^{-2}/\sum_{v \in V \setminus \{u\}} d(u,v)^{-2}$. In the lecture, we saw that with probability $\Omega(1/\log n)$, in each step we get to the next phase when we employ greedy routing. Hence, the expected number of steps is in $O(\log^2 n)$. Prove that the same bound on the number of steps holds w.h.p.!