

Network dynamics

The goal of this tutorial is to review the material on network dynamics and more particularly the notions of power law distribution and decentralised search.

Exercise 1: Power law

Consider an on-line news site, such as cnn.com or nytimes.com, which consists of a front page with links to many different articles. The people who operate such sites generally track the popularity of the various articles that get posted, asking questions like the ones that we've seen: "As a function of k , what fraction of all articles have been viewed by k people?" Let's call this the popularity distribution of the articles.

Now suppose that the operators of such a news site are considering changing the front page, so that next to each link is a counter showing how many people have clicked on the link. (E.g., next to each link it will say something like, "30,480 people have viewed this story," with the number getting updated over time.) First, what effect do you think this change will have on the behavior of people using the site? Second, do you expect that adding this feature will cause the popularity distribution of the articles to follow a power-law distribution more closely or less closely compared to the version of the site before these counters were added? Give an explanation for your answer.

Duration: 15 min

Exercise 2: P2P Overlays

In Figure 1, to implement the Chord DHT how many finger table entries should each node have? Draw the outgoing "finger" edges of nodes 5 and 14.

What "finger" edges of the overlay would be traversed during a lookup request issued by node 1 and targeting key 16?

How many messages maximum does a lookup request trigger?

Duration: 15 min

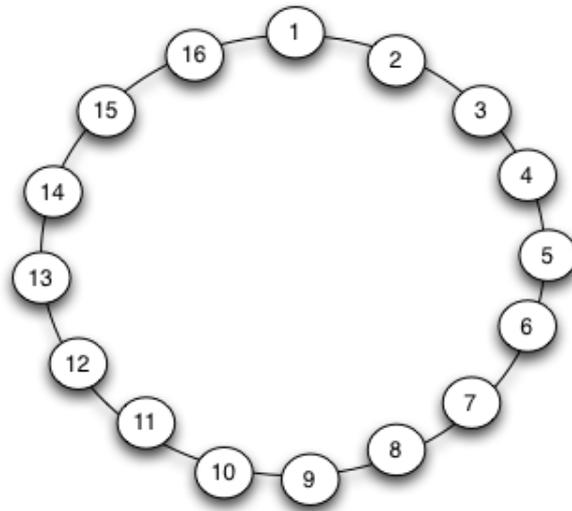


Figure 1: Peers organized in a ring overlay

Exercise 3: Small world

In the basic “six degrees of separation” question, one asks whether most pairs of people in the world are connected by a path of at most six edges in the social network, where an edge joins any two people who know each other on a first-name basis. Now let’s consider a variation on this question. For each person in the world, we ask them to rank the 30 people they know best, in descending order of how well they know them. (Let’s suppose for purposes of this question that each person is able to think of 30 people to list.) We then construct two different social networks:

- The “close-friend” network: from each person we create a directed edge only to their ten closest friends on the list.
- The “distant-friend” network: from each person we create a directed edge only to the ten people listed in positions 21 through 30 on their list.

In the close-friend network, is it possible that for each pair of people in the world, there is a path of at most six edges? Let’s think about how the small-world phenomenon might differ in these two networks. In particular, let C be the average number of people that a person can reach in six steps in the close-friend network, and let D be the average number of people that a person can reach in six steps in the distant-friend network (taking the average over all people in the world). When researchers have done empirical studies to compare these two types of networks (the exact details often differ from one study to another), they tend to find that one of C or D is consistently larger than the other. Which of the two quantities, C or D , do you expect to be larger? Give a brief explanation for your answer.

Duration: 20 min