Lecture 26: Small-World Peer-to-Peer Networks and Their Security Issues

Lecture Notes on "Computer and Network Security"

by Avi Kak (kak@purdue.edu)

April 27, 2017
11:42am

Goals:

1. Differences Between Structured P2P and Small-World P2P
2. Freenet as Originally Envisioned by Ian Clarke
3. The Small-World Phenomenon
4. Demonstration of the Small-World Phenomenon by Computer Simulation
5. Decentralized Routing in Small-World Networks
6. Small-World Based Examination of the Original Freenet
7. Sandberg's Decentralized Routing Algorithm for Freenet
8. Security Issues with the Freenet Routing Protocol
9. Gossiping in Small-World Networks

CONTENTS

	Section Title	*Page*
26.1	**Differences Between Structured P2P and Small-World P2P**	3
26.2	**Freenet as Originally Envisioned by Ian Clarke**	6
26.3	**The Small-World Phenomenon**	15
26.4	**Demonstration of the Small-World Phenomenon by Computer Simulation**	19
26.5	**Decentralized Routing in Small-World Networks**	41
26.6	**Small-World Based Examination of the Original Conceptualization of Freenet**	48
26.7	**Sandberg's Decentralized Routing Algorithm for Freenet**	50
26.8	**Security Issues with the Freenet Routing Protocol**	68
26.9	**Gossiping in Small-World Networks**	71
26.10	**For Further Reading**	76

26.1: DIFFERENCES BETWEEN STRUCTURED P2P AND SMALL-WORLD P2P

- First of all, both structured and small-world P2P networks are most commonly overlaid on top of the internet. So we can refer to them as **structured P2P overlays** and **small-world P2P overlays**.

- As we saw in Lecture 25, structured P2P overlays place topological constraints on what other nodes any given node is aware of for the purpose of data lookup or data retrieval. In a structured P2P overlay, a more-or-less uniformly distributed integer, $nodeID$, is assigned to each node. In the Chord protocol, for example, a node is aware of its immediate successor, which would be a node with the next larger value for $nodeID$. And, through its routing table, a node is also aware of a small number of additional nodes up ahead whose $nodeID$ values sample the node identifier space logarithmically.

- Structured P2P overlays of Lecture 25 are founded on the assumption that any node *can* exchange data with any other node

in the underlying network (meaning the internet). [Say that A and B are nodes in a structured P2P overlay. Let's say that at a given moment in time, B is not A's immediate neighbor in the P2P overlay and that B does not make any appearance at all in A's routing table. So A is not likely to forward its queries to B at this moment. But, after the addition of a few other nodes or departures thereof, it is entirely possible that B could become A's immediate successor (or predecessor) and/or that B would make an appearance in A's routing table. Should that happen, there would need to be a direct communication link in the underlying internet between A and B.]

- In small-world P2P overlays, on the other hand, it is the human owner of a node who decides which other nodes his/her node will communicate with directly. This feature of small-world P2P overlays could be used by a bunch of people to create their own private overlay network that would be invisible to the rest of the internet. **Such closed overlays are called** ***darknets*****.**

- In this lecture we will assume that it is NOT our intent to create a *closed* private overlay with a small-world P2P. Without requiring approval from all of the current participants, we want a human to be able to have his friends connect with his node — in the same manner that humans form and extend their friendships. In other words, in this lecture, we are interested in **open-ended small-world P2P overlays.**

- Such open-ended small-world P2P networks are also referred to as **unstructured** P2P networks.

- Considering the *ad hoc* nature of the connections in unstructured network overlays, we are interested in studying how messages are routed in such overlays and whether there exist any security problems with a given routing strategy.

- The best example of a small-world (unstructured) P2P overlay today is the Freenet that was proposed initially by Ian Clarke in a dissertation at the University of Edinburgh in 1999. Clarke's main focus was on creating a distributed system for key-indexed storage from where individuals could retrieve information while remaining anonymous. [As mentioned in Lecture 25, the system of web pages is an example of key-indexed storage in which the URLs are the keys and, for each key, the web page at that URL the corresponding value or data.] In other words, Clarke was interested in creating a "decentralized information distribution system" that would provide anonymity to both the providers and the consumers of information. [In Clarke's thinking, the regular internet is a highly centralized information system in which the routing is orchestrated by the DNS that directs an information consumer's query to the web pages of the information providers who stay at fixed locations. According to Clarke, the regular internet makes it all too easy to keep track of the information providers and and the information consumers.]

- The next section explains Clarke's original idea for the Freenet in greater detail.