582615 - Overlay and P2P Networks, Spring 2015

Exercise package 3 (20 points)

The exercises are intended to be done working in pairs. This package contains four exercises and an optional turbo challenge. During the course there will be three sets of exercises. The course book and the lectures contain some answers, but searching for outside sources too is strongly encouraged.

Schedule

There are two types of exercise sessions: Clarification sessions, where you can ask questions about the exercises or other matters about the course; and Answer sessions, where some answers to the returned exercises are presented and discussed.

- Clarification session: Tuesday 17.2. at 14:15
- **Exercise deadline**: Monday 23.2. at 12:00
- Answer session: Tuesday 24.2. at 14:15

Submission

Return your answers by email to juhani.toivonen@cs.helsinki.fi as an attached PDF or TXT document. Use “Overlay Exercise 3” as the subject line. The document should include:

- The title “Overlay exercise package 3”
- The name and student number of both writers
- The answers to the exercises

Assignments

Assignment 1 - Content Delivery Networks (4 points)

Content Delivery Networks (CDN) such as Akamai are powerful tools for distributing content from a publisher to users.

- What kinds of problems can be solved by using a CDN?
- Describe mechanisms used by CDNs to direct users to the content?
- Coral is an open CDN system. Describe how Coral delivers content to users.

References:

- https://www.usenix.org/legacy/events/nsdi04/tech/full_papers/freedman/freedman_html/

Assignment 2 - Kademlia (5 points)

Kademlia is a DHT algorithm which uses *bitwise exclusive or* (XOR) operation as a metric for node distance.

- What are the trade-offs in selecting k-bucket sizes for a Kademlia overlay?
- How much memory does the Kademlia routing table consume?
- How does Kademlia compare to other DHTs?

Kademlia is described in the article *"Kademlia: A Peer-to-peer Information System Based on the XOR Metric"* by Petar Maymounkov and David Mazières. It was published in the *1st International Workshop on Peer-to-peer Systems (IPTPS)* in 2002.

Reference: http://pdos.csail.mit.edu/~petar/papers/maymounkov-kademlia-lncs.pdf.

Assignment 3 - Dynamo (6 points)

Don't mix with DynamoDB; Dynamo is something different.

Dynamo is a cloud storage system used by Amazon. One of the use cases for Dynamo is to power Amazon's online shopping experiences. Dynamo is said to be a zero-hop DHT as all nodes contain enough information for routing directly to the content. Cassandra is a high performance distributed data store, an implementation of Dynamo, that is used in the back-end of several popular internet services worldwide.

- How does dynamo implement eventual consistency?
- What steps does Dynamo take to provide consistent shopping carts?
- How does Cassandra partition the data? What does Cassandra do when it notices that the data is distributed unequally among the nodes?
- What techniques does Cassandra use to achieve scalability and redundancy?

Dynamo is described in the article *"Dynamo: amazon's highly available key-value store"* by Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall and Werner Vogels. It was published at the *ACM Symposium on Operating Systems Principles (SOSP)* in 2007. Information about Cassandra can be found on the Apache Cassandra project home page and in *"Cassandra - A Decentralized Structured Storage System"* by Avinash Lakshman and Prashant Malik from Facebook.

References:

- http://www.allthingsdistributed.com/files/amazon-dynamo-sosp2007.pdf.
- http://cassandra.apache.org
- http://www.cs.cornell.edu/projects/ladis2009/papers/lakshman-ladis2009.pdf

Assignment 4 - Future of overlays and P2P (5 points)

Cloud computing pertains to the virtualization of hardware and software resources that they can be provided in an elastic and on-demand manner. Typically cloud platforms reside in large, centralized data centers. Some of them are federated over to multiple data centers in different geographical areas. To take the idea one step further, P2P and overlay technologies have been proposed to build a cloud platform on off-the-shelf desktop computers scattered around the office.

Software Defined Networking (SDN) introduces programmability into the network. It is based on the idea of separating the control- and data planes. The control plane is typically implemented as an overlay network that interacts with the network-level routers and switches.

- How can P2P help cloud computing?
- How can SDN and overlays be integrated and what benefits can be obtained from their integration?
- How could a P2P cloud benefit from SDN-compatible networking?

References:

- History of SDN: http://www.cs.princeton.edu/courses/archive/fall13/cos597E/papers/sdnhistory.pdf
- P2P Cloud Prototype: http://www.cs.unibo.it/babaoglu/papers/pdf/acm-sac-2012.pdf

Turbo challenge (optional) - Programming assignment

The turbo challenge allows you to recover lost points from other assignments, but will not increase the maximum points available. You can get full points from the exercise set without the turbo challenge.

Implement a BitTorrent tracker. The tracker does not need to seed the torrent itself, and it does not need to create torrent files, but it must track peers and announcements. Use of libraries, e.g. for Bencode formatting in the messages is allowed. When clients ask for peers, the tracker's response shoud prefer peers from the same Autonomous System (AS) as the requesting node.

A description about the BitTorrent Tracker protocol can be found here:
https://wiki.theory.org/BitTorrent_Tracker_Protocol

You can find out which AS an IP address belongs to by using, e.g., the JSON API from http://ipinfo.io/ (*try: 'curl http://ipinfo.io/8.8.8.8' and 'curl http://ipinfo.io/8.8.8.8/org'*)