



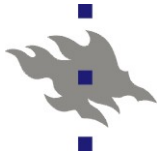
HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Overlay and P2P Networks

Introduction

Prof. Sasu Tarkoma

14.1.2013



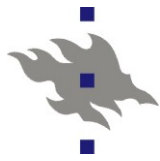
Contents

- Course Overview
- Lectures
- Assignments/Exercises



Course Overview

- Overlay networks and peer-to-peer technologies have become key components for building large scale distributed systems.
- This course will introduce overlay networks and peer-to-peer systems, discuss their general properties, and applications. The course will cover the following topics:
 - Currently deployed peer-to-peer systems and how they work
 - Distributed Hash Tables as a base for structured peer-to-peer systems
 - Peer-to-peer storage systems
 - Performance issues, legal aspects, and privacy issues
 - Peer-to-peer content distribution algorithms



General Info

Advanced course, 4 credits

The course replaces the P2P Networks course

You cannot take this course if you took the old course

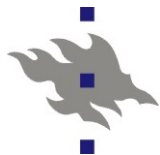
Requirements: basics of networking

Assignments/exercises done as group work (1-2 persons),
idea is to keep the same group structure (but do tell
about freeriders!)



Lectures

- Lectures
 - Monday 14-16 in D122 14.1.-21.2.
 - Thursday 12-14 in D122 14.1-21.2.
- Assignments
 - Friday 14-16 D122 21.09-22.2.
 - **First session on 18.1.**
 - Assignment topic given two weeks before deadline
- Course based on book
 - S. Tarkoma. Overlay Networks: Toward Information Networking. 260 pages. CRC Press / Auerbach, February 2010.



Overlay Networks Book

Introduction

- Overview
- Overlay Technology
- Applications
- Properties of Data
- Structure of the Book

Network Technologies

- Networking
- Firewalls and NATs
- Naming
- Addressing
- Routing
- Multicast
- Network Coordinates
- Network Metrics

Properties of Networks and Data

- Data on the Internet
- Zipf's Law
- Scale-free Networks
- Robustness
- Small Worlds

Unstructured Overlays

- Overview
- Early Systems
- Locating Data
- Napster
- Gnutella
- Skype
- BitTorrent
- Cross-ISP BitTorrent
- Freenet
- Comparison

Foundations of Structured Overlays

- Overview
- Geometries
- Consistent Hashing
- Distributed Data Structures for Clusters

Distributed Hash Tables

- Overview
- APIs
- Plaxton's Algorithm

- Chord
- Pastry
- Koorde
- Tapestry
- Kademlia
- Content Addressable Network
- Viceroy
- Skip Graph
- Comparison

Probabilistic Algorithms

- Overview of Bloom Filters
- Bloom Filters
- Bloom Filters in Distributed Computing
- Gossip Algorithms

Content-based Networking and Publish/Subscribe

- Overview
- DHT-based Data-centric Communications
- Content-based Routing
- Router Configurations
- Siena and Routing Structures
- Hermes
- Formal Specification of Content-based Routing Systems
- Pub/sub Mobility

Security

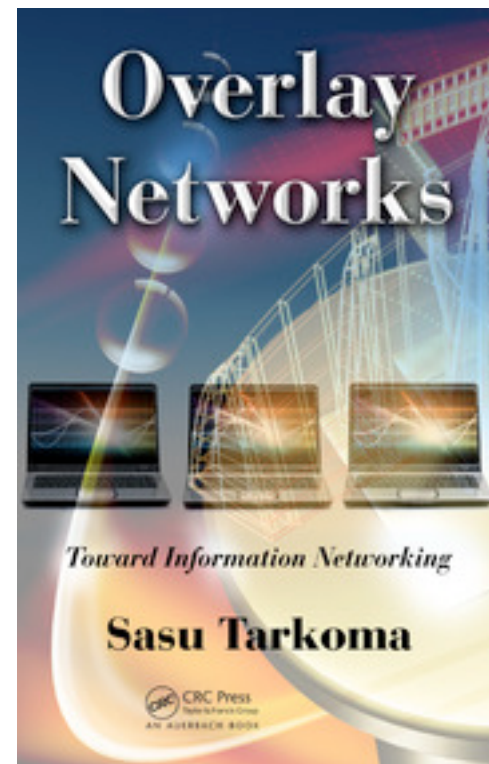
- Overview
- Attacks and Threats
- Securing Data
- Security Issues in P2P Networks
- Anonymous Routing
- Security Issues in Pub/Sub Networks

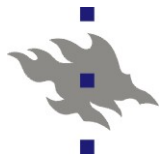
Applications

- Amazon Dynamo
- Overlay Video Delivery
- SIP and P2PSIP
- CDN Solutions

Conclusions

- References
- Index





Lectures

- 14.1. Introduction. Assignments
- 17.1. Unstructured networks I
- 21.1. BitTorrent and unstructured networks II
- 24.1. P2P Modelling (Petri Savolainen)
- 28.1. Power-law networks. Structured networks
- 31.1. Distributed Hash Tables (DHTs) I
- 4.2. DHTs II
- 7.2. Applications I
- 11.2. Applications II (also invited speakers)
- 14.2. Advanced topics
- 18.2. Conclusions and summary



Grading

Course grading will be based on the final exam and the assignments.

Course exam 27.2.2013 16:00 CK112

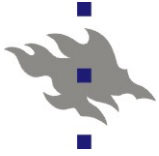
Final exam on 19.4.2013 16:00 in A111



Assignments/Excercises

- Assignments are given two weeks before the session, due date is the day before the assignment session 4pm
- Assignments are done in two-person groups (or alone), groups can change between assignments
- Assignments give bonus points for the exam
 - Max exam points 60, max bonus points 12
- Fridays 14-16 D122 last session 22.2.
 - 18.1. Reception on questions I
 - 25.1. Answers to questions I
 - 1.2. Reception on questions II
 - 8.2. Answers to questions II
 - 15.2. Reception on questions III
 - 22.2. Answers to questions III

Main theme	Prerequisites	Approaches learning goals	Meets learning goals	Deepens learning goals
Overlay and peer-to-peer networks: definitions and systems	Basics of data communications and distributed systems (Introduction to Data Communications, Distributed Systems)	Knowledge of how to define the concepts of overlay and peer-to-peer networks, and state their central features Ability to describe at least one system in detail	Ability of being able to compare different overlay and p2p networks in a qualitative manner Ability to assess the suitability of different systems to different use cases	Ability to give one's own definition of the central concepts and discuss the key design and deployment issues
Distributed hash tables	Basics of data communications and distributed systems (Introduction to Data Communications, Distributed Systems) Big-O-notation and basics of algorithmic complexity	Knowledge of the concepts of structured and unstructured networks and the ability to classify solutions into these two categories Knowledge of the basics of distributed hash tables Ability to describe at least one distributed hash table algorithm in detail	Ability of being able to compare different distributed hash table algorithms Ability of designing distributed hash table-based applications Knowledge of key performance issues of distributed hash table systems and the ability to analyze these systems	The knowledge of choosing a suitable distributed hash table design for a problem Familiarity with the state of the art
Reliability and performance modelling	Basics of probability theory Basics of reliability in distributed systems	Ability to model and assess the reliability of overlay and peer-to-peer networks by using probability theory Knowledge of the most important factors pertaining to reliability	Ability of analytically analyzing the reliability and performance of overlay and peer-to-peer networks Understanding of the design issues that are pertinent for reliable systems	Familiarity with the state of the art
Content distribution	Introduction to Data Communications	Knowledge of the basic content distribution solutions Ability to describe at least one overlay and p2p network based content distribution solution	Knowledge of different content distribution systems and the ability to compare them in detail Knowledge of several content distribution techniques	Familiarity with the state of the art
Security	Basics of computer security	Knowledge of the basic security issues with overlay and p2p networks Knowledge of the sybil attack concept	Ability to discuss how security problems and limitations can be solved Knowledge of how to prevent sybil attacks	Knowledge of how to prevent sybil attacks Familiarity with the state of the art



Contact information

Lecturer prof. Sasu Tarkoma (contact info on homepage)

Assignments: M.Sc Toni Ruottu(@cs.helsinki.fi)

Course homepage can be found: www.cs.helsinki.fi/courses



Questions?