

Exercise package 1 (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: Wednesday 18.1. at 12:15
- **Exercise deadline:** Tuesday 24.1. at 12:00
- Answer session: Wednesday 25.1. at 12:15

Submission

Return your answers by email to juhani.toivonen@cs.helsinki.fi as an attached PDF or TXT document. Use "Overlay Exercise 1" as the subject line. The document should include:

- The title "Overlay exercise package 1"
- The name and student number of the writer/writers
- The answers to the exercises

Assignments

Assignment 1 - Overlays and underlays (4 points)

TCP/IP is usually presented as a stack of layers, each with their own responsibility and assigned protocols. Overlay networks can be presented in a similar way.

- Generally speaking, what is the purpose of an overlay? And what is the purpose of an underlay?
- What kind of layers can be found in an overlay network stack and what responsibilities do they have?
- Do the layers of an overlay stack and those of the TCP/IP stack perform similar functions? What are they? Are they redundant?

Assignment 2 - Utility of overlays (5 points)

Overlay networks present applications with a different perspective to what the network looks like and what services it provides. In a sense they separate the physical and logical topologies of the network from each other, they may sector the network differently, etc.

- What kind of systems can use an overlay? What advantage would they get?
- Does using overlays sacrifice something? What are the trade-offs?
- Why would one use an overlay instead of simply changing the internet to do what they want?

Assignment 3 - Peer-to-Peer (P2P) (5 points)

Many overlay systems are either partially or entirely based on Peer-to-Peer technology.

- How is peer-to-peer different from the traditional client-server model?
- What is the main difference between a structured and unstructured P2P network?
- What strengths and what weaknesses does peer-to-peer have compared to the client-server model?
- What challenges do Network Address Translation (NAT) devices introduce to peer-to-peer? What ways are there to go around these challenges?

Assignment 4 - BitTorrent selection strategies (6 points)

At the heart of BitTorrent's strategic selections there are two key algorithms: the *Choke algorithm*, and the *Rarest first algorithm*. *Choke* is used to decide which of the peers that want to download are allowed to download and when. *Rarest first* is used to decide the order in which to request pieces of files for downloading. The algorithms are chosen and modified to make the network resilient even in the presence of misbehaving nodes. Descriptions for the algorithms can be found e.g. in "*Rarest first and choke algorithms are enough*" published by Arnaud Legout, G. Urvoy-Keller and P. Michiardi at ACM SIGCOMM '2006. (<http://conferences.sigcomm.org/imc/2006/papers/p20-legout.pdf>).

- What does the *Rarest first algorithm* do? How do the three added policies affect its operation?
- How fair is the original *Choke algorithm* towards the different peers? (Ones trying to seed the torrent, downloaders with fast or slow connections..)
- How does the *seed state modification* for the *Choke algorithm*, introduced in mainline BitTorrent 4.0.0 affect the overall performance?

Turbo challenge (optional, up to 4 points)

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. The turbo challenges may step outside the themes discussed on the course.

Traditionally BitTorrent coordinates lists of peers using central servers called Trackers. An alternative approach, an official extension to the BitTorrent protocol, uses something called a *Distributed Hash Table* (DHT) for the same purpose.

- What advantage does BitTorrent get from using a DHT?
- Explain how finding peers using the DHT works.