

CS-438: Decentralized Systems Engineering

Course Introduction and Basic Concepts

Hybrid local/remote teaching logistics

This course is trying the $\frac{1}{3}$ -local-at-once hybrid teaching mode

- Lectures will be both local and [concurrently broadcast live via Zoom](#)
- Local attendees: please join Zoom call but **keep audio muted & video off!**
- All attendees should use Zoom “raise hand” function for synchronous Q&A
- We’ll use [SpeakUp room number 63341](#) for asynchronous online Q&A
- Lectures (including Q&A) will be recorded and made available later

This plan is ***entirely provisional*** - anything may change depending on conditions

- If numbers keep getting worse too quickly, we will likely go online-only!

What is a Decentralized System?

An appealing but simplistic definition is a *system that works although no one's in charge*.

More precise definitions:

- *Distributed system*: a system of multiple computers (nodes) communicating over a network.
- *Decentralized system*: a distributed system in which different nodes or subsets of the network are owned or controlled by different people, organizations, or interests.

Some examples

- Centralized distributed systems
 - Google
 - Yahoo
 - Facebook
- Decentralized distributed systems
 - E-mail
 - UseNet
 - IRC
 - BitTorrent
 - Tor
 - Bitcoin
 - Ethereum

Course Goals

- Understand the fundamental challenges inherent in designing and building decentralized systems.
- Get a feel for the (limited) body of techniques and solutions we currently have for meeting these challenges.
- Examine a number of real, past and present systems, how they work, and how and why they succeeded or failed.
- Solidify this knowledge by applying it to the construction of a small, but working and usable, decentralized system.

Lecture 1: Course introduction

- Course goals
- **Course logistics**
 - Tentative schedule and topic outline
 - Course web sites and tools
 - Programming exercises overview

Tentative Course Syllabus

Week	Content	Week	Content
1	Course introduction, gossip and USENET	8	Consensus basics and algorithms
2	<i>National holiday, no lecture</i>	9	Blockchains and cryptocurrencies
3	Cryptography basics	10	Smart contracts
4	Identity and sybil attacks	11	Privacy and anonymity protection
5	Decentralized search	12	Cryptographic trust-splitting
6	File sharing and content delivery	13	Accountable computing
7	Content security attacks	14	Cloud, big data, and privacy-preserving data processing

Important course sites and tools

- Moodle for CS-438
 - <https://moodle.epfl.ch/course/view.php?id=15483>
 - Announcements and discussion forum
 - Slides, exercises, solutions
 - Project descriptions
- Lecture style
 - Most of the course material will be explained on the blackboard for better interactivity.
 - No slides usually. Instead, take notes, scribble copies of diagrams, etc.

Exercises and grading

- Weekly workload
 - Lectures: 2h - Mondays 10:15am-12:00pm in INF 1
 - Zoom link for remote participation: see Moodle
 - Self-guided work sessions: 2h - Mondays 1:15pm-3:00pm in INF 1
 - Room open to hack, discuss designs/problems, and interop-test with colleagues
 - *No sharing or copying code; everyone implements their own Peerster!*
 - TA-guided Q&A sessions: 2h - Fridays 3:15pm-5:00pm in INJ 218
 - Q&A, help with architecting your system and understanding the requirements
 - *The TAs will not debug your code for you.*
 - Homework: 3h-10h depending on your understanding of the concepts and programming skill
- Grading structure:
 - Labs: 60% of grade
 - Project: 40% of grade

Contact information

Please E-mail the instructors at: cs438@groupe.epfl.ch

Course instructors:

- Prof. Bryan Ford - bryan.ford@epfl.ch - BC210
- Teaching assistants:
 - Cristina Basescu - cristina.basescu@epfl.ch
 - Kirill Nikitin - kirill.nikitin@epfl.ch
 - Pasindu Tennage - pasindu.tennage@epfl.ch

Programming Exercise Structure

- Three exercise sets over the semester
 - Do not expect to do them last minute
 - Do not skip a part. Exercise build on each other so you will lose grades without skipping any amount of work
- TA-guided Q&A sessions
 - Discussion on design choices and how to architect your system
 - The TAs will not directly debug or fix your code for you!
- Main problem-solving and programming to be done “on your own”

End of semester project

- Idea will be developed during the semester
- Groups of 2 or 3
 - Every member must have a sub-project representing a readily-distinguishable contribution
 - Together the sub-projects should integrate into a running project
 - Each member will be asked question on his sub-project and evaluated independently
- Must be a project in the scope of the course
 - E.g. adding anonymity to peerster
 - Many potential topics will present themselves throughout the lectures and exercises
 - You're welcome to come up with your own; **ask** if you're uncertain it's in-scope

Why study decentralized systems?

- **Direct applicability:** only a decentralized system can solve computing problems in which there is no common authority everyone trusts.
- **Indirect applicability:** real life is full of decentralized systems — but unlike other relevant fields (econ, soc-sci, pol-sci), we get to build and not just study them.
- **Intellectual challenge:** making decentralized systems work reliably and securely is often fundamentally more difficult than centralized systems, because we get to make fewer simplifying assumptions — particularly about security.

Why Build a Decentralized System?

- Sometimes a basic requirement: e.g., federated or "business-to-business" systems in which there is no natural trusted authority.
- Decentralized system components/subsets may be more autonomous, and hence more available or reliable in the face of disconnection or network partition, if they aren't dependent on regular communication with a central authority.
- Successful decentralized systems can become vibrant, diverse evolutionary ecosystems when alternative, competing node implementations can coexist and interoperate but evolve independently (e.g., IRC clients, P2P clients).
- This diversity can in turn make a decentralized system less vulnerable to attack or infection: only a subset of the complete system is likely to be vulnerable to any given attack or infection.

Why might we prefer centralized systems?

- **Management simplicity:** it's much simpler to manage an army of obedient slaves; “managing” a decentralized system is akin to herding cats.
- **Security model simplicity:** a lot of hard security challenges become easy when you can just ask a central authority who we assume (rightfully or not) to be trustworthy.
- **Efficiency:** A centralized management and security model can allow for simpler and more efficient algorithms, less overhead due to crypto, etc.
- **Version compatibility:** one can limit the number/variety of software versions that exist in a centralized system, but a decentralized node might have to (try to) interoperate with any past, future, or independently developed software/hardware combination on another node.

Major General Topics and Applications

- Communication: messaging, chat, voice/video
- Decentralized search and data mining facilities
- Collaboration mechanisms (e.g., wiki)
- Social networking
- Deliberation, peer review (voting, reputation systems)
- Blockchains, cryptocurrencies, and smart contract systems

This course inherently has a heavy security component — which we approach not primarily from a formal or cryptographic perspective, but from a pragmatic systems perspective.

Recurrent issues and themes


- Identity (real, sybil) versus location
- Information integrity and privacy
- Behavior accountability
- Denial-of-service
- Protocol efficiency, in the normal case and under load or attack

Gossip and USENET



The Life and Death of USENET

Books published with "UseNet" in title by year

Source: amazon.com and Library of Congress listings as of 29-Aug-2010



Even inspired an “Intergalactic USENET”



Example USENET message

```
From: jerry@eagle.ATT.COM (Jerry Schwarz)
Path: cbosgd!mhuxj!mhuxt!eagle!jerry
Newsgroups: news.announce
Subject: Usenet Etiquette -- Please Read
Message-ID: <642@eagle.ATT.COM>
Date: Fri, 19 Nov 82 16:14:55 GMT
Followup-To: news.misc
Expires: Sat, 1 Jan 83 00:00:00 -0500
Organization: AT&T Bell Laboratories, Murray Hill
```

The body of the message comes here, after a blank line.