

CPSC 416 Distributed Systems

Winter 2023 Term 1 (October 24, 2023)

Tony Mason (fsgeek@cs.ubc.ca), Lecturer



Logistics



Teaching Assistants

Andy Hsu (andy.hsu@alumni.ubc.ca)

Hamid Ramezanikebrya (hamid@ece.ubc.ca)

Jonas Tai (jonastai@student.ubc.ca)

Cathy Yang (kaiqiany@student.ubc.ca)



Office Hours

Remember: Use Piazza for **all** official course-related communications

- Not on Piazza? Not official.
- Canvas “comments/messages” **are not monitored**



Office Hours:

Who	When	Where
Tony	Monday 14:00-15:00 Wednesday 16:00-17:00	Discord
Andy	Thursday 19:00-20:30	Discord
Hamid	Friday 16:30-18:00	Kaiser 4075
Jonas	Wednesday 09:00-10:00	Discord
Cathy	Friday 09:00-10:30	X237

Self-Assessment

This week

- Design Project 2 Code Submission Due Thursday (October 26 @ 17:00)
 - **Note for credit the code must compile**
- Design Project 2 Implementation Report Due Thursday (October 26 @ 23:59)



Next Week

- Capstone Project: Weekly Report 2 Due Tuesday (October 31 @ 17:00)
- Design Project 3 Peer Review Due Tuesday (October 31 @ 17:00)
- Self-Assessment (Kleppmann Chapter 5) Due Tuesday (October 31 @ 17:00)
- Design Project 2 Peer Review Implementation Report Due Tuesday November 2 @ 17:00)

Note:

- You are strongly encouraged to collaborate with others on this
- You should use tools at your disposal to answer these questions
- **Do not forget to submit it.**

Capstone Project Overview



Learning Goals (Capstone Project)

Problem Identification and Analysis

Research Literature Review

Design & Planning

Implementation

Self-Reflection and Critical Thinking

Collaboration and Teamwork

Communication Skills

Ethical Considerations

Continuous Learning & Adaptability

Feedback & Iteration



Capstone Project: Key Goals

You need to identify a project that interests you

You need to decide if you will work solo, or with a team

You need to decide what *you* want to learn from this project

- Note: the **rubric** is focused around your ability to learn
 - Design
 - Self-reflection
 - Communications/interaction
 - Problem identification, scoping, and evaluation

It is *deliberately* open-ended because the emphasis is *learning how to learn*.



What Makes a Good Project?

Relevance & Impact

Originality & Creativity

Depth of Understanding

Iterative Process

Self-Reflection

Effective Communication

Collaboration & Team Dynamics

Self-reflection & Growth

Feedback & Openness

Passion & Enthusiasm

Notice what *isn't* listed?



Generic Project Ideas

Distributed File System

Distributed Task Scheduler

Distributed Key-Value Store

Distributed Consensus Algorithm Comparison

Distributed Caching System

Distributed Lock Manager

Distributed Stream Processing System

Distributed Machine Learning

Distributed Blockchain System

Distributed Simulation Framework



Advanced Generic Project Ideas

Geo-Replicated Databases

Distributed Tracing & Monitoring

Serverless Architectures

Distributed Graph Processing

Edge Computing

Distributed Security Protocols

Distributed Real-time Collaboration Tools

Distributed Quantum Computing

Distributed Virtual Reality (VR)

Research Replication: Reproducibility



Research Project Ideas

Distributed Systems for Decentralized Finance (DeFi)

- Smart Contracts
- De-centralized exchanges
- Lending Platforms

Hybrid MRDT-CRDT Systems

- Design a system combining both MRDTs and CRDTs.
- Explore trade-offs and benefits

Distributed Systems in Healthcare

- Patient Data Sharing (security, authorization, access control)
- Telemedicine (scheduling, messaging, tracking)
- Real-time data monitoring/alerting



Distributed Databases & NewSQL Systems

Example Systems: [Google Spanner](#), [CockroachDB](#), [TiDB](#)

Key Paper: “[One Size Fits All: An Idea Whose Time Has Come and Gone](#)”



Distributed Machine Learning

Key Paper: [Large Scale Distributed Deep Networks](#)

Examples: [TensorFlow](#), [Horovod](#)



Key Paper: [Kafka: a Distributed Messaging System for Log Processing](#)

Examples: [Apache Kafka](#), [Apache Flink](#), [Apache Storm](#)



Serverless Architectures & FaaS

Key Paper: Cloud Programming Simplified: [A Berkeley view on serverless computing.](#)

Example Systems: [AWS Lambda](#), [Azure Functions](#), [Google Cloud Functions](#)



Distributed Tracing & Monitoring

Key Paper: [Dapper, a Large-Scale Distributed Systems Tracing Infrastructure](#)

Examples: [Zipkin](#), [Jaeger](#)



Distributed Ledger & Blockchain

Key Papers:

- [Bitcoin: A Peer-to-Peer Electronic Cash System](#)
- [Algorand: Scaling Byzantine Agreements for Cryptocurrencies](#)
- [Blockchain: Digitally Rebuilding the real estate industry](#)



Examples: [Ethereum](#), [Hyperledger Fabric](#)

Note: I couldn't find a credible open-source land title registry implementation

Geo-distributed System and Edge Computing

Key Paper: [The Emergence of Edge Computing](#)

Examples: [AWS Wavelength](#), [Azure Edge Zones](#)



Distributed Graph Processing

Key Paper: [Pregel: a system for large-scale graph processing.](#)

Examples: [Apache Giraph](#), [Neo4j](#)



Distributed Multi-agent Systems

Key Paper: [An Introduction to Multi-Agent Systems](#)

Example: [JADE \(Java Agent DEvelopment Framework\)](#)



Distributed Virtual Reality & Gaming

Key Paper: On Consistency and Network Latency in Distributed Interactive Applications

- [Part 1](#)
- [Part 2](#)

Examples: [Unity Multiplayer](#), [Epic's Unreal Engine](#)



Distributed File Systems

[Network File System \(NFS\)](#)

[Andrew File System \(AFS\)](#)

[DCE/DFS](#)

[CIFS](#)

[Frangipani](#)

[IPFS](#)



Distributed Hash Tables

[Chord](#)

[Pastry](#)

[Tapestry](#)



Challenge

Pick a project

Pick a focus

Pick a team

Remember: the **primary evaluation** is based upon:

- Your ability to motivate your project
- Your ability to design your solution
- Your ability to evaluate your design
- Your ability to *communicate* your results effectively

Notice what isn't on this list?



Questions?





THE UNIVERSITY OF BRITISH COLUMBIA

