

# CPSC 416 Distributed Systems

Winter 2023 Term 1 (September 12, 2023)

Tony Mason ([fsgeek@cs.ubc.ca](mailto:fsgeek@cs.ubc.ca)), Lecturer



# Logistics



# Teaching Assistants

Andy Hsu ([andy.hsu@alumni.ubc.ca](mailto:andy.hsu@alumni.ubc.ca))

Hamidreza Ramezanikebrya ([hamid@ece.ubc.ca](mailto:hamid@ece.ubc.ca))

Jonas Tai ([jonastai@student.ubc.ca](mailto:jonastai@student.ubc.ca))

Cathy Yang ([kaiqiany@student.ubc.ca](mailto: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	Thursday 11:00-12:30	X150, Table 1&2
Cathy	Friday 09:00-10:30 (Starting Sep. 22)	X237

# Self-Assessment

## This week

- Post-lecture review assignment (Useful?) – Due before *next* lecture
- Distributed Systems Design Recipe assignment – Due Friday (Sep 15 @ 23:59)



## Note:

- You are strongly encouraged to collaborate with others on this
- You should use tools at your disposal to answer these questions
  - **Including your favourite large language model** (e.g., ChatGPT, your older sibling, or your pet.)
- As previously noted, you get full credit if you submit. **Do not forget to submit it.**

# Failure



# Today's Failure

Source: [Microsoft had three staff at Australian data centre campus when Azure went out - Cloud - Storage – iTnews](#)

When: 08:41 UTC on 30 August 2023 to 06:40 UTC on 1 September 2023

Duration: 46:59 (e.g, just short of **two days**)

Effect: “Issues affecting access to” Azure, Microsoft 365 and Power Platform.

TL;DR version: power sag caused chiller failure caused thermal failure of storage devices. Secondary problems (“bugs”) caused approval of stale requests and marked operational nodes as “not healthy” (degraded or failed). Recovery tools didn’t work well (“SQL did not have well tested tools on hand that were built to move databases when the source ring was in [a] degraded health scenario”)

**Cascading failures are common.** Recovery is hard – seldom used paths, staff training, etc.



MURPHY  
was an optimist





# Learning Goals



# Learning Goals

## Chapter 3: File Formats

- How you organize data matters (performance, recovery)
- Details matter (Understanding trade-offs)
  - Media characteristics
  - Usage patterns



## Chapter 4: Implementing B-Trees (Organizing persistent storage)

- How do we handle mutability efficiently?
- How do we access data quickly?
- How do we deal with space reclamation?

Note: these issues are recurring themes in systems.

# Discussion



# Discussion

Too much logistics and failure... Let's talk.

Petrov, Chapters 3 & 4



# Chapter 3: High Level

Cell Layout and Data Types

Slotted Pages

Variable Size Data

B-Tree Layout and Pages

Versioning and File formats

Checksumming and Data Integrity



# Chapter 4: High Level

B-Trees: key organizational structure for storage

Search

Splitting and Merging

Maintaining Balance

Optimized Data Loading

Space Efficiency

- Compression
- Cleaning/Compaction



# Questions?





THE UNIVERSITY OF BRITISH COLUMBIA

