ETH zürich

Week 2 - Cloud Storage

Aljaž Medič

Plan for today

1. Exercise Sheet 1 Questions
2. Quiz
3. HTTP
4. Vector Clocks
5. Merkle Trees

Exercise Sheet 1

Any questions about the exercise sheet solutions?

QUIZ

Name the mandatory components of a URI plus other optional ones.

Scheme, Path (mandatory); Authority, Query, Fragment (optional).

List the four basic HTTP methods.

GET, PUT, POST, DELETE.

What are the four design principles of Amazon Dynamo?

Incremental stability, Symmetry, Decentralization, Heterogeneity.

QUIZ

What is meant by heterogeneous data?

Data that doesn't fulfil domain or relational integrity

A data lake can only store structured data after schema enforcement.

False. (Data lakes can store unstructured data in raw form.)

How does the CAP theorem define the limitations of distributed storage systems?

Distributed systems can only achieve two of: Consistency, Availability, Partition Tolerance

QUIZ

What are the main advantages of key-value stores over RDBMS?

Simplicity, scalability, performance at scale

What is the difference between an object and a value, in terms of size?

Object ~5TB, Value ~400KB

Object stores like Amazon S3 primarily store structured data.

False. (They store unstructured or semi-structured data.)

QUIZ

Key-value stores usually guarantee strong (atomic) consistency.

False. (They usually offer eventual consistency.)

Object storage supports a physical file hierarchy of directories.

False. (Hierarchy is only logical; physically flat.)

PUT and POST are idempotent HTTP methods.

False. (PUT is idempotent, POST is not.)

QUIZ

GET is an idempotent HTTP method.

True.

REST APIs are widely used for cloud storage because they provide a standardized HTTP-based interface.

True.

Why does it become desirable to denormalize data at large scales?

Performance reasons (avoids joins)

QUIZ

Static websites (no server-side logic) can be hosted directly on S3.

True.

A system that is AP and has eventual consistency will be consistent.

False. (In practice, such a system is rarely ever consistent. (Marketing trick))

Key-value stores store data as records with predefined schemas, similar to relational databases.

False. (They store flexible key-value pairs without predefined schema.)

QUIZ

The use of virtual nodes in Dynamo helps balance the load.

True.

Give all three paradigms for scaling when one machine isn’t enough.

Scaling up, Scaling out, Optimizing code.

Name the two main paradigms for storing data.

Import into database (ETL), Query in place (Data lake).

QUIZ

List the SLA availability levels for 4 and 7 nines, with their approximate maximum downtime (round up on the last non-zero time unit).

99.99% = 53m/year; 99.99999% = 4s/year.

What are the three operations in a basic key-value store API?

get(key), put(key, value), delete(key).

URI and HTTP

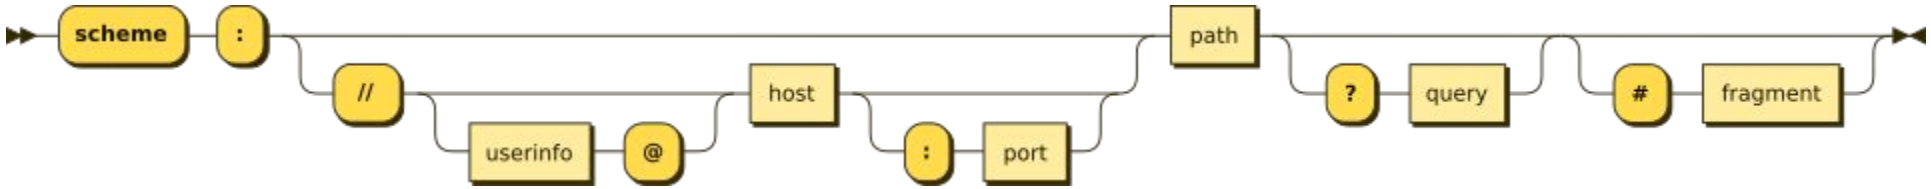

Safe, idempotent, and cacheable request methods

The following table lists HTTP request methods and their categorization in terms of safety, cacheability, and idempotency.

Method	Safe	Idempotent	Cacheable
`GET`	Yes	Yes	Yes
`HEAD`	Yes	Yes	Yes
`OPTIONS`	Yes	Yes	No
`TRACE`	Yes	Yes	No
`PUT`	No	Yes	No
`DELETE`	No	Yes	No
`POST`	No	No	Conditional*
`PATCH`	No	No	Conditional*
`CONNECT`	No	No	No

* `POST` and `PATCH` are cacheable when responses explicitly include freshness information and a matching `Content-Location` header.

HTTP response status codes

HTTP response status codes indicate whether a specific HTTP request has been successfully completed. Responses are grouped in five classes:

1. Informational responses (`100` – `199`)
2. Successful responses (`200` – `299`)
3. Redirection messages (`300` – `399`)
4. Client error responses (`400` – `499`)
5. Server error responses (`500` – `599`)

The status codes listed below are defined by RFC 9110.

Sources:
- Wikipedia (URI)
- Mozilla HTTP Docs (Request Methods)

Vector Clocks

Dynamo paper: **xx ([S_2, 1])** or **yy ([S_1, 1], [S_3, 1], [S_0, 1])**

(they have to denote writing node explicitly)

Our notation: **S_2 : xx ([0,1,0,0,0])** or **S_0 : yy ([1,1,0,1,0])**

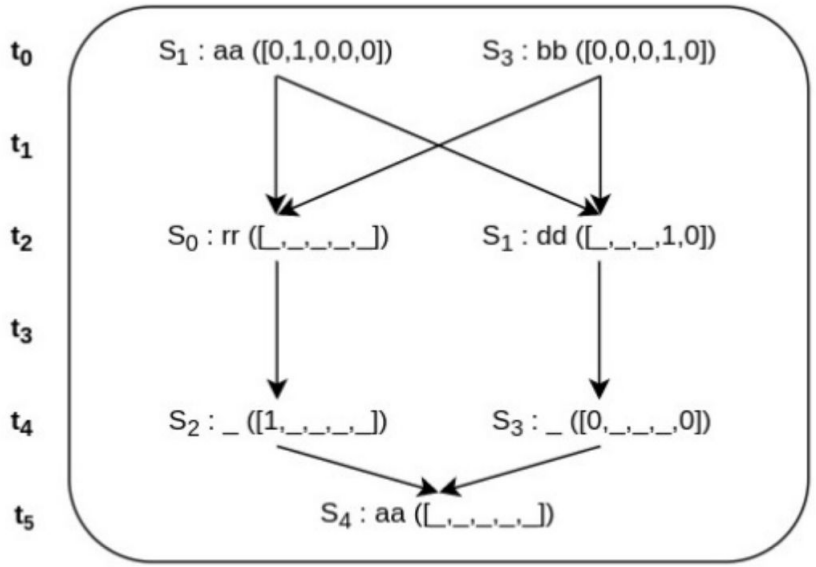

Obj. 2

Diagram source: HonyuHe

Merkle Trees

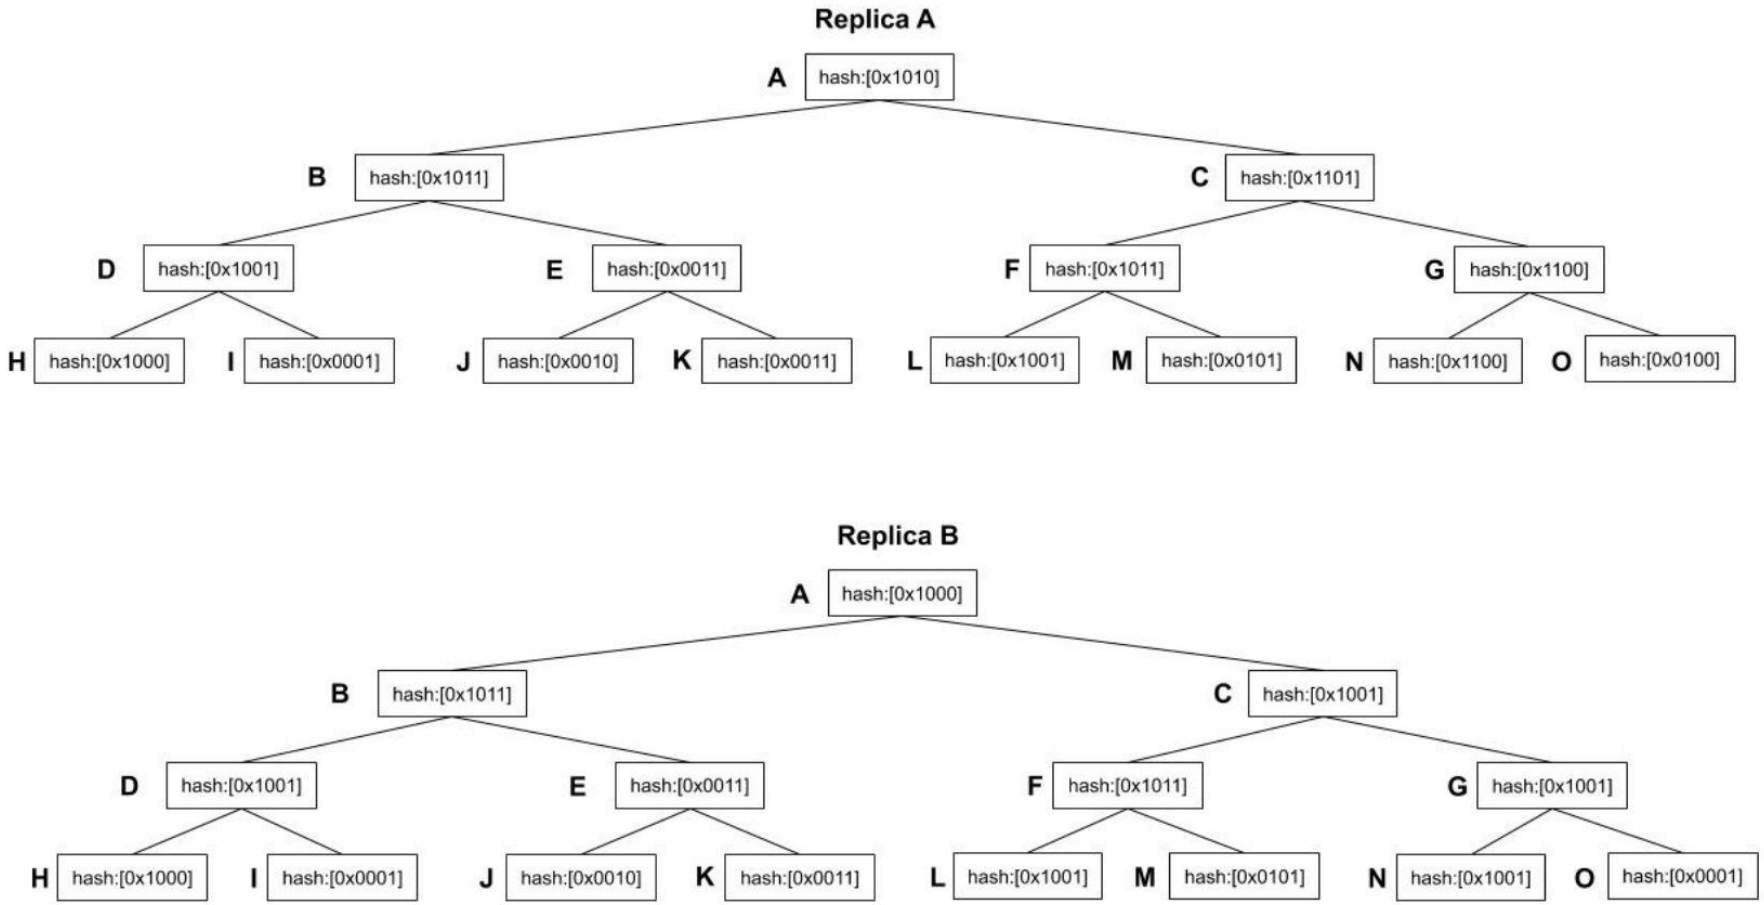

ETH zürich

See you next week!

Aljaž Medič
amedic@ethz.ch

Slides

Suggestions