

TDT4225 Store, distribuerte datamengder (Very large, distributed data volumes)

- Data models and query languages
- Storage and retrieval
- Encoding and evolution
- Replication
- Partitioning
- Transactions
- Trouble with distributed systems
- Consistency and consensus
- SSDs and LSM trees
- Time and clocks on distributed systems
- Examples of distributed databases
 - o Amazon Dynamo
 - o Google Spanner
 - o Dostoevsky (Harvard DASLAB)
- Programming exercises in groups (4 students) using virtual machines
 - o Programming in Python with MySQL
 - o Programming in Python with MongoDB