

CS-300

Data-intensive systems

Ailamaki Anastasia, Kashyap Sanidhya

Cursus	Sem.	Type
Communication systems	BA6	Opt.
Computer science minor	E	Opt.
Computer science	BA6	Obl.
Data science minor	E	Opt.
HES - IC	E	Opt.

Language of teaching	English
Credits	6
Session	Summer
Semester	Spring
Exam	Written
Workload	180h
Weeks	14
Hours	5 weekly
Lecture	2 weekly
Exercises	1 weekly
Project	2 weekly
Number of positions	

Summary

The purpose of this course is to discuss the design of database and operating systems concepts using a hands-on approach.

Content

Topics we will cover include database and operating system architecture, data models, query optimization and planning, query engine and system programming, networked and distributed systems, and storage systems.

- Entity-relationship and relational model
- Relational algebra and Calculus
- The SQL query language
- Storage and indexing
- Query operators and optimization
- Transaction Management and Concurrency control
- Isolation mechanisms: process abstraction, system calls, and interrupts
- Virtual memory: MMU, TLB, paging, on-demand paging, working set
- Task scheduling: Context switch, scheduling algorithms, coroutines
- Basic synchronization mechanisms: locking and latching, task coordination
- File system: UNIX file representation, inodes, crash recovery, buffer cache, logging mechanisms

Learning Prerequisites**Required courses**

CS-202 Computer Systems

CS-302 Parallelism and concurrency in software (from 2024-2025)

Recommended courses

CS-200 Computer architecture

Important concepts to start the course

Algorithms and data structures, programming, parallelism

Learning Outcomes

By the end of the course, the student must be able to:

- Identify and manage key components of database and operating systems
- Choose or critique design choices for DB and OS system software
- Express application information requirements and model the data of an application
- Create and design a database with a practical application in mind while justifying choices
- Explore how a DBMS and/or an OS performs work
- Report performance and possible optimizations for applications

Teaching methods

Lectures, exercises, and projects

Expected student activities

- Attend the lectures in order to ask questions and interact with the professor
- Attend the exercises session to solve and discuss exercises about the recently taught material
- Work on a team project which covers the practical side of the course, e.g. build a key component of an OS or build an application for using a database system
- Take a midterm and a final exam

Assessment methods

Labs (30%)

Midterm (25%)

Final exam (45%)

Supervision

Office hours	Yes
Assistants	Yes

Resources

Moodle Link

- <https://go.epfl.ch/CS-300>