

- Faculté des sciences
- www.unine.ch/sciences

Large-scale Distributed Systems (3IN2032)

Filières concernées	Nombre d'heures	Validation	Crédits ECTS
Master en informatique	Cours: 2 ph Exercice: 2 ph	écrit: 2 h	5

ph=période hebdomadaire, pg=période globale, j=jour, dj=demi-jour, h=heure, min=minute

Période d'enseignement:

- Semestre Automne

Equipe enseignante:

Dr. Etienne Riviere

Objectifs:

This course presents the paradigms, algorithmic techniques and design principles for large-scale distributed systems (LSDS) design. The focus of the course is on decentralized, autonomous and self-organizing systems. It favors practical case studies of state-of-the-art system designs from academia and industry.

The course covers: LSDS characteristics and requirements, gossip-based protocols, DHTs and other routing/querying substrates, distributed file systems and distributed data management, large-scale processing & MapReduce, multicast, publish-subscribe, and content-based routing. A unique feature for a course on LSDS is the use of hands-on experimentation of a set of representative algorithms on a distributed testbed using real protocols deployments and no simulation. This task is made simple thanks to a specialized framework developed at the university and used by several research groups worldwide.

Evaluation: small projects reports, and final exam based in part on an industry system example.

Contenu:

On successful completion of this course, you will:

- understand the key paradigms and services that one can use to build scalable applications;
- understand the objectives of fault tolerance, load balancing and efficiency and what techniques are appropriate to achieve them;
- have a general knowledge of the existing solutions used at major companies such as Amazon, Yahoo! or Google, and understand why the traditional solutions are not effective anymore;
- have a practical experience in writing and evaluating code on a real large-scale testbeds, processing, presenting and discussing the results of your observation

Forme de l'évaluation:

The evaluation is on the final exam (50% of the grade) and the grades of the project assignments (50%). There are 3 project assignments. They are mandatory and are due on fixed dates announced at the first lecture. Upon failure at the exam, the grade for the assignments will be kept when the student passes the exam another time (note that it is not possible to secure a 4 by passing only the exam due to the 50%/50% rule).

Documentation:

no mandatory book. some research paper reading may be required during the course.

a publication about the course in a conference dedicated to distributed systems teaching can be found here:
<http://cs.gsu.edu/~tcpp/curriculum/?q=advanced-technical-program>

Pré-requis:

Mandatory: knowledge of imperative programming (e.g., C, Java, Python, Perl, etc.);

Recommended: prior course or seminar on small-scale distributed systems (distributed clocks, message passing, RPC, etc.)

Contact the instructor if you would like to know if your background is sufficient for the course.

Forme de l'enseignement:

Lessons and practical sessions using a dedicated systems for easy prototyping and evaluation