

Large-Scale Distributed Systems and Networks

Programme course

11 credits

Storskaliga distribuerade system och nätverk

TDDE35

Valid from: 2018 Spring semester

Determined by

Board of Studies for Computer Science
and Media Technology

Date determined

Main field of study

Computer Science

Course level

First cycle

Advancement level

G1X

Course offered for

- Computer Science and Software Engineering, M Sc in Engineering

Specific information

This course is not available for exchange students

Entry requirements

Note: Admission requirements for non-programme students usually also include admission requirements for the programme and threshold requirements for progression within the programme, or corresponding.

Prerequisites

Calculus, statistics, and programming knowledge (preferably both in java and C).

Intended learning outcomes

Computer networks have become an indispensable part of the infrastructure of our modern society. With billions of people and devices being connected and using critical distributed services implemented over the Internet, for example, it is becoming very important to understand how these networks, as well as the distributed systems and services operating over these networks, are designed to scale to large number of machines and users. Also, at the level of individual machines, it is important to know how to build applications and services that effectively scale with the resources (e.g., the number of cores, processors, etc.). In this course, we will use a combination of theory and practice (including exploration of real data) to gain a deeper understanding of modern large-scale systems and services.

Within the area of computer networks, participants successfully completing the course are expected to be able to:

- Explain, describe, and analyze a typical network architecture, including arguing regarding the importance of network layers and encapsulation
- Explain the different basic types of protocols, communication channels, and

network types

- Design, implement, verify, and test your own protocols
- Explain fundamental performance tradeoffs, including showing an understanding of where delays can occur in a network, what different types of delay that exist, the impact of packet losses and jitter on various protocols
- Using concrete examples, and in detail, describe the interaction between the different protocols in the network architecture, and the protocols associated with the different layers
- Describe and analyze the most common application architectures in the Internet, how the most important application-layer protocols work, the service they provide, as well as have the ability to design and implement their own application-layer protocols
- Analyze and explain important design considerations at the transport layer, including hands-on knowledge of how flow control and congestion control works, and how reliable data transfer is implemented
- Motivate and explain how routing and forwarding is implemented on the Internet, including basic design and implementation principles of network-layer protocols used to ensure scalability
- Describe and explain different link-layer technologies and how they work
- Exemplify how different types of security services can be implemented in different layers with the help of different standards
- Analyze and exemplify some of the unique challenges as we are moving towards increasingly mobile users
- Explain and discuss the fundamentals of how multimedia services are provided over the Internet

The students are also expected to understand how distributed systems can be built on-top of the network architecture to provide scalable services, as well as how multicore systems and embedded systems can be used to further enhance services. More specifically, after successfully completing the course the student is expected to be able to:

- Define what a distributed system is and its most important goals
- Explain the relationship between architectures, processes and communication
- Exemplify different types of transparency, scaling techniques
- Analyze and explain some of the fundamental differences in different system architectures
- Describe and explain how to achieve synchronization, consistency and replication
- Motivate and explain the design of various types of distributed system architectures, including object-based distributed systems and Web-based distributed systems (including how a proxy cache works)
- Understand fundamental homogeneous and heterogeneous multicore architecture concepts and their performance implications; basic techniques for multicore programming with threads and tasks; and some techniques to design parallel algorithms and analyze their complexity, including parallel scalability
- Understand system-level methods and tools for the design of real-time

embedded systems; understand basic tradeoffs and design implications that need to be taken into consideration when making system-level design decisions; and place the design in a bigger context (including in the context of the hardware architecture and software implementation).

By introducing general design concepts, some basic scientific methodologies (such as basic systems performance modeling), exploration of real-world data, and a general systems thinking, with scale and performance as important aspects, used throughout the course, we also expect the student to be able to:

- Explain using concrete examples fundamental network design principles and scalability tradeoffs
- Design and perform targeted experiments to critically evaluate network and distributed systems technologies
- Apply basic system models and analysis methods to analyze systems and networks
- As a team, plan and conduct a project study of an identified problem within a selected technology area, including experiments using real data sources (in some cases collected by the students themselves)
- Based on a project study, present and explain (both written and orally) findings within a selected area of technology, to an audience with similar general knowledge of computer networks
- Give/receive constructive feedback to/from other students

Course content

Computer networking: The fundamental design principles of computer networks, their protocols, and the Internet stack. Application layer protocols (e.g., HTTP), transport layer protocols (e.g. TCP), network layer protocols (e.g., IP and BGP), link layer protocols (e.g., Ethernet, WiFi, and Bluetooth). Introduction to multimedia applications, wireless networking, and network security for each of the layers.

Distributed systems, multicore systems, and embedded systems: Basic distributed architectures and their processes and communication. Synchronization, replication, and consistency issues and tradeoffs. Object-based and Web-based systems. Multicore architectures, their opportunities, and basic challenges they present. Embedded systems and their integration into a wide range of modern systems.

Project: The precise topics for the project will vary slightly from year to year, to keep the projects exciting and up-to-date with developments in the areas.

Recurring topics include: Fundamental properties of large-scale systems (e.g., power laws, rich-gets-richer); Scalable systems and designs (e.g., hierarchical vs. flat designs; layered designs); Measurement, modeling and analysis methods using real network data; Important modern distributed systems such as cloud-based services (e.g., EC2), CDNs, the Internet routing architecture itself, and social networks.

Teaching and working methods

The course consists of both theory (lectures and assignments) and practical hands-on training and exploration (lab assignments and project). The course has two written exams. The first on networking and the second on introductory material about distributed systems, multicore systems, embedded systems, and basic system science methodologies. The project should result in a written report, should be presented in a seminar during which the students will act as both presenters and opponents (evaluating and providing each other with feedback, such as to improve the reports and projects). Assignments should be done in pairs. The projects should be done in groups of three-to-four students. The course runs over the entire spring semester.

Examination

PRA1	Project work	2 credits	U, G
LAB1	Laboratory work	3 credits	U, G
TEN2	Written examination	3 credits	U, 3, 4, 5
TEN1	Written examination	3 credits	U, 3, 4, 5

For a pass grade in the course, at least a pass grade is needed for all components. The final grade will be calculated using the average from the individual exam grades.

Grades

Four-grade scale, LiU, U, 3, 4, 5

Other information

Supplementary courses:
Distributed Systems; Advanced Networking (TDTS21); Individual projects

Department

Institutionen för datavetenskap

Director of Studies or equivalent

Patrick Lambrix

Examiner

Niklas Carlsson

Course website and other links

Education components

Preliminary scheduled hours: 80 h

Recommended self-study hours: 213 h

Course literature

Books

J. F. Kurose and K. W. Ross, (2017) *Computer networking: A top-down approach*
7th Pearson

Other

Various articles and online resources

Common rules

Course syllabus

A syllabus has been established for each course. The syllabus specifies the aim and contents of the course, and the prior knowledge that a student must have in order to be able to benefit from the course.

Timetabling

Courses are timetabled after a decision has been made for this course concerning its assignment to a timetable module. A central timetable is not drawn up for courses with fewer than five participants. Most project courses do not have a central timetable.

Interrupting a course

The vice-chancellor's decision concerning regulations for registration, deregistration and reporting results (Dnr LiU-2015-01241) states that interruptions in study are to be recorded in Ladok. Thus, all students who do not participate in a course for which they have registered must record the interruption, such that the registration on the course can be removed. Deregistration from a course is carried out using a web-based form: www.lith.liu.se/for-studenter/kurskomplettering?l=sv.

Cancelled courses

Courses with few participants (fewer than 10) may be cancelled or organised in a manner that differs from that stated in the course syllabus. The board of studies is to deliberate and decide whether a course is to be cancelled or changed from the course syllabus.

Regulations relating to examinations and examiners

Details are given in a decision in the university's rule book:
<http://styrdokument.liu.se/Regelsamling/VisaBeslut/622678>.

Forms of examination

Examination

Written and oral examinations are held at least three times a year: once immediately after the end of the course, once in August, and once (usually) in one of the re-examination periods. Examinations held at other times are to follow a decision of the board of studies.

Principles for examination scheduling for courses that follow the study periods:

- courses given in VT1 are examined for the first time in March, with re-

examination in June and August

- courses given in VT2 are examined for the first time in May, with re-examination in August and October
- courses given in HT1 are examined for the first time in October, with re-examination in January and August
- courses given in HT2 are examined for the first time in January, with re-examination at Easter and in August.

The examination schedule is based on the structure of timetable modules, but there may be deviations from this, mainly in the case of courses that are studied and examined for several programmes and in lower grades (i.e. 1 and 2).

- Examinations for courses that the board of studies has decided are to be held in alternate years are held only three times during the year in which the course is given.
- Examinations for courses that are cancelled or rescheduled such that they are not given in one or several years are held three times during the year that immediately follows the course, with examination scheduling that corresponds to the scheduling that was in force before the course was cancelled or rescheduled.
- If teaching is no longer given for a course, three examination occurrences are held during the immediately subsequent year, while examinations are at the same time held for any replacement course that is given, or alternatively in association with other re-examination opportunities. Furthermore, an examination is held on one further occasion during the next subsequent year, unless the board of studies determines otherwise.
- If a course is given during several periods of the year (for programmes, or on different occasions for different programmes) the board or boards of studies determine together the scheduling and frequency of re-examination occasions.

Registration for examination

In order to take an examination, a student must register in advance at the Student Portal during the registration period, which opens 30 days before the date of the examination and closes 10 days before it. Candidates are informed of the location of the examination by email, four days in advance. Students who have not registered for an examination run the risk of being refused admittance to the examination, if space is not available.

Symbols used in the examination registration system:

** denotes that the examination is being given for the penultimate time.

* denotes that the examination is being given for the last time.

Code of conduct for students during examinations

Details are given in a decision in the university's rule book:
<http://styrdokument.liu.se/Regelsamling/VisaBeslut/622682>.

Retakes for higher grade

Students at the Institute of Technology at LiU have the right to retake written examinations and computer-based examinations in an attempt to achieve a higher grade. This is valid for all examination components with code "TEN" and "DAT". The same right may not be exercised for other examination components, unless otherwise specified in the course syllabus.

Retakes of other forms of examination

Regulations concerning retakes of other forms of examination than written examinations and computer-based examinations are given in the LiU regulations for examinations and examiners,

<http://stydokument.liu.se/Regelsamling/VisaBeslut/622678>.

Plagiarism

For examinations that involve the writing of reports, in cases in which it can be assumed that the student has had access to other sources (such as during project work, writing essays, etc.), the material submitted must be prepared in accordance with principles for acceptable practice when referring to sources (references or quotations for which the source is specified) when the text, images, ideas, data, etc. of other people are used. It is also to be made clear whether the author has reused his or her own text, images, ideas, data, etc. from previous examinations.

A failure to specify such sources may be regarded as attempted deception during examination.

Attempts to cheat

In the event of a suspected attempt by a student to cheat during an examination, or when study performance is to be assessed as specified in Chapter 10 of the Higher Education Ordinance, the examiner is to report this to the disciplinary board of the university. Possible consequences for the student are suspension from study and a formal warning. More information is available at <https://www.student.liu.se/studenttjanster/lagar-regler-rattigheter?l=sv>.

Grades

The grades that are preferably to be used are Fail (U), Pass (3), Pass not without distinction (4) and Pass with distinction (5). Courses under the auspices of the faculty board of the Faculty of Science and Engineering (Institute of Technology) are to be given special attention in this regard.

1. Grades U, 3, 4, 5 are to be awarded for courses that have written examinations.
2. Grades Fail (U) and Pass (G) may be awarded for courses with a large degree of practical components such as laboratory work, project work and group work.

Examination components

1. Grades U, 3, 4, 5 are to be awarded for written examinations (TEN).
2. Grades Fail (U) and Pass (G) are to be used for undergraduate projects and other independent work.

3. Examination components for which the grades Fail (U) and Pass (G) may be awarded are laboratory work (LAB), project work (PRA), preparatory written examination (KTR), oral examination (MUN), computer-based examination (DAT), home assignment (HEM), and assignment (UPG).
4. Students receive grades either Fail (U) or Pass (G) for other examination components in which the examination criteria are satisfied principally through active attendance such as other examination (ANN), tutorial group (BAS) or examination item (MOM).

The examination results for a student are reported at the relevant department.

Regulations (apply to LiU in its entirety)

The university is a government agency whose operations are regulated by legislation and ordinances, which include the Higher Education Act and the Higher Education Ordinance. In addition to legislation and ordinances, operations are subject to several policy documents. The Linköping University rule book collects currently valid decisions of a regulatory nature taken by the university board, the vice-chancellor and faculty/department boards.

LiU's rule book for education at first-cycle and second-cycle levels is available at http://styrdokument.liu.se/Regelsamling/Innehall/Utbildning_pa_grund-_och_avancerad_niva.