

Parallel and Distributed Software Systems (E017930)

Course size (nominal values; actual values may depend on programme)
 Credits 6.0 Study time 180 h Contact hrs 60.0 h

Course offerings and teaching methods in academic year 2017-2018

A (semester 1)	English	practicum	5.0 h
		seminar: coached	10.0 h
		exercises	
		lecture	30.0 h
		self-reliant study	15.0 h
B (semester 1)	Dutch	activities	
		guided self-study	30.0 h
		seminar: coached	10.0 h
		exercises	
		self-reliant study	15.0 h
		activities	
		practicum	5.0 h

Lecturers in academic year 2017-2018

De Turck, Filip	TW05	lecturer-in-charge
Fostier, Jan	TW05	co-lecturer

Offered in the following programmes in 2017-2018

	crdts	offering
Bridging Programme Master of Science in Electrical Engineering (main subject Communication and Information Technology)	6	A
Brugprogramma Master of Science in Bioinformatics (main subject Engineering)	6	A
Bridging Programme Master of Science in Computer Science Engineering	6	B
Bridging Programme Master of Science in Computer Science Engineering	6	A
Master of Science in Electrical Engineering (main subject Communication and Information Technology)	6	A
Master of Science in Electromechanical Engineering (main subject Control Engineering and Automation)	6	A
Master of Science in Electromechanical Engineering (main subject Electrical Power Engineering)	6	A
Master of Science in Bioinformatics (main subject Engineering)	6	A
Master of Science in Electromechanical Engineering (main subject Maritime Engineering)	6	A
Master of Science in Electromechanical Engineering (main subject Mechanical Construction)	6	A
Master of Science in Electromechanical Engineering (main subject Mechanical Energy Engineering)	6	A
Master of Science in Computer Science	6	A
Master of Science in Computer Science Engineering	6	B
Master of Science in Computer Science Engineering	6	A
Exchange Programme in Bioinformatics (master's level)	6	A
Exchange Programme in Computer Science (master's level)	6	A

Teaching languages

Dutch, English

Keywords

Parallel and distributed software, high performance computing, cloud computing, big data processing, communication, coordination, synchronization, efficiency, programming models, fault tolerance.

Position of the course

This course will teach students the concepts regarding the different aspects of the design and implementation of distributed software. The course will provide the students with a state-of-the-art overview of parallel and cloud computing systems, design of parallel algorithms, software engineering specifically for these applications, and management of high performance and cloud-based systems. The emphasis is on the software side and on the different programming models. Hardware/architecture aspects are assumed to be covered in other courses and are only used to the extent necessary to understand the impact on the software performance.

Contents

- Fundamentals, definitions & terminology, classification.
- Performance metrics & limiting factors: speedup, efficiency, scalability (strong & weak), high availability, Amdahl's and Gustafson's law, CAP theorem, network cost modeling, failure.
- Distributed software: message passing, remote procedure call (RPC), distributed objects, remote message invocation (RMI), request-reply protocol, marshalling. Message oriented middleware and services.
- Cloud computing models: service models, deployment models, payment models; cloud platforms and programming models.
- Data-driven architectures: scale out vs. scale up, move processing power to data, avoid random access, scalability; MapReduce; Spark; NoSQL; Big data for enterprise applications.
- High performance distributed-memory computing: MPI, point-to-point communication, collective communication, problem decomposition, case studies.
- Shared-memory programming: data protection, achieving concurrency, mutexes, semaphores, condition variables, deadlocks, false-sharing, thread-safety, atomic operations, lock-free programming, programming models, case studies.
- Monitoring and management of large-scale parallel and distributed software systems: large-scale measurements and monitoring, autonomic and control theory based management.
- GPU programming: kernels, thread hierarchy, memory hierarchy, control flow.

Initial competences

Basic programming skills in C/C++ and Java. Basic knowledge of datastructures and algorithms. Basic knowledge of operating systems.

Final competences

- 1 To know and understand the most important characteristics of distributed systems and the associated algorithmic problems when designing parallel and distributed software.
- 2 To know and understand the most important architectures to realize parallel and distributed systems and the most important software technologies to realize parallel and distributed applications.
- 3 To know and understand the basic techniques to solve algorithmic problems associated with parallel and distributed systems.
- 4 To be able to design parallel and distributed algorithms using an appropriate programming model.
- 5 To be able to realize and validate an implementation of parallel and distributed software.
- 6 To be able to assess the performance and scalability of parallel and distributed software.

Conditions for credit contract

Access to this course unit via a credit contract is determined after successful competences assessment

Conditions for exam contract

This course unit cannot be taken via an exam contract

Teaching methods

Guided self-study, lecture, practicum, self-reliant study activities, seminar: coached exercises

Learning materials and price

- Syllabus via VTK
- Slides on Minerva
- Source code examples via Minerva

References

- George Coulouris, Jean Dollimore, Tim Kindberg, Gordon Blair, "Distributed Systems: Concepts and Design (5th Edition)", Pearson Publishers.
- Rajkumar Buyya, James Broberg, Andrzej M. Goscinski, "Cloud Computing: Principles and Paradigms", Wiley Publishers.
- Cloud programming (online references)
- MPI, The complete reference (online)
- Peter Pacheco: "An Introduction to Parallel Programming", Morgan Kaufmann.
- Ian Foster, "Designing and Building Parallel Programs", Addison-Wesley Inc.
- Jimmy Lin, Chris Dyer, "Data-intensive Text Processing with MapReduce".

Course content-related study coaching

- Practicals are supervised by assistants.
- Additional information via Minerva.

Evaluation methods

end-of-term evaluation and continuous assessment

Examination methods in case of periodic evaluation during the first examination period

Written examination

Examination methods in case of periodic evaluation during the second examination period

Written examination

Examination methods in case of permanent evaluation

Skills test, report

Possibilities of retake in case of permanent evaluation

examination during the second examination period is not possible

Extra information on the examination methods

- During examination period: written closed-book exam; written open-book exam.
- During semester: graded practicals and homework assignments.

Calculation of the examination mark

- 75% exam
- 25% practicals and homework assignment

In case the score for either the examination or the practicals and homework assignment is less than 8/20, the final score for this course will be limited to 8/20.