Master of Science
Parallel and Distributed Computer Systems

Courses description

Advanced Techniques in Computer and Network Security

The course is focused on current issues in computer and network security. There are discussed security issues of the electronic information from the perspective of ensuring the confidentiality, integrity, authenticity and nonrepudiation of the information. The course is a natural continuation of the discipline of “Cryptography” and addresses from practical perspective of the cryptographic techniques and protocols (authentication, secret sharing, key exchange, subliminal channels), with applications in e-mail security, IP security, Web security and security of computer network (protected by firewall, IDS, IPS devices and antivirus protection techniques). Also, the course presents principles used in security evaluation of cryptographic modules, products and applications derived from industrial security standards: Security Requirements for Cryptographic Modules (ISO 19790) and Common Criteria for Information Technology Security Evaluation (ISO 15408).

Cluster and Grid Computing

The course covers the specific concepts of Cluster and GRID computing. It is presented the basic concepts of cluster, as a new approach of parallel and distributed processing system, which consists of a collection of interconnected stand-alone heterogeneous systems cooperatively working together as a single, integrated computing resource. In the course are presented the type of clusters, cluster architecture, new concept in OS services for distributed processing , physical cluster interconnections and interconnect support, cluster programming environments, monitoring and performance analysis tools.
The course presents the essence of Grids how to utilize highly flexible network architectures, and how to sharing of all computing resources, not just data. Are presented the grid technologies, an extensible and open Grid architecture, general aspects of basic components that enable interoperability among different Grid resources. There are presentend the principal Grid characteristics: Wide geographical distribution,

Heterogeneous, Resource sharing, Multiple admin policies, Resource coordination, Transparent access, Dependable, Consistent, Pervasive. It is presented a Sample Grid Computing Environment: Resource Sharing & Aggregation and Grid Architecture for Computational Economy. The Layered Grid Architecture is presented.
In the second part of the course the students have to choose a topic of application of the Grid computing and they have to present an essay about it.

Computer and Network Security

Computer and Network Security class covers the topic of computer and network security, from advanced cryptographic techniques, innovative algorithms, security protocols and policies, trust ensuring infrastructures and specific technologies for securing networks.
At the end of this class, the students should be able to: develop a complex security policy for ensuring the safety of the provided services, identify threats and vulnerabilities at network and operating system level, identify types of attacks and mitigate them, identify the vulnerabilities of layer 2 and 3 equipments, implement a solution for ensuring a high level of security for the equipments, monitor the network, servers and workstations for identifying and stopping attacks, configuring a firewall, an IDS/IPS and an AAA system.
This class also has an important practical component, which includes: configuring routers that define the network perimeter with the use of specific instruction set for ensuring a high level of security, configuring layer 2 and 3 equipments in a safe manner, configuring a firewall for allowing basic network operations, configuring an IDS/IPS, configuring AAA both on routers and switches.

Distributed algorithms

Distributed algorithms are a very actual and important area in Computer Science. This is a master level class, in which we assume the students are already familiar with the fundamental notions and models of distributed algorithms. First, we will discuss about causal order and total order. We will present the problems and discuss why traditional solutions will not work in all cases. We will then discuss the possible solutions, both in academia and industry and study some of the most important research papers that appeared on these topics. A particular example of a distributed application, a collaborative text editor, will be discussed and we will see how we can relax the strong conditions from total order for this particular case by introducing the operational transformation technique. Finally, we will discuss about distributed consensus and peer to peer systems.

Distributed Systems

The course aims at teaching and integrating the main concepts, principles, models, and techniques for the development of distributed system software. Another objective is increasing students' capacity to use the acquired knowledge in modeling, design of software components for distributed systems, implementing middleware programs based on modern technologies, evaluating the implemented systems against specification and design criteria, and optimizing the performance of solutions with engineering methods and instruments.
Course Content: Models and architectures of large-scale dynamic distributed systems. P2P systems for content distribution: structured and unstructured architectures, search methods; content replication techniques; anonymity and reputation management. Event-driven distributed systems: architecture, components; complex event processing, ECA (event, condition, action) model and finite-state automata with interval timestamps; intelligent engines for event processing. Cloud computing: resource provisioning, efficient use of resources; energy management; data security; data storage and management techniques. Compute and data intensive models, MapReduce. Context-based distributed systems: context sensing, communication and processing; ontology based context models; security and privacy.

Parallel Programming

The Parallel Programming lecture is outlining a series of programming paradigms in the context of modern parallel computer architectures. First an overview of parallel programming models is given considering various programming models, issues such as productivity, performance, and portability and presenting a number of models for communication, synchronization, memory consistency and runtime systems. In this context, in the second part of the lecture, a series of parallel programming paradigms with shared (OpenMP, PThreads, Cilk, TBB, HPF, Chapel, Fortress, Stapl), and distributed memory (MPI, Charm++, Stapl), parallel global address shared space (UPC, X10) as well as other atypical paradigms are presented (Linda, MapReduce, MATLAB DCE). A basic understanding of computer architectures, operating systems, parallel and distributed algorithms and compilers is assumed.

Scheduling Methods and Algorithms in Distributed Systems

The objectives of this course are: Classification of scheduling problems; Specific attributes and models of scheduling algorithms; The complexity of scheduling algorithms; Scheduling and resource management; Advance reservation and co-allocation of resources; Scheduling policies; Algorithms and classical methods of scheduling: scheduling on a single processor; Scheduling for parallel systems; Scheduling in distributed systems (Grid, P2P); Scheduling with dependencies and workflows; Methods of optimization and performance analysis; Fault tolerance and rescheduling. Applications have as main objectives analysis, implementation and testing of presented algorithms. The project aims to create a context for the implementation of scheduling policies. Optimization of scheduling by implementing optimization techniques will be another focus of practical applications (projects). It will also study the applications of scheduling algorithms: courses and exams planning, land vehicles planning, aircraft schedule.

Operating Systems

The Operating Systems lecture goal is to stimulate research activities and develop student abilities to implement advanced concepts by presenting and offering for debate research papers addressing advanced topics in the operating systems space. During the lecture, an open debate is started to discuss the research paper(s). The scientific topics include virtualization, advanced synchronization mechanisms, network deep packet inspection, operating systems security, vulnerability analysis, mobile OS topics.
The lecture is doubled by a practical project realized in a team of 2-3 students based on a novel idea. The purpose of the project is to implement a software component and to write a scientific research paper that can be presented at a conference.

Advanced Topics in Distributed Systems

Large-scale distributed systems are at the core of our daily life: they enable applications such as web search, online shops and social networking.

These highly-optimized applications run on thousands of machines inside datacenters, processing petabytes of data to answer user queries in well under one second. This course tackles all aspects related to building such large-scale applications, discussing recent research advances the field.

The course first discusses infrastructure topics including virtualization, novel networking protocols and topologies, distributed file systems, distributed caching and shared memory. Modern frameworks for building distributed processing are discussed next - these include the popular Map/Reduce paradigm, as well as more recent advances such as in-memory processing (Spark). Finally, large scale distributed systems are discussed including web-search, key-value stores and distributed databases. A special emphasis is placed on practical skills: the course is accompanied by lab sessions where all the topics discussed in the course are made available to students for experimentation.

Dependable Systems

The course presents advanced system architectures for dependable systems. Students have a chance to gain a solid foundation (theory and practical) on dependability as a property of systems capable to provide high reliability, safety, availability, security and ease of maintenance. During this class students will familiarize with modern paradigms related to the development of reliable systems, characteristics of reliable systems, models and architectures specific to these types of systems, analytical techniques for addressing the confidence of systems, error analysis, management of risk factors, methods for increasing reliable systems, detection and recovery, evaluation and reconfiguration in software architectures. After going through this course students will be able to: 1. understand the principles and models of developing reliable systems, 2. to critically analyze system reliability and 3. to be able to understand architectural solutions to increase the dependability provided by systems and applications.

Research activities

The research activity in Parallel and Distributed Systems is focused on topics with high complexity, usually related to the research projects of the teaching staff involved in this program, and to projects developed together with partner companies and institutions. It targets the study of the state of the art and the development of original solutions of problems, based on instruments that are specific to parallel and distributed systems. Also, the development of creativity, personal initiative and pro-activity, and the update of professional, scientific, and technical knowledge are among the most important objectives. The activity will include: discussion and selection of the research subjects, documentary research and elaboration of the state of the art, problem identification, solution development, solution evaluation and presentation of results, elaboration of the final research report.