LINKÖPING UNIVERSITY
Department of Computer and Information Science (IDA)
Prof. Christoph Kessler

TDDD25 Distributed Systems / *Distribuerade System*
Reading directions / *Läsanvisningar*

The materials discussed at the lectures will be *directly* covered by the written examination; this material you have to understand and, at the same time, know how to apply to solve problems.

In order to prepare for the exam, you have to study:

1. **Lecture notes**: *all the material* presented in the lecture notes may appear in the examination.

2. **Textbook**: you find below chapters and paragraphs from the two recommended course books which are related to the examination topics and serve for a better understanding of the material.

Notice: for each of these textbooks, there are several topics discussed at the lectures, which are ***not*** covered in the book. The lecture notes should be sufficiently explicit to understand them.

- **2.A** if you use the book by **van Steen and Tanenbaum, "*Distributed Systems*", 4th edition 2023**:
  (recommended from 2026)

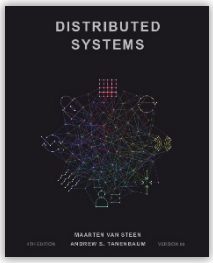

Chapter 1. Introduction
1.1 From networked systems to distributed systems
1.2 Design goals (except the part on cryptography pp. 22-23)
1.3 A simple classification of distributed systems (especially pp. 40-42 on enterprise application integration and 1.3.3 Pervasive systems)
1.4 Pitfalls

Chapter 2. Architectures
2.1 Architectural styles
2.2 Middleware and distributed systems
2.3.1 Simple client-server architecture
2.3.2 Multitiered architectures
2.4 Symmetrically distributed system architectures (peer-to-peer systems)
2.5.1 Cloud computing
2.5.2 Edge-cloud architecture

Chapter 3. Processes
(3.1 and 3.2 should be known from the Operating Systems course already.)
3.2.4 Application of virtual machines to distributed systems
3.3.3 Client side software for replication transparency
3.4 Servers
3.5 Code migration: 3.5.1-3.5.2

Chapter 4. Communication
(4.1.1, 4.3.1 and 4.3.3 may be useful background reading if you have not taken a computer networks course yet)
4.1.2 Types of communication
4.2 Remote procedure call
4.3.2 Advanced transient messaging: ZeroMQ (until p. 218). Not MPI, this will be covered in great detail in TDDE65.
4.3.3 part on Message brokers (pp.224-226)
4.4 Multicast communication

Chapter 5. Coordination
5.1 Clock synchronization
5.2 Logical clocks
5.3 Mutual exclusion 5.3.1-5.3.4 and Note 5.5 (pp. 278-280)
5.4 Election algorithms: 5.4.1 and 5.4.2

Chapter 6. Naming
6.2.3 Distributed hash tables

Chapter 7. Consistency and Replication
7.1 Introduction
7.2.1 Data-centric consistency models: sequential and causal consistency
7.5 Consistency protocols: 7.5.1 – 7.5.2

Chapter 8. Fault Tolerance
8.1 Introduction to fault tolerance, esp. failure models, synchronous/asynchronous systems, failure masking by redundancy
8.2 Process resilience 8.2.1-8.2.2 and 8.2.5 (Byzantine agreement)
8.3 Reliable client-server communication
p.524-526 Message ordering

- **2.B** if you use the book by **Coulouris et al., "*Distributed Systems – Concepts and Design*" (5th edition, 2011)** (recommended until 2025, still useful in 2026 as alternative to van Steen and Tanenbaum):

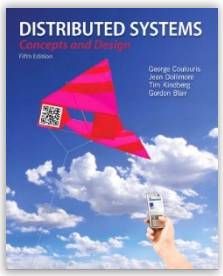

Chapter 1. Characterization of Distributed Systems
1.1 Introduction
1.2 Examples of Distributed Systems
1.5 Challenges

Chapter 2. System Models
2.1 Introduction
2.3 Architectural models
2.4 Fundamental Models (without security model)

Chapter 3. Networking and Internetworking
3.4.6 TCP and UDP

Chapter 4. Interprocess Communication
4.3.0 External data representation and marshalling,
4.3.1 CORBA common data representation and marshalling

4.3.4 Remote object references
4.5 Network virtualization: Overlay networks

Chapter 5. Remote Invocation
5.1 Introduction
5.2 Request-reply Protocols
5.3 Remote Procedure Call (without Sun RPC case study)
5.4 Remote Method Invocation

Chapter 6. Indirect Communication
6.1 Introduction
6.2.2 Group Communication; Implementation Issues
6.3 Publish-subscribe systems

Chapter 8. Distributed Objects and Components
8.1 Introduction
8.3 Case Study: CORBA

Chapter 10. Peer-to-Peer Systems
10.1 Introduction
10.2 Napster and its Legacy
10.3 Peer-to-Peer Middleware
10.4 Routing overlays

Chapter 14. Time and Global States
14.1 Introduction
14.2 Clocks, Events, and Process States
14.3 Synchronizing Physical Clocks
14.4 Logical Time and Logical Clocks
14.5 Global States

Chapter 15. Coordination and Agreement
15.1 Introduction
15.2 Distributed Mutual Exclusion (without Maekawa's algorithm)
15.3 Elections
15.4.3 Ordered multicast (without implementing causal ordering, overlapping groups, multicast in synchronous and asynchronous systems)
15.5.3 The Byzantine Generals Problem in Synchronous Systems

Chapter 18. Replication
18.1 Introduction
18.5 Transactions with Replicated Data (without virtual partition algorithm)

Chapter 20. Distributed Multimedia Systems
20.6.2 BitTorrent

Some other material related to the course topic:

* http://www.omg.org/ (on OMG and CORBA).