

# Distributed Systems II

TDA297(CTH), DIT290 (GU)

---

LP3 2012

7.5 hec

<http://www.cse.chalmers.se/edu/course/TDA297>



# The Teachers

---

- Instructor: Philippas Tsigas
  - E-mail: [tsigas@chalmers.se](mailto:tsigas@chalmers.se)
  - Phone: +46-31-772 5409
  - Office: 5106
  - Office hours: By appointment
- Assistants:
  - Farnaz Moradi – [moradi@chalmers.se](mailto:moradi@chalmers.se)
  - Ioannis Nikolakopoulos – [ioaniko@chalmers.se](mailto:ioaniko@chalmers.se)
  - Office hours: To be announced

# The Teachers

- Farnaz Moradi  
[moradi@chalmers.se](mailto:moradi@chalmers.se)  
Phone: 772 1045  
Room: 5111
- Ioannis  
Nikolakopoulos  
[ioaniko@chalmers.se](mailto:ioaniko@chalmers.se)  
Phone: 772 5720  
Room: 5107
- Office hours: To be announced





# Course Representatives

---

Johansson	William	<a href="mailto:wiljoh@student.chalmers.se">wiljoh@student.chalmers.se</a>
Nowdehi	Nasser	<a href="mailto:nowdehi@student.chalmers.se">nowdehi@student.chalmers.se</a>
Svensson	Martin	<a href="mailto:martsven@student.chalmers.se">martsven@student.chalmers.se</a>
Söderberg	Mathias	<a href="mailto:soderbem@student.chalmers.se">soderbem@student.chalmers.se</a>
Walulya	Ivan	<a href="mailto:walulya@student.chalmers.se">walulya@student.chalmers.se</a>



# Reading

---

- 4th Edition of the book: "Distributed Systems: Concepts and Design"
  - written by: George Coulouris, Jean Dollimore and Tim Kindberg
  - published by Addison-Wesley, ISBN 0-321-26354-5
- Some extra material: Distributed Algorithms (Notes) + papers.



# Description

---

Distributed systems are **popular and powerful computing paradigms**. Their importance increases as networked computers become more common than freestanding ones, especially since many different types of computers can be found in networks. In this course we will see the points of **inherent difference and strength of distributed systems** compared with sequential or strongly-coupled systems; consequently, we will also **study the issues and problems** that have to be addressed and solved efficiently for these differences to be taken advantage of, so that the system retains its strength and high potential.



# Goals

---

- 1. Replication; The advantages and costs of replicating data:
  - Potential improvement in response times and reliability
  - Extra communication costs involved in keeping data consistent.
- 2. Fault-tolerant Agreement in Distributed Systems:
  - (a very special and significant problem, since it is a key issue in most synchronisation and coordination problems in distributed systems)
  - Study of the basic impossibility results and discuss their implications
  - Proceed with solutions and protocols for systems with certain strengths and design structures.



# Goals cont.

---

- 3. Resource Allocation.
- 4. Distributed algorithms: How to design and analyse distributed algorithms?
- 5. Sensor Networks.
- 6. Denial of Service Attacks.





# This is an advanced course

---

- No list of things to learn
  - The book and assignments cover most exam material
- Find out what you need
  - Lectures
  - The book
  - Internet
- Lectures and assignments in parallel
  - Don't wait for the lectures before you start with the assignments

# Approximate Course Schedule

Week (Chalmers week/LV)	Monday	Wednesday	Friday
<b>3 (1)</b> (Jan 17-21)	<b>Lecture 1:</b> Introduction	<b>Lecture 2:</b> Broadcasting	<b>Assignment Lecture 1:</b> Introduction to Lab 1
<b>4 (2)</b> (Jan 24-Jan 28)	<b>Lecture 3:</b> Broadcasting & Replication	<b>Lecture 4:</b> Replication	<b>Lecture 5:</b> Replication & Quorum Consensus
<b>5 (3)</b> (Mar 31-Feb 4)	<b>Assignment Lecture 2:</b> Introduction to Lab 2 + Questions and answers for all labs.	No Lecture – Work on assignments	<b>Lecture 6:</b> Distributed denial of service attacks
<b>6 (4)</b> (Feb 7-11)	<b>Lecture 7:</b> Distributed Transactions & Concurrency Control in Distributed Transactions <b>Deadline for lab 1 at 23:59</b>	CHARM – No Lecture	<b>Assignment Lecture 3:</b> Questions and answers for all labs.
<b>7 (5)</b> (Feb 14-18)	<b>Lecture 8:</b> Atomic Commit protocols & Byzantine General	<b>Lecture 9:</b> Distributed Algorithms	<b>Assignment Lecture 4:</b> Sensor Networks + Introduction to Lab 3 + Questions and answers for all labs.
<b>8 (6)</b> (Feb 21-25)	<b>Lecture 10:</b> Grouping and routing in sensor networks <b>Deadline for lab 2 at 23:59</b>	<b>Assignment Lecture 5:</b> Questions and answers for all labs.	<b>Lecture 11:</b> Mutual Exclusion and Resource Allocation. Dining Philosophers
<b>9 (7)</b> (Feb 28-Mar 4)	<b>Lecture 12:</b> Generalization of the Dining Philosophers	<b>Lecture 13:</b> Drinking Philosophers and Efficient Resource Allocation	No Lecture <b>Deadline for lab 3 at 23:59</b>
<b>10 (8)</b> (Mar 7-11)	<b>Assignment presentations:</b> Presentations and Demos of the 3rd lab.	<b>Lecture 14:</b> Efficient Resource Allocation continued	<b>Lecture 15:</b> Closing
	<b>EXAM:</b> Monday the 14th of March 2011 in the V building 14:00-18:00		



# Examination

---

- 1st Lab due: around 10th of February
- 2nd Lab due: around 17th of February
- 3rd Lab due: around 1st of March
- Final examination: 05th of March



# Resources

---

## **Full support page for the Coulouris' book:**

- <http://www.cdk4.net/>
- <http://www.cdk5.net/>

## **Slides:**

- At homepage after lecture
- Last years slides available, use them as a reference point before the lecture:  
<http://www.cse.chalmers.se/edu/year/2010/course/TDA297/>



# Resources cont.

---

## **1st Assignment:**

- Distributed bulletin board
- <http://www.cse.chalmers.se/edu/course/TDA297/labs/lab1.html>

# Resources cont(2).



## 2nd Assignment:

- Reliable and ordered multicast
- <http://www.cse.chalmers.se/edu/course/TDA297/labs/lab2.html>

## 3rd Assignment:

- Routing in Sensor Networks
- <http://www.cse.chalmers.se/edu/course/TDA297/labs/lab3.html>







# Why Distributed Systems?

---

# Suggestions?





# Because they are there

---

## Distributed Applications

- Automated Banking Systems
- Tracking Roaming Cellular Telephones
- Air-Traffic Control Systems
- Fly-by-wire Systems
- The World-Wide- Web



# Because they have many +

---

## Network vs. Centralised Systems

- Cycles always available
- Incremental growth
- Independent failure
- Heterogeneity
- Increased Autonomy
  - purchasing
  - management
  - software



# Ingredients of D.S.

---

- Multiple computers
- Interconnections

Shared states

# Asynchronous Systems

- Communication is reliable but incurs potentially unbounded delays





# Distributed Computation

---

A distributed computation is a *single execution* of a *distributed program* by a collection of *processes*. Each sequential process generates a *sequence of events* that are either *internal events*, or *communication events*

The *local history* of process  $p_i$  during a computation is a (possibly infinite) sequence of events  $h = e_i^1, e_i^2, \dots$

The *global history* of a computation is a set  $H = \bigcup_{i=1}^n h_i$



# State of the computation

---

- Imagine stopping a distributed computation by **stopping all** of its **processes simultaneously**.
- The combined **states of** each of the **processes, plus** the contents of the **messages in transit** between processes will then tell us the exact **global state of the computation**.



# State of the computation

---

- The problem with our **asynchronous system** is that there is **no such thing as simultaneity**.
- The concept of a system state is still used in analysis

The local state of process  $p_i$  after executing event  $e_i^k$  is denoted  $\sigma_i^k$ . The initial local state of  $p_i$  is  $\sigma_i^0$ .

A **global state** is a collection of local states  $\Sigma = \bigcup_{i=1}^n \sigma_i$ .



# Cause and Effect

---

- No simultaneity in an asynchronous system  
=> we need something else in order to study the events in such a system.
- We do have the notion of cause and effect, however.
- If one event  $e_i$  caused another event  $e_j$  to happen, then  $e_i$  and  $e_j$  could never have happened simultaneously:

$e_i$  happened before  $e_j$





# Causal Orders

---

We define a binary relation  $\rightarrow$  over events, such that

1. If  $e_i^k, e_i^l \in h_i$  and  $k < l$ , then  $e_i^k \rightarrow e_i^l$  ■
2. If  $e_i = \text{send}(m)$  and  $e_j = \text{receive}(m)$ , then  $e_i \rightarrow e_j$  ■
3. If  $e \rightarrow e'$  and  $e' \rightarrow e''$ , then  $e \rightarrow e''$



When  $e \rightarrow e'$ , we say  $e$  *causally precedes*  $e'$  or  $e$  *happened before*  $e'$ .

We define *concurrent* as  $e \parallel e' \equiv \neg(e \rightarrow e' \vee e' \rightarrow e)$

# Space Time Diagram

A distributed computation is a *partially ordered set* (poset)  $\gamma = (H, -)$ .

A distributed computation can be depicted in a *space-time diagram*:

