DTU Compute
Department of Applied Mathematics and Computer Science

Distributed Systems (02220)

Course Presentation

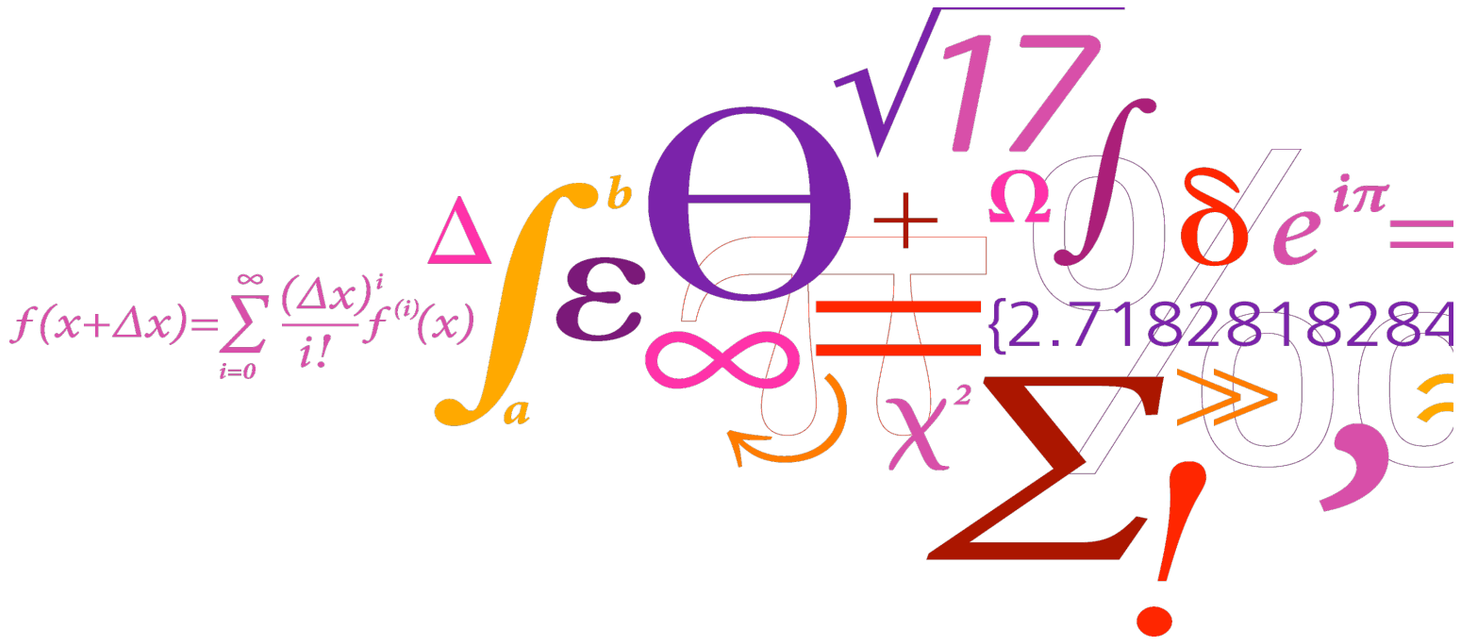

Nicola Dragoni

- Associate Professor in Distributed Systems and Security in Distributed Systems

  DTU Compute

  Technical University of Denmark (DTU), Denmark

- Professor in Computer Engineering

  Centre for Applied Autonomous Sensor Systems (AASS)

  Örebro University, Sweden

“Salient Ingredient”: Communication

- Do NOT let people guess!
- **Meet**
- **Talk**
- **Share**
- **Listen**
- **Ask**
- **Help**

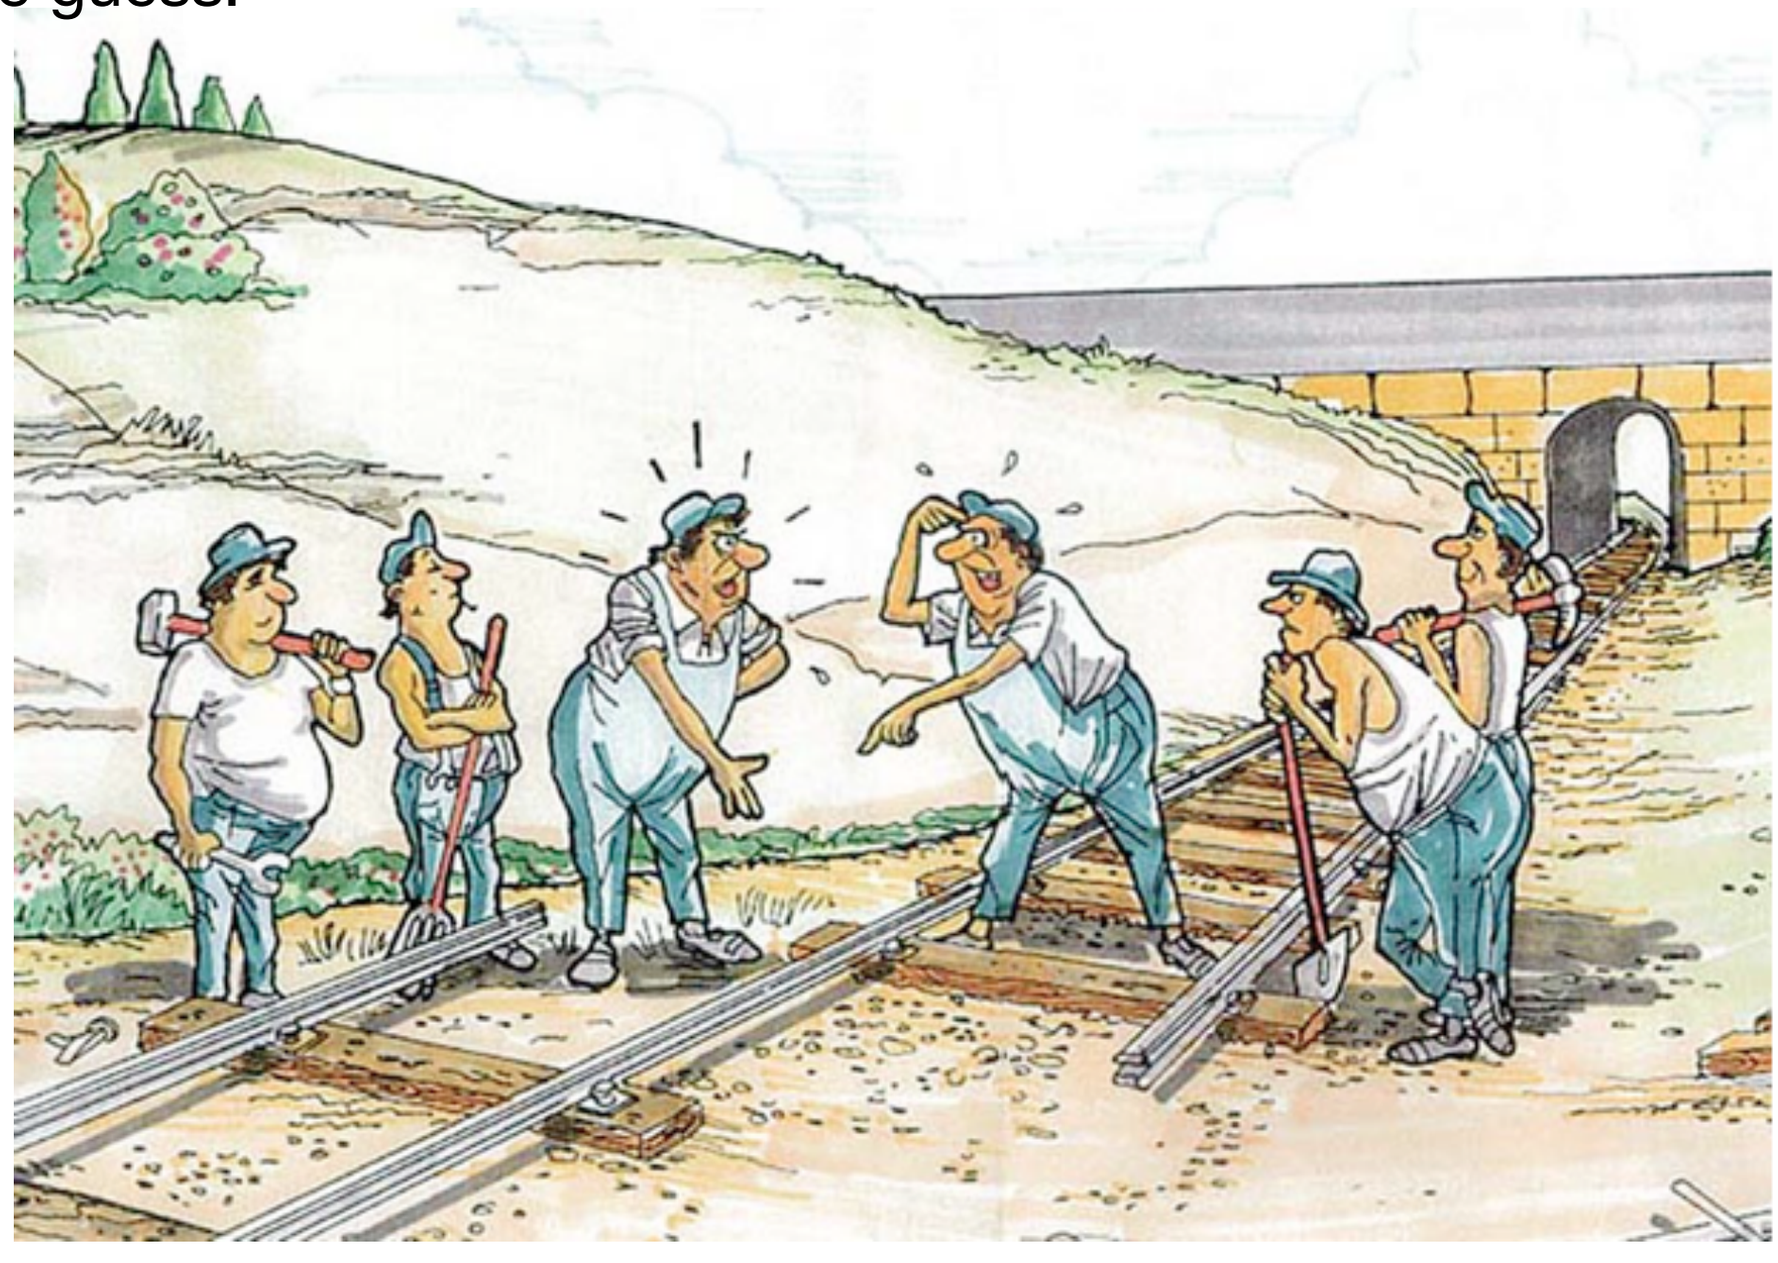

Communication: How and When

- **Email**: ndra@dtu.dk
  - For a **quick reply**, include **[02220]** in the **SUBJECT** of your email!
  - If **time-consuming issue**, ask for a **meeting**

- **AFTER** each lecture

Distributed Systems?

- **Networks of computing devices** are everywhere:
  - mobile phone networks
  - social networks
  - campus networks
  - home networks
  - Internet
  - ...

Distributed System

- A possible definition: *a distributed system is a system in which hardware or software components located at networked **devices** communicate and coordinate their actions only by passing messages*

- Networked devices (i.e., devices that are connected by a network) may be spatially separated by any distance:
  - separate continents
  - same building
  - same room
  - ...

> *"You know you have a distributed system when the crash of a computer you've never heard of stops you from getting any work done."*
>
> [Leslie Lamport]

Why Distributed Systems?

- The motivation for constructing and using distributed systems stems from a desire to share resources

- Resource = abstract term that characterises the range of things that can be usefully be shared in a networked computer system:

  - Hardware components: disks, printers, ...

  - Software entities: files, databases, and data objects of all kinds

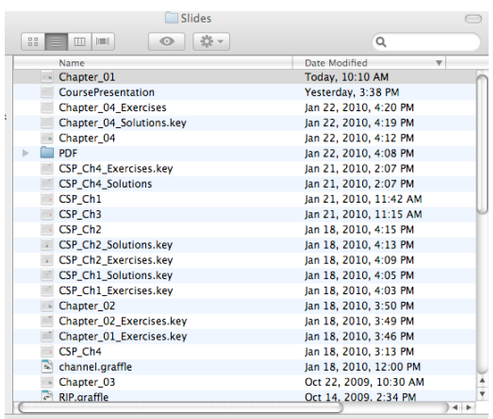

Selected Application Domains and Associated Networked Applications

Finance and commerce	eCommerce e.g. Amazon and eBay, PayPal, online banking and trading
The information society	Web information and search engines, ebooks, Wikipedia; social networking: Facebook and MySpace.
Creative industries and entertainment	online gaming, music and film in the home, user-generated content, e.g. YouTube, Flickr
Healthcare	health informatics, on online patient records, monitoring patients
Education	e-learning, virtual learning environments; distance learning
Transport and logistics	GPS in route finding systems, map services: Google Maps, Google Earth
Science	The Grid as an enabling technology for collaboration between scientists
Environmental management	sensor technology to monitor earthquakes, floods or tsunamis

Example: The Internet

- A vast interconnected collection of computer networks of many different types
  - Programs running on the computers connected to it interact by passing messages, employing a common means of communication (Internet protocols)

- A very large distributed system
  - It enables users, *wherever they are*, to make use of open-ended services (WWW, email, file transfer, multimedia services, ...)

A Map of the First Internet (ARPANET, ~1971)

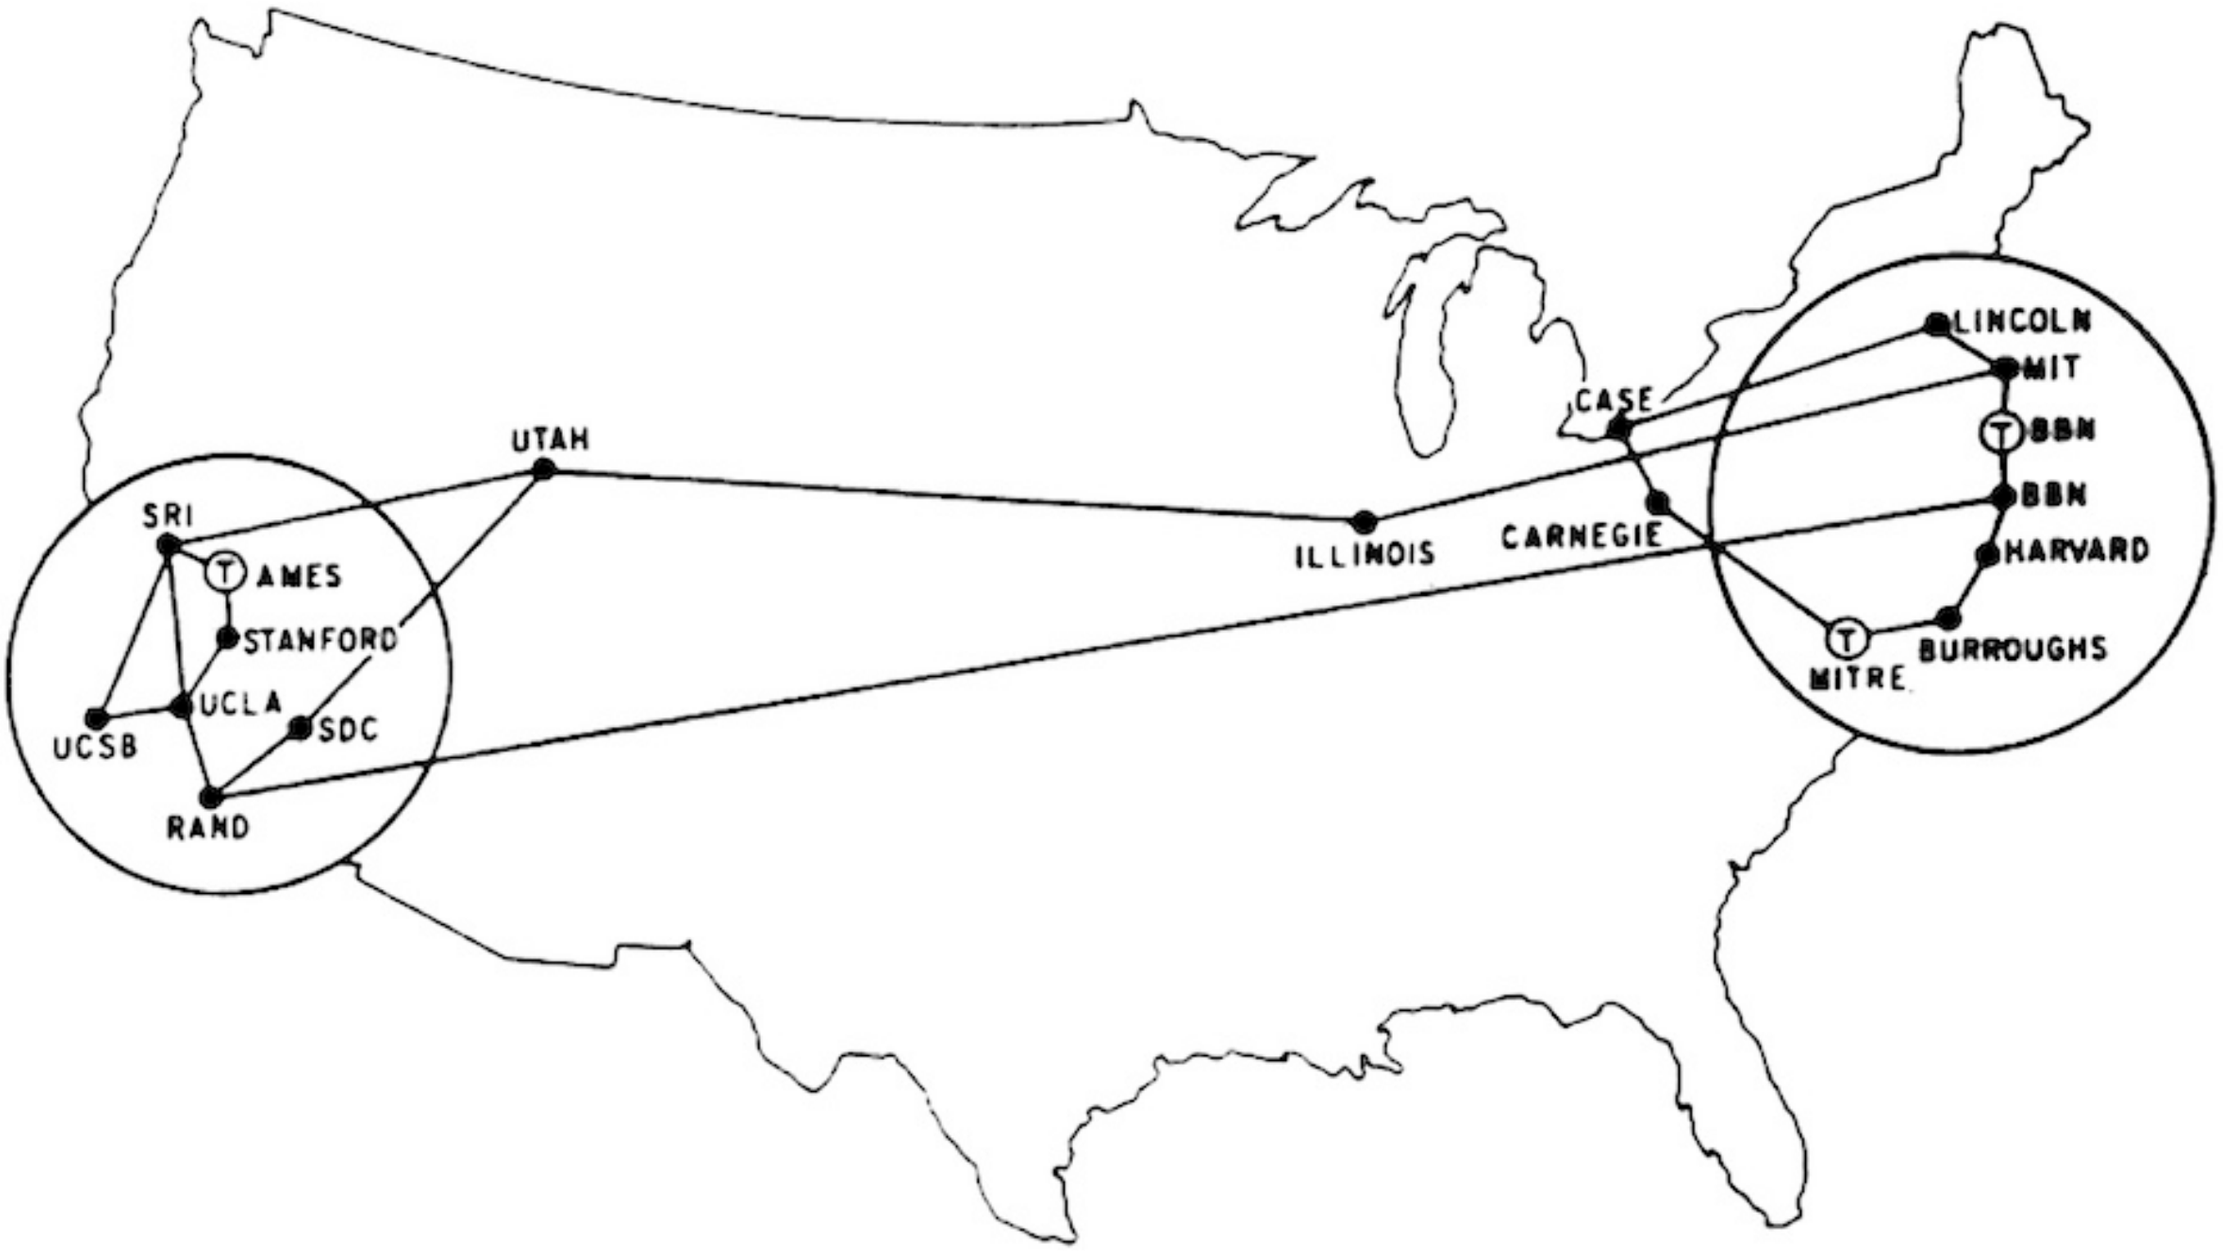

Web (November 2003)

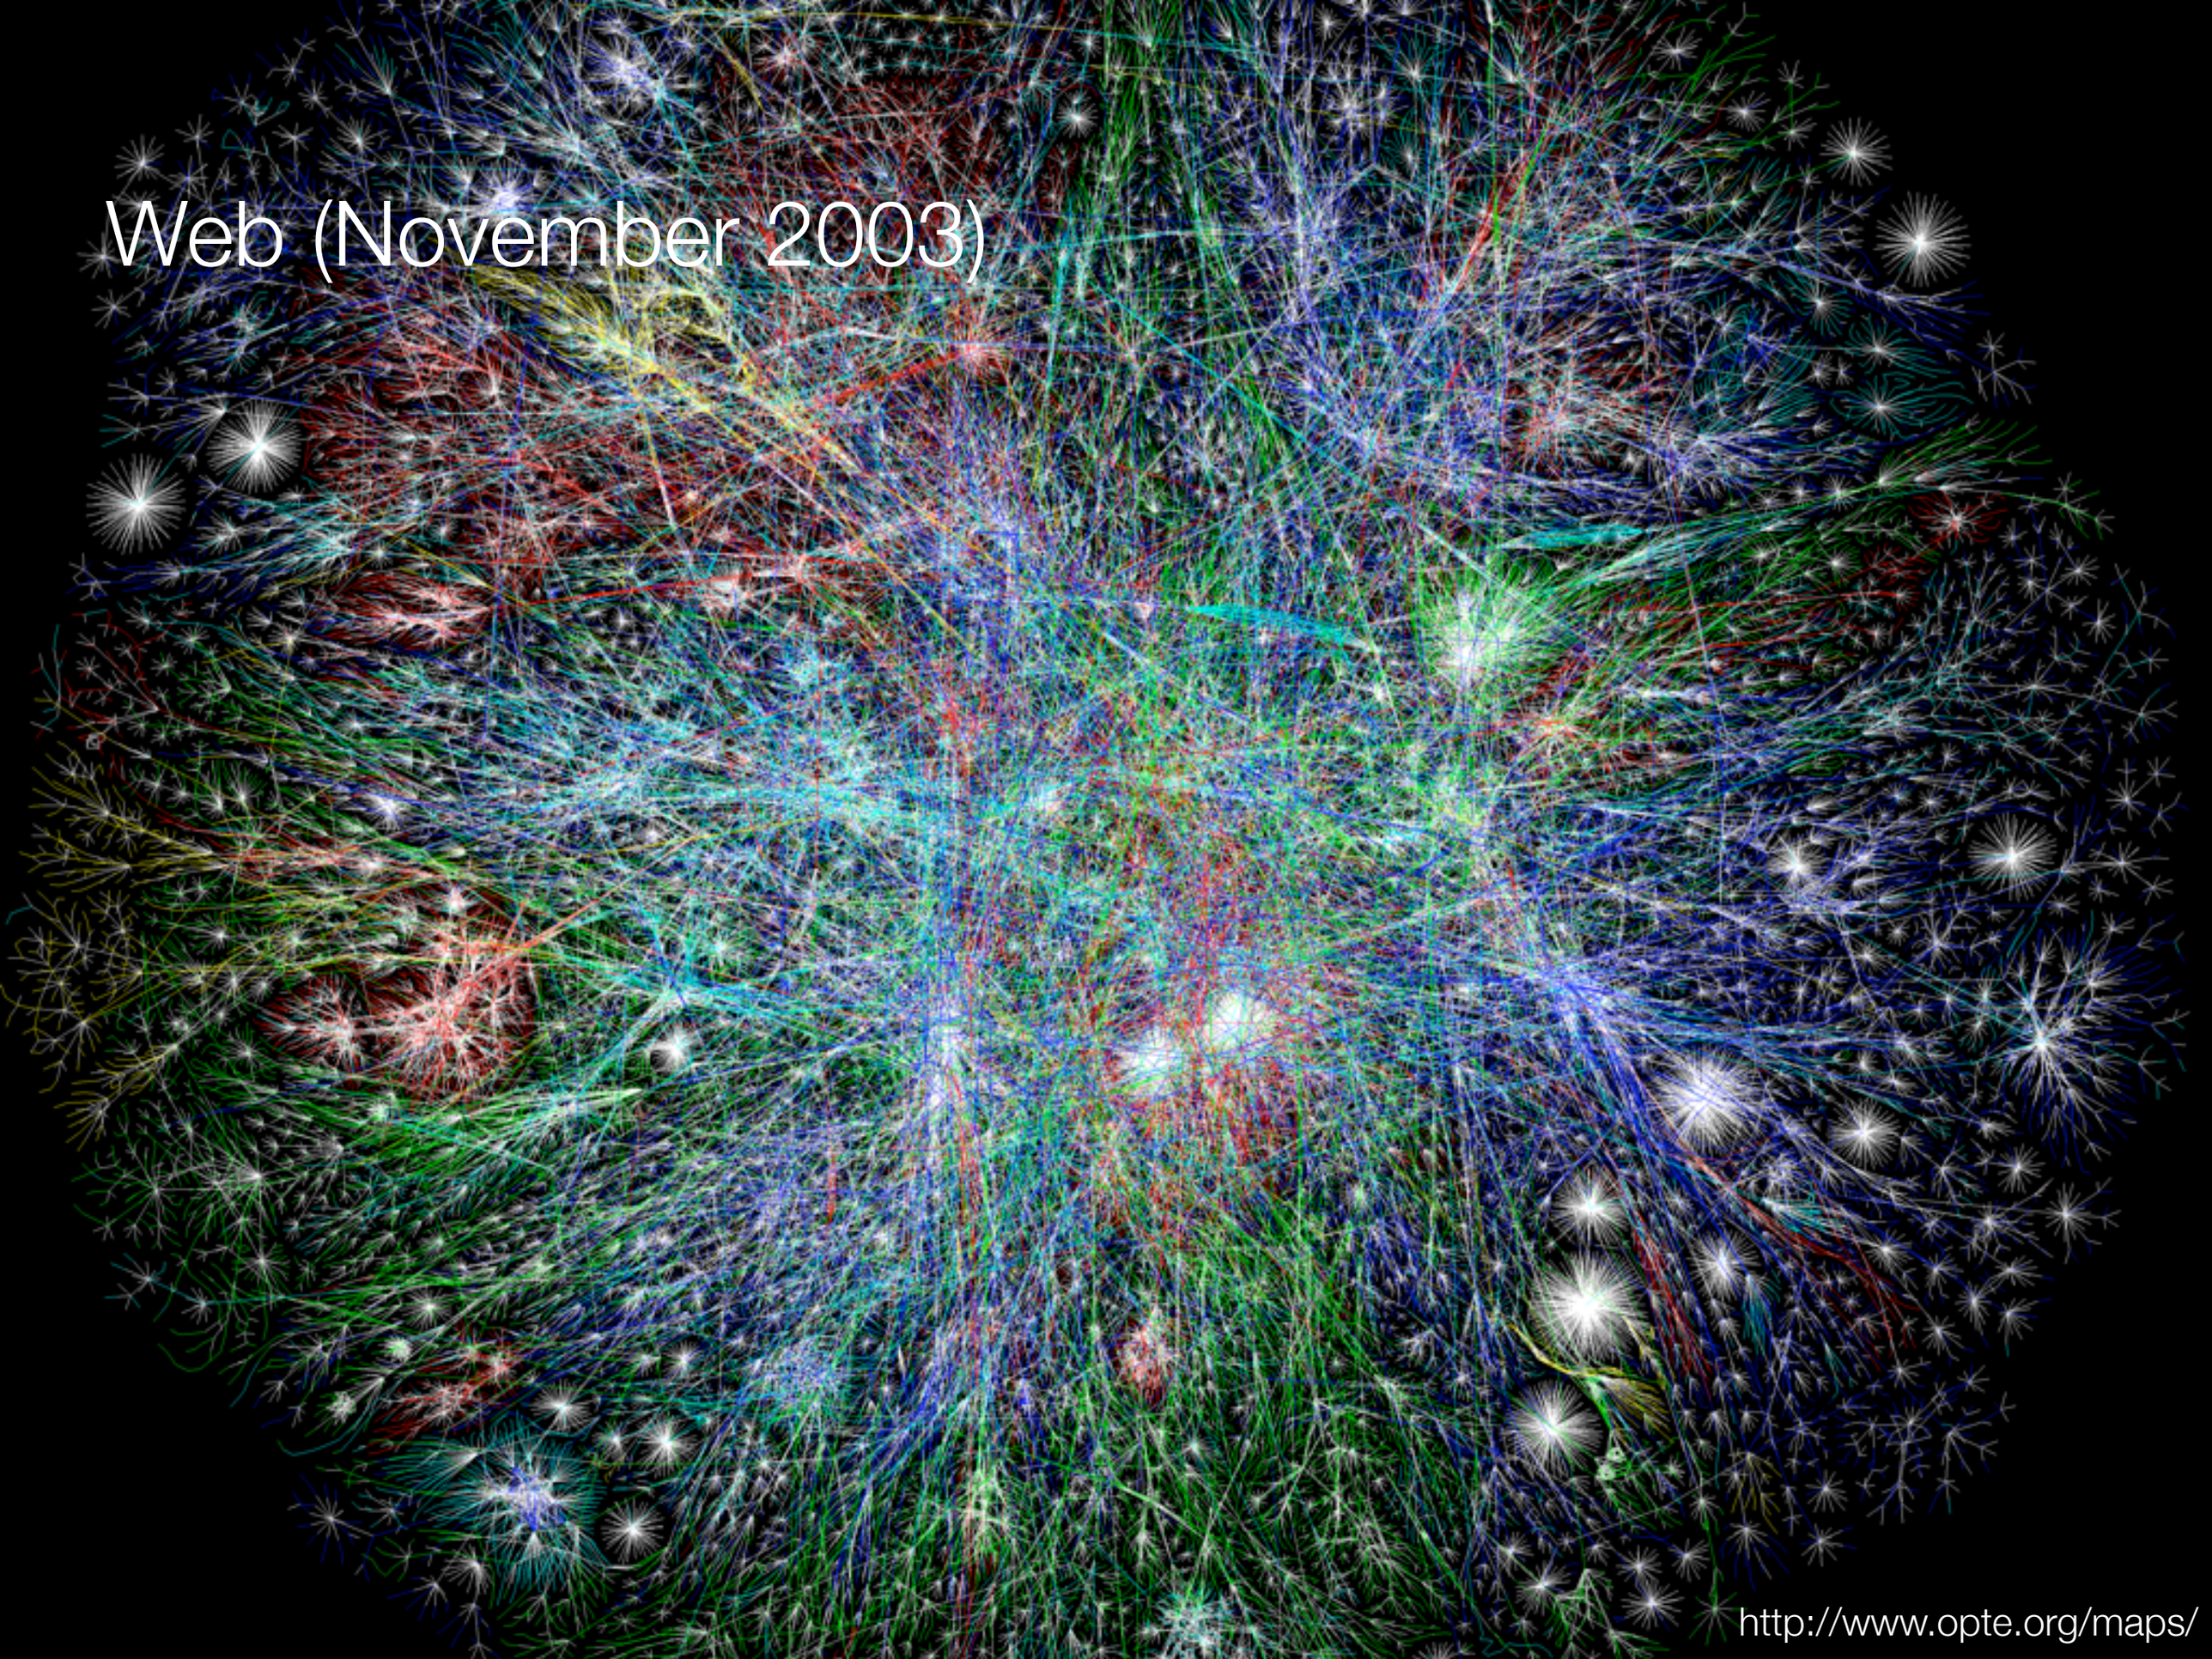

http://www.opte.org/maps/

Facebook (December 2010)

Facebook (January 2014)

Distributed Systems

02220

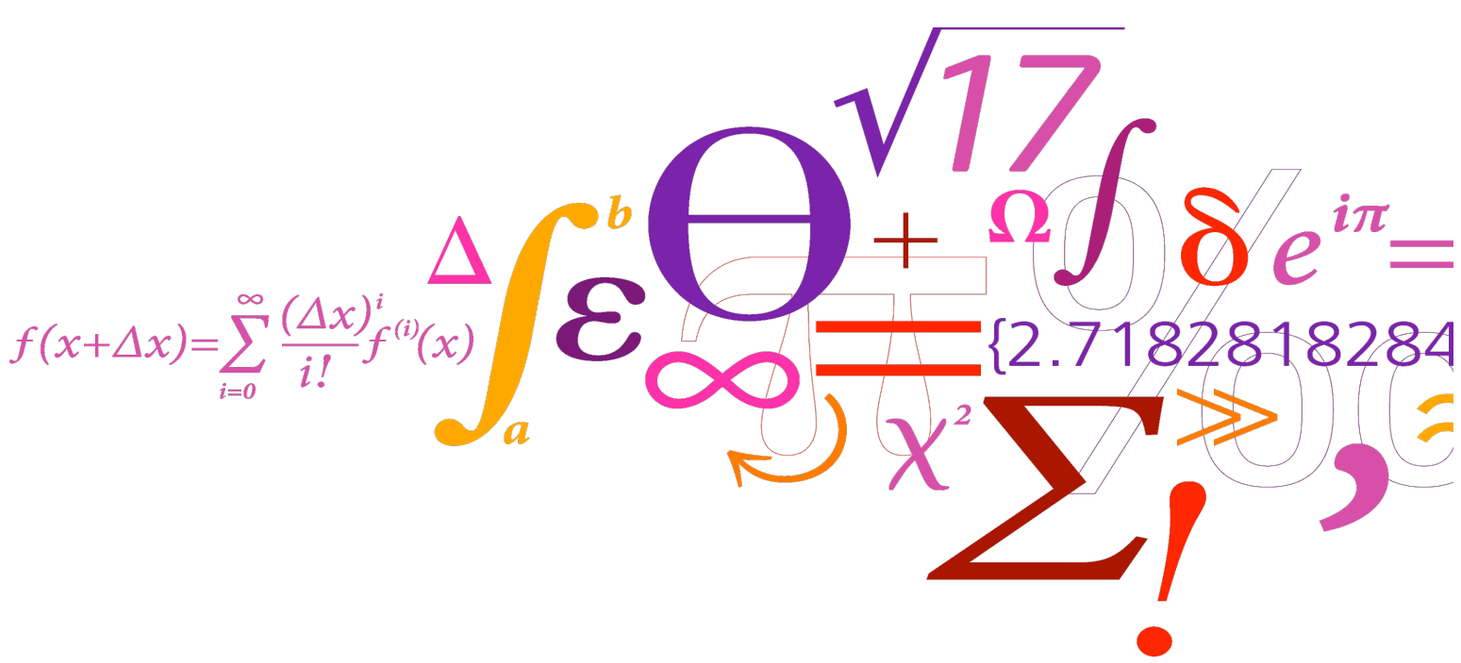

Web Site: http://www2.imm.dtu.dk/courses/02220

02220 - DISTRIBUTED SYSTEMS

General course objectives: to illustrate the types of problems which arise and methods used in the design and analysis of systems of interconnected computing devices

Responsible: Nicola Dragoni, 322, 118, ndra@dtu.dk

Prerequisites:

- basic understanding of concurrent systems and the problems which arise in systems with concurrent activities
- basic knowledge of computer networks

Some of these may be attained through course 02158 - Concurrent programming (formerly 02152 - Concurrent systems)

Point (ECTS): 7.5

Course type: BSc/MSc - Advanced Course

Not applicable together with: 02221, 02222

Duration of course: 13 weeks

Activity plan (2017): a *tentative* activity plan is available here. *N.B.: the plan is preliminary and may be changed at any time. Check it regularly*!

Type of assessment: the grade is based on a written report and a final written exam; the partial grade for the report can be carried over to the following (autumn) term

Exam duration: 4 hours

Aid: no aid

Evaluation: 7 step scale, internal examiner

Textbook:
(main reference)
George Coulouris, Jean Dollimore, Tim Kindberg, Gordon Blair
Distributed Systems - Concepts and Design
Addison Wesley, 2011 (Fifth Edition)
Available from Polyteknisk Boghandel (or any other bookshop...)

Aim and Prerequisites

- Aim of the course:

To illustrate the types of problems which arise and methods used in the design and analysis of systems of interconnected computing devices

- Prerequisites:

  - basic understanding of concurrent systems and the problems which arise in systems with concurrent activities

  - basic knowledge of computer networks

Organization (Main Topics)

- **Foundations** (3 lectures)
  - Introduction (today)
  - Models
  - Basic Protocols

- **Guest Lecture** (1 lecture)
  - netcompany

- **Communication** (1 lecture)
  - Interprocess Communication
  - Remote Invocation

- **Middleware** (1 lecture)
  - P2P computing

- **Distributed Algorithms** (5 lectures)
  - Logical Time
  - Global States
  - Coordination and Agreement

- **Hot Topics** (1 lecture)
  - TBA

- **Final Thoughts** (1 lecture)
  - exam, feedback, project, …

Activity Plan

- The activity plan is available on the 02220 Web site:
  - www2.imm.dtu.dk/courses/02220/2017/DS_scheduling_2017.html

This plan is preliminary and may be changed at any time.

Week	Day			Lecturer	Lecture Contents	Where	Comments and Material
					Presentation of the Course		
01	Thursday	02	Feb.	Nicola	- 02220: -- activity plan -- exam -- project -- DTU green challenge -- Introduction to pervasive computing	341, aud. 21	Introduction to the course, including information about the course activity plan, book, teaching assistants, project and final assessment
					Introduction to Distributed Systems		
					- Architectural Models		This lecture is based on **chapters 1 and 2** of the book

- **N.B.: the plan is preliminary and may be changed at any time**

Textbook

- G. Coulouris, J. Dollimore, T. Kindberg, G. Blair
  Distributed Systems - Concepts and Design
  Pearson Education, 2012 (Fifth Edition)

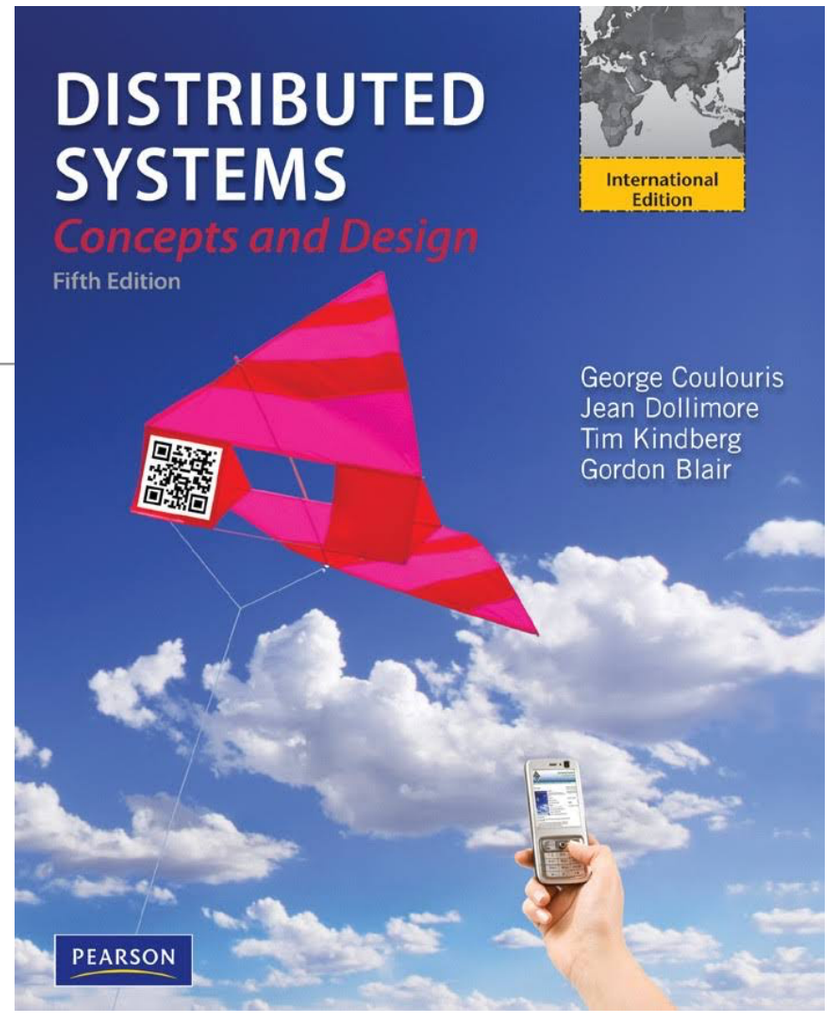

- Other books:

- A.D. Kshemkalyani, M. Singhal
  Distributed Computing: Principles, Algorithms, and Systems
  Cambridge University Press, 2011

- Robin Sharp
  Principles of Protocol Design
  Springer, 2008, ISBN 978-3-540-77540-9

Teaching Assistants

- **Adriana Messalli** (s151347@student.dtu.dk) —> projects
- **Michele De Donno** (s166213@student.dtu.dk) —> exercises

Assessment

- The **final grade** is based on
  - a **project** (**not mandatory, but <u>IT COUNTS</u> for the final grade**)
  - a final **written exam** (**mandatory**)

- **Overall assessment**: only **ONE final grade** (exam **+** project)
  - **<u>NO</u> official grades for the project**
- The partial grade for the project can be carried over to the following term

Calculation of the Final Grade

- **Final grade =**
  - IF **grade of written exam** ∈ **{-3, 0}** THEN **fail**
  - IF **grade of written exam** ∈ **{2, 4, 7, 10, 12}** THEN **exam + project**

+
  - if **no project**, **-2 steps** (e.g., 7 —> **2**, 12 —> **7**, 4 —> **0**)
  - if **bad project**, **-1 step** (e.g., 7 —> **4**, 12 —> **10**, 2 —> **0**)
  - if **ok/sufficient project**, **+0** (e.g., 7 —> **7**, 12 —> **12**, 2 —> **2**)
  - if **good project**, **+1 step** (e.g., 7 —> **10**, 10 —> **12**, 2 —> **4**)

Written Exam

- **Written** exam, **4 hours**, **no aid** (but don't forget to bring the brain.. :)

- **Any possible question**

  on any topic

  covered in the lectures

  (see activity plan)

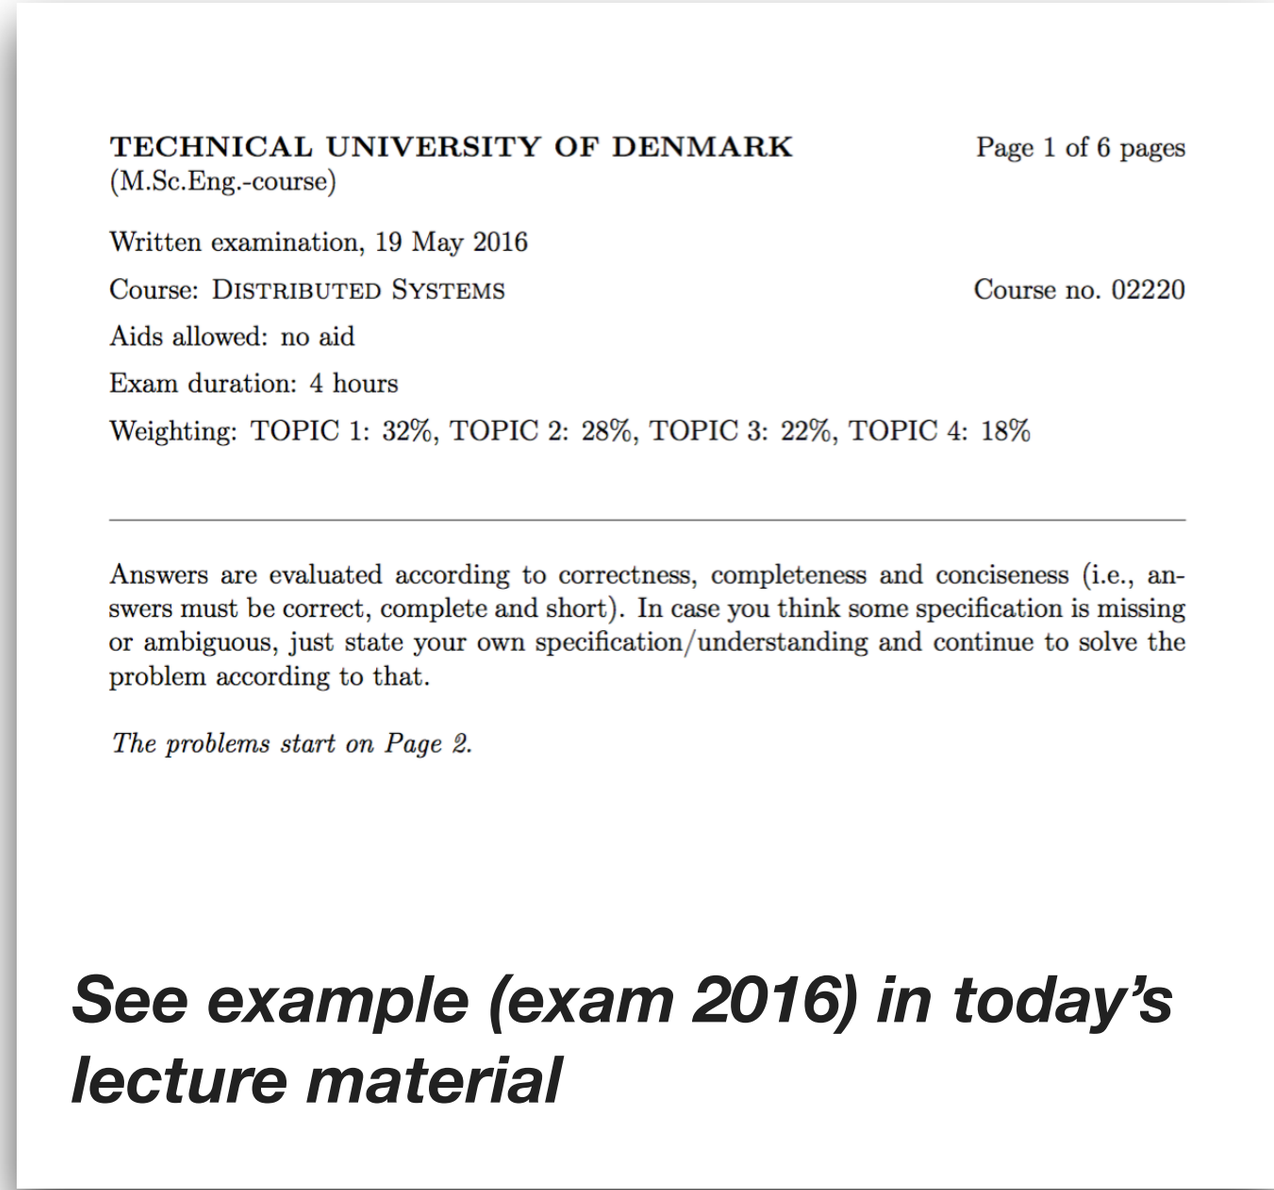

TECHNICAL UNIVERSITY OF DENMARK Page 1 of 6 pages
(M.Sc.Eng.-course)

Written examination, 19 May 2016

Course: DISTRIBUTED SYSTEMS Course no. 02220

Aids allowed: no aid

Exam duration: 4 hours

Weighting: TOPIC 1: 32%, TOPIC 2: 28%, TOPIC 3: 22%, TOPIC 4: 18%

Answers are evaluated according to correctness, completeness and conciseness (i.e., answers must be correct, complete and short). In case you think some specification is missing or ambiguous, just state your own specification/understanding and continue to solve the problem according to that.

The problems start on Page 2.

See example (exam 2016) in today's lecture material

Project

- Goal: **design** (and implementation) of a **distributed system**
- Groups of **2** or **3 students**
  - Self-organisation! If you need HELP to find project mates, send an email to **Adriana** (s151347@student.dtu.dk, include **[02220]** in the **SUBJECT** of your email!)
- **VERY IMPORTANT:** The project must be **OFFICIALLY ACCEPTED by me**
  - document (max 1 page) describing your **IDEA**
  - yes/no/yes_with_feedback from me

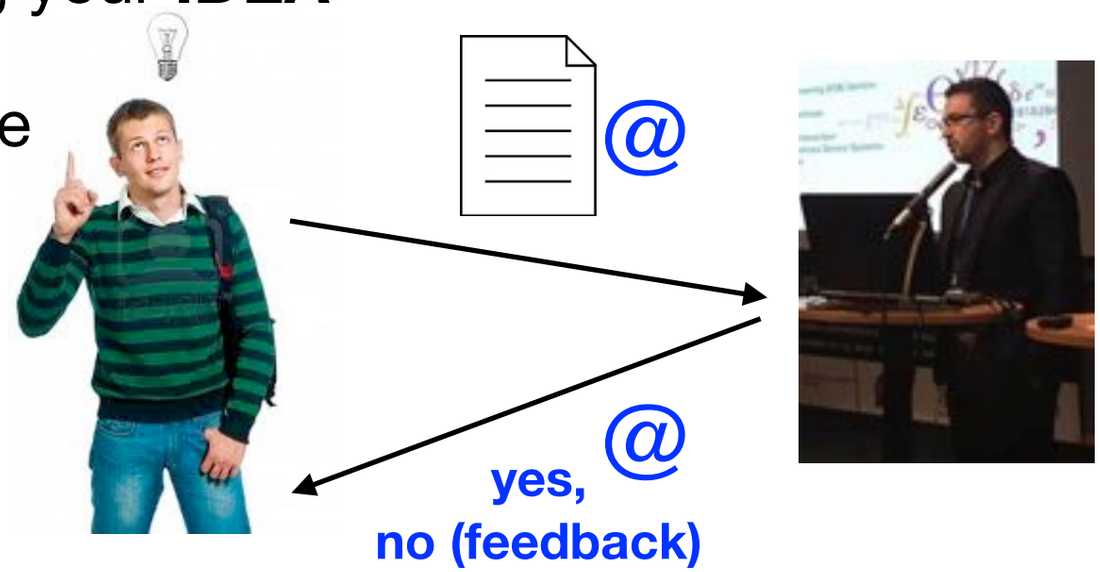

What is a **Good Project**?

- Clear **PROBLEM** definition
  - **WHAT** is the problem? **WHY** is it significative?
- **DESIGN** of a solution (**HOW** do you plan to solve the problem? What other solutions have been proposed? Focus on the **idea**, no technical details)
- **IMPLEMENTATION** details (**HOW** do you plan to implement the solution? Focus on **technology**, but no need to develop a working system)

- For EACH STEP you take in the design and implementation of the system:
  - Consider and evaluate alternatives
  - Motivate WHY you have decided for that step (in terms of the alternatives)
- Examples to illustrate problem and solution

Project Requirements

- **One fundamental document** (see material in activity plan):
  - Project **requirements** (i.e., ***rules you have to follow***)
- The requirements fall into **three categories**:

Legal Requirements [requirements which are *dictated by DTU's rules* and which have the status of *legal requirements*]

Mandatory Requirements for Technical/Scientific Reports [requirements which are *mandatory for any good technical/scientific report*, such as the ones you produce in 02220]

Recommended Requirements [requirements which *we advise you to follow* in order to present the results of the lab project in a good way]

What Topic for My Project?

WHAT

YOU LIKE?

(…provided that is compliant with the “Distributed Systems” context…)

What Topic for My Project?

DO WHAT

YOU LIKE!

(…provided that is compliant with the “Distributed Systems” context…)

Examples of (Good) Projects (2016)

- **Distributed File Sharing System**
  [Deals with the problem of sharing files over the Internet in an easy and secure manner, by establishing a direct connection to avoid storing the file on a third-party server]
- **Gas Payment from the Car**
  [… in particular how the communication between the car and the gas station, as well as with the cloud, should happen in order to build a system that is secure, fault-tolerant and scalable, and how this information should be stored]
- **FixIt! Notify Damages in Your City**
  [FixIt is an application where people can post new incidences in order to notify urban damages to the maintenance service, so they can be quickly repaired]
- **Programming Language-Independent Multiplayer Communication**
  [… to facilitate communication between different players in multiplayer games, regardless of the game engine (and of whether a game engine was even used) and regardless of the programming language used to develop the game]
- **Mobile-Based Application of the Assassin Game**
- **Automated Parking Guidance System using Vehicular Ad-hoc Networks**

Deadlines

- **Monday, May 08**

  PROJECT

  report (+ source code) have to be submitted electronically through **Campusnet**

- **Monday, May 15** (please, check the date on DTU portalen!)

  WRITTEN EXAM (GOOD LUCK! :-)

14	Thursday	04	May	Nicola	**Final Thoughts + Feedback**		
					- Feedback on the course, how to pass the 02220 exam, ... - Solutions to (requested) exercises - Feedback on projects	341, aud. 21	...
/	**Monday**	**08**	**May**	/	**Project Deadline: report (+ source code) have to be submitted electronically through Campusnet before midnight.**		
/	**Monday**	**15**	**May**	/	**WRITTEN EXAM: GOOD LUCK! :-)**		

Why Not... in This Course?

- Why not **advanced, new and cool technologies**?
- Why not **Web Services** and/or **Cloud Computing** and/or **Big Data**?
- Why not **security in distributed systems**?
- ... in other words...

WHY (mostly) FOUNDATIONS of DISTRIBUTED COMPUTING?

- ... because there are specific DTU courses on above topics...
- ... because **we live in exponential times**...!

Did you know? (Watch Youtube video…)

Key Facts

- **Technology** is growing at an **exponential rate**

- From the video:

  - The amount of **new technical information** is **doubling every 2 years**

    - For students starting a 4-years technical or college-degree, this means…

      … that **half of the TECHNOLOGY they learn in their first year of study might be outdated by their third year of study**

- **Goal is to focus on foundations in order to be ready for the technological changes!**