Networks and Distributed Systems

Olaf Landsiedel

Definition

- I believe you know what a network is…
- But, what is a **Distributed System?**
  - Have you ever seen one?
  - Have you ever used one?
- A Distributed System is characterized by?
  - Multiple devices
  - Connected by a network
  - Cooperating on some task

Examples

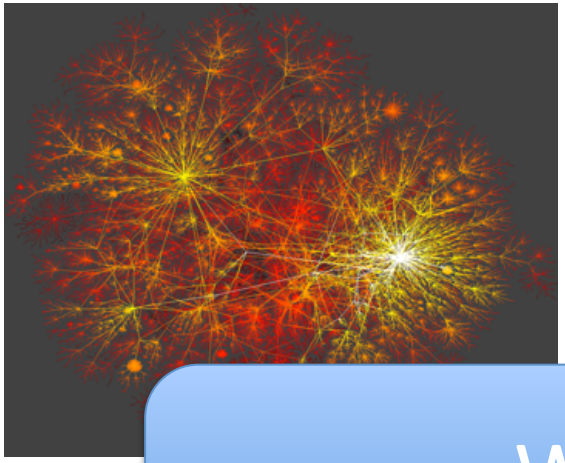

We teach you how to
build large-scale systems

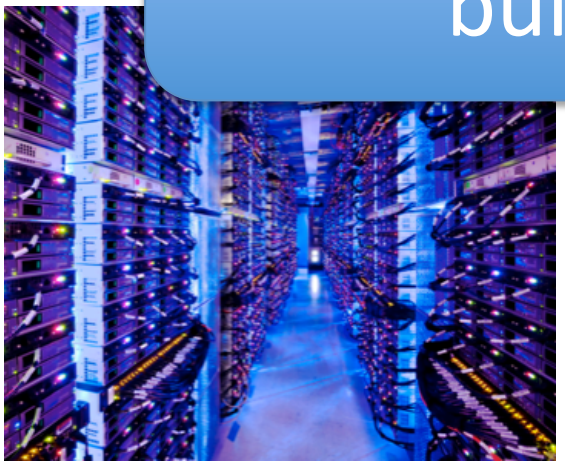

Cloud Computing / Data Center

Phone Network

Power Grid

One more Example

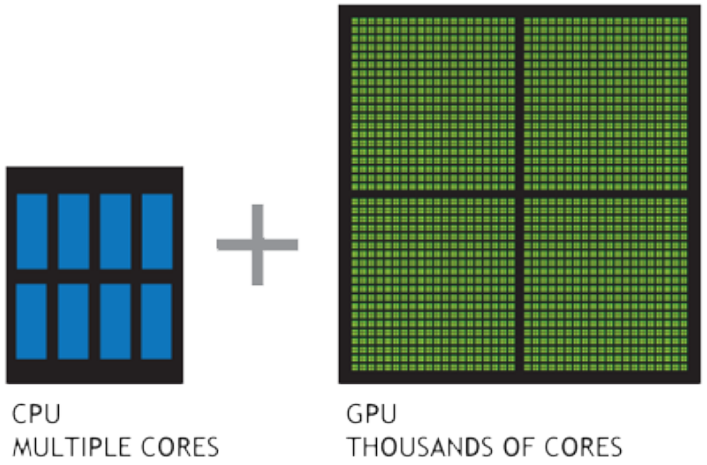

- A modern computer is a distributed system
  - Multi-core CPU
  - Multi-core GPU
  - …
- Actually
  - Even a modern cell phone

Note

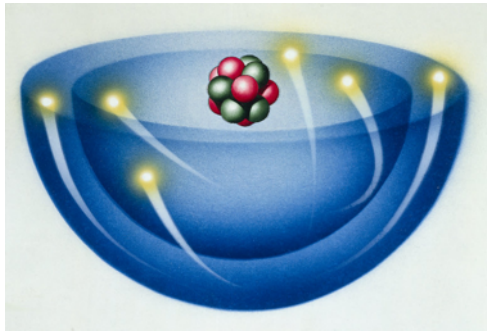

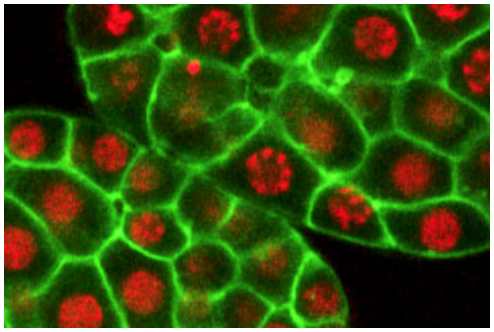

- Non computer-driven "distributed systems"
  - Atoms
  - Molecules
  - Society
  - Animals (ants, bees, …)
  - …
- Not topic of our lectures

Distributed Systems vs. Networks

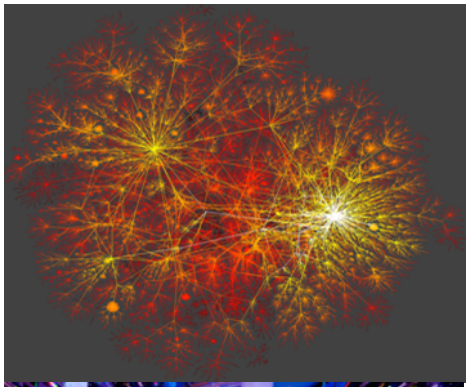

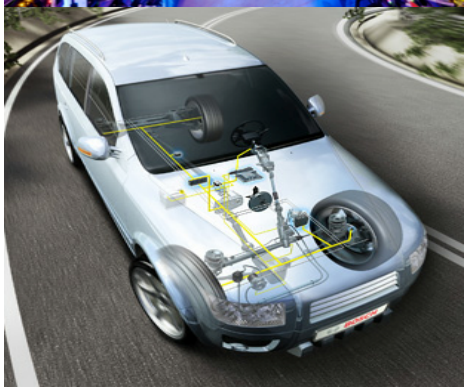

- Networking is worried about
  - Sending a message from here to there
  - Not what you do with the message
- Distributed Systems
  - Assume:
    There is a way to send messages
  - Focus: How you build a system using those messages
  - **Teach you what things to do with a network**

COURSES

Computer Systems and Networks

120 credits (MSc, 2 years)

Distributed Systems Profile

https://www.chalmers.se/en/education/programmes/masters-info/Pages/Computer-systems-and-networks.aspx

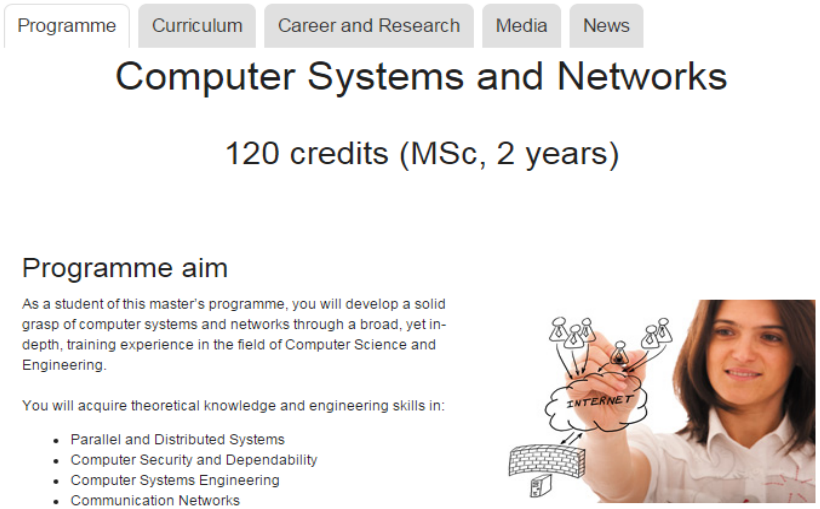

Programme | Curriculum | Career and Research | Media | News

Computer Systems and Networks

120 credits (MSc, 2 years)

Programme aim

As a student of this master's programme, you will develop a solid grasp of computer systems and networks through a broad, yet in-depth, training experience in the field of Computer Science and Engineering.

You will acquire theoretical knowledge and engineering skills in:

- Parallel and Distributed Systems
- Computer Security and Dependability
- Computer Systems Engineering
- Communication Networks

Computer Systems and Networks

120 credits (MSc, 2 years)

Distributed Systems Profile

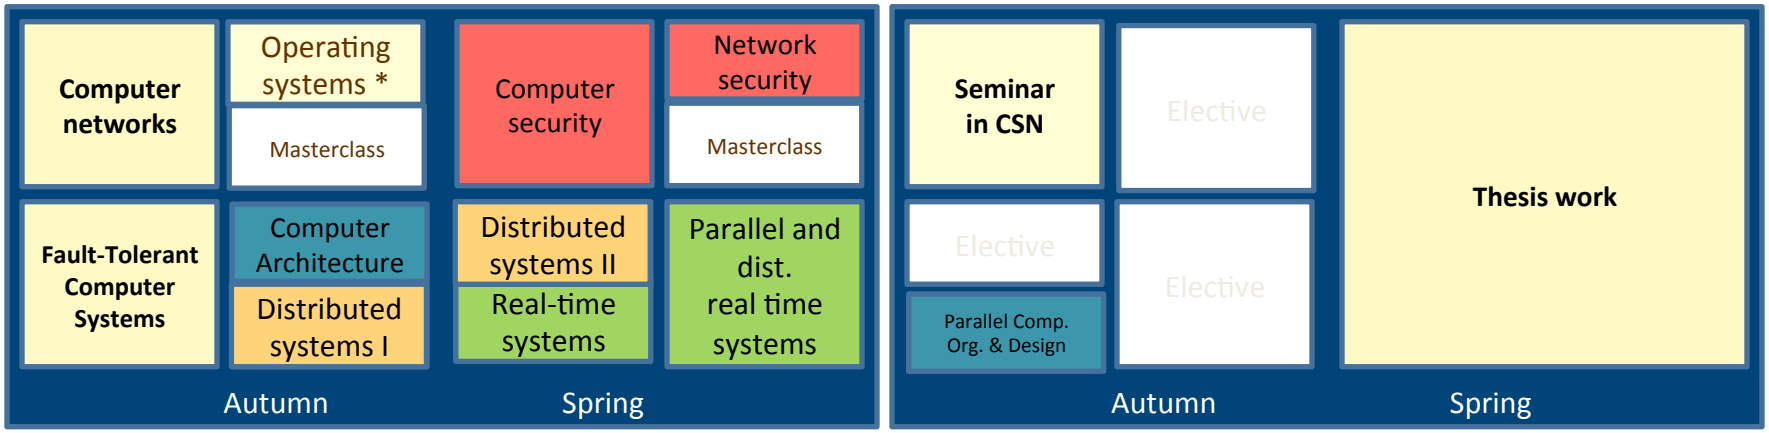

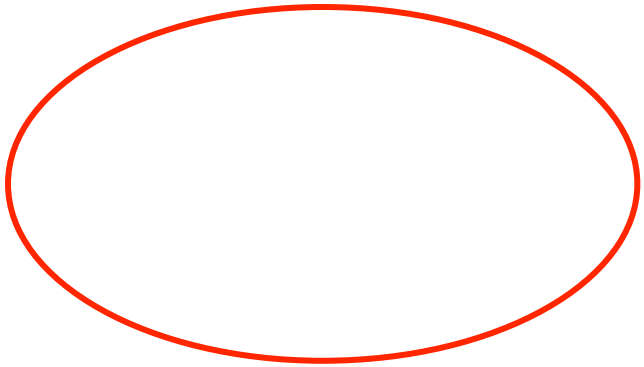

Also available as elective courses from other programs

Course Goals in a Nutshell

- Lectures: Teach you Distributed Systems
  - What do they do?
  - How do they work?
- Labs: Give you hands-on experience
  - Feel the challenges
  - Master the techniques
- Have some fun!
  - Optional: you can pass without it

Will I learn something useful?

- We hope so!
  - This our key goal
- From an email we got from a former student
  - ”[…] I'm […] making a living out of building distributed systems, […] rest assured I've been finding the contents of your course very useful. :)”
    - Started working at Spotify
- We hope you will have a similar experience

More courses

- **Distributed Systems:**
  - **Distributed Systems**, LP2, 7.5hec(hp), TDA596 (Chalmers), DIT240 (GU)
  - **Distributed Systems advanced (Distribuerade system fk.),** LP3 – 7.5 hec (hp), TDA297 (CTH), DIT290 (GU)
- Project Courses
  - DAT295 - **Autonomous and Cooperative Vehicular Systems**, Lp2, 7.5hec
  - DAT300 - **ICT support for adaptiveness and security in the smart grid**, LP4, 7.5hec
- Broader Field
  - EDA387 - **Computer networks**, LP1, 7.5 hec
  - EDA343, EDA344, LEU061 **Datakommunikation**, LP1, LP3, LP4. 7.5 hec
  - EDA491 - **Network security**, LP4, 7.5 hec

HISTORY

History

- In the examples
  - Many different distributed systems
- How did we get here
  - Where do all these DSs come from?
  - What is the trend?
    - Will their number increase even more?

1943

I think there is a world market for maybe five computers

Thomas J. Watson, 1943; Chairman and CEO of International Business Machines (IBM)

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1969

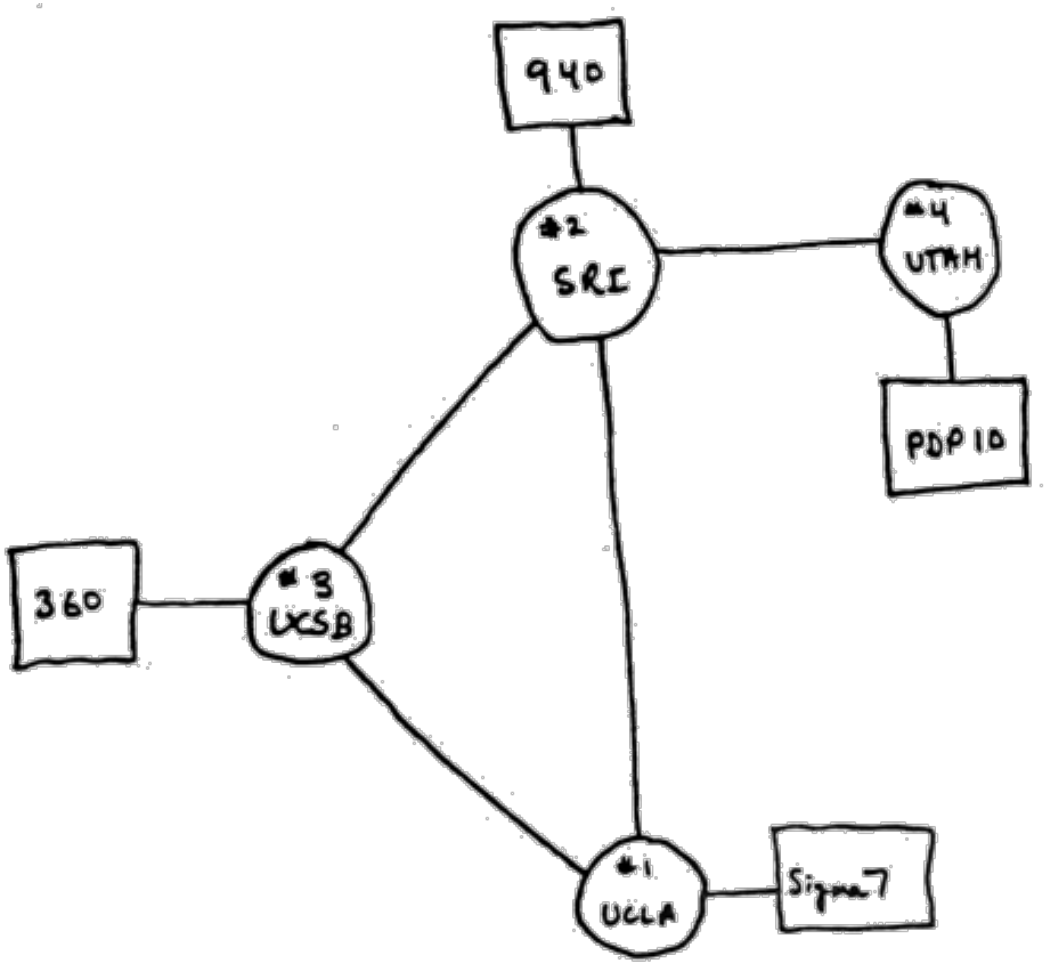

ARPANET begins…with a deployment at UCLA, Stanford, UCSB, and Utah (one computer per site)

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1969, 29 Oct, 22:30:
First data on the Internet

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992

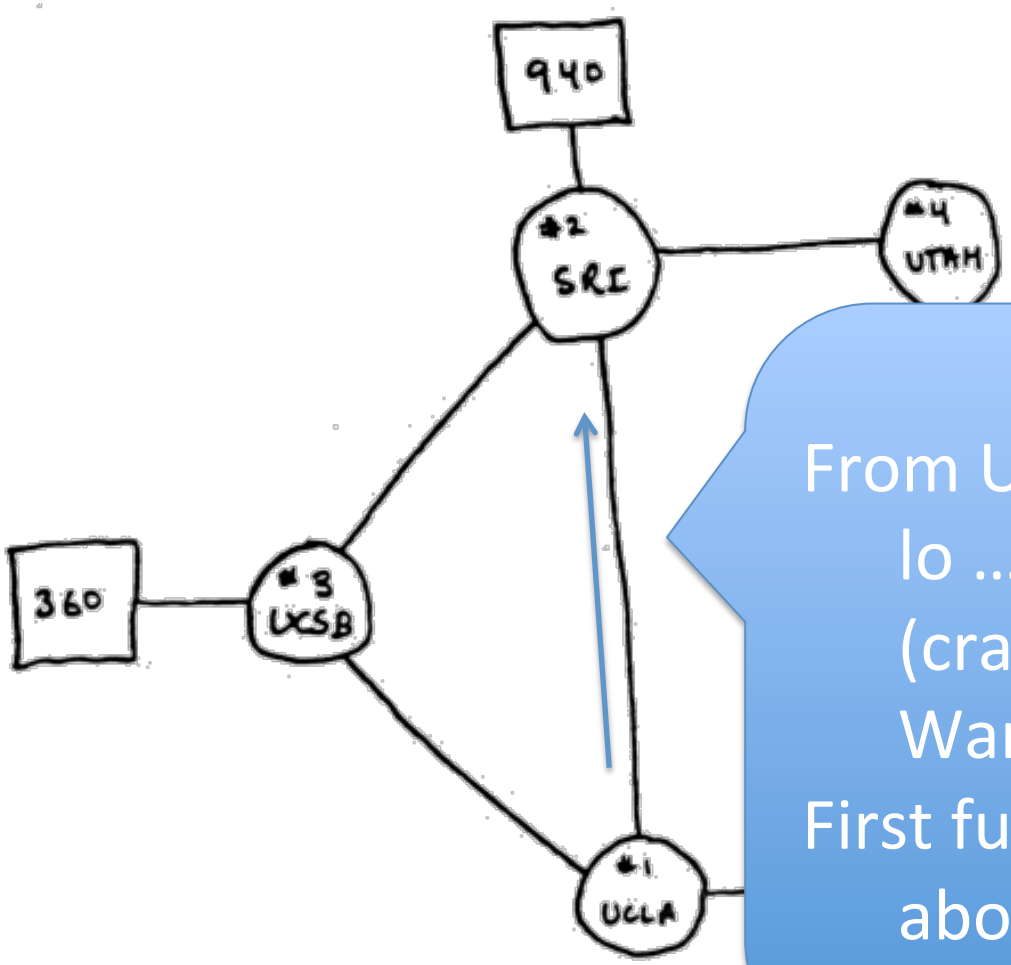

From UCLA to SRI:

lo ….

(crash of SRI machine)!

Wanted to send “login”

First full-login:

about one hour later

1969, 29 Oct, 22:30:
First data on the Internet

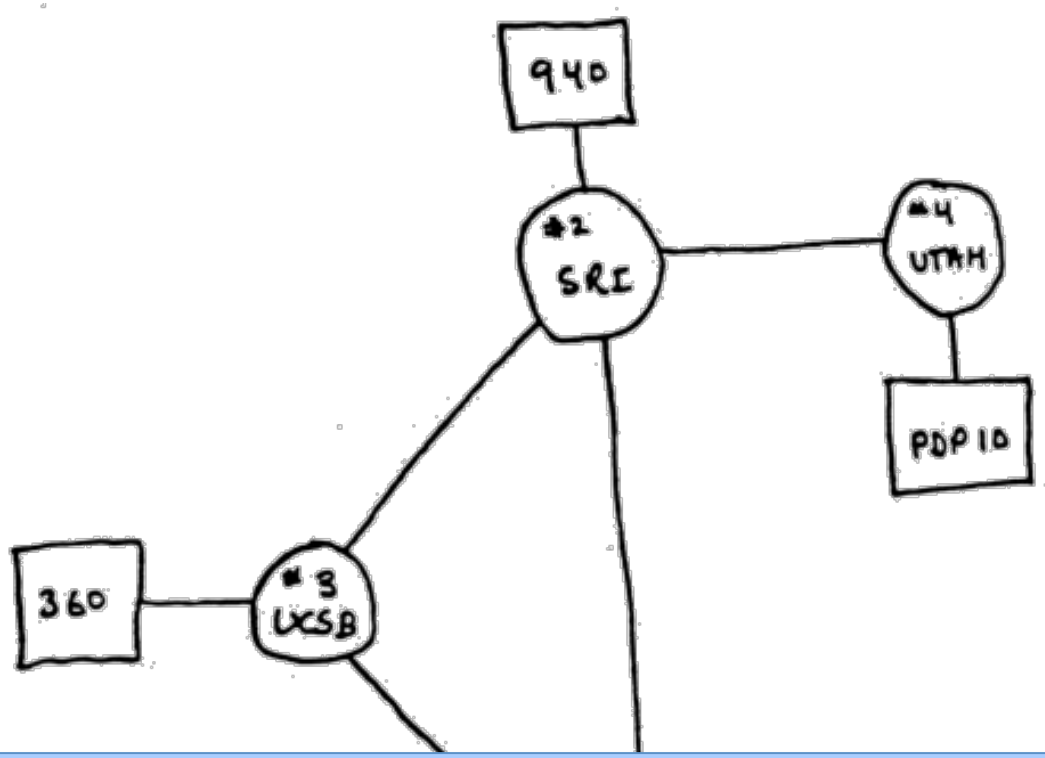

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

Lessons Learned:

1. First words/letters on the Internet: “lo”
2. Not many things in the Internet work on the first try

1977

1969→

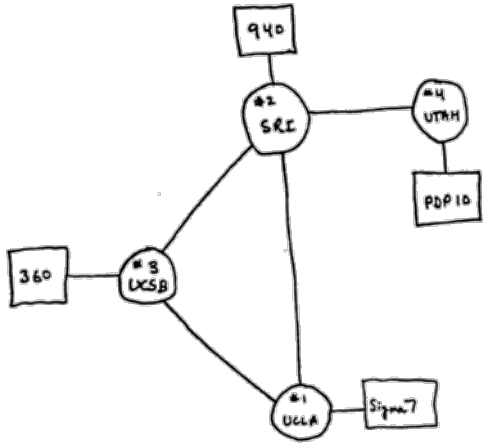

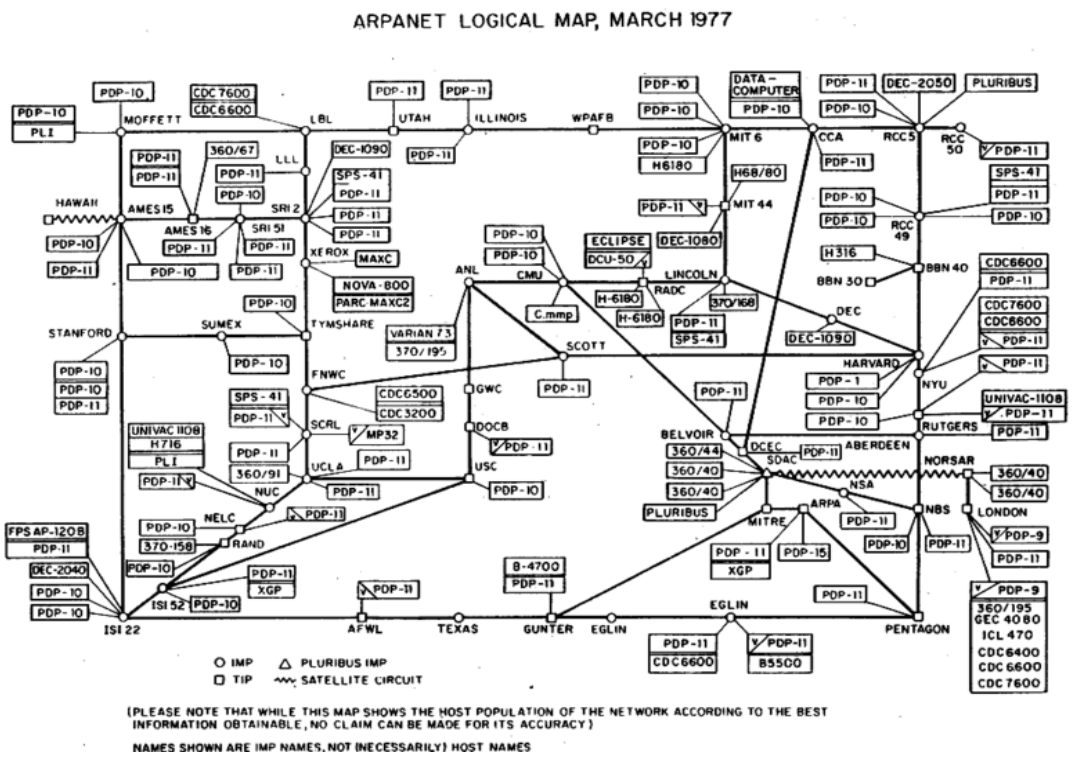

←1977

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

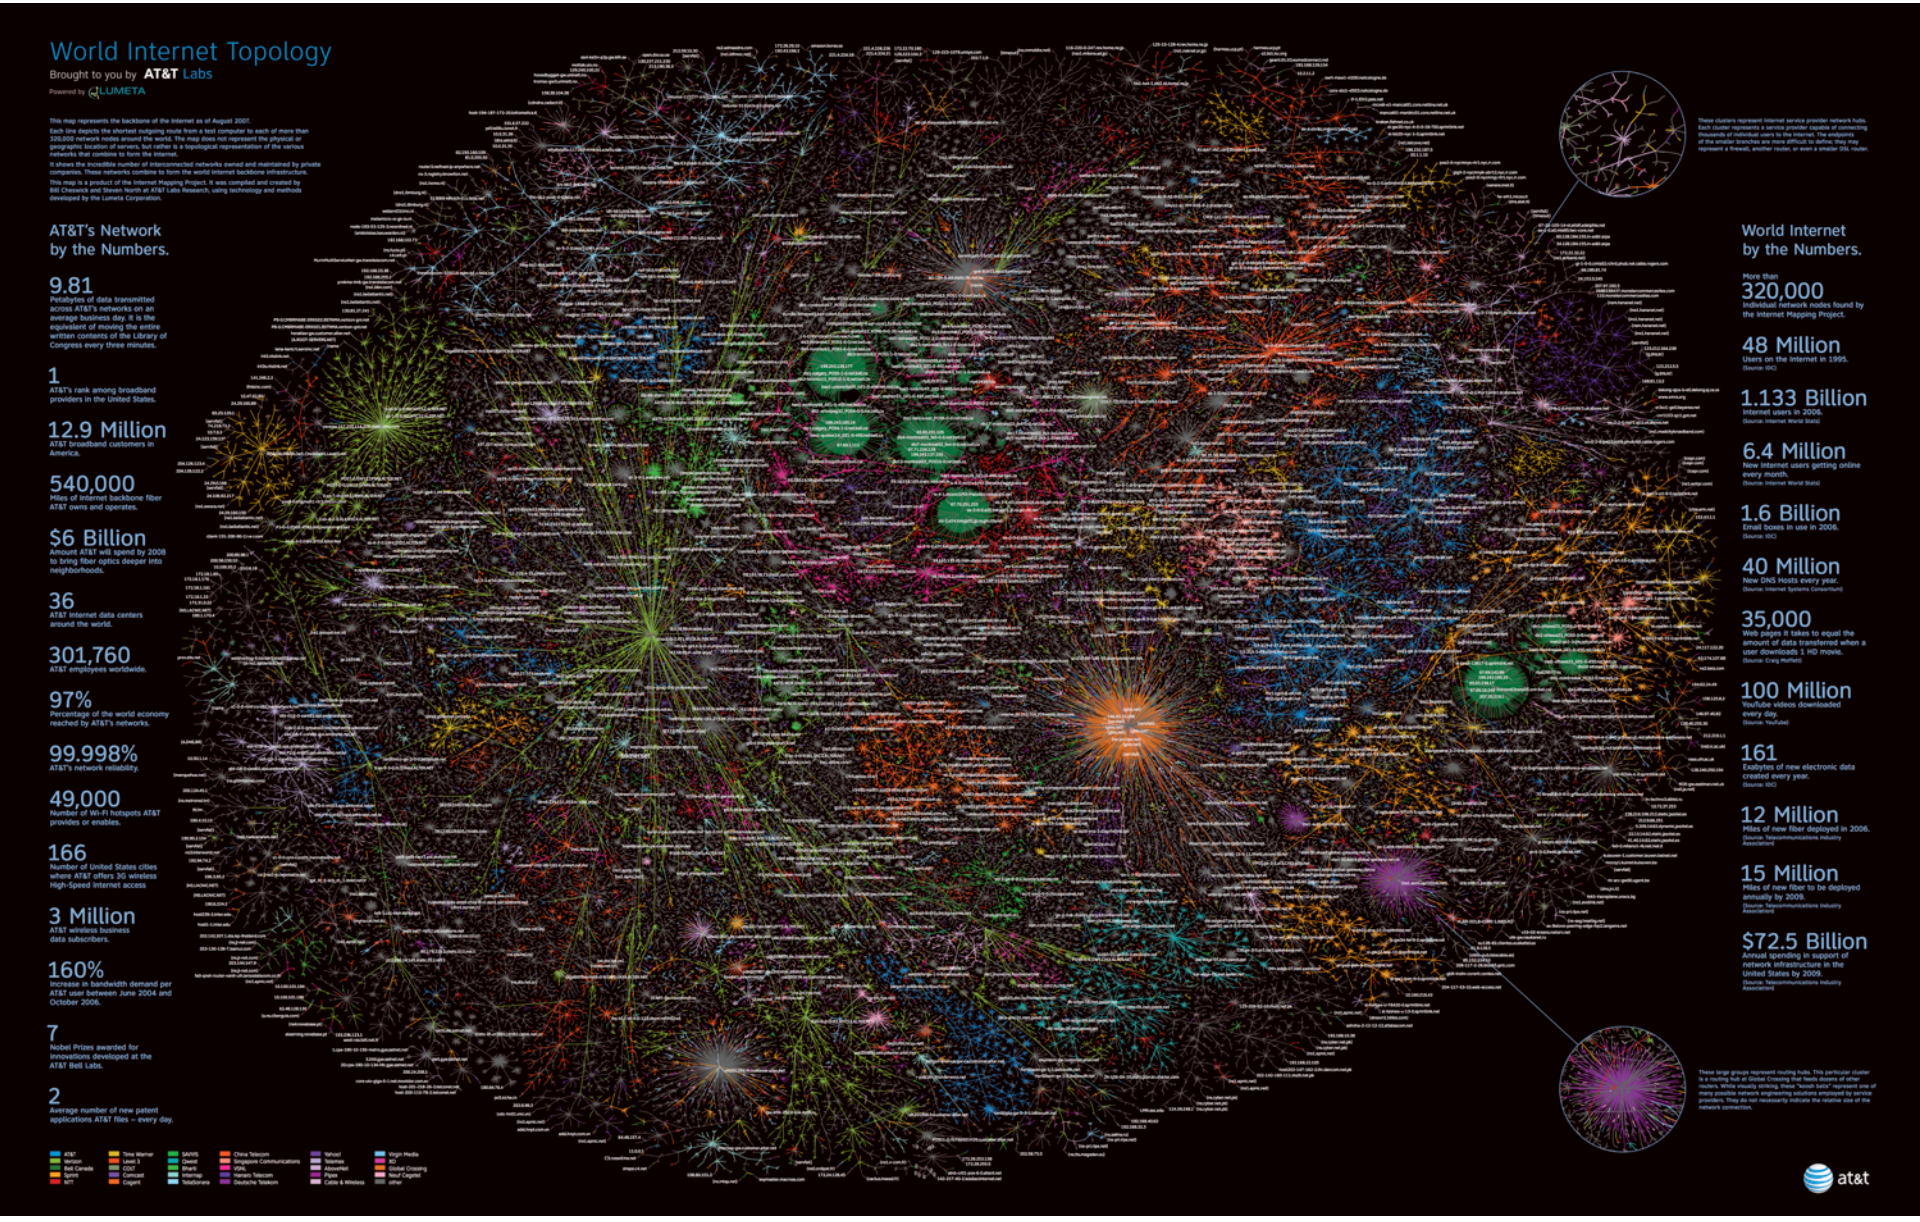

Internet 2007 (just the backbone)

www2.research.att.com/~north/news/img/ATT_Labs_InternetMap_0730_10.pdf

1971

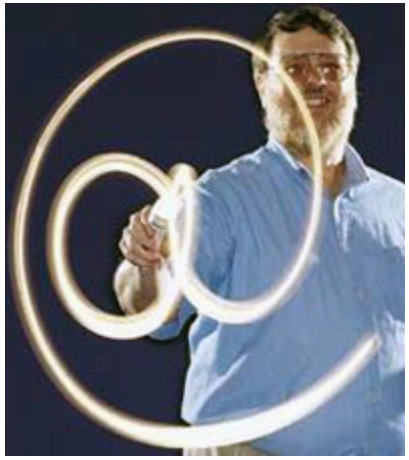

Ray Tomlinson creates first email program

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1974

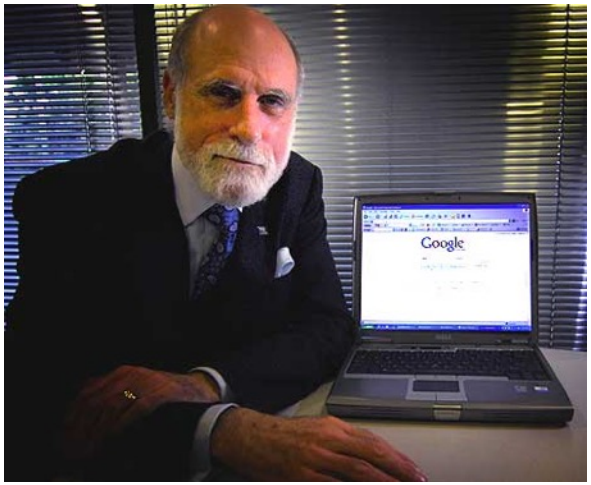

TCP / IP defined by Vint Cerf & Bob Kahn

2004: both received the Turing Award

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1984

Paul Mockapetris introduces DNS

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1989 – The Web Emerges

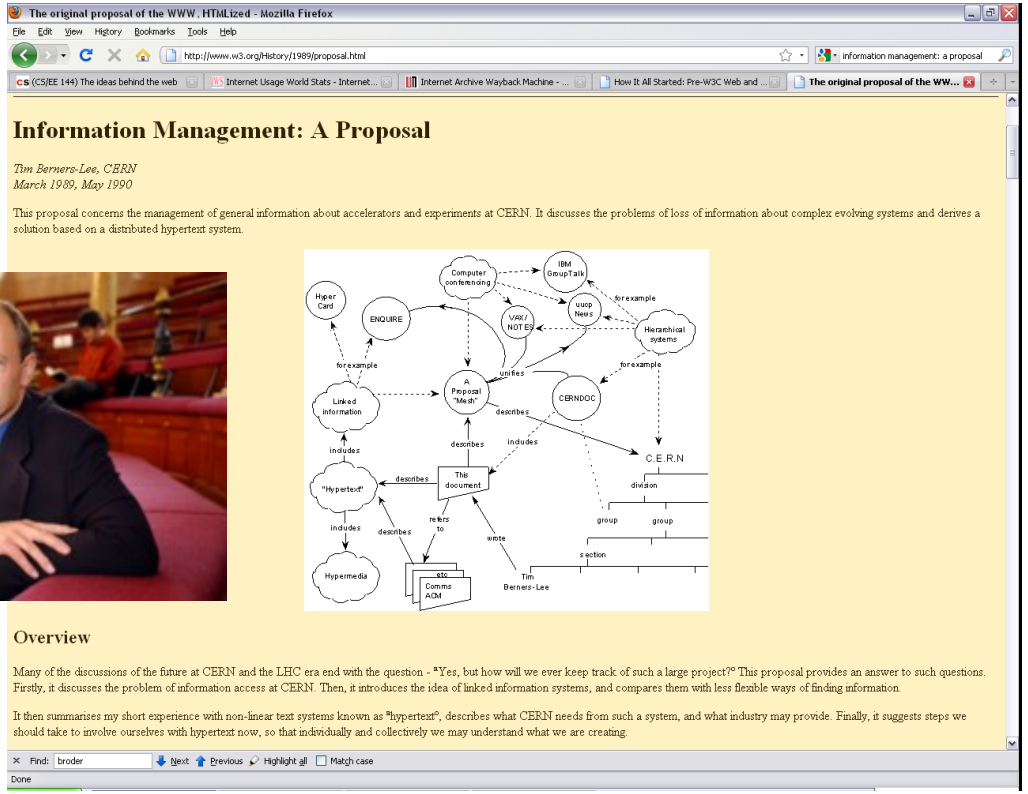

Tim Berners-Lee writes "Information Management: A proposal" at CERN

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1990

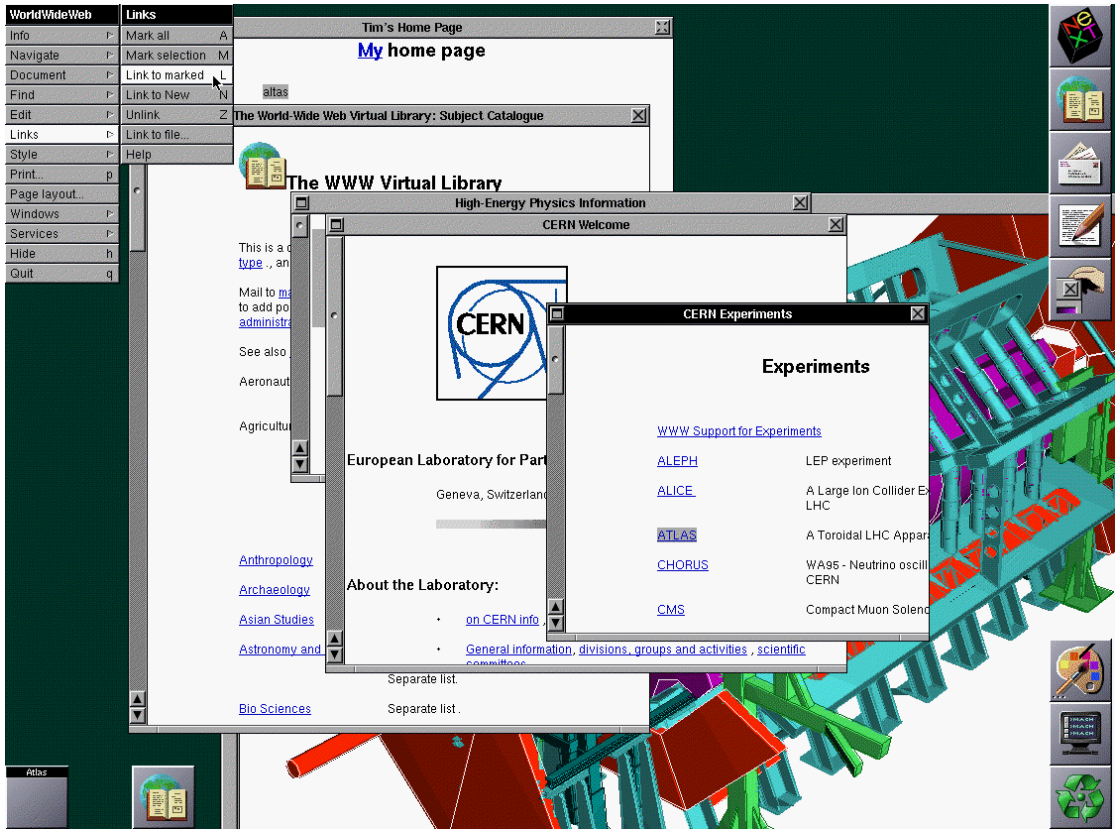

First browser developed at CERN

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1991

First paper appears on the project at Hypertext conference

→ Only accepted as a poster!

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1993

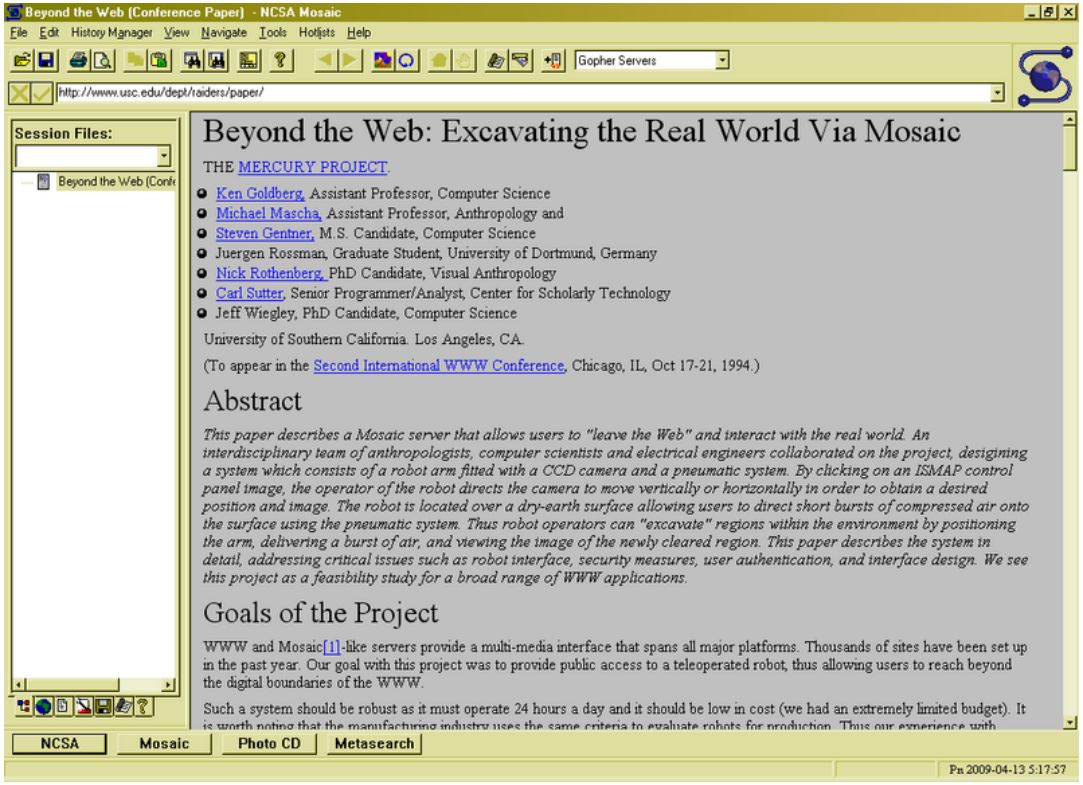

Mosaic became the first graphical browser

CERN agrees to allow public use of web protocol royalty-free!

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1994

→ Mosaic goes commercial
(later becomes Netscape)

→ Traditional dialups (AOL, CompuServe, Prodigy) begin to sell Internet access.

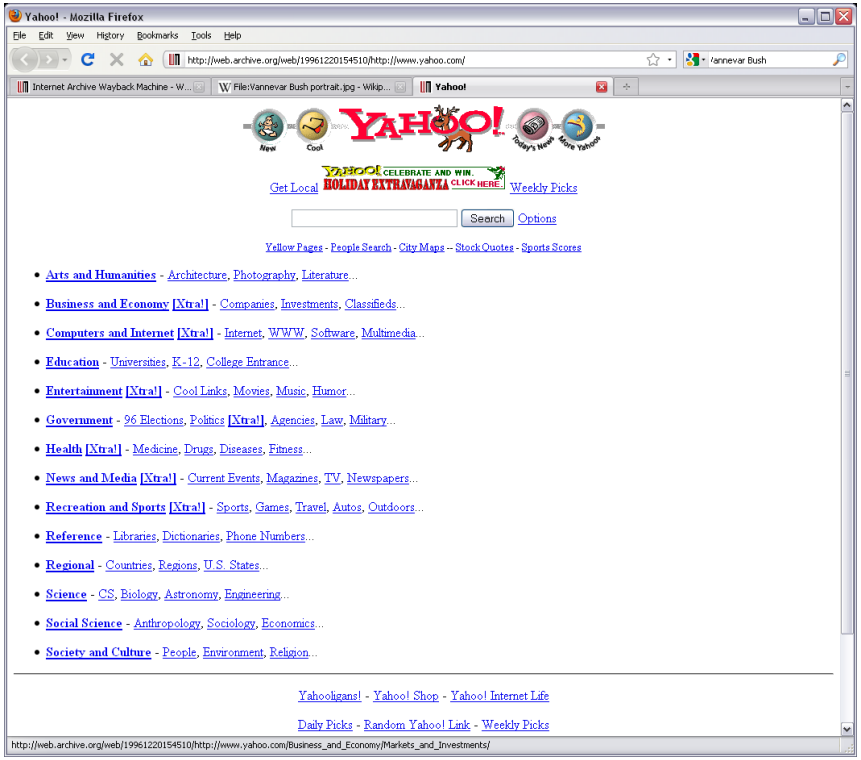

Yahoo
circa
1996

"Jerry's Guide to the world wide web" started …
it eventually became Yahoo

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

1995+

Amazon arrives and the commercialization of the web begins

Amazon circa 1999

- Pre-me (<1979)
- Pre-you (<1989)
- 1989
- 1990
- 1991
- 1992
- 1993
- 1994

Today

- How many connected devices do you have?

- Many!
  - Desktop
  - Laptop
  - (Smart)phone
  - Tablet
  - TV / gaming console
  - …

Summary: A bit of History

Mainframe age (60's & 70's):
One computer for many

PC age (80's & 90's):
One computer for each, partially networked

Cloud computing
Mobile, ubiquitous computing
(Today, > 2000):
Many computers for each, networked

Tomorrow?

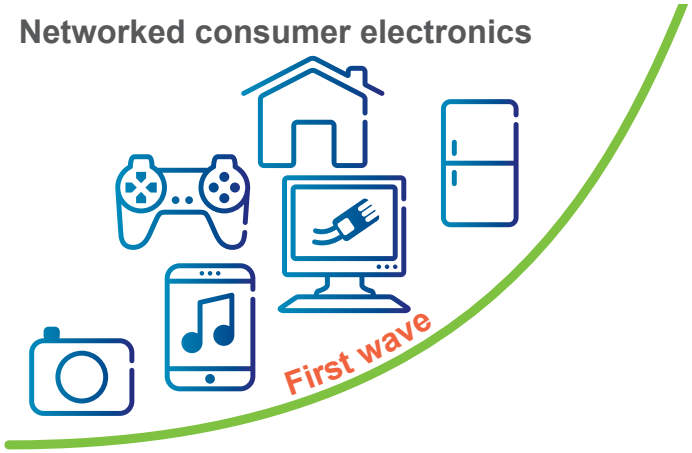

Today

Tomorrow?

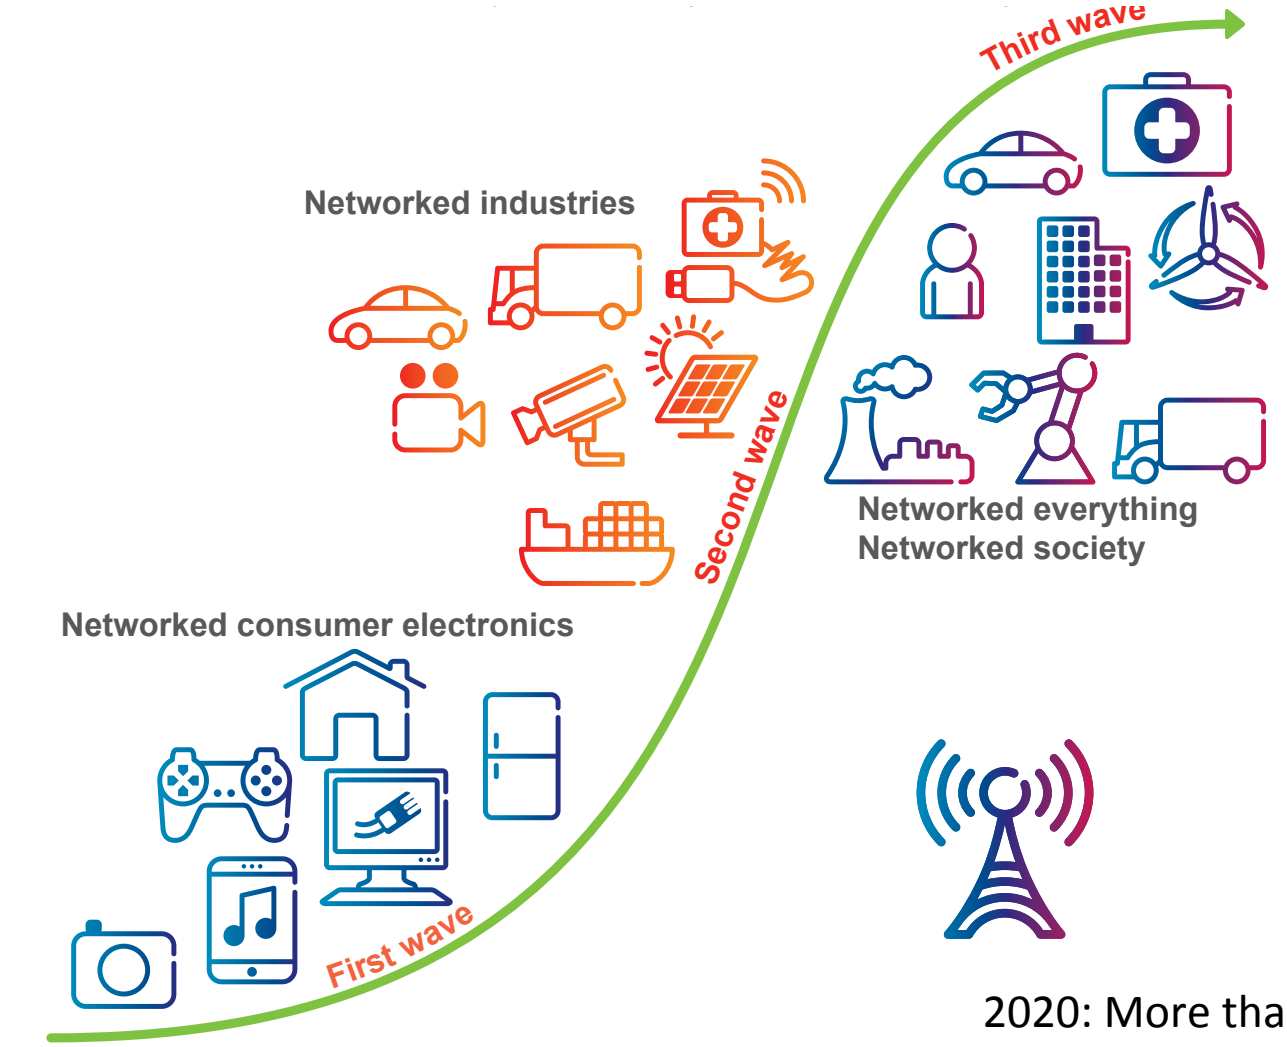

Today

2020: More than 50 Billion connected devices (Ericsson)

Tomorrow? Networked Society!

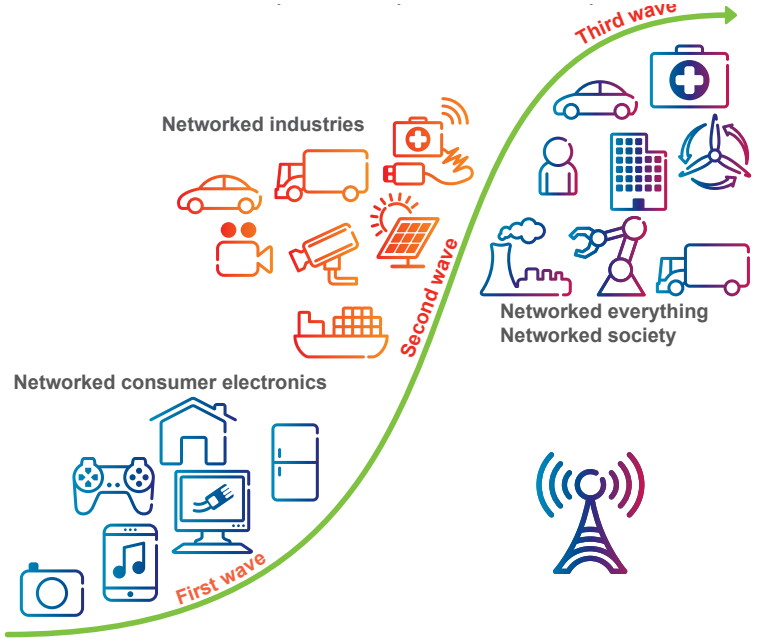

- Distributed Systems touch all aspects of daily life!
  - Integral building block for our networked society
  - Strongly increasing in numbers
  - Result: Very good topic to study ;-)

COURSE TOPICS

Course Topics: Motivation

- Assume: your task is to build
  - Facebook or
  - Amazon or
  - just a simple web application

- What challenges do you face?

The Eight Fallacies of Distributed Systems

- The network is reliable
- Latency is zero
- Bandwidth is infinite
- The network is secure
- Topology doesn't change
- There is one administrator
- Transport cost is zero
- The network is homogeneous

Mechanisms

- This course
  - Mechanisms to deal with these challenges
  - Generic mechanisms
    - Not bound to the Internet
    - But: Examples mostly Internet bound
      - Easier to understand for most students
      - Compared to power grids, cars, …

Course Content: Mechanisms

- Architectures & Processes
- Mutual exclusion & Election
- Naming
- Clocks and Time
- Consistency & replication
- Fault tolerance

Computer Systems and Networks

120 credits (MSc, 2 years)

Distributed Systems Profile

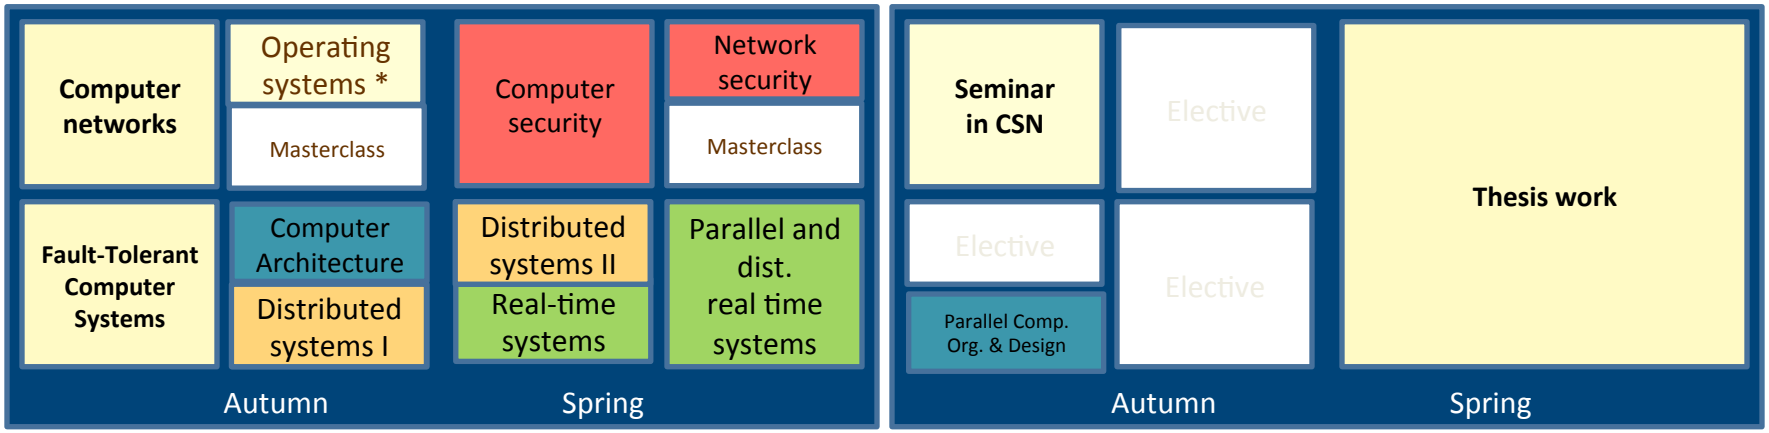

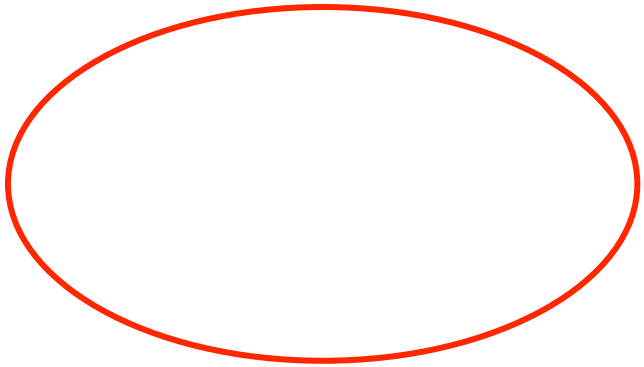

Also available as elective courses from other programs

Questions