



Networks and Distributed Systems

Olaf Landsiedel


Networks and Distributed Systems

- What is...
 - A computer network?
 - Have you ever seen one?
 - Have you ever used one?
 - A distributed system?
 - Have you ever seen one?
 - Have you ever used one?

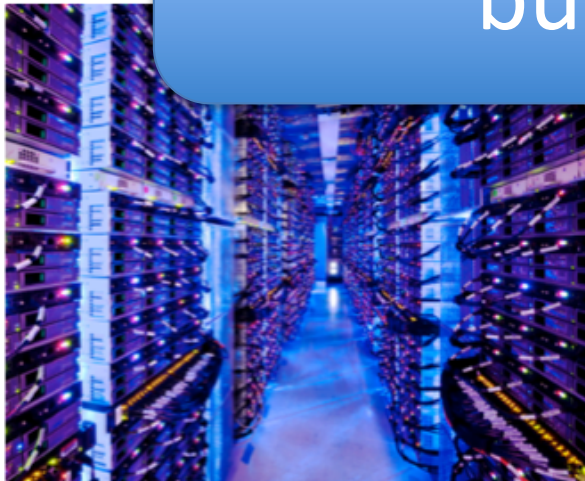
Networks and Distributed Systems

- A Computer Network is characterized by
 - Bring data from A and B
 - By exchanging messages
- 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

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**



Networks and Distributed Systems

HISTORY

History

- In the examples
 - Many different distributed systems and networks
- How did we get here
 - Where do all these networks and distributed systems 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


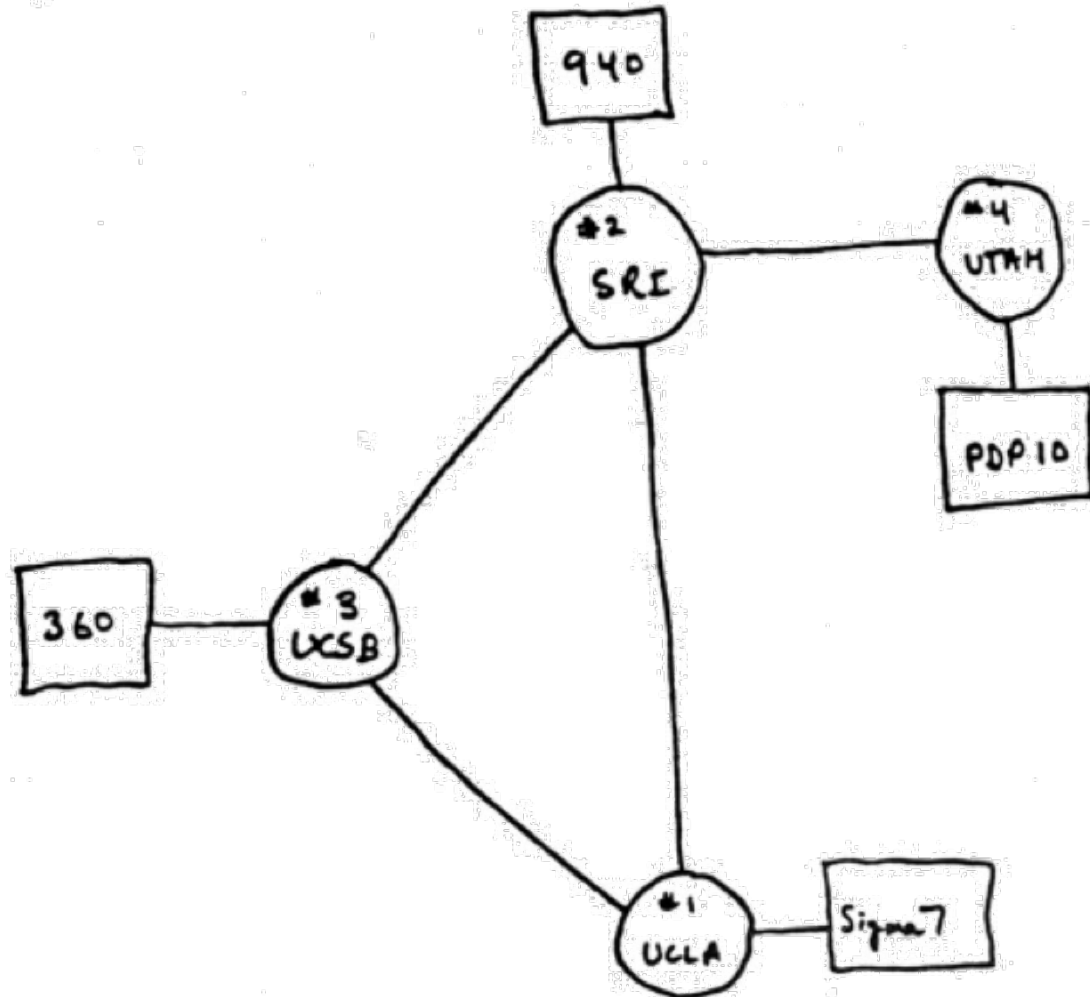
1993

1994

If this statement had been correct, we would not teach:


- Computer Networks, Distributed systems, ...
- Or, more precisely no Computer Science

1969



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


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"

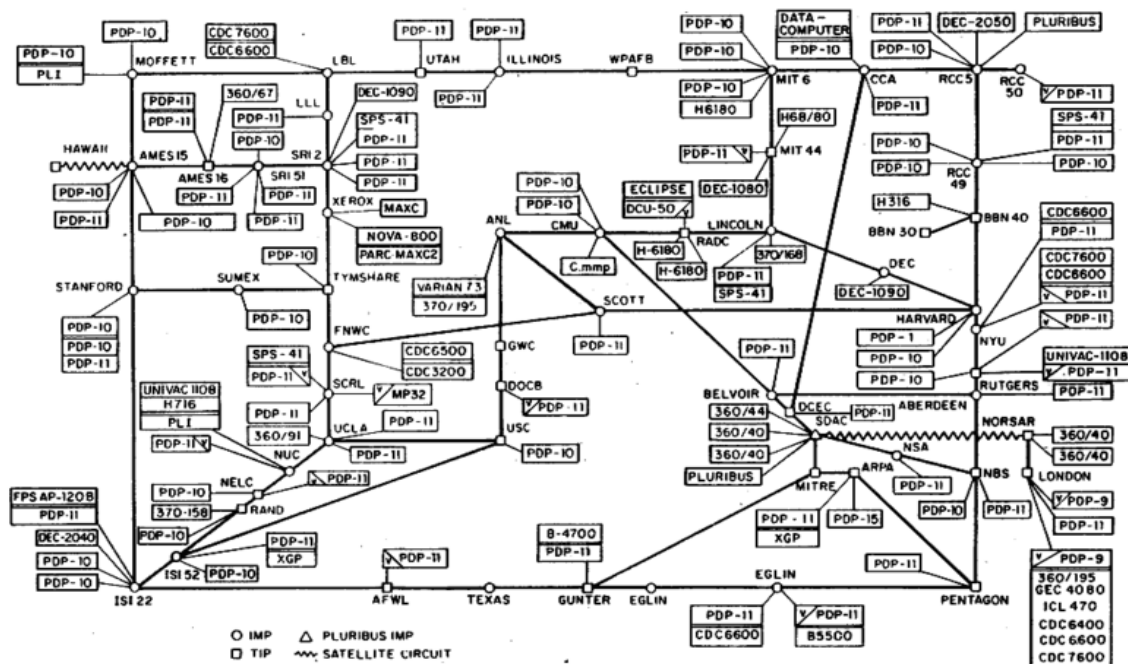
2. Not many things in the Internet work on the first try

1977

1969 →



ARPANET LOGICAL MAP, MARCH 1977



(PLEASE NOTE THAT WHILE THIS MAP SHOWS THE MOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY.)
 NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

Pre-me (<1979)

Pre-you (<1989)

1989

1990

1991

1992

1993

1994

← 1977



World Internet Topology

Brought to you by AT&T Labs

Powered by LUMETA

This map represents the backbone of the Internet as of August 2007.

Each line depicts the physical outgoing route from a host computer to each of more than 100,000 servers around the world. The map does not represent the physical or geographic location of servers, but rather is a topological representation of the various networks that comprise the Internet.

It shows the available number of interconnected networks owned and maintained by private companies. These networks continue to form the world's Internet backbone structure.

This map is a product of the Internet Mapping Project. It was compiled and created by 300 researchers and team leads at AT&T Labs Research, using technology and methods developed by the LUMETA Corporation.

AT&T's Network by the Numbers.

9.81

Petabytes of data transmitted across AT&T's network in an average business day is the equivalent of moving the entire written contents of the Library of Congress every three minutes.

1

AT&T's rank among broadband providers in the United States.

12.9 Million

AT&T broadband customers in America.

540,000

Miles of Internet backbone fiber AT&T owns and operates.

\$6 Billion

Amount AT&T will spend by 2008 to bring fiber optics closer into neighborhoods.

36

AT&T Internet data centers around the world.

301,760

AT&T employees worldwide.

97%

Percentage of the world economy reached by AT&T's network.

99.998%

AT&T's network reliability.

49,000

Number of Wi-Fi hotspots AT&T provides or enables.

166

Number of United States cities where AT&T offers 3G wireless High-Speed Internet access.

3 Million

AT&T wireless business data subscribers.

160%


Increase in bandwidth demand per AT&T user between June 2004 and October 2006.

7

Nobel Prizes awarded for inventions developed at the AT&T Bell Labs.

2

Average number of new patent applications AT&T files - every day.



These nodes represent Internet service providers, search engines, and other major entities in the network. The number of connections to these nodes is a measure of their importance in the network.

World Internet by the Numbers.

More than **320,000**

autonomous systems found by the Internet Mapping Project.

48 Million

users on the Internet in 2005.

1.133 Billion

Internet users in 2006.

6.4 Million

New Internet users getting online every month.

1.6 Billion

Small files in use in 2006.

40 Million

New DNS records every day.

35,000

Web pages it takes to equal the amount of data transferred when a user downloads a 100 movie.

100 Million

YouTube videos downloaded every day.

161

Terabytes of new electronic data created every year.

12 Million

Miles of new fiber deployed in 2006.


15 Million

Miles of new fiber to be deployed annually by 2008.

\$72.5 Billion

Annual spending in support of network infrastructure in the United States by 2008.

Source: AT&T



These large groups represent routing hubs. The particular color is a routing ID or other identifying the hubs. Each of these hubs is a major node in the network, and the number of connections to these hubs is a measure of their importance in the network.




Internet 2007 (just the backbone)


1971



Ray Tomlinson creates first email program



1974



Pre-me (<1979)

Pre-you (<1989)

1989

1990

1991

1992

1993

1994

TCP / IP defined by Vint Cerf & Bob Kahn




2004: both received the Turing Award



1984



Paul Mockapetris introduces DNS



1989 – The Web Emerges

Pre-me (<1979)

Pre-you (<1989)

1989

1990

1991

1992

1993

1994

Information Management: A Proposal

Tim Berners-Lee, CERN
March 1989, May 1990

This proposal concerns the management of general information about accelerators and experiments at CERN. It discusses the problems of loss of information about complex evolving systems and derives a solution based on a distributed hypertext system.

Overview

Many of the discussions of the future at CERN and the LHC era end with the question - "Yes, but how will we ever keep track of such a large project?" This proposal provides an answer to such questions. Firstly, it discusses the problem of information access at CERN. Then, it introduces the idea of linked information systems, and compares them with less flexible ways of finding information.

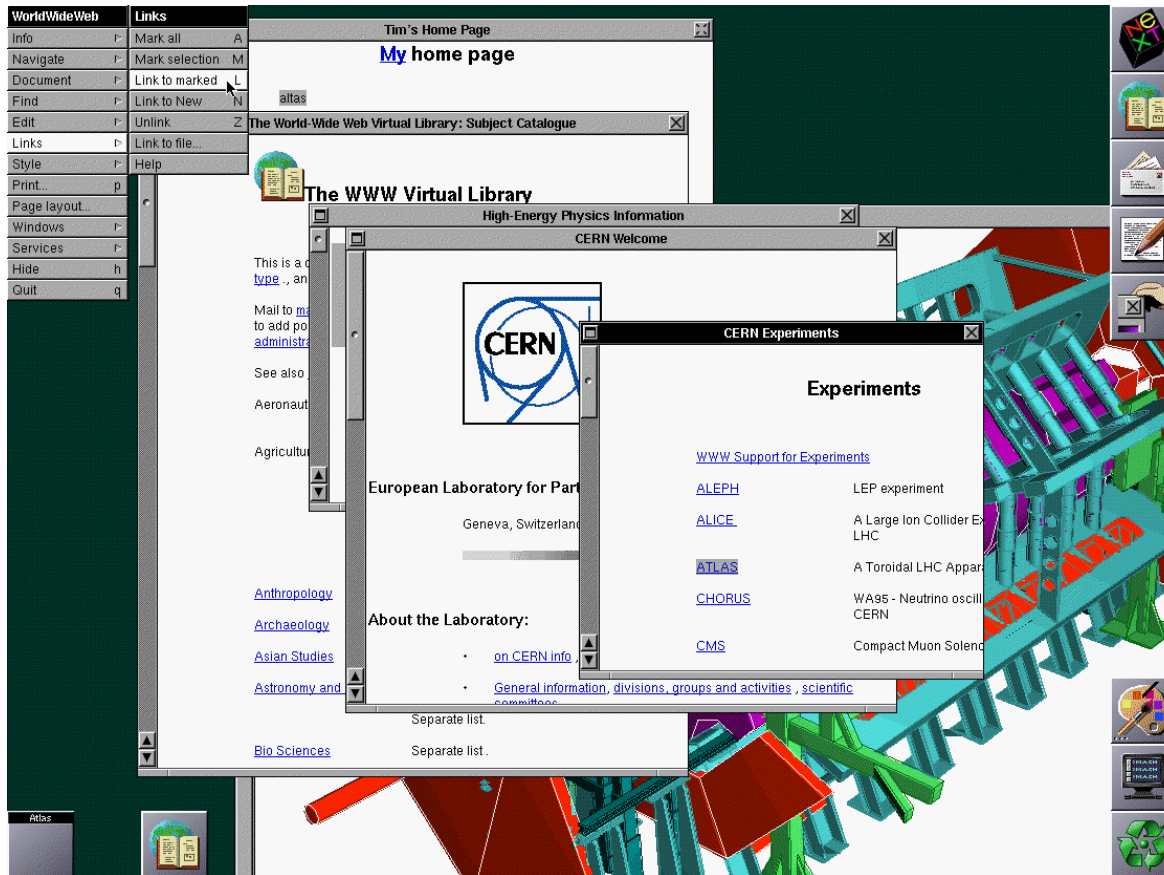
It then summarises my short experience with non-linear text systems known as "hypertext", describes what CERN needs from such a system, and what industry may provide. Finally, it suggests steps we should take to involve ourselves with hypertext now, so that individually and collectively we may understand what we are creating.



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



1990



First browser developed at CERN

Pre-me (<1979)

Pre-you (<1989)

1989

1990

1991


1992

1993


1994



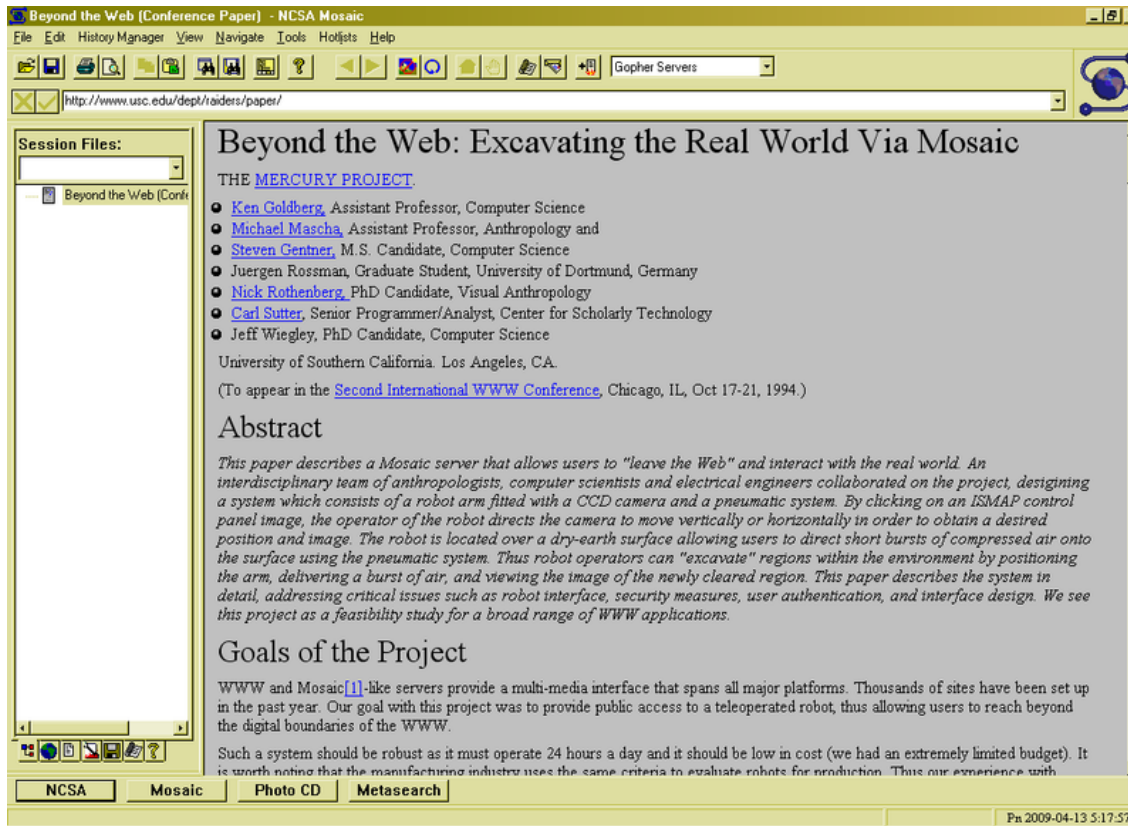
1991



First paper appears on the project at Hypertext conference
→ Only accepted as a poster!



1993



Pre-me (<1979)

Pre-you (<1989)

1989

1990

1991

1992

1993

1994

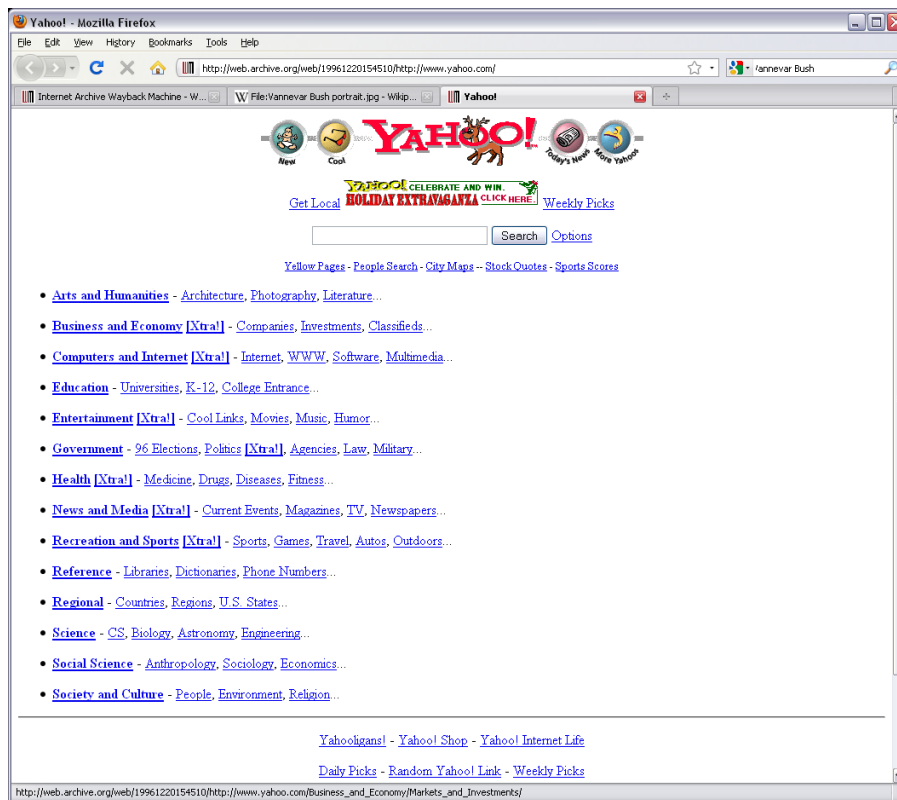
Mosaic became the first graphical browser

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



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

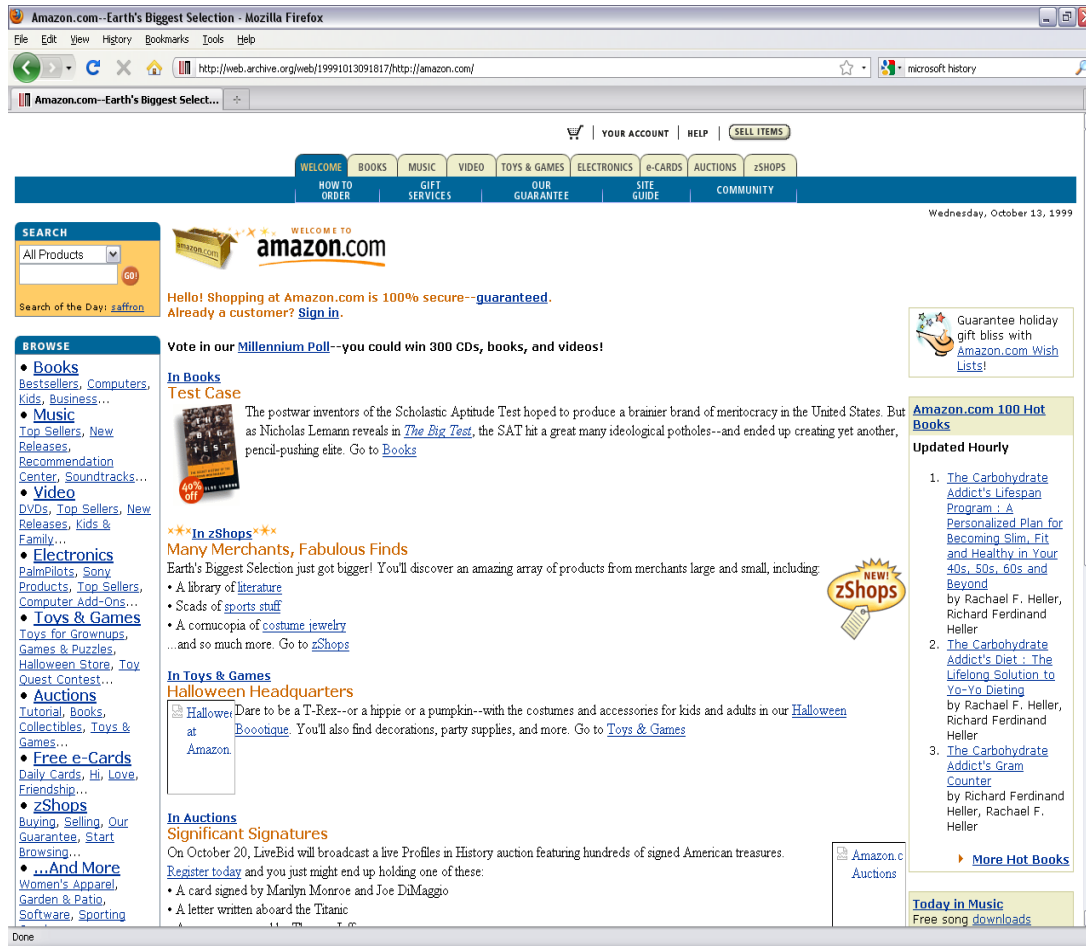
1993

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

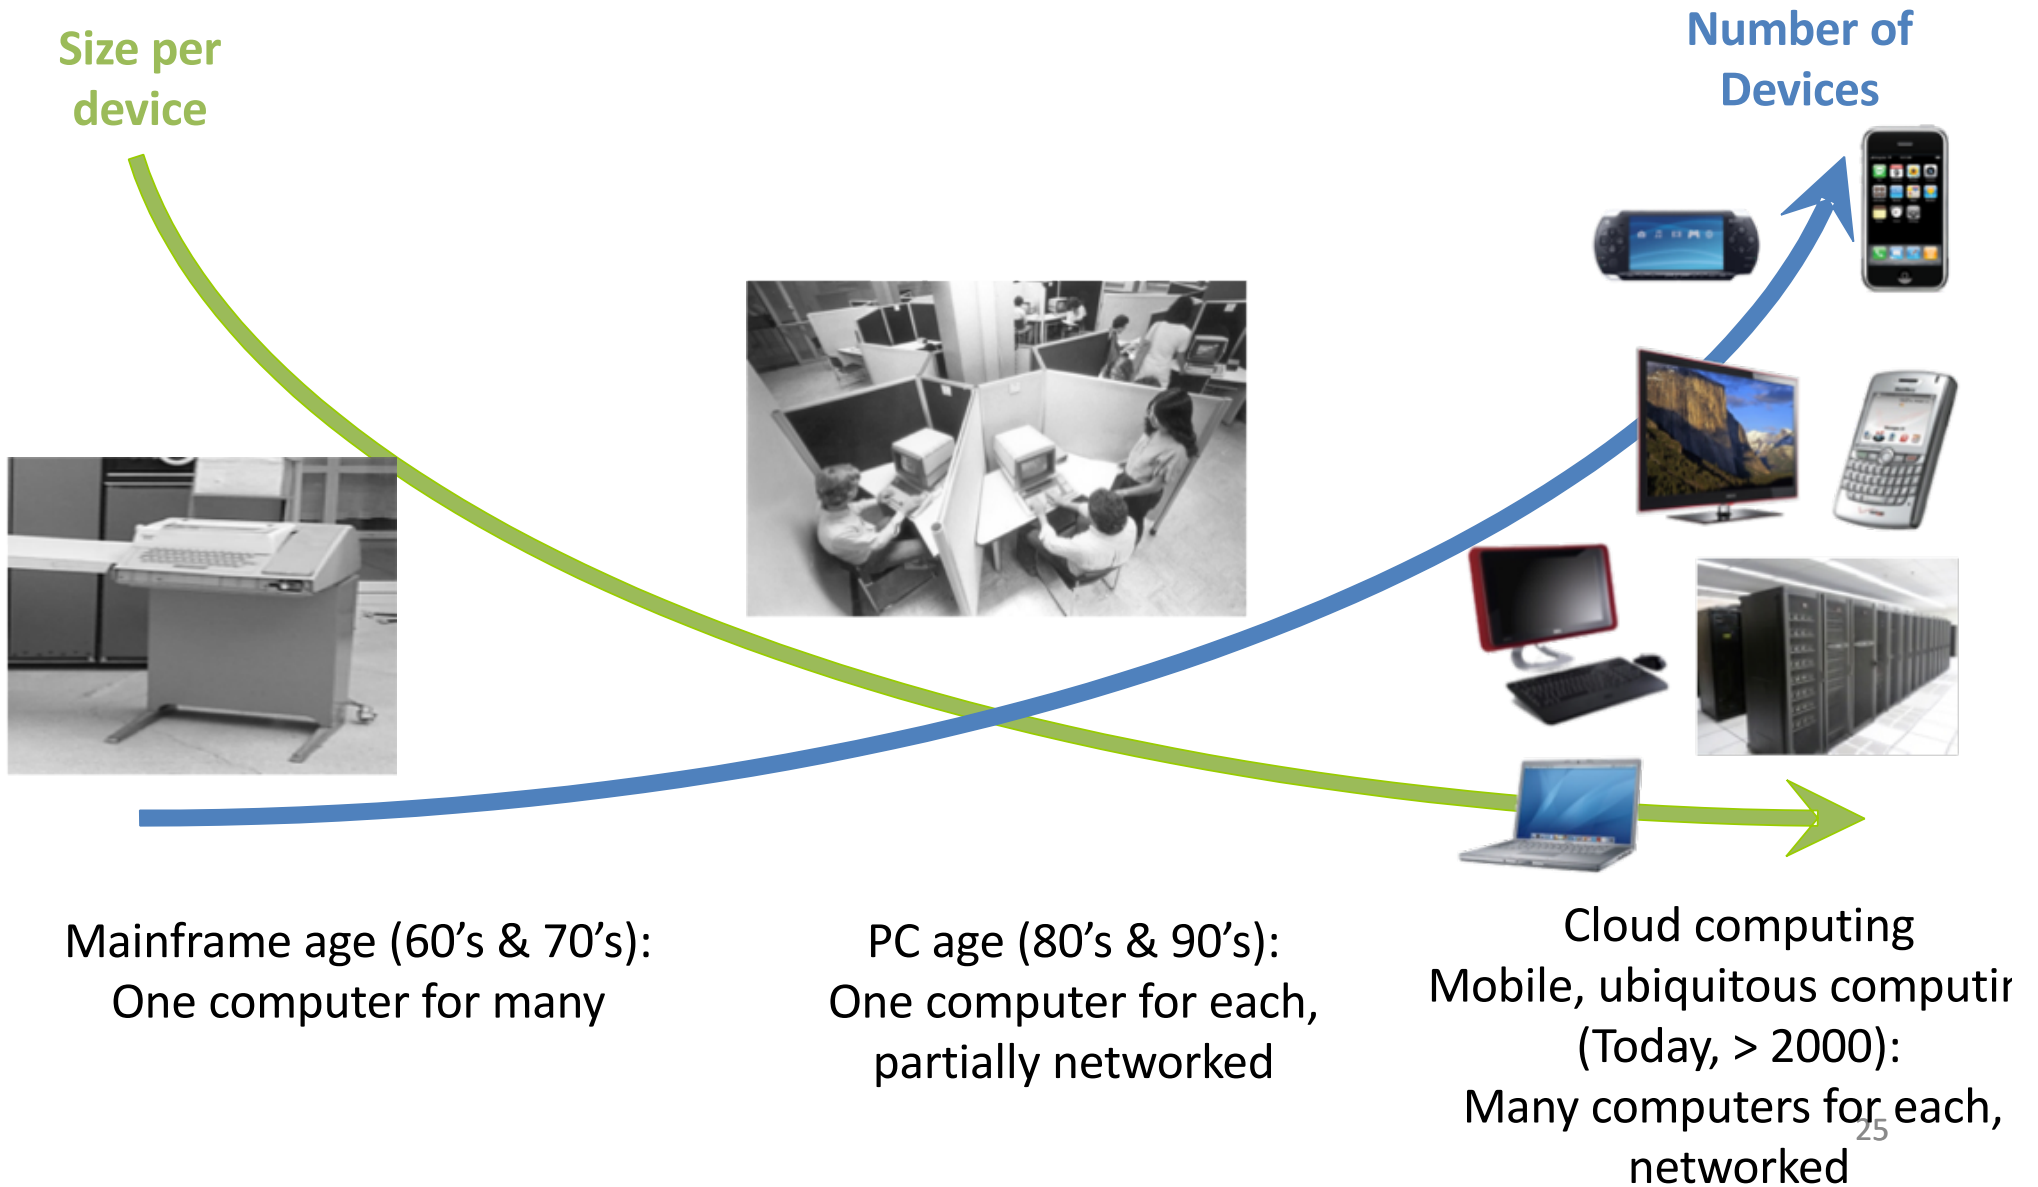
1994




Today

- How many connected devices do you have?
- Many!
 - Desktop
 - Laptop
 - (Smart)phone
 - Tablet
 - TV / gaming console
 - ...

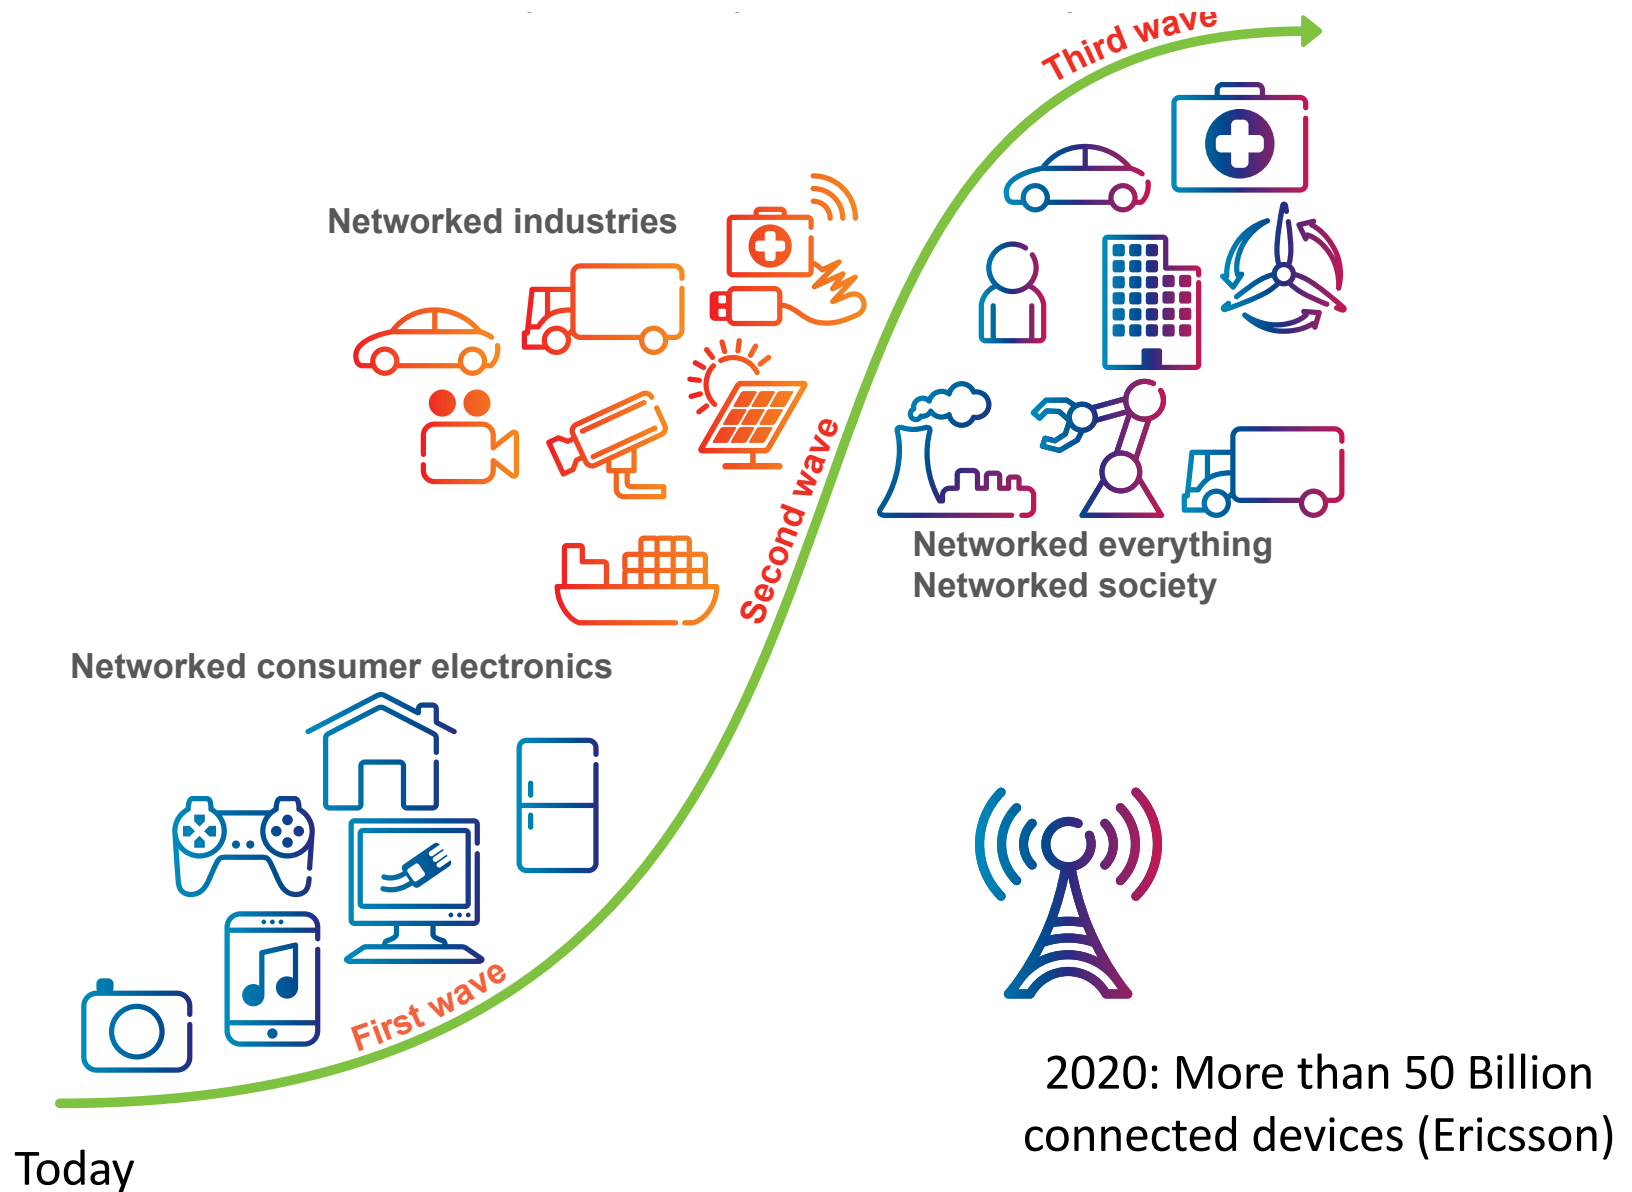
Summary: A bit of History



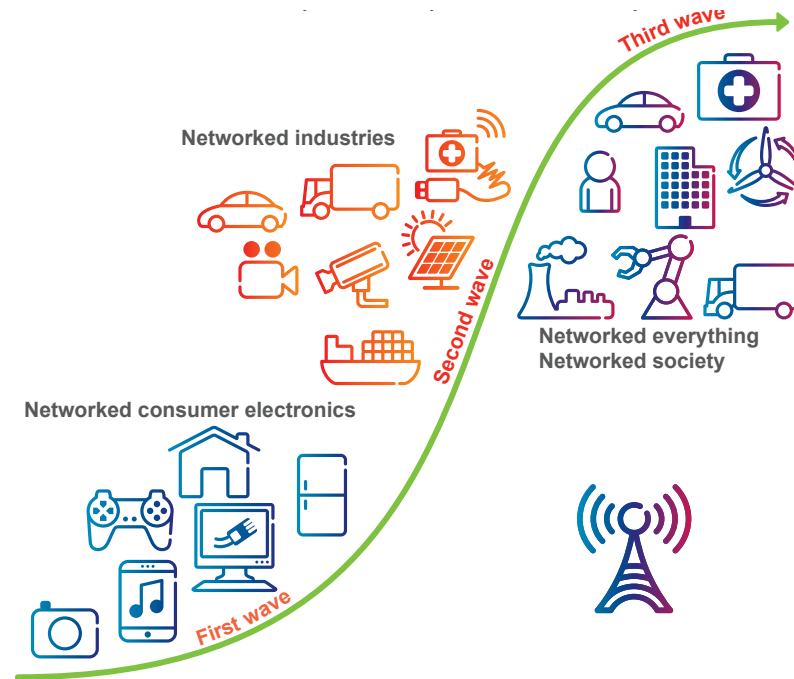
Tomorrow?



Tomorrow?



Tomorrow? Networked Society!



- Networks and 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 ;-)

Computer Systems and Networks

MASTER PROGRAM

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



- Compulsory courses
- Elective courses
- Tracks

* Mandatory only if not taken before

** Recommended elective project courses: Autonomous and Cooperative Vehicular Systems (second study period), OCT Support for adaptiveness and Security in the smart grid (fourth study period)

*** Choose two out of these course tracks: Computer security, Real-time systems, Distributed systems and Computer architecture

Networks and Distributed Systems

COURSES

Courses

- **Networks:**
 - EDA387 - **Computer networks**, LP1, 7.5 hec
 - EDA343, EDA344, LEU061 **Datakommunikation**, LP1, LP3, LP4. 7.5 hec (Bachelor)
- **Operating Systems:**
 - EDA092/DIT400 **Operating Systems**, LP1, 7.5 hec
- **Distributed Systems:**
 - **Distributed Systems**, LP2, 7.5hec, TDA596 (Chalmers), DIT240 (GU)
 - **Distributed Systems advanced**, LP3 – 7.5 hec, 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

[Data Communication and later Computer networks]

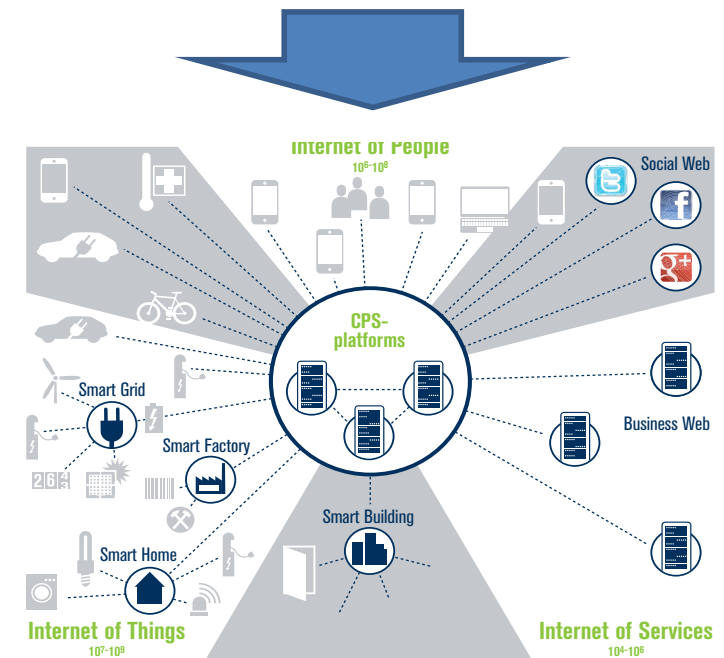

Internet & its context + evolution

Course aims

- Learn well the basic data-networking principles and methods, to follow **constant change** in the field
- Learn to deal with bigger problems by breaking into small ones

After completion of the course, you are able to

- distinguish **network services, related protocols, new systems** relating with IoT, varying data flows and virtualization (Software Defined Networks)
- Understand and think **possibilities and constraints** in the existing systems
- build and configure a **working network**





Source: Bosch Software Innovations 2012

continuous evolution

Course Operating Systems

- Course covering how operating systems bridge hardware / software and users.
- Broad spectrum, from:

Threads management
in multicore CPUs




Virtual memory



Security

Courses Distributed Systems


- Learn to build large-scale distributed systems
 - And the associated challenges



Internet



Facebook, etc.

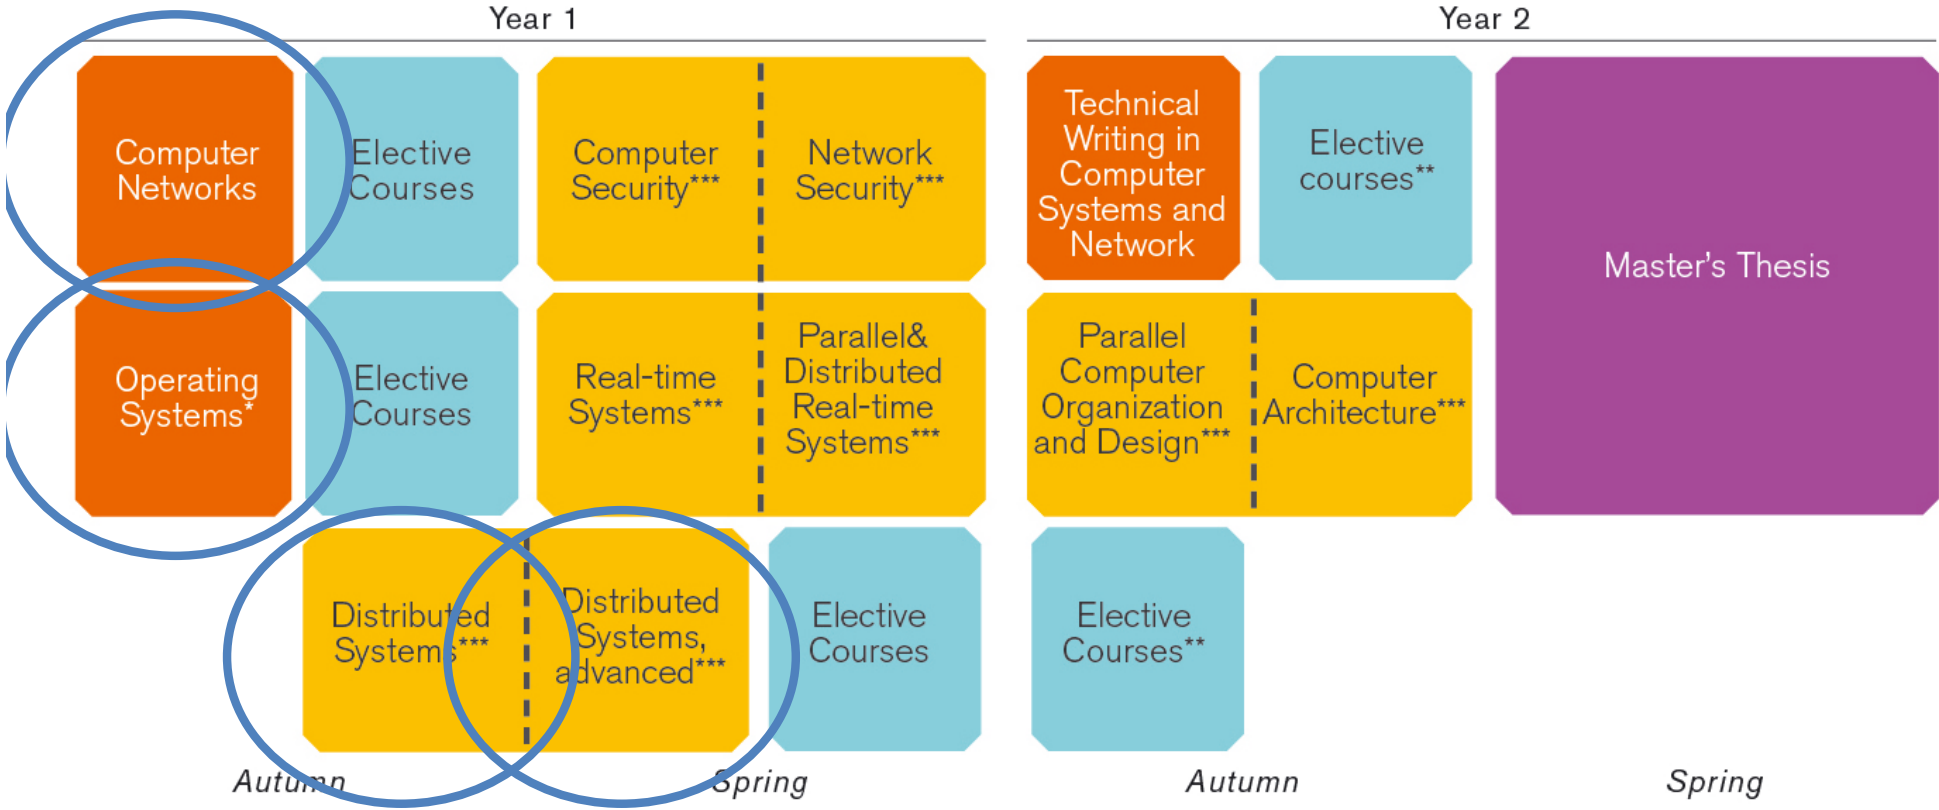


Modern Cars

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

Computer Systems and Networks



- Compulsory courses
- Elective courses
- Tracks

* Mandatory only if not taken before

** Recommended elective project courses: Autonomous and Cooperative Vehicular Systems (second study period), OCT Support for adaptiveness and Security in the smart grid (fourth study period)

*** Choose two out of these course tracks: Computer security, Real-time systems, Distributed systems and Computer architecture

Questions