CSE Seminar Series

Networks and Distributed Systems
(+ Data as ”bonus” - mainly Big)
2019-03-07

Marina Papatriantafilou
Networks and Systems Division
CSE Department
Chalmers & Gothenburg Un.

What is a Distributed System?

"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

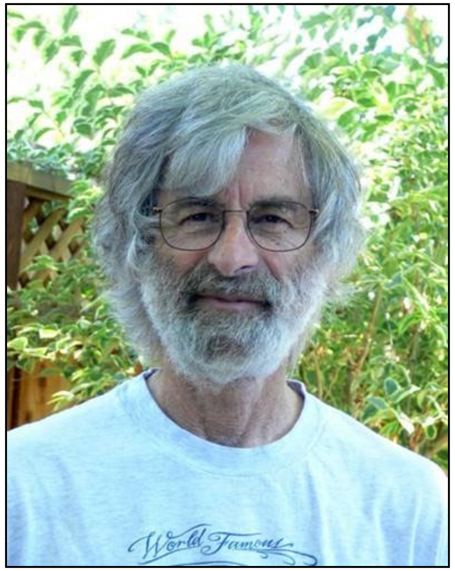

Leslie Lamport

A Distributed System

A set of computing&communicating processes,
collaborating for acheiving local and/or global goals

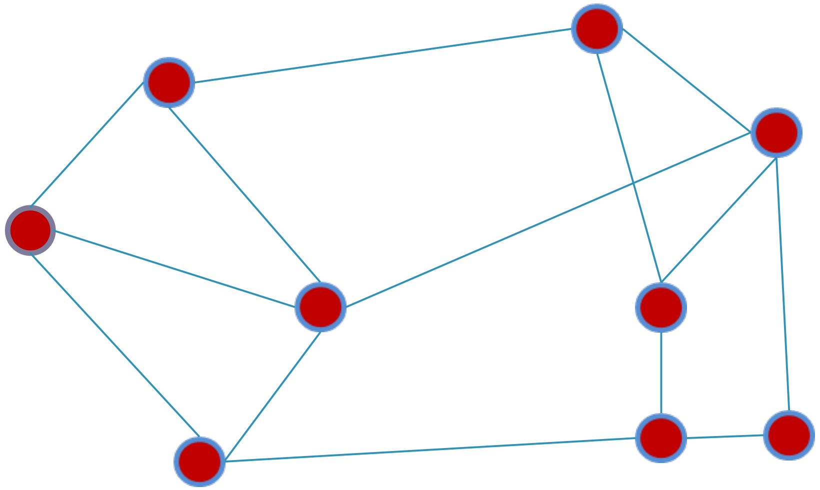

Distributed Systems?

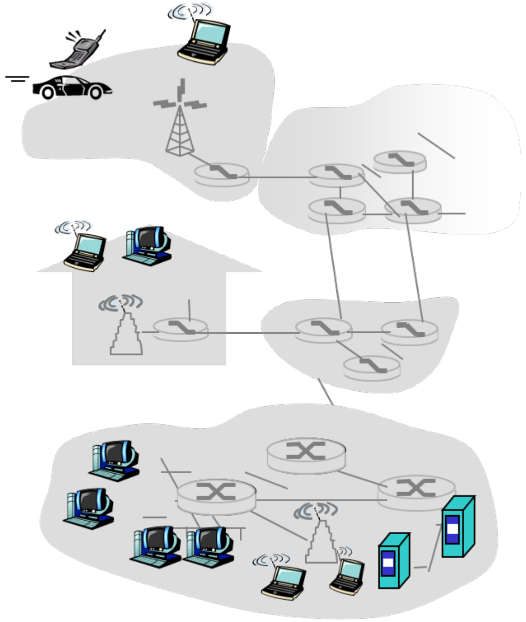

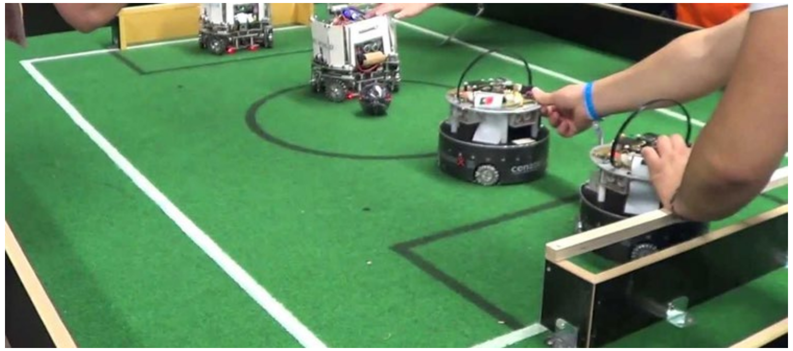

Figs: Computer Networking: A Top Down Approach by Kurose&Ross; robocup.org; Chalmers Gulliver prj by E. Schilller; ebgames.ca

Layered system perception

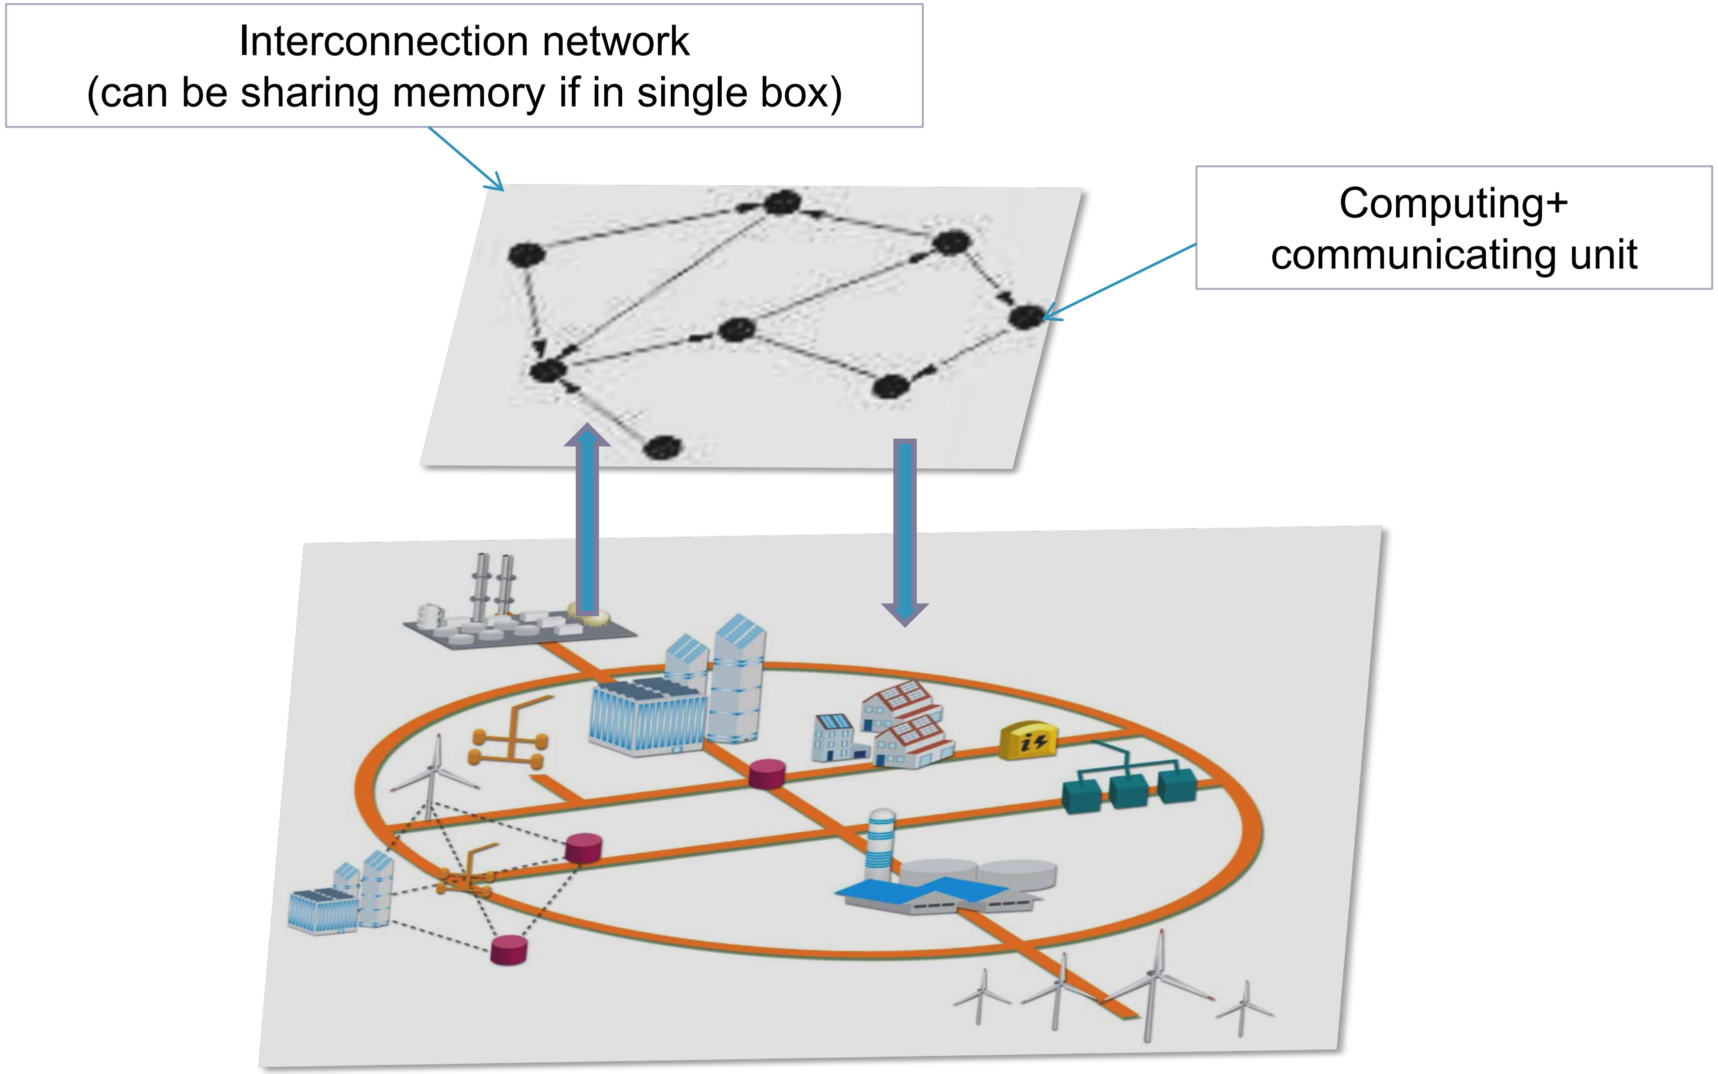

Distributed Systems vs. Networks

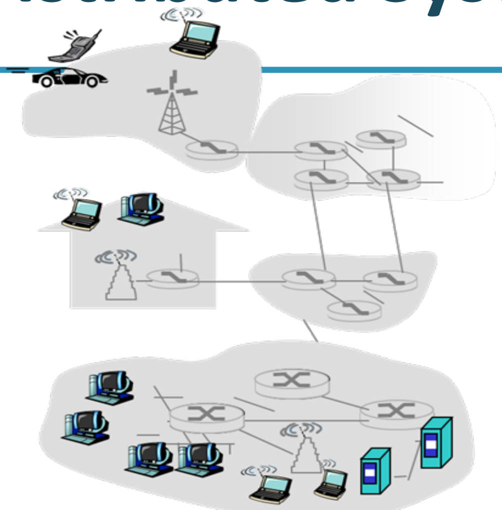

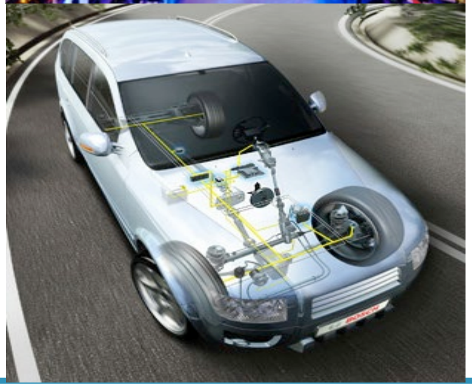

- Networking is worried about
  - Sending a message from here to there
  - Not what you do with the message
  - **We teach you how networks are built and how they function**

- Distributed Systems
  - Assume: There is a way to communicate
  - Focus: How you build a system using those messages
  - **We teach you what things to do with a network**

Inter-net-working, Data processing and Distributed Computing in interplay in IoT

A lot of data to be communicated, distributed, processed

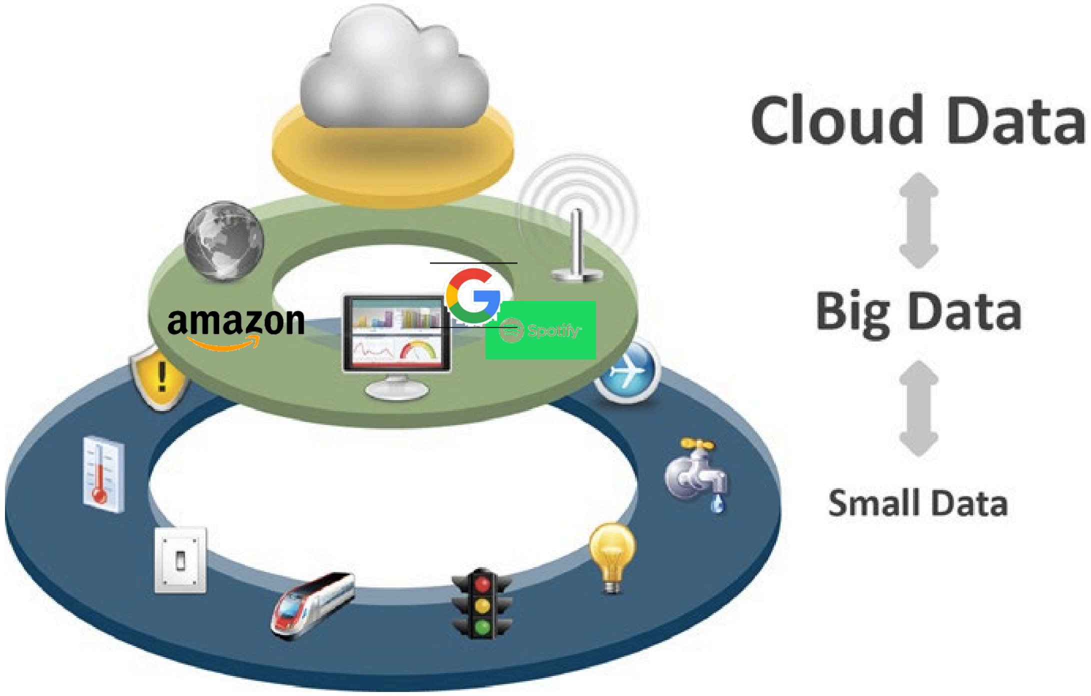

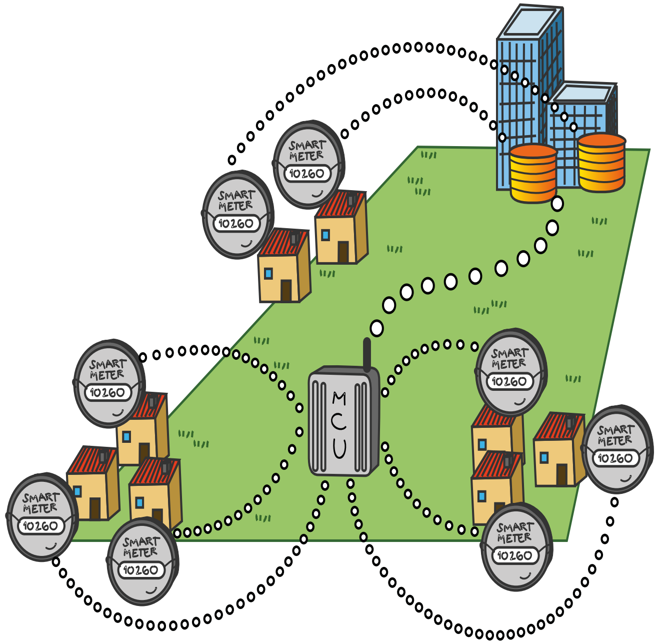

Example study topics in these domains

- Send, share data
- Aggregate-data/monitor @local-level
- Learn data-patterns @data-center, @local-level
- Ensure consistency/synchronization among copies @updates

Figs::/[/www.iebmedia.com;](http://www.iebmedia.com;)
Vincenzo Gulisano / Rocio Rodriguez

Let’s hit the road

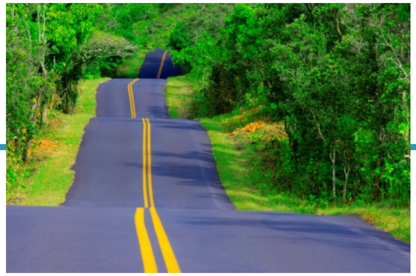

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects

Distributed system synchronization: once upon a time…

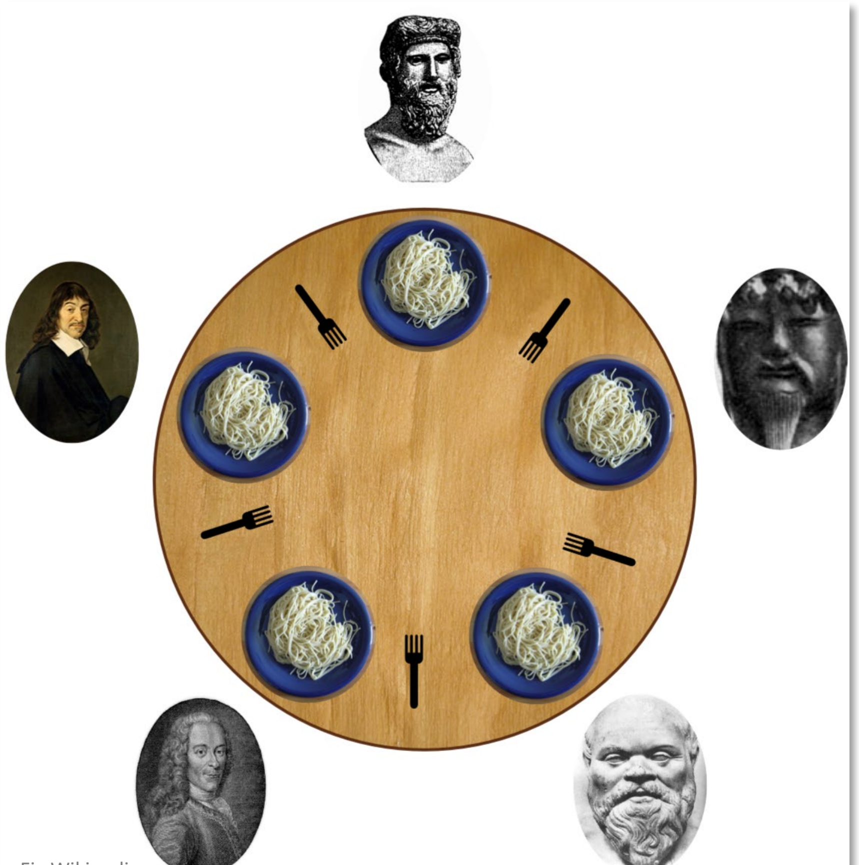
Fig Wikipedia

- [Dijkstra 1965]: **Dining philosophers:** example problem in concurrent algorithms&systems to illustrate synchronization issues and techniques for resolving them
- exam exercise ☺, presented in terms of computers competing for access to tape drive peripherals

”Internet”: once upon a time ….

1969

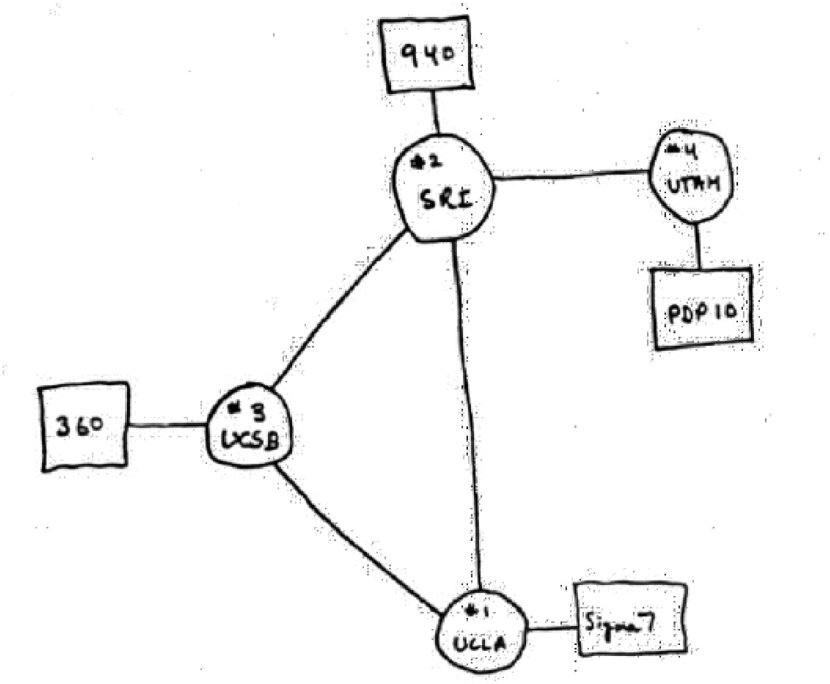

ARPANET begins…with a deployment at UCLA, Stanford, UCSB, and Utah (one computer per site)

Leonard Kleinrock (now prof Emeritus, UCLA) about the Internet:

What was going through your mind when you sent the first host-to-host message (from UCLA to the Stanford Research Institute)?

Frankly, we had no idea of the importance of that event. We had not prepared a special message of historic significance, as did so many inventors of the past (Samuel Morse with "What hath God wrought." or Alexander Graham Bell with "Watson, come here! I want you." or Neal Amstrong with "That's one small step for a man, one giant leap for mankind.") Those guys were *smart*! They understood media and public relations. All we wanted to do was to login to the SRI computer. So we typed the "L", which was correctly received, we typed the "o" which was received, and then we typed the "g" which caused the SRI host computer to crash! So, it turned out that our message was the shortest and perhaps the most prophetic message ever, namely "Lo!" as in "Lo and behold!"

Earlier that year, I was quoted in a UCLA press release saying that once the network was up and running, it would be possible to gain access to computer utilities from our homes and offices as easily as we gain access to electricity and telephone connectivity. So my vision at that time was that the Internet would be ubiquitous, always on, always available, anyone with any device could connect from any location, and it would be invisible. However, I never anticipated that my 99-year-old mother would use the Internet—and indeed she did!

and later …

1969

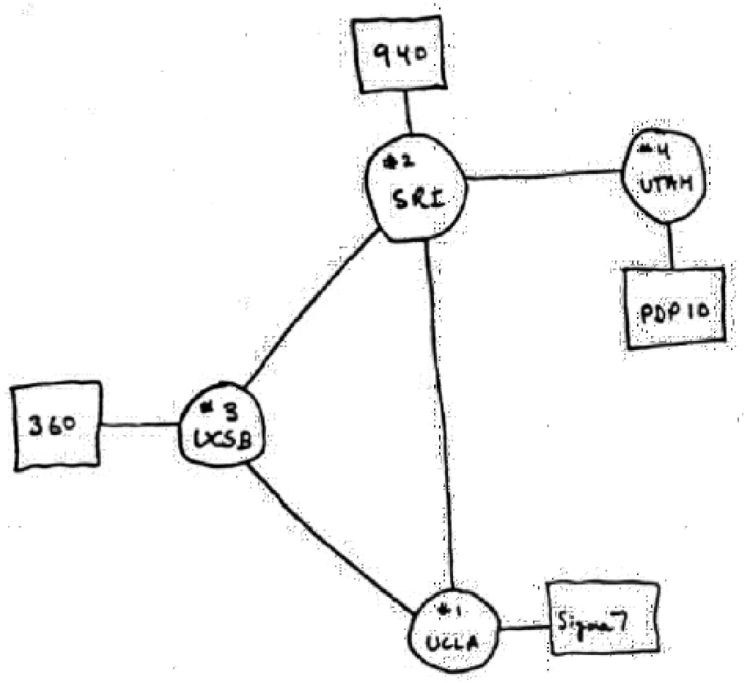

ARPANET begins…with a deployment at UCLA, Stanford, UCSB, and Utah (one computer per site)

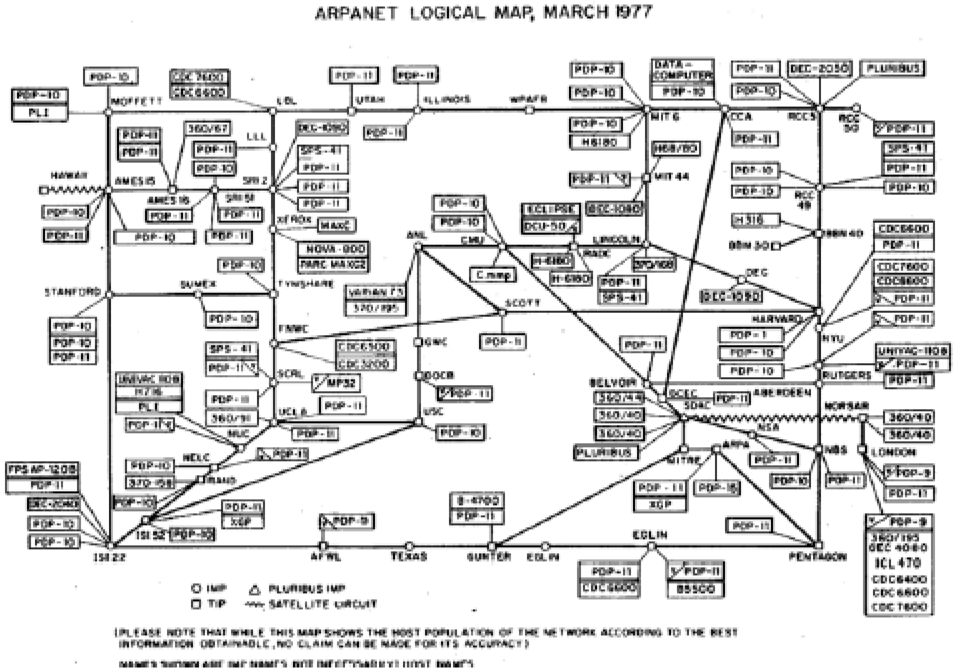

←1977

Adapted from slides on the
Computer Systems and Networks Masters program
by O. Landsieldel

… and later …

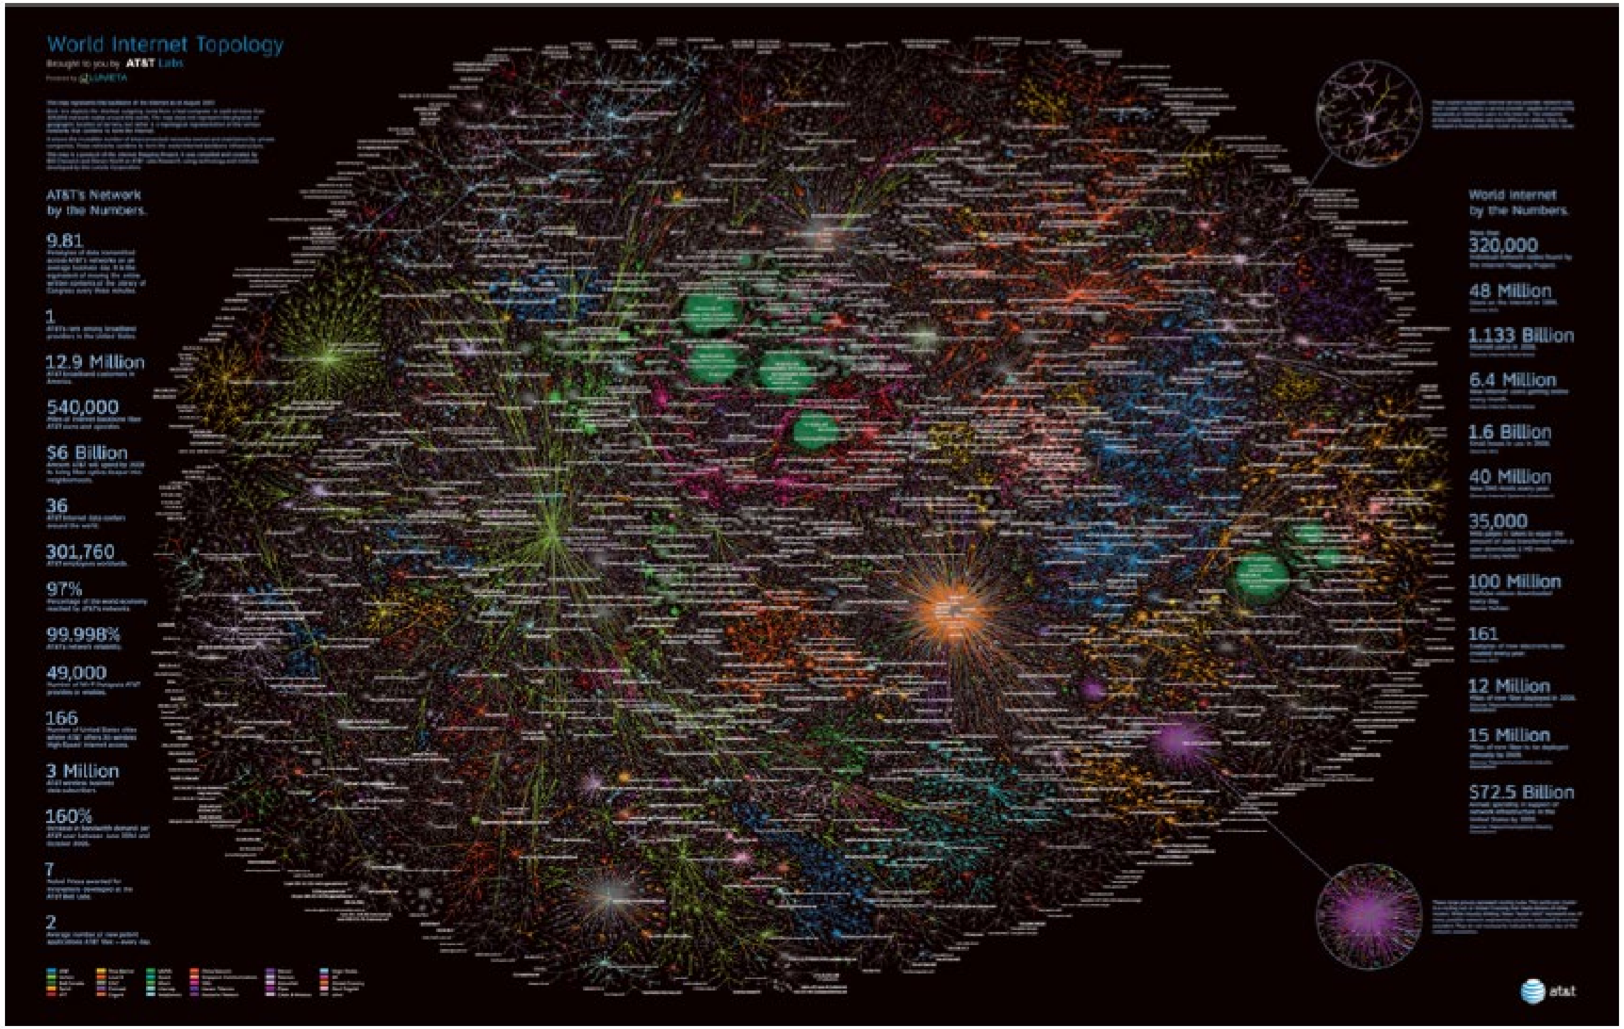

Internet 2007 (just the backbone)

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

How was this enabled? (examples)

1974

TCP / IP defined by Vint Cerf & Bob Kahn

2004: both received the Turing Award

1971

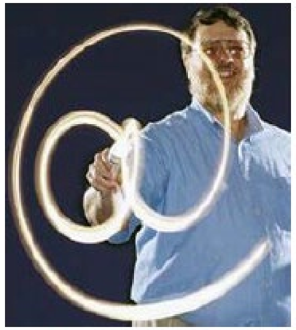

Ray Tomlinson creates first email program

1984

Paul Mockapetris introduces DNS

How was this enabled? (examples cont)

70-80's: foundations about time and coordination in distributed systems; *concurrent R/W shared data; wait/lock-free algorithms* [Courtois, Heymans, Parnas] [Misra] [Lamport] : asynchronous HW?!

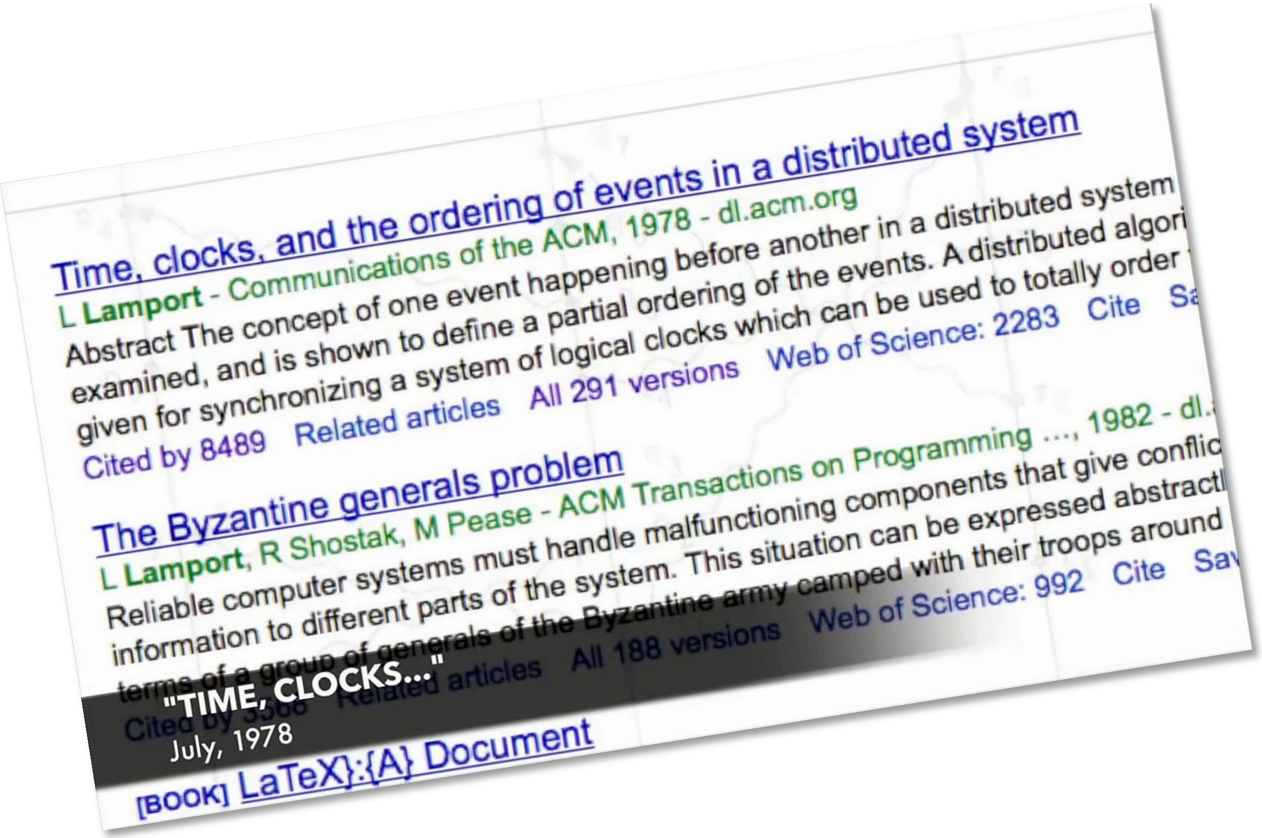

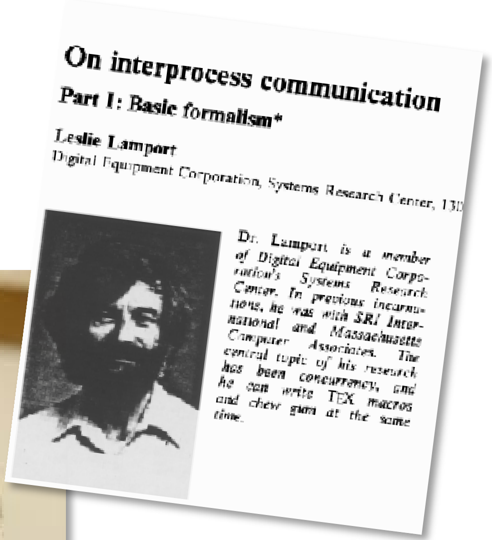

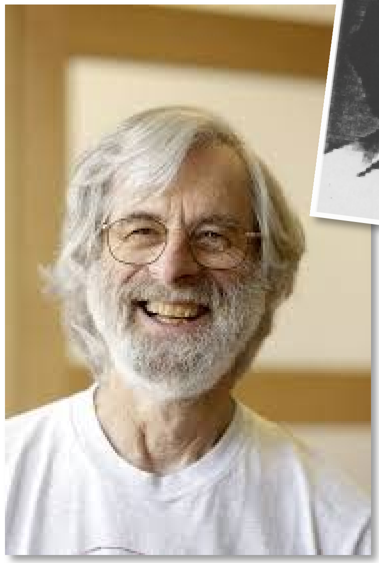

Leslie Lamport:
Turing award winner 2013 for his work on distributed systems synchronization, consistency, robustness

How was this enabled (examples cont):

1989 – The Web Emerges

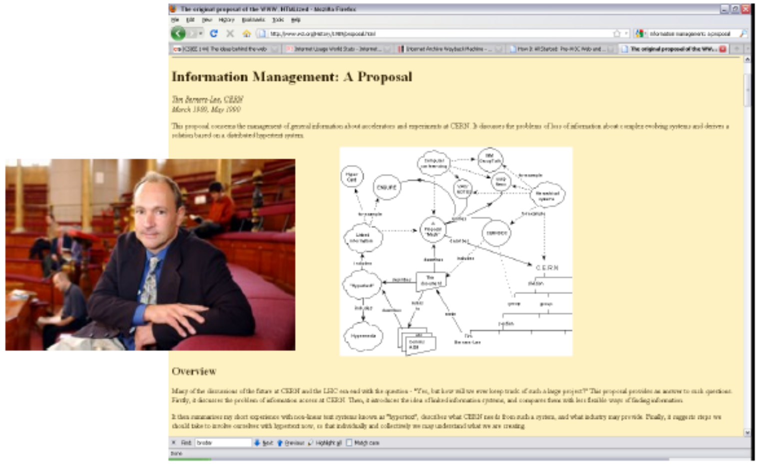

Tim Berners-Lee writes “Information Management: A proposal” at CERN

1990

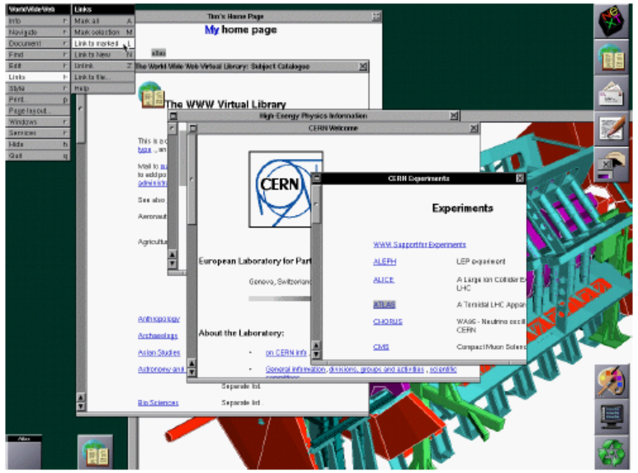

First browser developed at CERN

1993

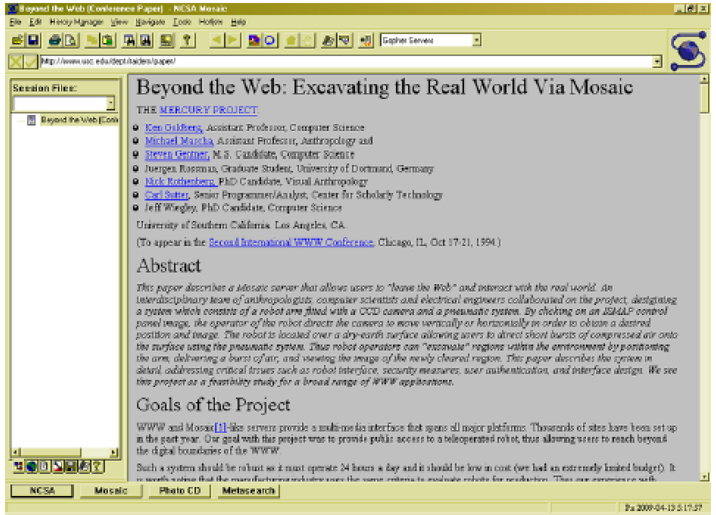

Mosaic became the first graphical browser

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

1995+

Amazon arrives and the commercialization of the web begins

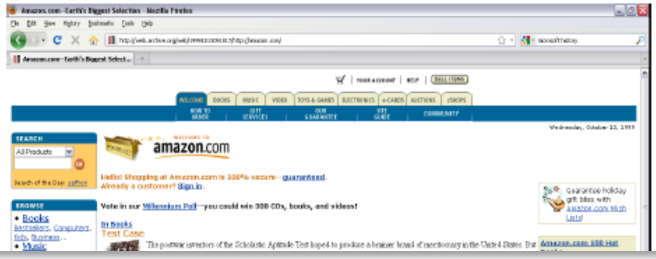

Amazon circa 1999

Adapted from slides CSE Seminar 2018 by O. Landsieldel

How was this enabled? (examples cont)

2000's: p2p applications, social networks, Content Distribution Networks, … ;
multi/many-core data processing; asynchronous hardware!

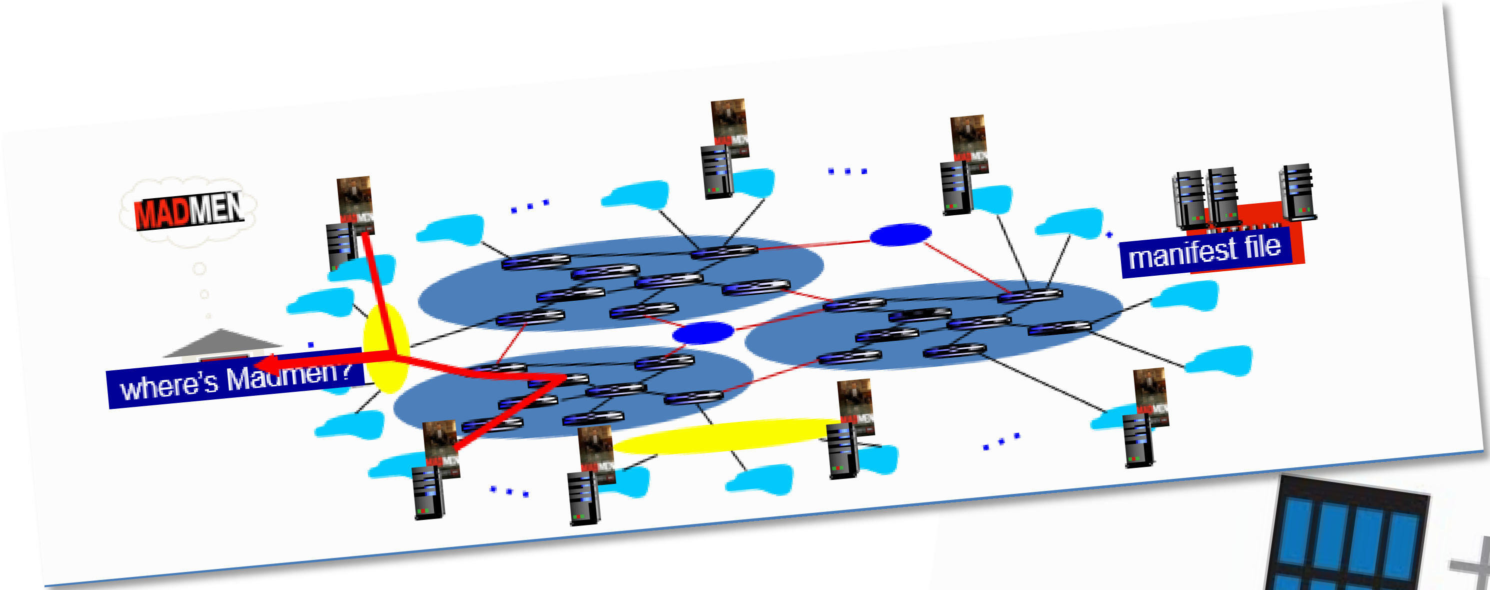

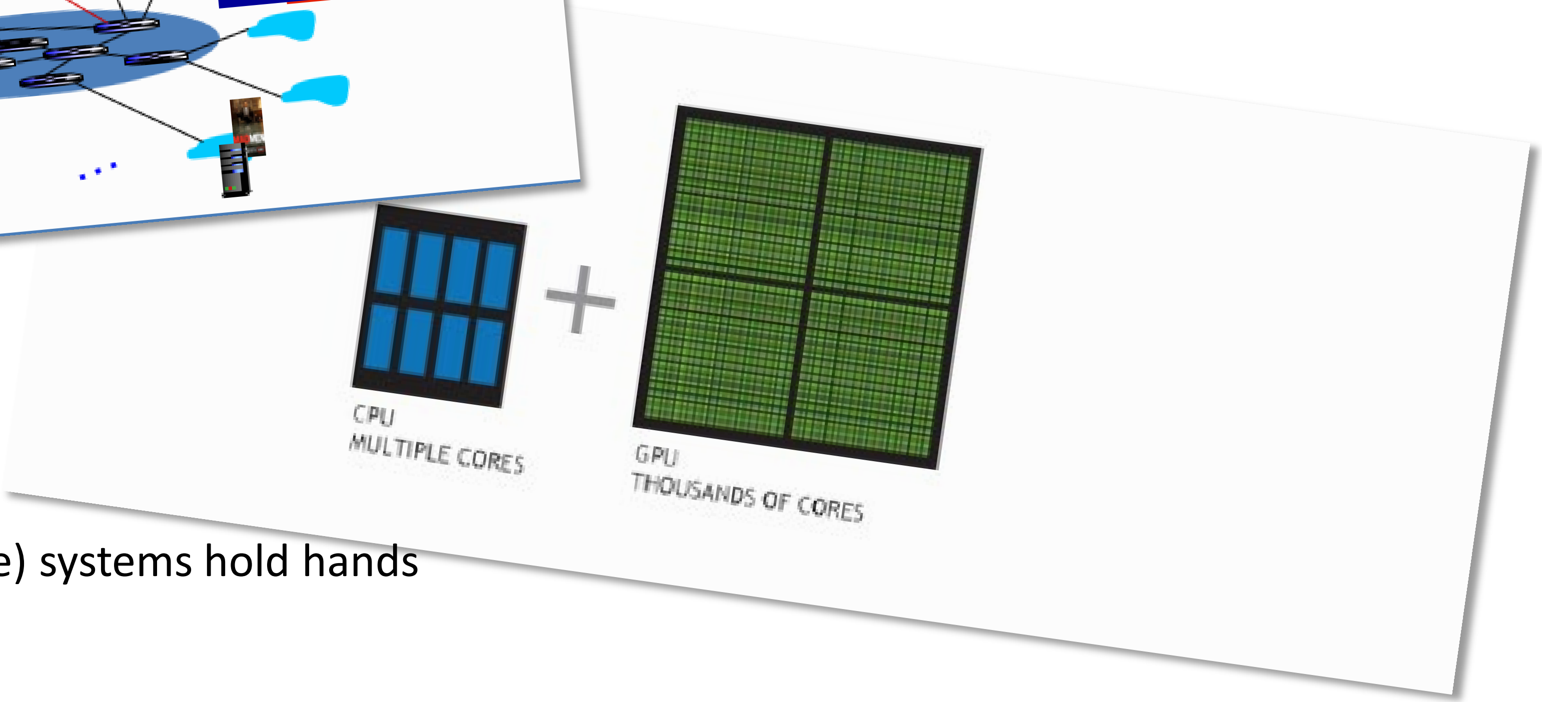

Distributed (inclulind parallel, multicore) systems hold hands with Networks

Roadmap

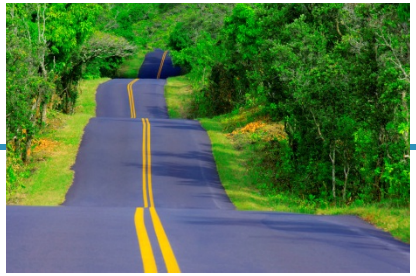

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects

The Future is Distributed

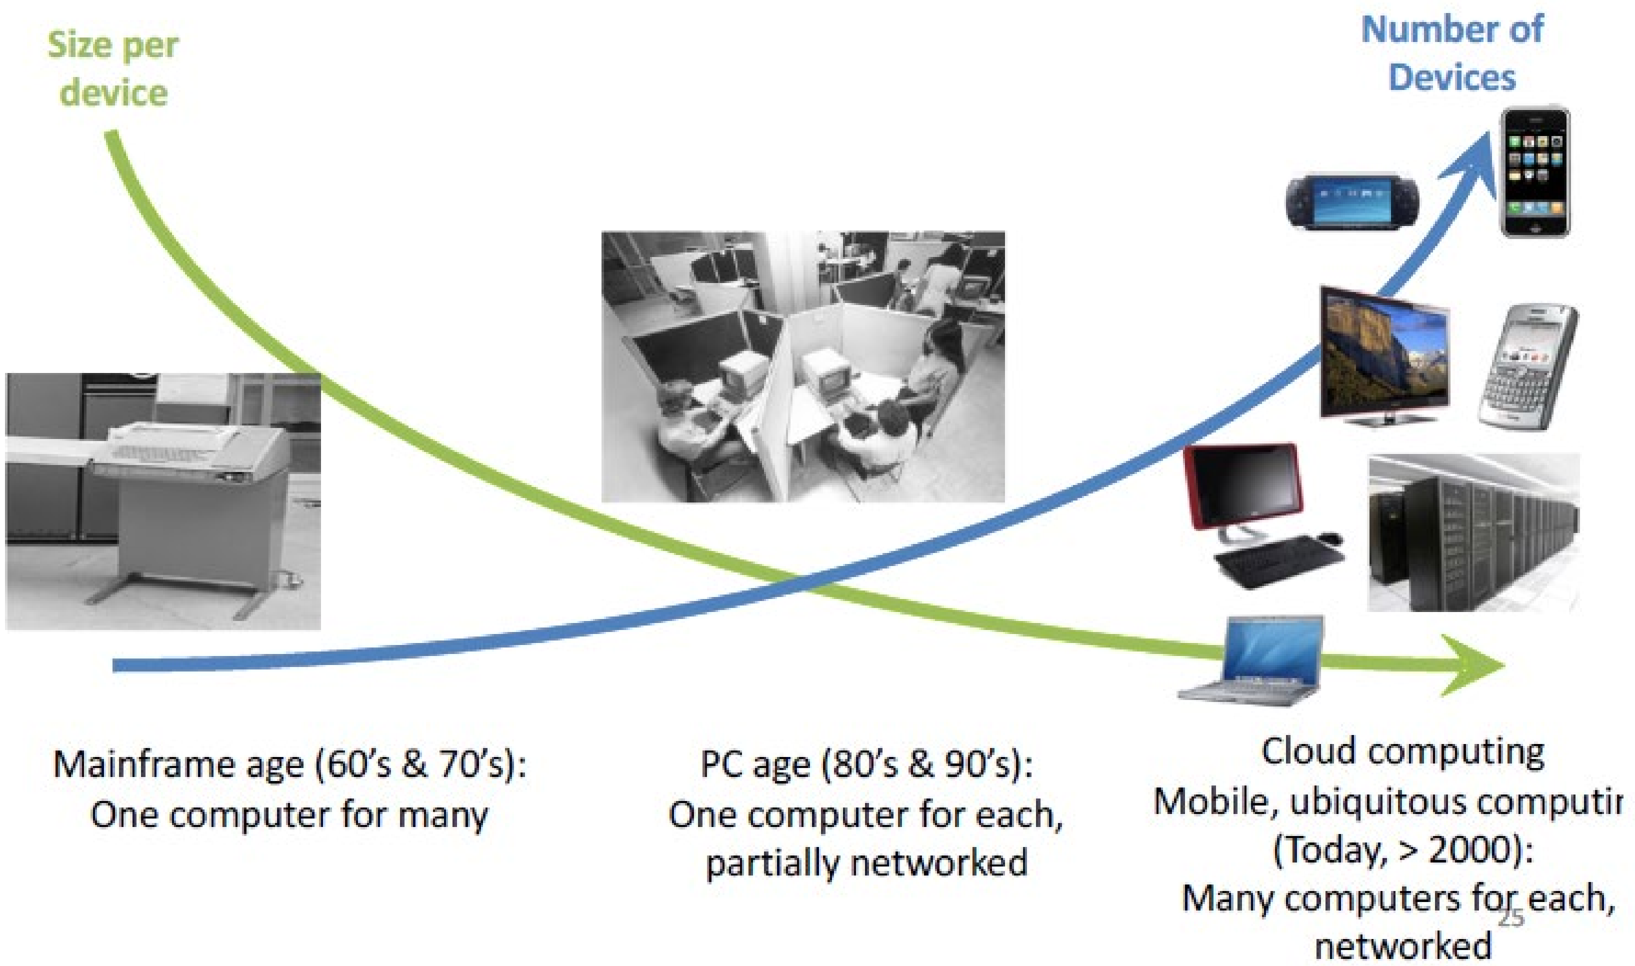

Manycores,synchronous hardware

Adapted from slides CSE seminar O. Landsiedel

The Future is Distributed

- Networks and Distributed Systems touch significant aspects of daily life!
  - Integral building block for our networked society

What Makes a Smart City?

Multiple Applications Create Big Data

Application	Data
Connected Plane	40 TB per day (0.1% transmitted)
Connected Factory	1 PB per day (0.2% transmitted)
Public Safety	50 PB per day (<0.1% transmitted)
Weather Sensors	10 MB per day (5% transmitted)
Intelligent Building	275 GB per day (1% transmitted)
Smart Hospital	5 TB per day (0.1% transmitted)
Smart Car	70 GB per day (0.1% transmitted)
Smart Grid	5 GB per day (1% transmitted)

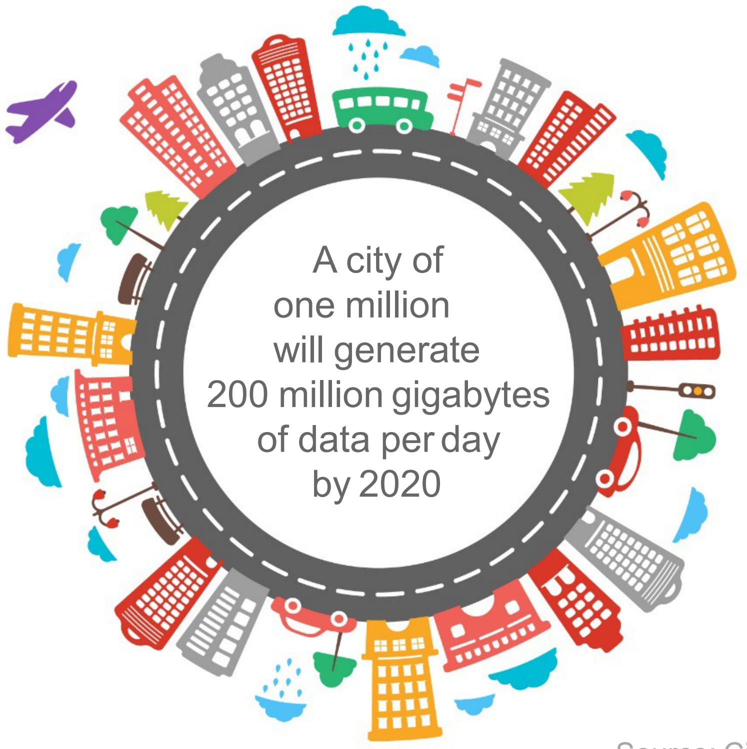

Back to Index

Source: Cisco Global Cloud Index, 2015–2020

 45

Data Created vs. Data Center Traffic

Data Created Outpaced

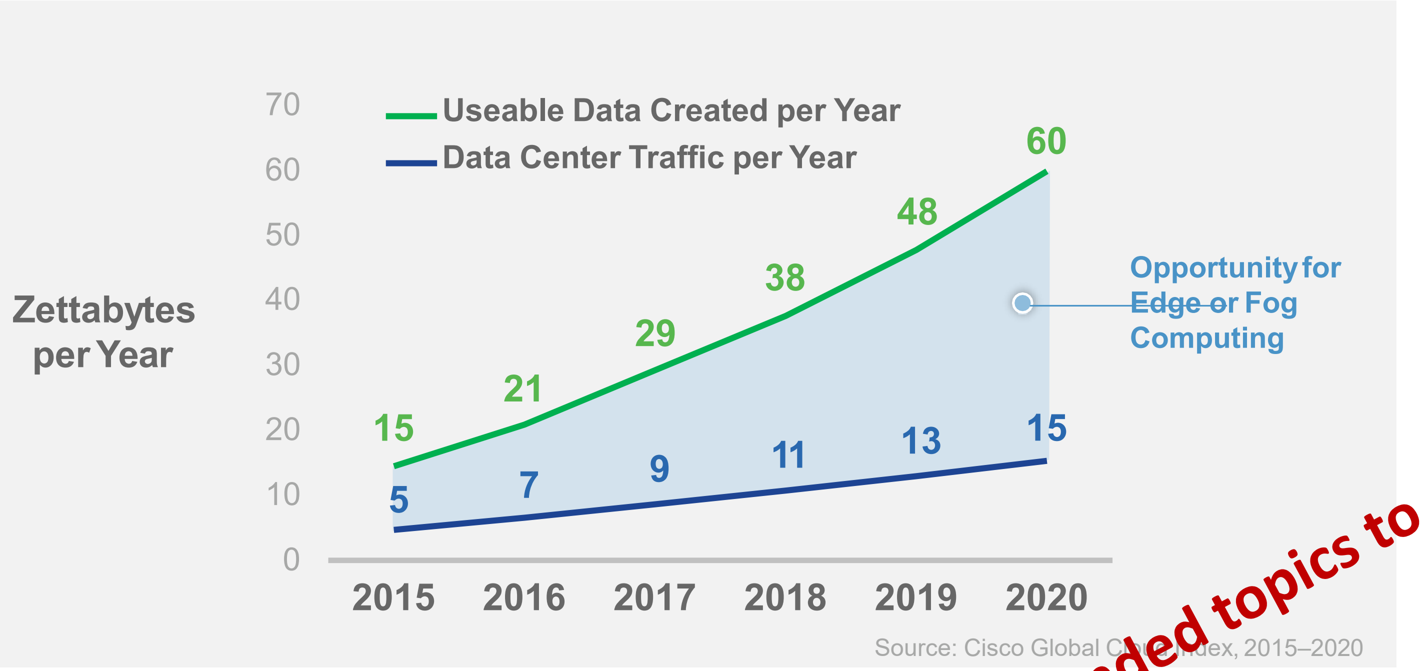

Recommended topics to study ;)

Michael Stonebraker:
Turing award winner 2014 for his work on stream-data procsessing, enabling <u>*in-network data processing*</u> and revolutionizing database systems

Summary:

Networks & Data: Big =>

Distributed Computing and Systems: ”break” big problems into smaller, local ones

Roadmap

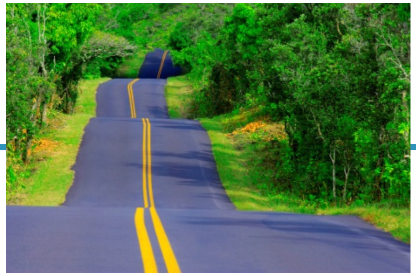

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects

Putting things together

Offererd in our curriculum:

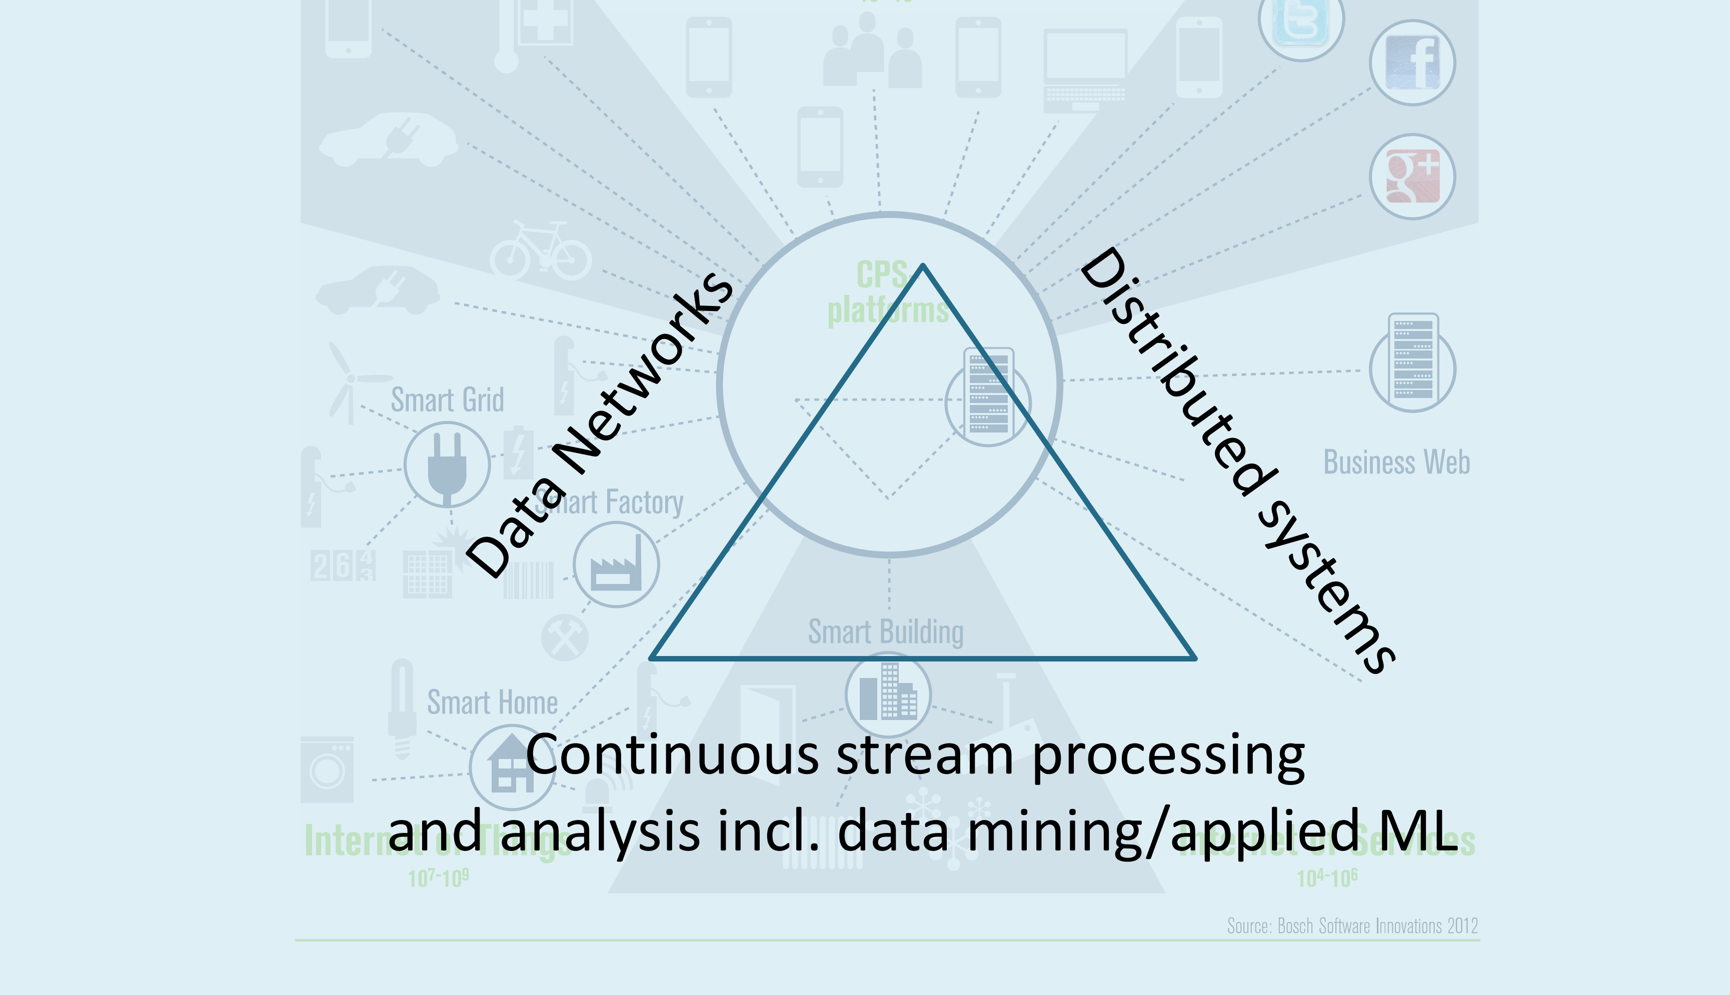

e.g.: MS prorgam @CTH, Specialization options @GU

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

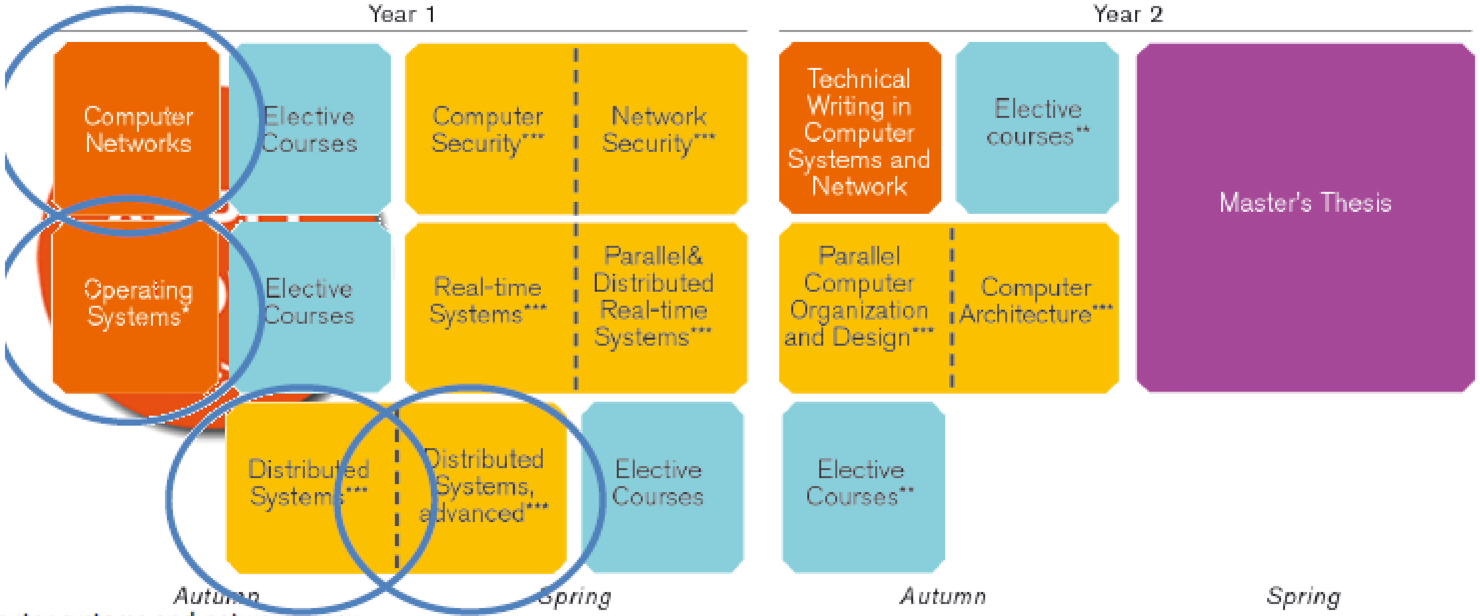

https://www.chalmers.se/en/education/programmes/masters-info/Pages/Computer-systems-and-networks.aspx

Compulsory courses

Elective courses

Tracks

* Mandatory only if not taken before

** Recommended elective project courses: Autonomous and Cooperative Vehicluar 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

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, LP1, 7.5hec

Roadmap

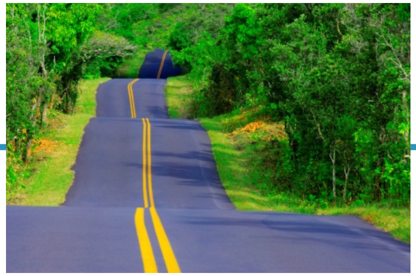

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects

Operating Systems Course

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

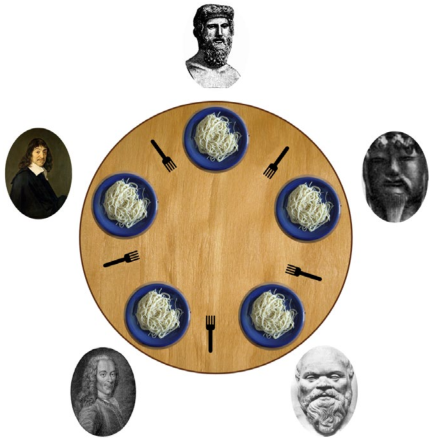

Threads management in multicore CPUs

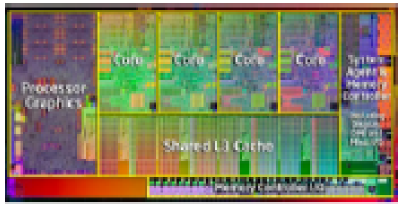

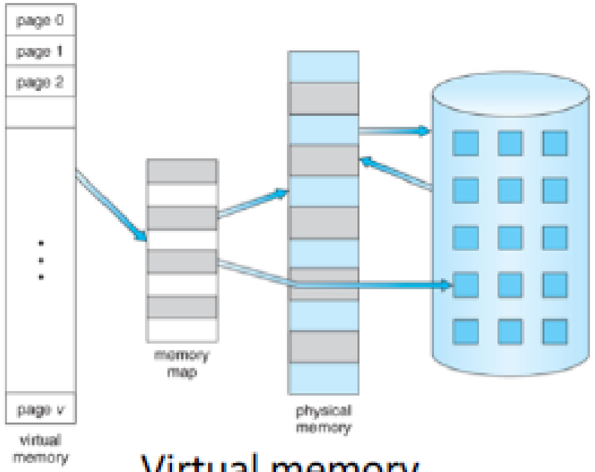

Virtual memory

Security

Courses Distributed Systems

You learn: How to build *large-scale distributed systems* and the associated challenges, *principles & practice*

- eg CAP thm [Brewer's conjecture 1998; Gilbert&Lynch2002 proof]

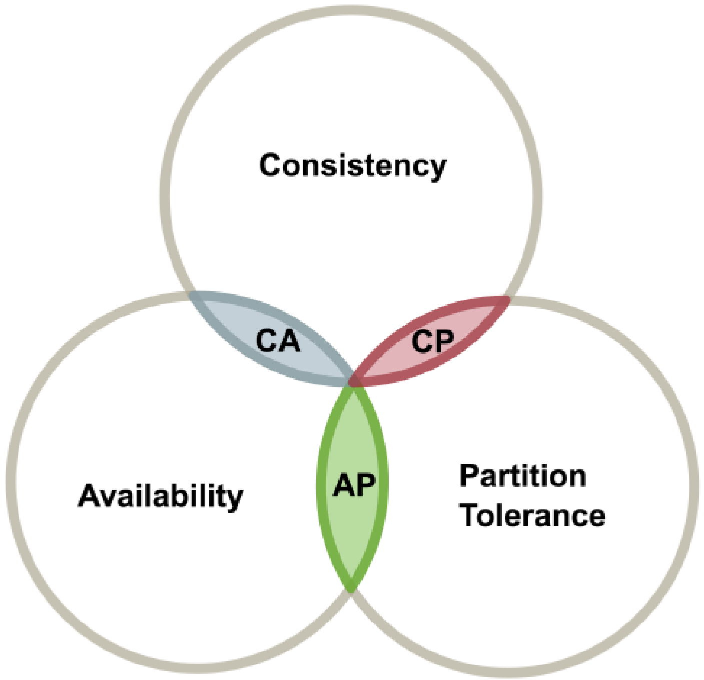

amazon

Spotify

Eg. *applied in Spotify's, Amazon's systems*: partitioning of servers happens!
=> eventual consistency in distributed state [CRDTs: Shapiro et-al]

Courses Computer Communication

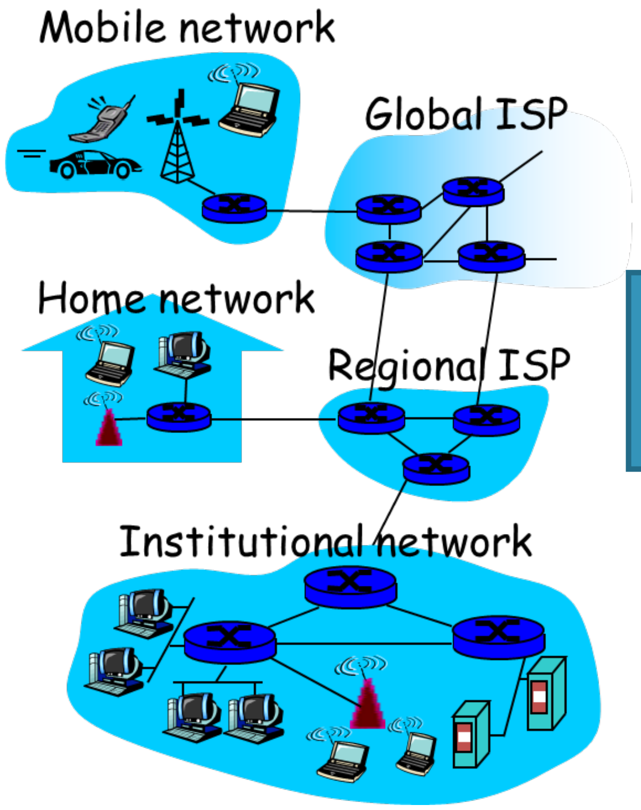

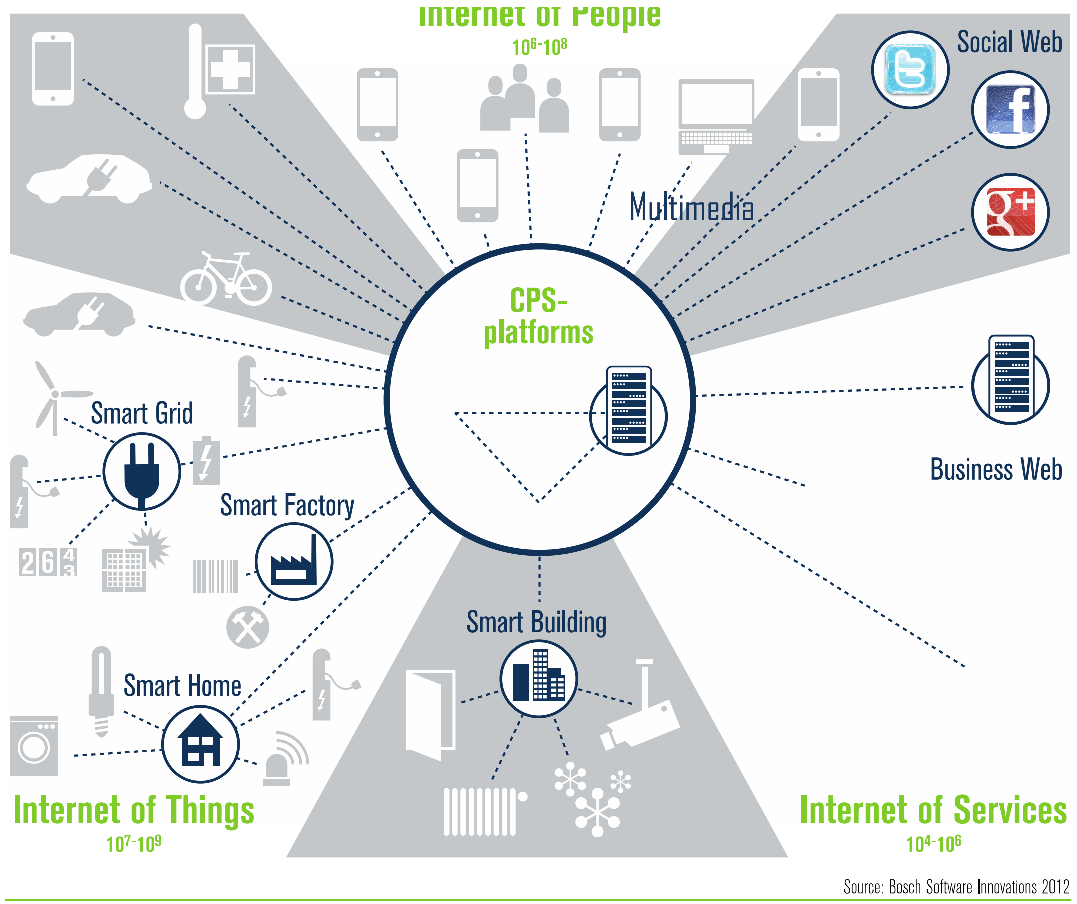

You get knowledge to build the basis …

… to follow continuous evolution ….

Courses Computer Communication, example content

Software-Defined Networks: logically separated control plane

i.e. you learn how to support
Networks with Distributed systems
and Data Processing

compute tables separately,
in data-center/distributed system
and distribute

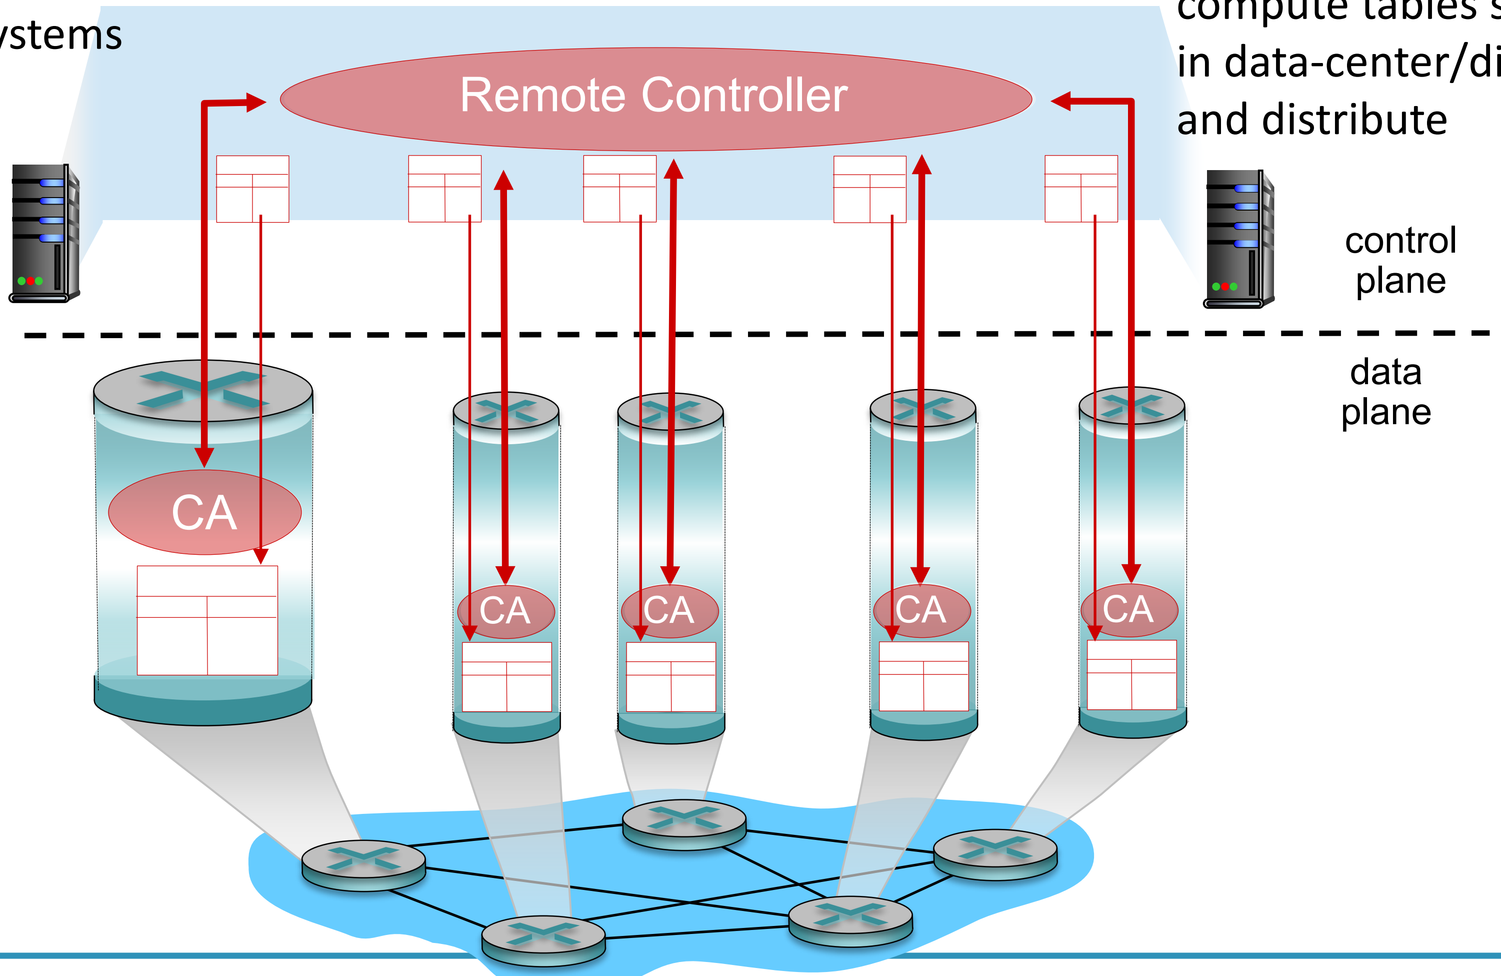

Project-based courses: eg ICT in data-driven cyberphysical systems

Example projects: e. in the context of smart grid systems

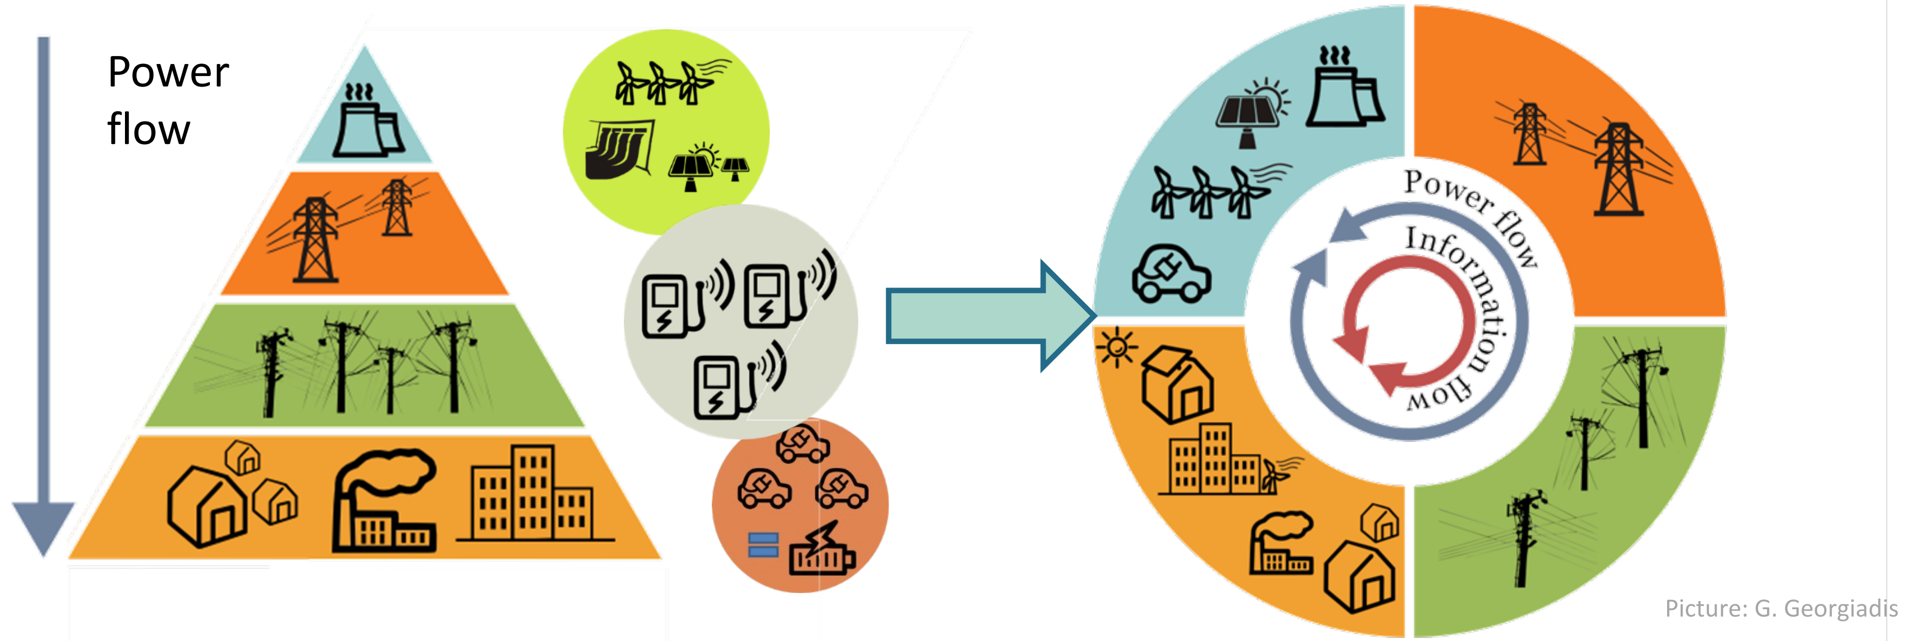

Context: paradigm shift
from “adapt generation to demand “to “adapt consumption to availability +”

Project-based courses: eg ICT in data-driven cyberphysical systems

Example projects:

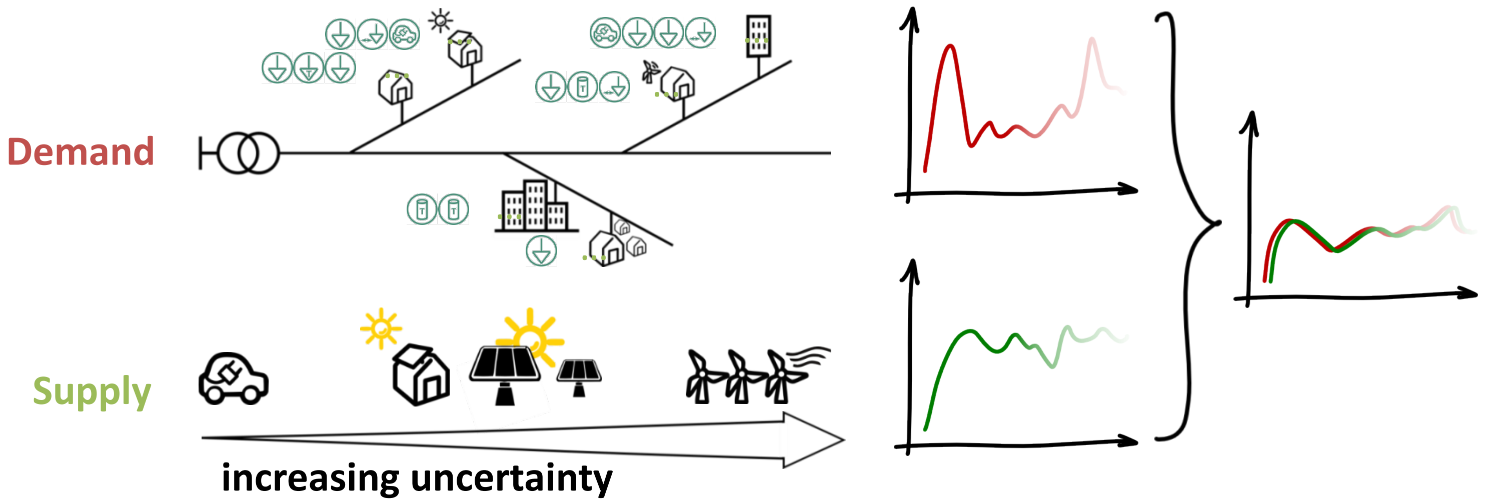

The project: **Adaptive, autonomous and collective load balancing**

The goal: Shaping streaming demands to streaming supply, taking into account energy storage options and consumption/generation data

Project-based courses: eg ICT in data-driven cyberphysical systems
Example projects:

The project: Reliability of RT object detection:

Goal: understand limits of ML-processing with noisy data on embedded GPU platforms

Figure 9: Odroid XU4

(a) Cloudy image with YOLOv3 benchmark

(b) Same image with YOLOv3-tiny and random noise

(c) Same image with YOLOv3-tiny and gaussian noise

(d) Same image with YOLOv3-tiny and artificial rain noise

Figure 13: Comparison of YOLOv3 and YOLOv3-tiny annotation on a cloudy image with different noises

Figs: report A. Mosshammer, C vRosen Johansson, M. Romain

Highlights

Guest lectures by: Ericsson research, Volvo, Spotify, FlexLink, ABB, Göteborg Energi, Svenska Kraftnät and more

Example empolyments of graduates:

Spotify, Volvo, Zenuity, Ericsson research, RISE, ABB, academic institutes internationally (Max Plank Inst, Purdue Un, EPFL, …)

- Masters projects with relevance for industry and academia, including publishable work
- Comments of appreciaton in course evaluations
- Continued contacts after graduation

Roadmap

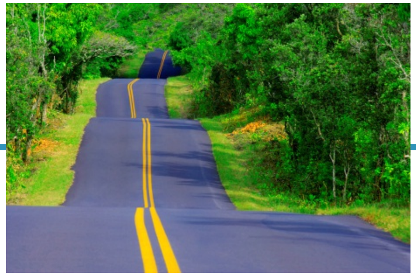

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects

Our research DCS @NS division (approx 30 pers):

Strong publications records (incl. PODC, ICDE where Turing award talks have been delivered

Parallel &stream computing

Distributed systems, IoT

Security, dependa bility

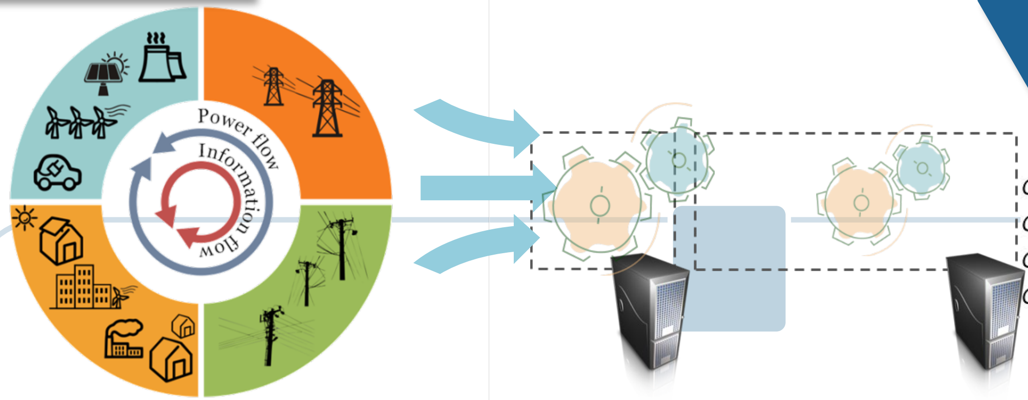

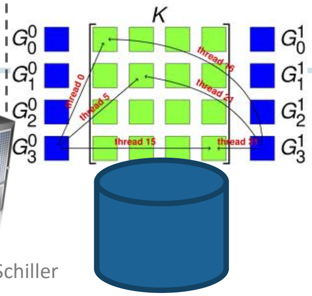

Fig. Giorgos Georgiadis/ Vincenzo Gulisano / Rocio Rodriguez / Chalmers Magazine Gulliver project/Elad Schiller

Energy, production, transport

- data-driven distributed monitoring, resource matching

- Microgrids demo work

On-the-fly/stream data processing & analysis

-Data validation, monitoring, …

- -Security, privacy

Energy/efficient computation

parallel/multicore/ processing, incl. on embedded processors

Vehicular systems

- data-driven situation-awareness
- communication

&coordination, e.g. virtual traffic-lights

- Gulliver testbed

2 research awards by ACM DEBS Distributed Event Processing to DCS@Chalmers

Results adopted by nVIDIA, Intel, C++

Example projects/results: Geospatial monitoring

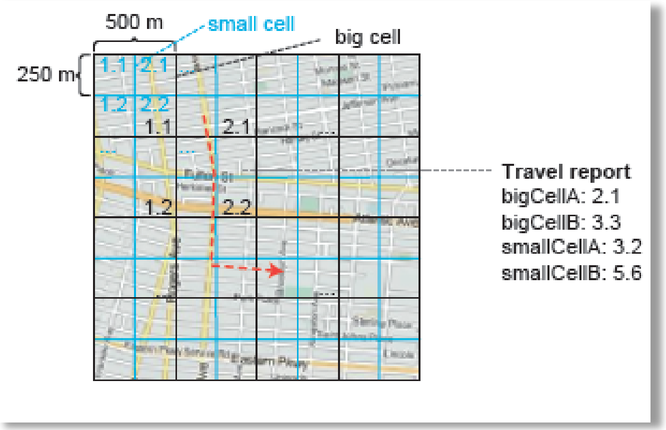

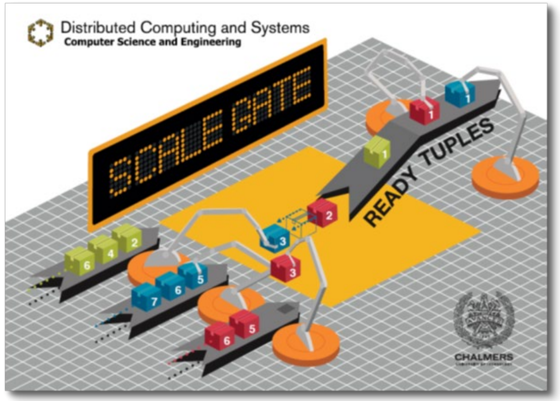

Fig. Vincenzo Gulisano / Rocio Rodriguez

2015: DETERMINISTIC REAL-TIME ANALYTICS OF GEOSPATIAL DATA STREAMS THROUGH SCALEGATE OBJECTS
http://www.chalmers.se/en/departments/cse/news/Pages/debs2015.aspx
ACM DEBS 2015 Grand Challenge best solution award

- Top k frequent routes @NY, profitable cells (near-real time window-based streaming)
- > 110,000 tuples/sec throughput, < 46 msec latency, 1yr data processed in 11 min

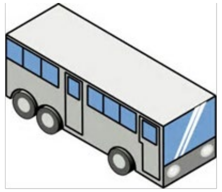

Join Traffic Streams from Dublin

2017: Maximizing Determinism in Stream Processing Under Latency Constraints
ACM DEBS 2017 best paper award
[collab of our team with Athens Uni. of Business]

Fig. Yiannis Nikolakopoulos

Example research projects/results: Shared Data Objects for ultra-efficient processing on many-core embedded systems

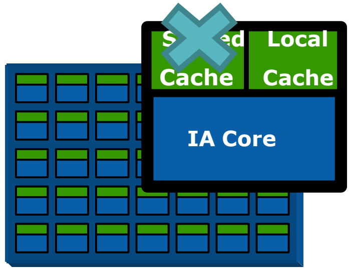

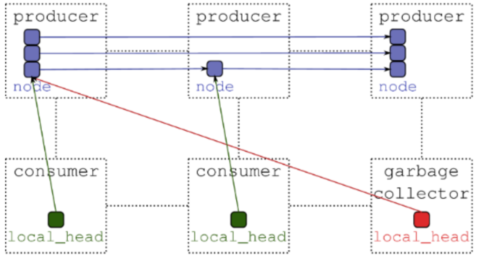

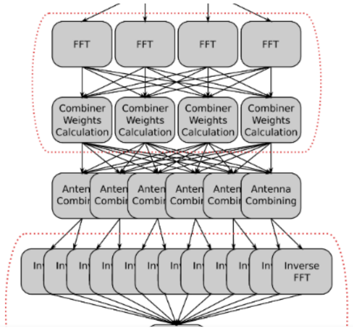

H/W Evolves

- Networks on chip (NoC)
- Limited support for synchronization primitives

New Software for Concurrent access data objects

Eg ScaleGate [CGNPT14]

LTE (FFT) Uplink benchmark

- Focus on communication patterns: n-to-1, n-to-n
- Up to **42% improvement** in execution cycles

ERICSSON

STAMINA: Processing & analysis of data STreams in AMI for Awareness and Adaptiveness in electricity grids

WASP WALLENBERG AUTONOMOUS SYSTEMS PROGRAM

Possibile quality correlation from distribution and AMI

Distributed data pre-processing and validation

Fig. Vincenzo Gulisano / Rocio Rodriguez

Data-driven alerts analysis NTL / risks

Facilitate automated funtionality; ”lower-voltage SCADA”

AMI: Advanced Metering Infrastructure

Example recent work/results:
Continuous parallel processing of LIDAR data

- Lidar provides cloud of 3D points: high rate sensor (MBps)
  - It requires different processing tools (filtering, clustering, segmentation, …)

- New: We enabled possibilities to extract useful information from raw data, in real-time streaming, even on embedded hardware [ICDCS2018]

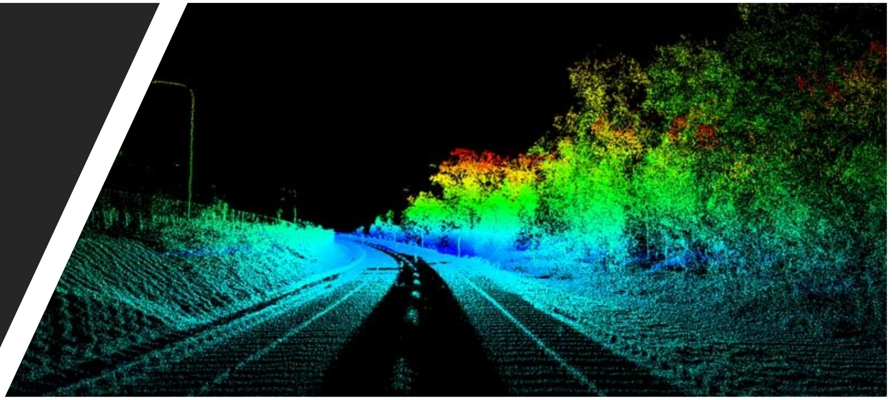

Fig: Hannah Najdataei

Recent & current related research projects of the team

1. **EXAMINE:** Extracting information out of big data in Advanced Metering Infrastructure (Göteborg Energii);
2. **FiC:** Future Factories in the Cloud (SSF);
3. **EU/FP7 EXCESS** on energy-efficient computation in embedded devices (EU; Chalmers' team coordinator)
4. **EU/FP7 KARYON** on safety Kernels in vehicular systems;
5. **EU/FP7 CRISALIS** on security of critical infrastructures;
6. **RICS:** on resilient information and control systems (MSB);
7. **EU Horizon 2020 United Grid** - with CTH Elteknik
8. **Gulliver**: a test-bed for developing, demonstrating and prototyping vehicular systems (SAFER, AoA Trasnsport);
9. **Scheme:** on Software abstractions for heterogeneous multi-core computers (SSF);
10. **EPOC** (Energy on Campus, Chalmers Area of Advance Energy and Building Future)
11. **iTRANSIT** on intelligent traffic management systems (FFI);
12. Fine-grain synchronization and memory consistency in parallel programming (VR);
13. Big Data and IoT for Sustainable living (Sw. Energy Agency);
14. Data-driven and Distributed Algorithms for Safe and Sustainable Traffic (SAFER and Chalmers Transport);
15. Adaptive energy dispatch in Smart Grids (E.ON and Swedish Energy Agency);
16. Concurrent Data Structures for Heterogeneous Parallel Programming (VR);
17. **D-SAS**: Data summaries and stream processing for autonomous systems (Wallenburg Autonomous Systems and Software Programme - WASP);
18. **STAMINA:** Processing & analysis of data STreams in AMI for Awareness and Adaptiveness in electricity grids (GE & WASP);
19. **OODIDA**: On-board Off-board Distributed Data Analytics (Vinnova FFI);
20. **HARE**: Self-deploying and Adaptive Data Streaming Analytics in Fog Architectures (VR);
21. Models and Techniques for Energy-Efficient Concurrent D[illegible]igns (VR)

Systems Security

Distri-buted systems, IoT

Parallel &stream computing

Faculty members responsible/involved:

Magnus Almgren

Vincenzo Gulisano

Marina Papatriantafilou

Elad Schiller

Philippas Tsigas

CHALMERS AoA Building Futures, Energy, ICT, Transport

Roadmap

- Overview
- Some history
- Present and projection to the future
- Possibilities in our curriculum
- Some course-related info
- Our research team and highlights of results & projects
- ***Next: See you at the courses!*** ☺