Storskaliga Distribuerade System och Nätverk (Large-scale Distributed Systems and Networks)

Slides by Niklas Carlsson

Systems thinking

- We want to understand the full system and the ecosystem it operates within; e.g.,
  - Understanding the full system
  - Looking at the parts and how they interact
- This course provide many examples ...

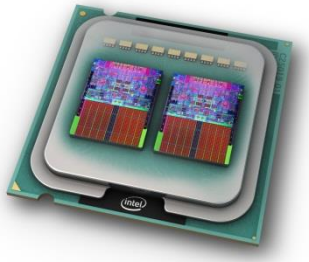

Components, overall system, service(s)

- Components together provide some service(s)
  - Typically want good overall performance
- Data storage and processing often distributed
  - Data/information stored and moving between components
  - Processing in different components
- How do we best design, utilize, and run systems?

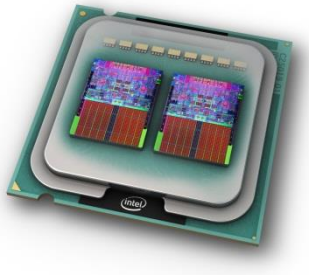

Embedded Systems

- We find microprocessors everywhere …

Multicore and Manycore Systems

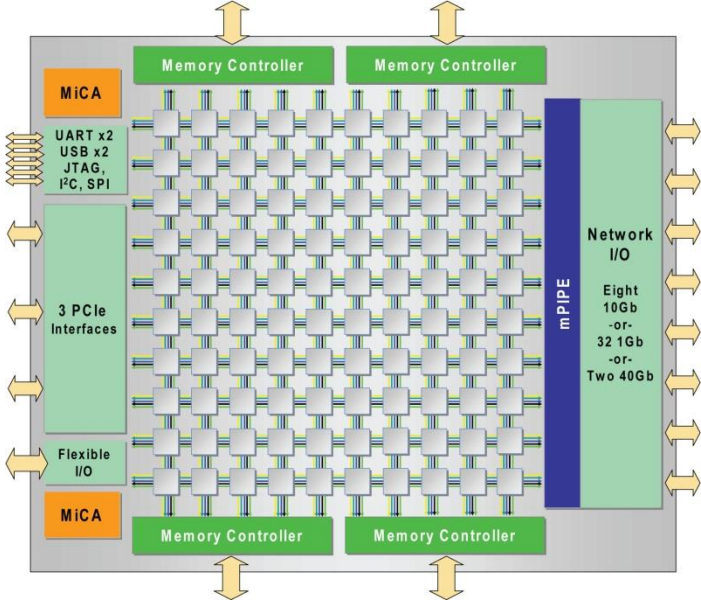

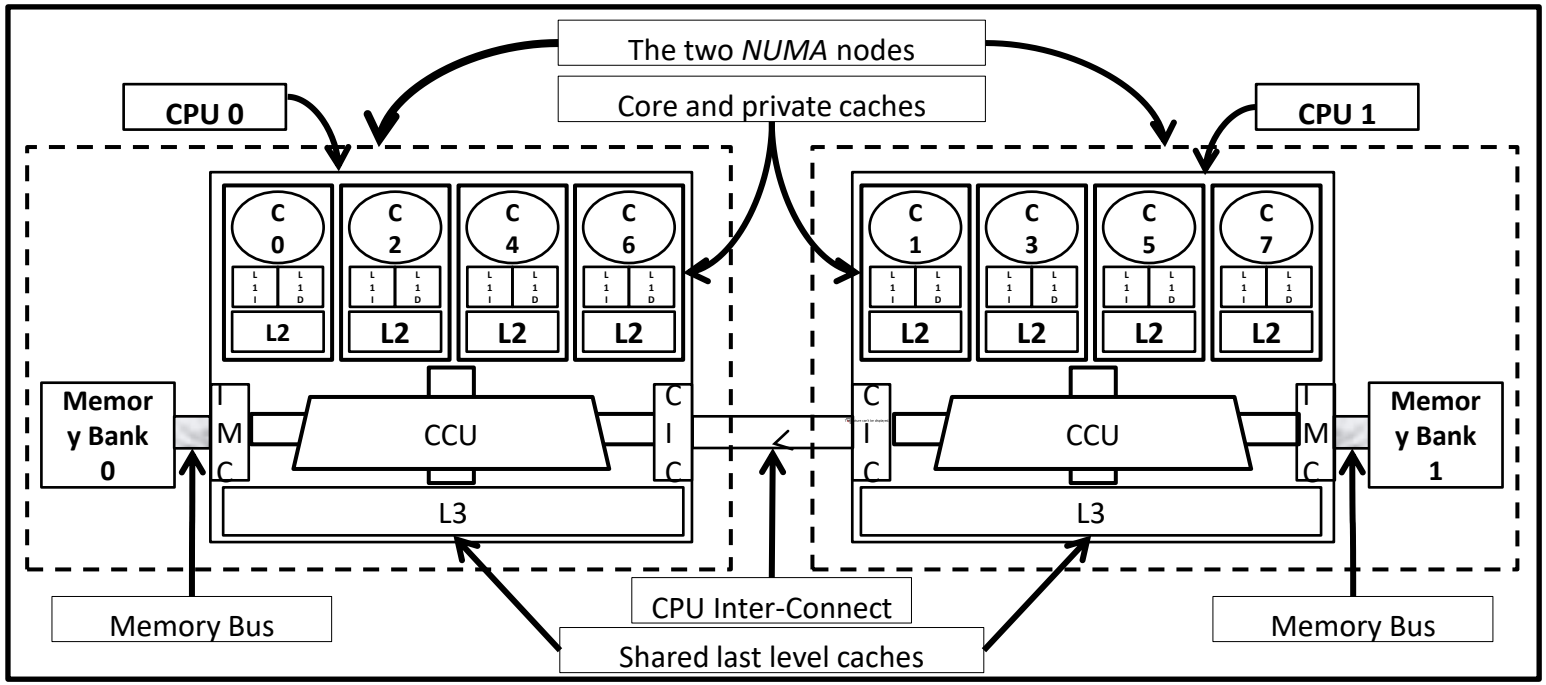

Distributed Systems

- Everything is connected …

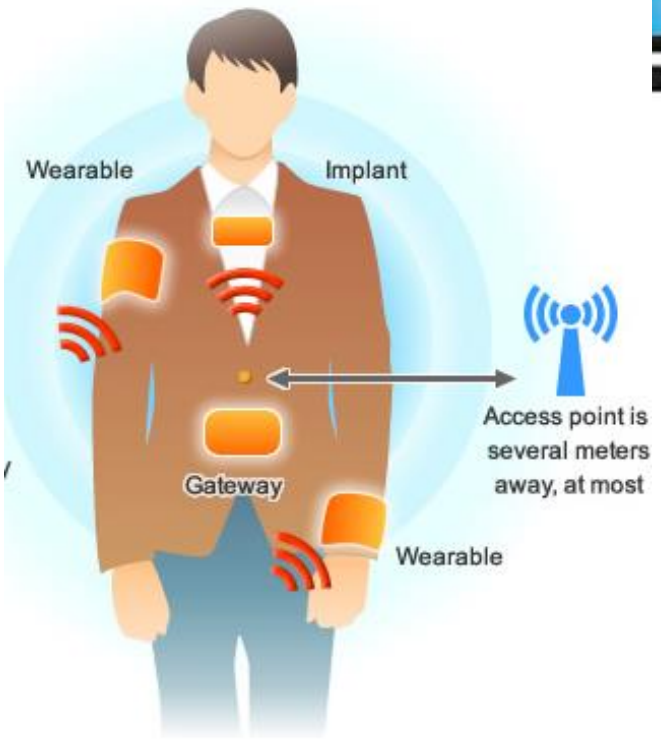

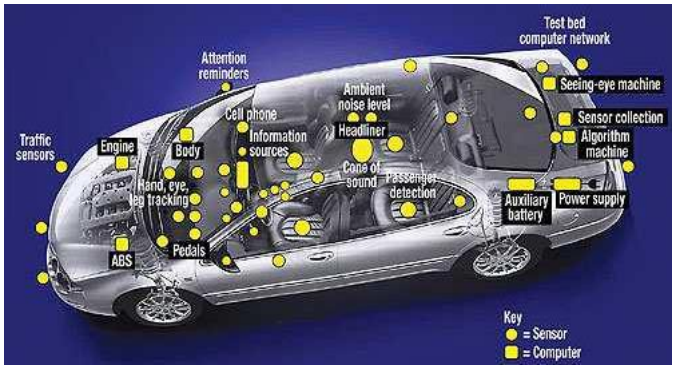

… looking towards the future …

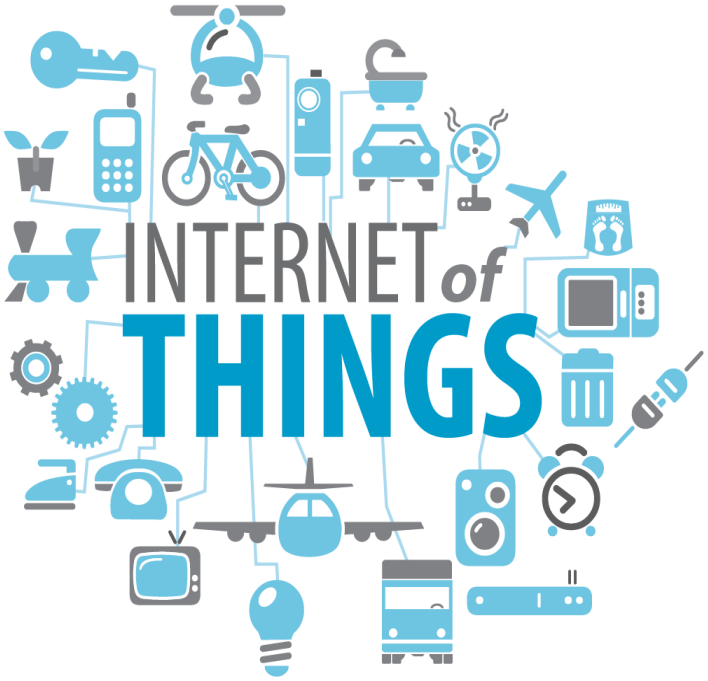

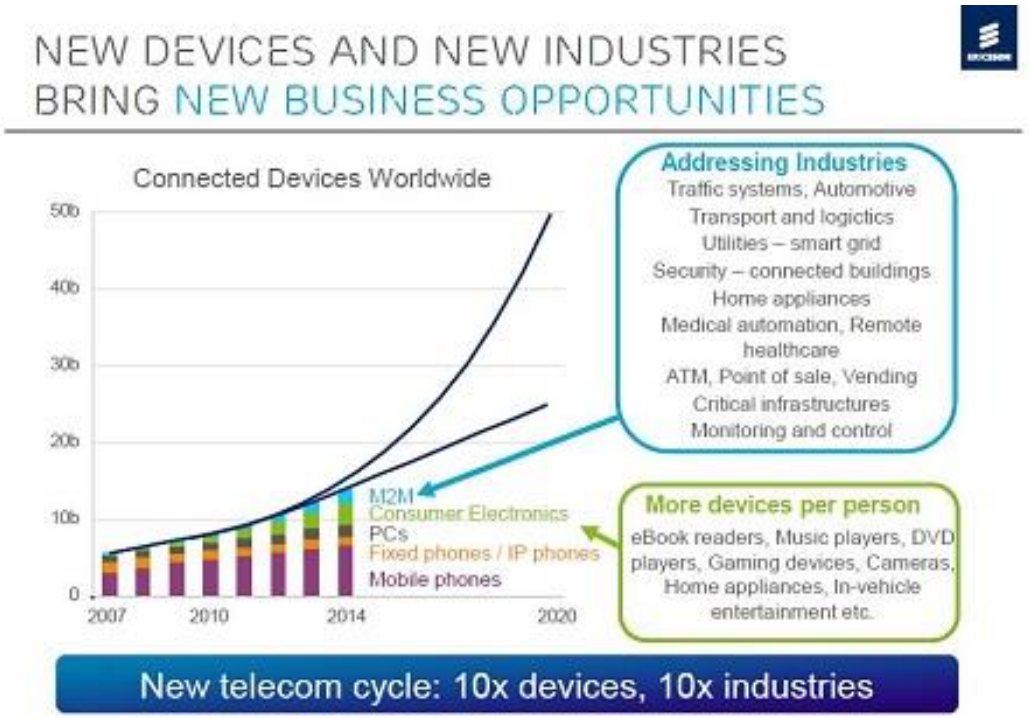

Trends and visions

- Everything that can be connected will be connected
  - 50B devices (perhaps more like 500B ...)
- IoT and smart cities
  - Machine-to-machine
- High-definition 3D streaming/games to heterogeneous clients

Personalized service and personal footprints in a connected world …

Scalability

- Typically want solutions that “scales”
  - Ability of a system, network, or process to handle a growing amount of work effectively
  - Capability to increase its total output under an increased load when resources are added
- Typically want the costs or resource capacity needed to scale sub-linearly with demand OR the performance to improve at least proportionally to the capacity added

Scalability examples

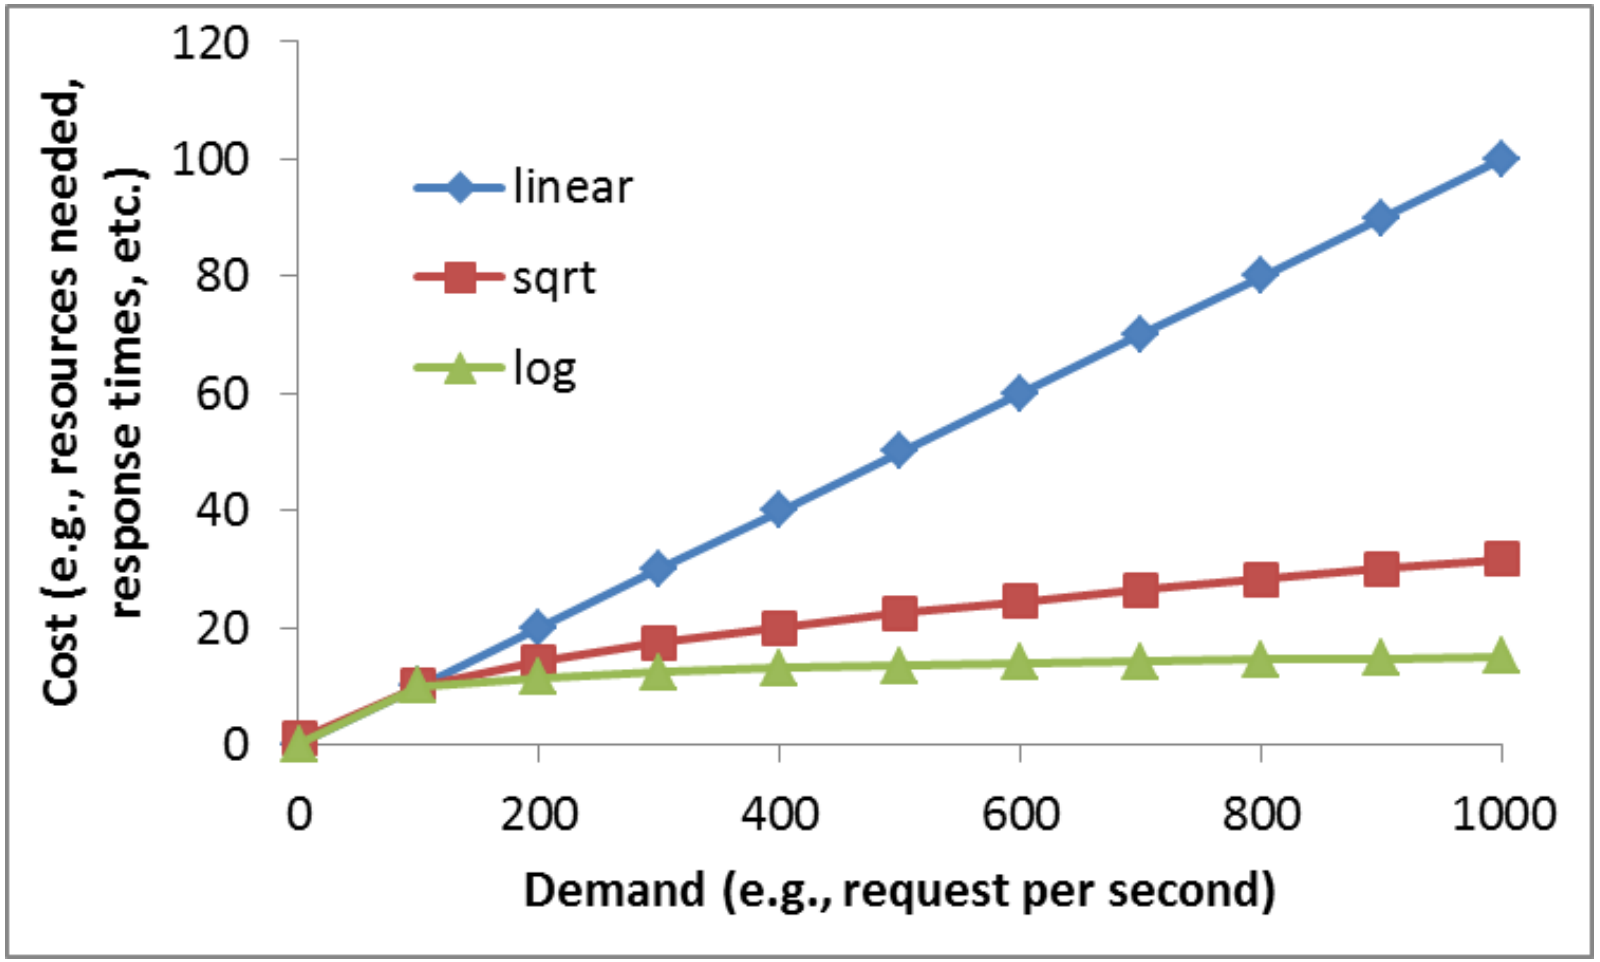

Server-client vs. P2P: example

Time to distribute F to N clients, using client/server approach

$$d_{cs} = \max\{ NF/u_s, F/\min(d_i) \}$$

… and using a P2P approach

$$d_{P2P} = \max\{ F/u_s, F/\min(d_i) , NF/(u_s + \sum u_i) \}$$

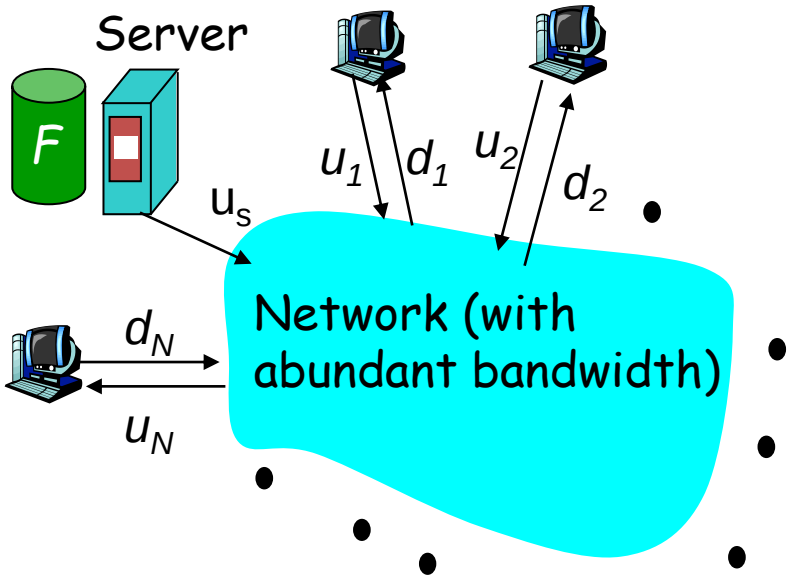

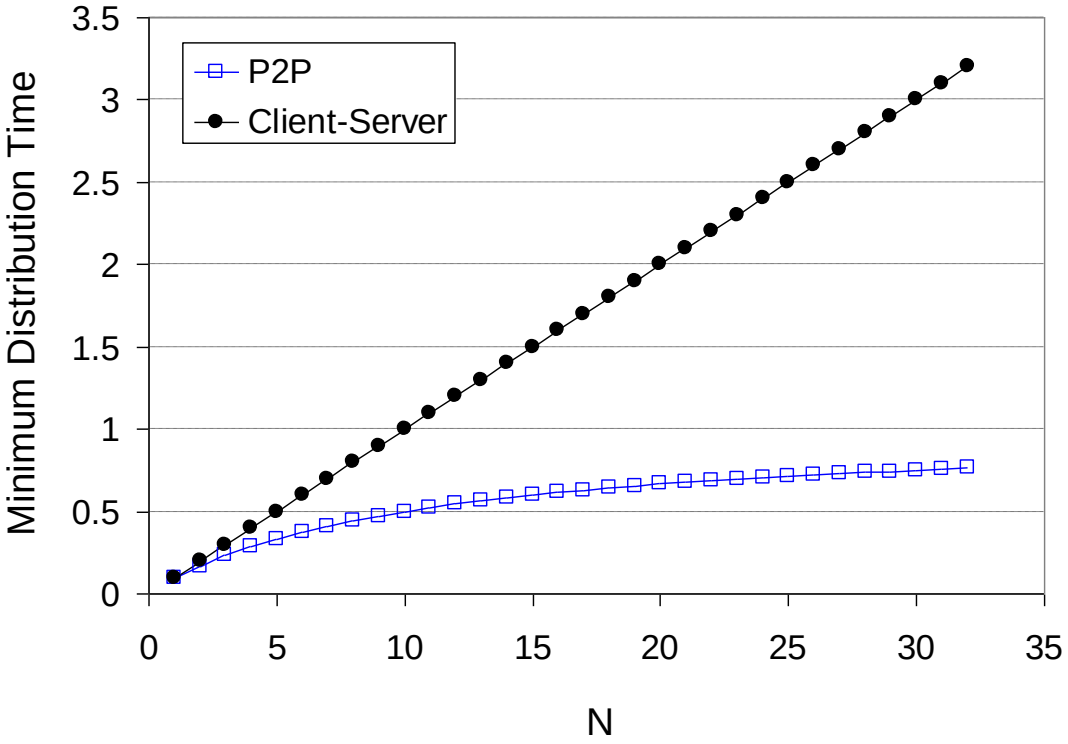

Client upload rate = u,
F/u = 1 hour, u_s = 10u, $d_{min} \geq u_s$

Parallelism

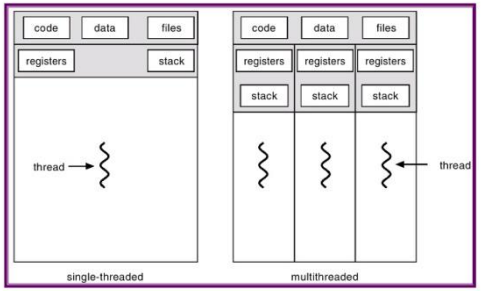

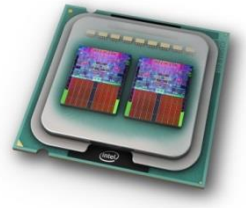

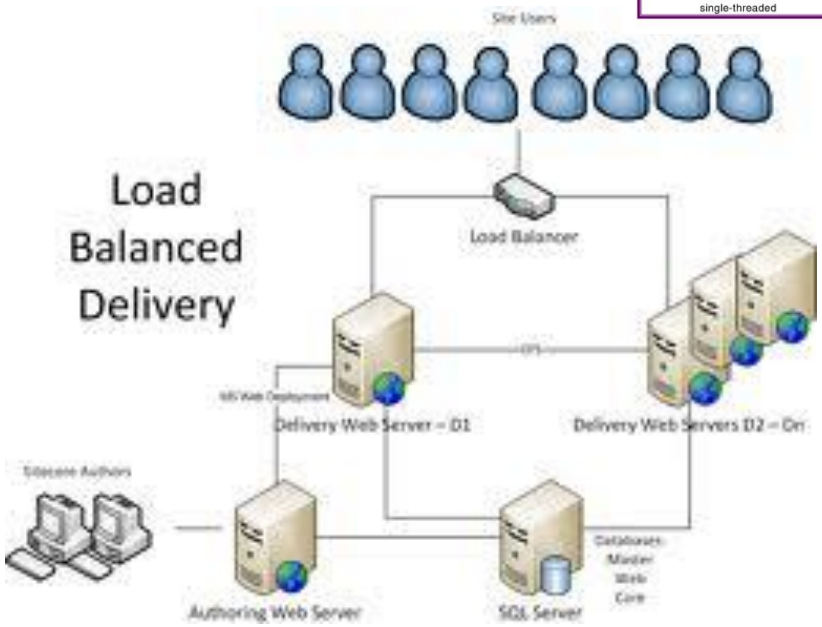

Power/energy

- The systems and their individual parts consume energy and cost money …

The Power Issue

- Power = Static (leakage) power + Dynamic (switching) power
- Dynamic power ~ Voltage^2 * Clock frequency
  where Clock frequency approx. ~ voltage
  → Dynamic power ~ Frequency^3
- Total power ~ #processors

Processor architecture	#cores	Voltage	Frequency	Performance	Power	Power efficiency [Gflops/W]
Classical superscalar	1x	1x	1x	1x	1x	1x
"Faster" superscalar	1x	1.5x	1.5x	1.5x	3.3x	0.45x
Multi-core	2x	0.75x	0.75x	1.5x	0.8x	1.88x

→ Preferable to use multiple slower processors than one superfast processor

… PROVIDED THAT the application can be parallelized efficiently!

Dynamic Power management (cont'd)

Hardware Support (e.g. Intel Xscale Processor)

- RUN: operational
- IDLE: Clocks to the CPU are disabled; recovery is through interrupt.
- SLEEP: Mainly powered off; recovery through wake-up event.
- Other intermediate states: DEEP IDLE, STANDBY, DEEP SLEEP

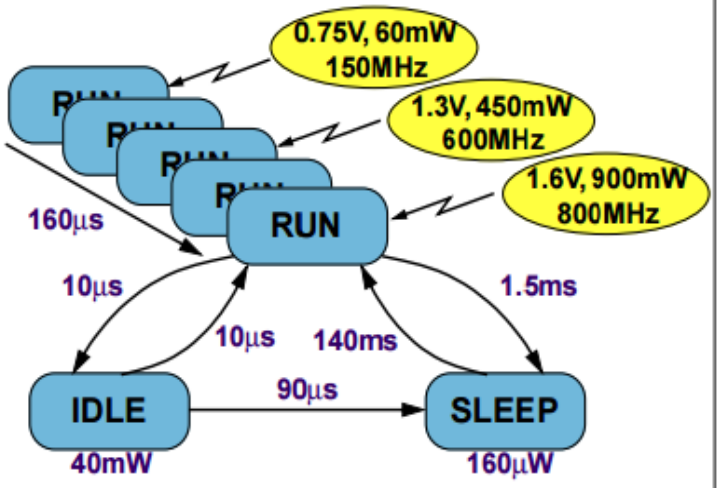

How do we understand these tradeoffs?

- There are three main methods used in the design of performance evaluation studies:
- Analytic approaches
  - the use of mathematics, Markov chains, queueing theory, Petri Nets, abstract models…
- Simulation approaches
  - design and use of computer simulations and simplified models to assess performance
- Experimental approaches
  - measurement and use of a real system

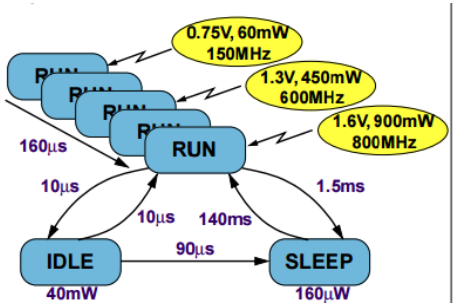

Recommended follow up courses

- **TDTS21 Advanced Networking (6hp)**
- TDDD25 Distributed Systems (6hp)
- TDDI08 Embedded Systems Design (4hp)
- TDDI11 Embedded Software, 6 credits (6hp)
- TDDC78 Programming Parallel Computers (6hp)
- TDDD56 Multicore and GPU Programming (6hp)

- … and lots of interesting thesis projects, of course!

(Other projects are also possible; e.g., on individual basis …)

The exams 2020

- Exam dates:
  - TEN2: June 4 (8-12)
  - TEN1: June 10 (8-12)
- Please remember to register
  - LiU will not allow late registrations for exams
  - No last-minute participation is allowed, pre-registration is required
- Format TEN2
  - Four parts: each worth a 1/4th of the points …
  - Bonus question (only on original exam: June 2020)
  - Bonus points from project (only June 2020 exam): Will not be informed of these until AFTER the exam (but will be decided before marking)
- Format TEN1
  - Networking …
  - Bonus question (only on original exam: June 2020)

… general on exams …

- Read all instructions carefully
- Please explain how you derived your answers. Your final answers should be clearly stated.
- Where a discourse or discussion is called for, be concise and precise.
- If necessary, state any assumptions you made in answering a question. However, remember to read the instructions for each question carefully and answer the questions as precisely as possible.
- Solving the *wrong question may result in* deductions! It is better to solve the *right question incorrectly, than the wrong* question correctly.
- The grading is U, 3, 4, 5 based on the correctness and clarity of answers, displayed understanding of the material, and how close are the answers to indicated length limit of each question.

… yet more on 2020 exam …

- Questions will be available via Lisam (https://lisam.liu.se/) room "Submissions" after the exam starts (and on course website as backup).
- Duration of the exam is 4 hours, and results must be submitted in the course Lisam room before the deadline.
- Answers should be submitted as a single pdf file.
- Please type answers in English. No handwritten answers are allowed.
- Figures can be drawn by hand or by computer. If you draw by hand, please scan (or take a photo) and insert into the file you use to generate your pdf.
- The exam is not anonymous due to the exceptional situation of COVID-19. Please show understanding and tolerance of potential grading delays and technical issues as we switch to the remote examinations.
- You must answer all questions yourself. No group work or outside help is allowed.
- Do not copy-paste text from other sources. Answers will likely be checked through Urkund for plagiarism.
- By taking the exam, you accept a solemn obligation not to cheat.

Good luck on the exam!!