T-110.5121 Mobile Cloud Computing Summary 28.11.2012

Teacher: Yrjö Raivio
Assistant: Eetu Jalonen
Aalto University, School of Science
Department of Computer Science and Engineering
Data Communications Software
Email: yrjo.raivio(at)aalto.fi
Course email: t-110.5121(at)tkk.fi

Social Networks and Operators

The Facebook Mobile Plot Thickens: Orange To Offer 'Party Call', A Social Calling Service Starting Summer 2013

INGRID LUNDEN

Wednesday, November 21st, 2012 — 29 Comments

Facebook Phone it may not be, but today we have one more example of how users will be able to use **Facebook** on their phones — and in place of traditional phone services. Today, multinational French carrier **Orange** announced that it would offer "Party Call," a "social calling" service on the social network. Available on both mobile and desktop versions of the social network, users will be able to call each other, and make "group calls", without needing to know or use each other's mobile numbers to do it.

The deal was announced at **Hello**, and innovation event organized by Orange owner France Telecom. Party Call was described by the carrier today as "A partnership with Facebook, to introduce more

Source: http://techcrunch.com/2012/11/21/the-facebook-mobile-plot-thickens-orange-inks-deal-for-social-calling-service-starting-summer-2013/

Targets

- "You have a clear view of the advantages of ***distributed computing technologies***, such as ***cloud computing***, in the ***mobile*** space. You understand the ***core properties*** as well as the ***strengths*** and ***limitations*** of mobile cloud computing. You have a vision how clouds will ***change*** the mobile ecosystem, reviewed both from ***technology*** and ***business*** angles."
- (Mobile) Disruption
- Neutral, scientific and critical view, over the hype
- Technology, Business, Theoretical and Practical approach

Requirements

- **5 ECTS: 24 + 0 (2 + 0), not applicable to post graduate studies**
  - Lectures 24 h, Lecture preparation 24 h, Assignments 48 h, Exam preparation 36 h, Exam 3 h
  - Lectures are not obligatory but highly recommended
- **Exam**
  - Tue 18.12.2012, 9-12, Tue 08.01.2013, 9-12 or Tue 28.05.2013, 9-12 (**remember to register in Oodi**)
  - Structure:
    - 1 obligatory question: 6 definitions
    - 3 questions, 2 must be answered
    - 1 obligatory essay
    - 6 points from each: 0-11=0, 12-13=1, 14-15=2, 16-18=3, 19-21=4, 22-24=5
- **Assignments in pairs (possible alone with good reasons)**
- **Course feedback!!!**

Grading

- **Exam 50% + Both assignments together 50%**
- **Both assignments have the same weight, e.g. 25%**
- **Exam and assignments evaluated 0-5**
- **To pass the whole course, each component must be passed at least with grade 1**
- **Example:**
  - Exam: 5
  - Assignment 1: 3
  - Assignment 2: 5
  - Total: 50% x 5 + 25% x 3 + 25% x 5 = 4.5 = grade 5 (rounded to closest integer)

Exam reading material

1. Armbrust, Michael, Fox, Armando, Griffith, Rean, Joseph, Anthony D., Above the Clouds: A Berkeley View of Cloud Computing, Feb. 10, 2009. Available at: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf
2. Lee Badger, Tim Grance, Robert Patt-Corner and Jeff Voas: Draft Cloud Computing Synopsis and Recommendations, Recommendations of the National Institute of Standards and Technology, May 2011, available at: http://csrc.nist.gov/publications/drafts/800-146/Draft-NIST-SP800-146.pdf
3. Rolf Harms and Michael Yamartino: The Economics of the Cloud, Nov. 2010, available at: http://www.microsoft.com/presspass/presskits/cloud/docs/The-Economics-of-the-Cloud.pdf
4. Greenberg, Albert, Hamilton, James, Maltz, David A., Patel, Parveen (2009) The cost of a cloud: research problems in data center networks, SIGCOMM Comput. Commun. Rev., Vol. 39, No. 1, pp. 68-73.
5. Mohammad Hajjat, Xin Sun, Yu-Wei Eric Sung, David Maltz, Sanjay Rao, Kunwadee Sripanidkulchai, and Mohit Tawarmalani, Cloudward Bound: Planning for Beneficial Migration of Enterprise Applications to the Cloud, ACM SIGCOMM'10, (Sections 1-2, 2 pages)
6. Ross Anderson: Security Engineering, 1st Edition (2001), Chapter 1, available from http://www.cl.cam.ac.uk/~rja14/Papers/SE-01.pdf (12 pages); Cloud Security Alliance: Top Threats to Cloud Computing V1.0, March 2010, (14 pages), available from https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf; Security Guidance for Critical Areas of Focus in Cloud Computing, v. 2.1, Dec 2009, Section II/Domain 5, (6 pages), available from https://cloudsecurityalliance.org/csaguide.pdf
7. *Lutz Schaubert and Keith Jeffery: Research in Future Cloud Computing, May 2012, Section II/Domain 5, (6 pages), available at: http://cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf (NOT REQUIRED)*
8. Antero Juntunen, Eetu Jalonen and Sakari Luukkainen: HTML5 in Mobile Devices: to be published in HICSS-46 in Jan 2013. (10 pages)
9. Matti Kemppainen: Mobile Computation Offloading: A Context Driven Approach, IWORK paper from 2011. (10 pages)

Lecture schedule

12.09 Introduction, Basics & Assignment 1, Yrjö Raivio and Eetu Jalonen
19.09 Mobile Networks, Jukka K. Nurminen
26.09 Programming on open APIs, Olli Rinne/Apps4Finland
03.10 Mobile Cloud Business, Antero Juntunen
10.10 Scalable Cloud Computing, Keijo Heljanko
17.10 Mobile Cloud, Yrjö Raivio and HTML5, Eetu Jalonen
24.10 No lecture, exam week
31.10 Green Cloud Computing, Tommi Mäkelä & Assignment 2, Yrjö Raivio
07.11 Cloud Computing in Data Centres, Jarno Laitinen, Risto Laurikainen and Peter Jenkins/CSC
16.11 Cloud Security, Abu Shohel Ahmed/Ericsson
21.11 Industry keynote, Markku Lepistö/NSN
28.11 Resource Provisioning & Summary, Yrjö Raivio and Eetu Jalonen
05.12 Spare

Massive data volumes

	WAN bandwidth/mo.	CPU hours (all cores)	disk storage
Item in 2003	1 Mbps WAN link	2 GHz CPU, 2 GB DRAM	200 GB disk, 50 Mb/s transfer rate
Cost in 2003	$100/mo.	$2000	$200
$1 buys in 2003	1 GB	8 CPU hours	1 GB
Item in 2008	100 Mbps WAN link	2 GHz, 2 sockets, 4 cores/socket, 4 GB DRAM	1 TB disk, 115 MB/s sustained transfer
Cost in 2008	$3600/mo.	$1000	$100
$1 buys in 2008	2.7 GB	128 CPU hours	10 GB
cost/performance improvement	2.7x	16x	10x
Cost to rent $1 worth on AWS in 2008	**$0.27–$0.40** ($0.10–$0.15/GB × 3 GB)	**$2.56** (128× 2 VM's@$0.10 each)	**$1.20–$1.50** ($0.12–$0.15/GB-month × 10 GB)

Bottleneck: Bandwidth
Example: Facebook 1 PB totally, 2-3 TB added each day
1 TB drive, 1 Gbit/s I/O = 2 h 13 min

$T = 10^{12}$
$P = 10^{15}$

Source: Ambrust et al, Above the Clouds: A Berkeley View of Cloud Computing, 2009

Some numbers

- **Global ICT business size**
  - 2008: $ 383 B, 4% cloud
  - 2012: $494 B, 9% cloud
- **Largest growth in storage**
- **SMEs have best opportunities to adapt**
- **Also small countries with good infrastructure**
- **Startup costs for SMEs dropped dramatically**

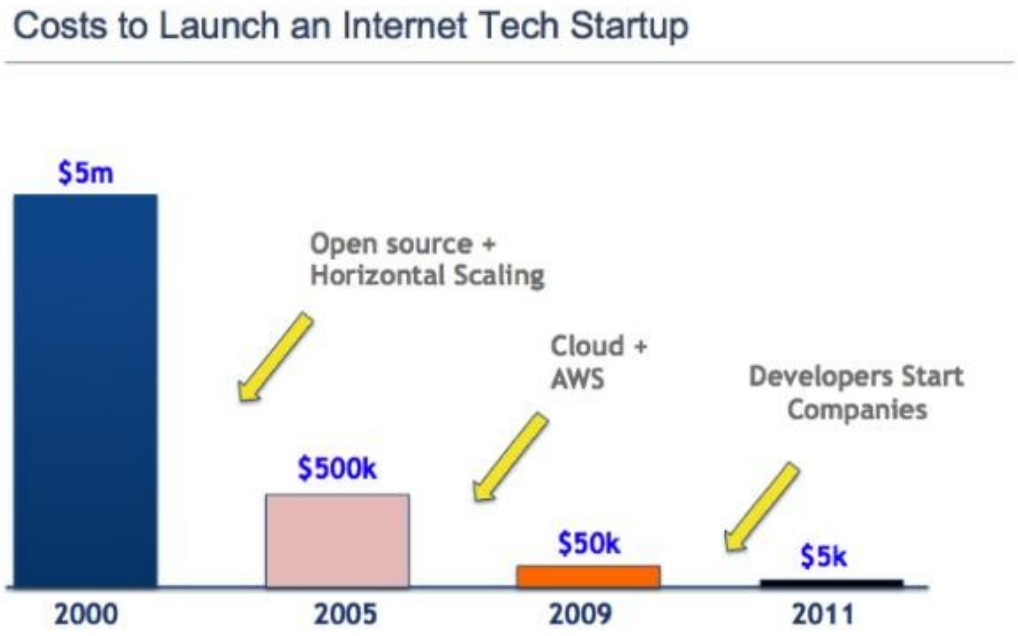

Source: F. Etro, The Economic Impact of Cloud Computing on Business Creation, Employment and Output in Europe, 2009

Source: M. Suster, It's Morning in Venture Capital, blog, May 23, 2012

Mobile internet access doubled in a year

Proportion of global Web pageviews from mobile devices, by region, in May 2011 and May 2012							
	Global	**Africa**	**Asia**	**Europe**	**North America**	**Oceana**	**South America**
Mobile pageviews in May 2012	10.1%	12.9%	18.0%	5.1%	8.6%	7.5%	2.8%
Mobile pageviews in May 2011	5.8%	6.7%	8.3%	2.7%	7.8%	4.8%%	2.8%
Top mobile operating system for pageviews, by region, in May 2012							
1	Android 23.8%	Symbian 45.6%	Symbian 33.3%	iOS 41.2%	iOS 47.2	iOS 70.5%	Android 24.9%
2	iOS 23.0%	Series 40 23.1%	Series 40 18.8%	Android 33.4%	Android 38.3%	Android 23.7%	Series 40 15.1%
3	Symbian 20.3%	Android 6.8%	Android 15.5%	BlackBerry 11.7%	BlackBerry 6.1%	Symbian 2.0%	iOS 13.6%
4	Series 40 11.8%	Samsung 4.6%	Samsung 12.4%	Symbian 5.63%	Windows 1.0%	Windows 0.8%	Symbian 13.6%
Source: StatCounter (May 2012)						**via: mobiThinking**	

Source: http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats/b#mobilepageviews

Mobile capabilities are improving but battery capacity is still a bottleneck

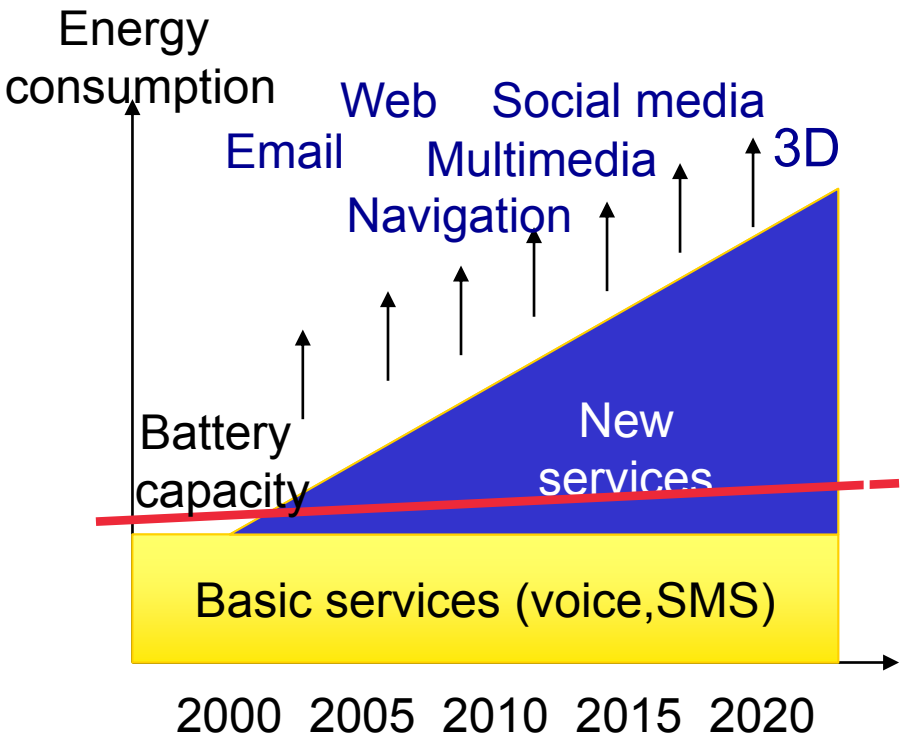

- Less new services **NOK**
- More frequent charging **NOK**
- Physically larger battery **NOK**
- More energy efficient chips **OK**
- Intelligent methods **OK**
- Radical battery inventions **?**
- HTML5 **?**

Source: professor Jukka K. Nurminen

(Mobile) Cloud Computing

"Mobile Cloud computing is a model for enabling convenient, on-demand mobile network access to a shared pool of configurable mobile computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

Adapted from: P. Mell and T. Grance, The NIST Definition of Cloud Computing, 2009

Definition 2

1. The illusion of ***infinite computing resources*** available *on* ***demand***, thereby ***eliminating*** the need for Cloud Computing users to plan ***far ahead*** for provisioning.
2. The ***elimination*** of an ***up-front*** commitment by Cloud users, thereby allowing companies to start ***small*** and increase hardware resources only when there is an ***increase*** in their needs.
3. The ability to ***pay for use*** of computing resources on a short-term basis as needed (e.g., processors by the hour and storage by the day) and ***release*** them as needed, thereby rewarding conservation by letting machines and storage go when they are no longer useful.

Source: Ambrust et al, Above the Clouds: A Berkeley View of Cloud Computing, 2009

Fixed vs. Mobile

- **Cloud performance**
  - Processing time at the data center
  - Processing time on the device
  - Network latency
  - Data transport time
- **Mobile challenges**
  - Variable reliability, less throughput, longer latency
  - Energy
  - Limited resources, example Dell Desktop vs. iPhone 4
    - 3 times less processing power
    - 8 times less memory
    - 5 times less storage capacity
    - 10 times less network bandwidth
- **But mobiles are always with you and provide context data (location, sensor data, camera)**

Source: http://www2.alcatel-lucent.com/blogs/techzine/2010/mobile-cloud-computing-challenges/

What is Mobile Cloud Computing (MCC)

- ***Mobile Cloud Computing:*** **mobile apps are processed and data stored in a cloud rather than on the native handset, referred as *Mobile Computing***
- ***Mobile Offloading*****: moving data, code block or virtual machine from mobile to cloud or vv.**
- **Using mobile context to enhance cloud based services**
- **Cloud within mobile (Hyrax, Cuckoo, ThinkAir, MAUI, CloneCloud..)**

Key issues

- **Deployment models (Public, Private, Hybrid, Community Cloud)**
- **Service models (SaaS, PaaS, IaaS)**
- **Key Benefits**
  - Economies of scale
  - Elasticity
  - Resource planning
  - Pay-as-you-go
  - Always available
- **Technology**
  - Virtualization, Storage, SLA, Provisioning, Energy, Security, Mobility

Source: Armbrust, Michael, Fox, Armando, Griffith, Rean, Joseph, Anthony D., ”Above the Clouds: A Berkeley View of Cloud Computing”, Feb. 10, 2009.

Deployment models

- **Public cloud**
  - Resources made available to the general public via the Internet
  - Scalable
  - Pay for what you use
- **Private cloud**
  - Host own resources
  - Provide to internal customers only
  - Provision with cloud interfaces
- **Hybrid cloud**
  - Some resources provided internally and some outsourced
- **Community cloud**
  - Address special needs of a community

Main benefits - Service models

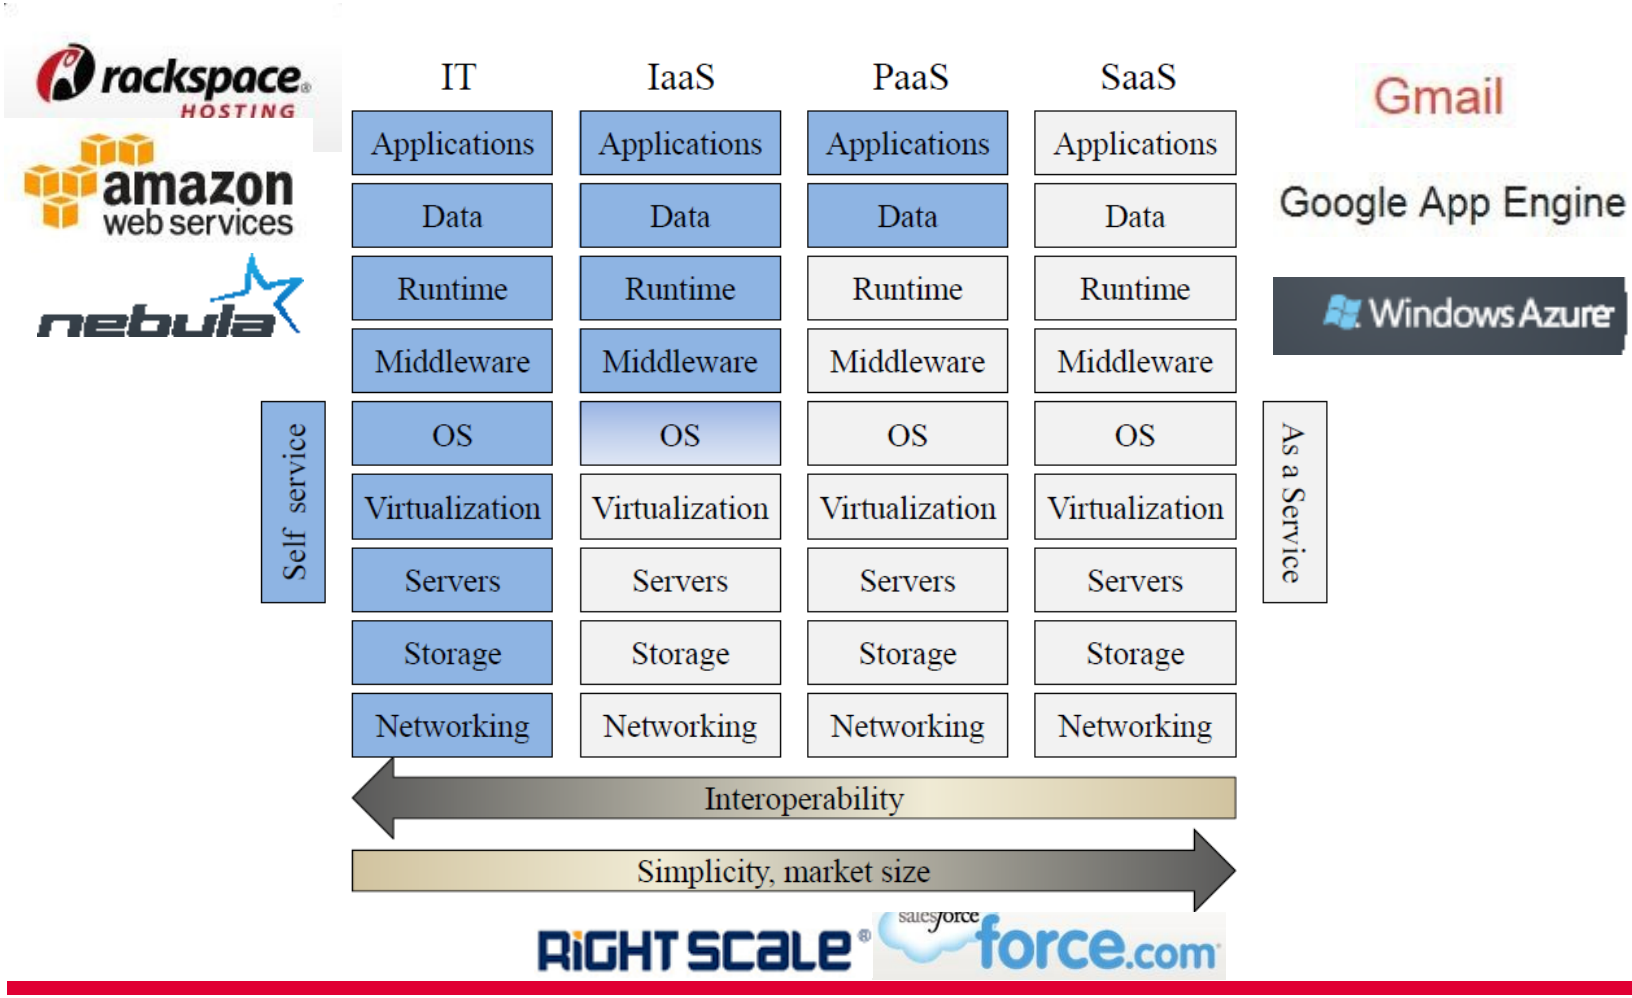

Economies of scale

- Cheaper MIPS (5-7 times)
- Better utilization of computing resources (5-10% to 60-80%)
- Multi-tenancy: one instance can serve several customers
- Less admin people per server (from 1:100 up to 1:10 000)
- Worth 1$ IT requires 8$ admin costs

FIG. 4: ECONOMIES OF SCALE (ILLUSTRATIVE)

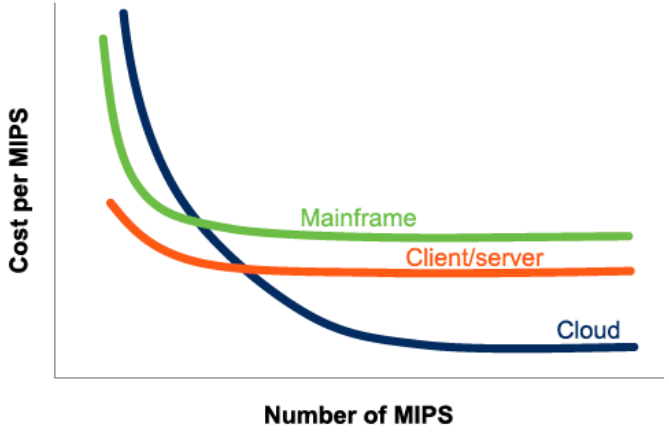

Source: Microsoft.

FIG. 15: ECONOMIES OF SCALE IN THE CLOUD

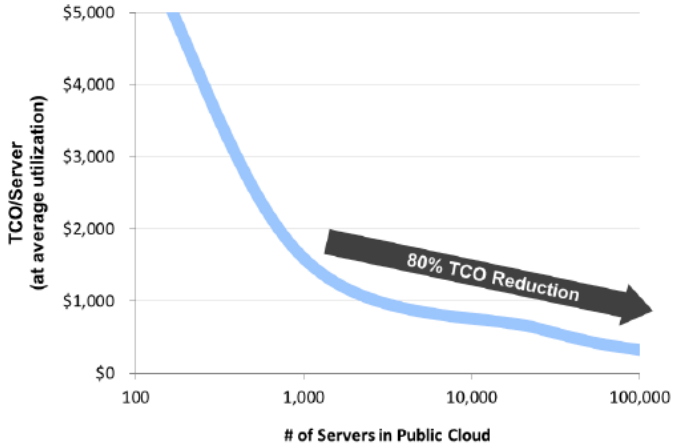

Source: Microsoft.

Source: Rolf Harms and Michael Yamartino: The Economics of the Cloud, Nov. 2010.

Statistical multiplexing

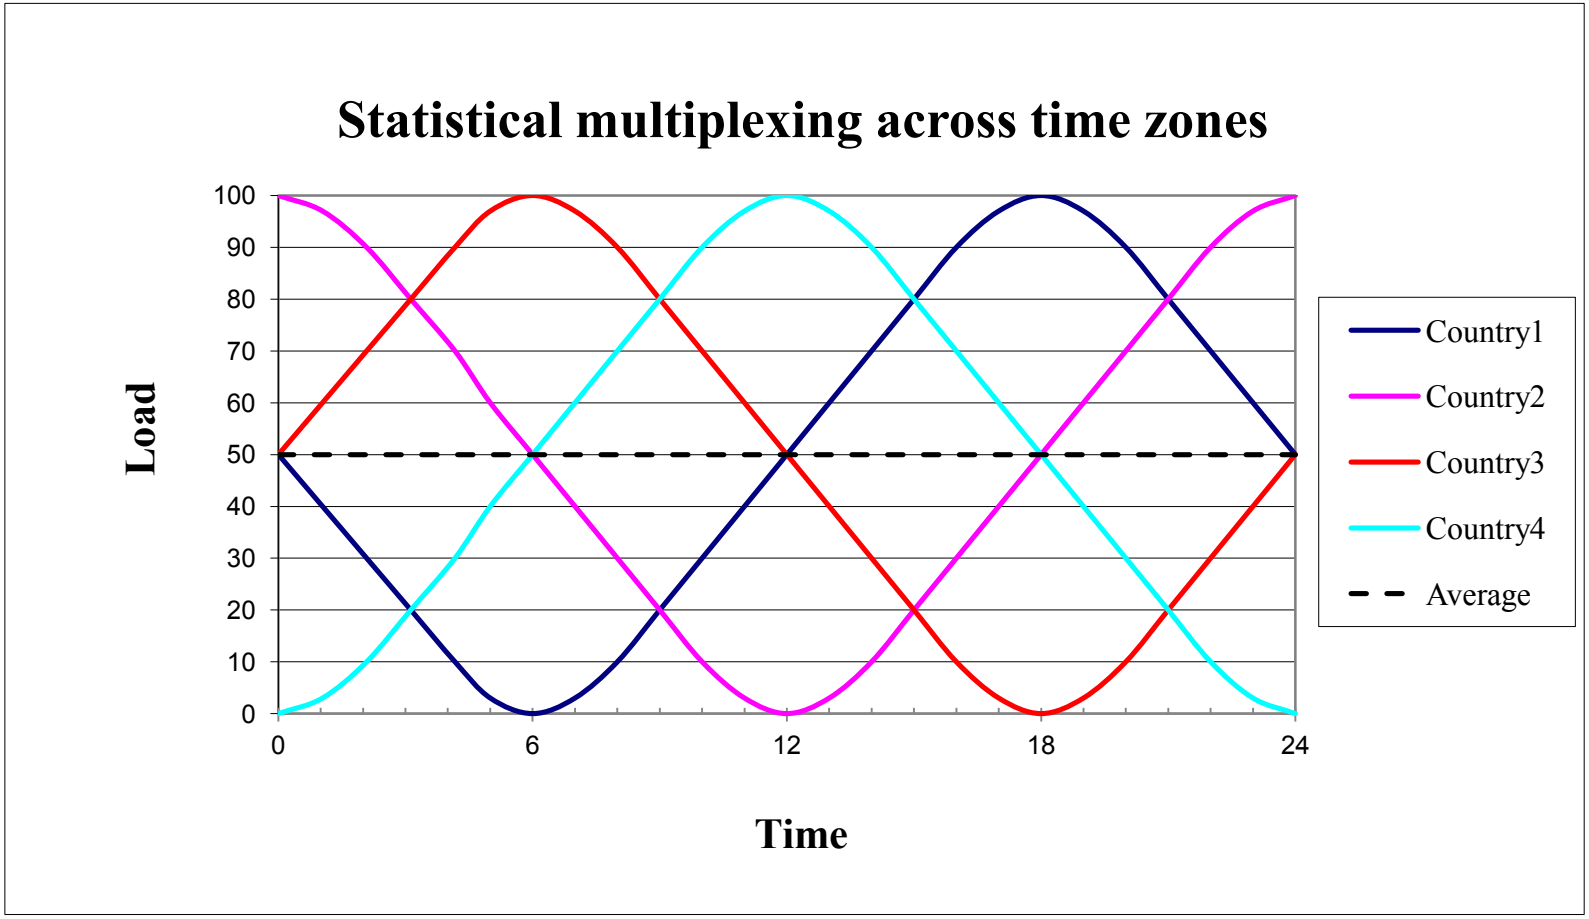

Elasticity – pay-as-you-go

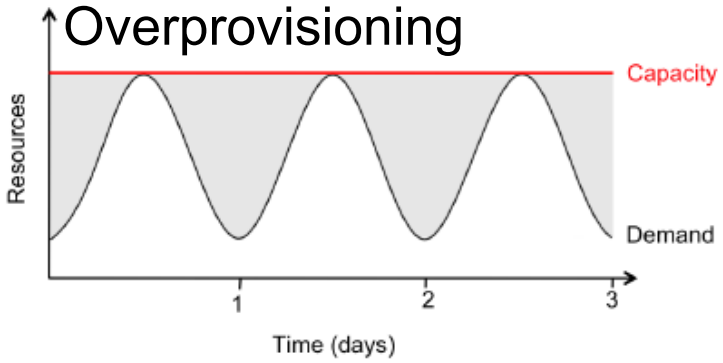

(a) Provisioning for peak load

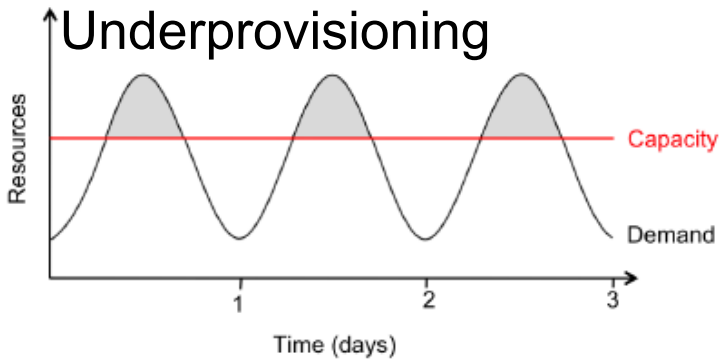

(b) Underprovisioning 1

- **Avoid high upfront investment, avoid risk**
- **Adapt to changing business**
- **Buy or lease**
- **Amortizise value to investment period**

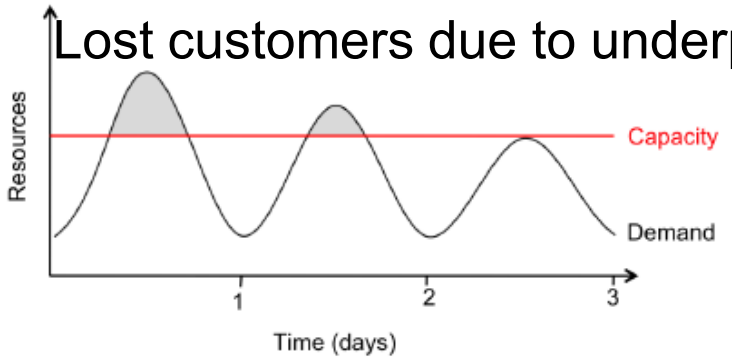

(c) Underprovisioning 2

Source: Ambrust et al, Above the Clouds: A Berkeley View of Cloud Computing, Feb 2009

Amortization

- **Hybrid cloud**
  - Public vs. Private
  - www.cloudonomics.com
- **Cost structure by Greenberg (2009):**
  - ~45% Servers CPU, memory, storage
  - ~25% Infrastructure
  - ~15% Power draw
  - ~15% Network Links
- **Staff? 1$ IT : 8$ Admin!**
- **Equated Monthly Installment**

$$A_m = E \frac{\frac{r}{12}(1+\frac{r}{12})^m}{(1+\frac{r}{12})^m - 1}$$

- **Net Present Value**

$$NPV = \sum_0^N \frac{P_T - C_T}{(1+r)^T} + \frac{S}{(1+r)^N} - E$$

Case SMSC

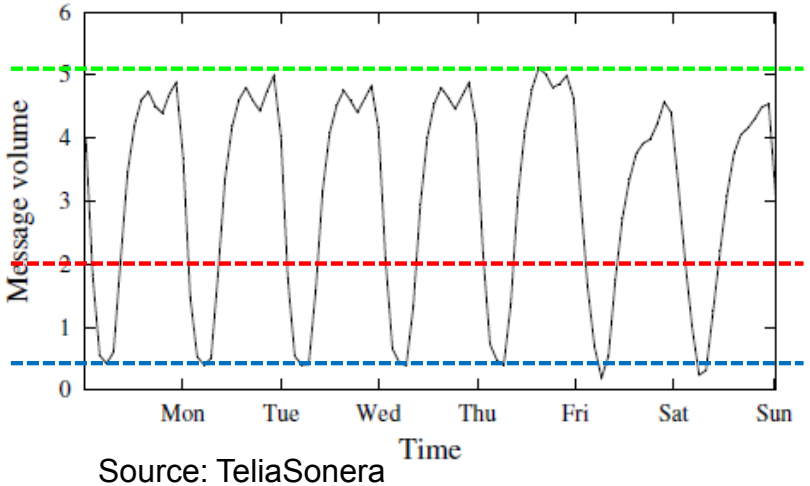

Source: TeliaSonera

E = basic investment
N, m = length of the investment
R = rate of interest
Pt = annual revenue
Ct = annual cost
S = residual value

Always available

- **Anyone, anytime, anywhere**
- **High availability?**
- **Typical public cloud SLA promise**
  - 99.95% = max 4 h 23 min down time per year
- **Telecom**
  - 99.999% = 5 min
- **Availability Zone, fully (?) independent computing systems**
- **Using two Availability Zones**

$$P_P = 1 - P_F{}^2 = 1 - (1 - A)^2 = 99.9999\%$$

Resource planning

- **Resources can be optimized to meet service needs**
- **Service integration time can be shortened, example Short Message Service Center setup from 2 weeks to 4 minutes**

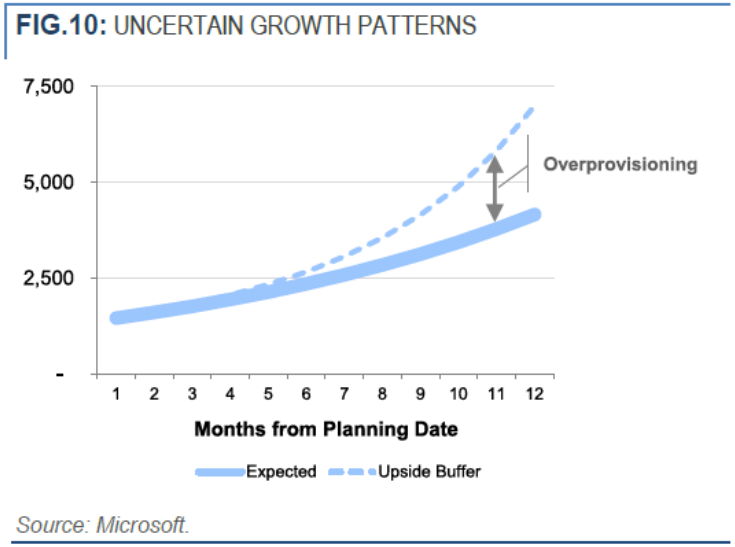

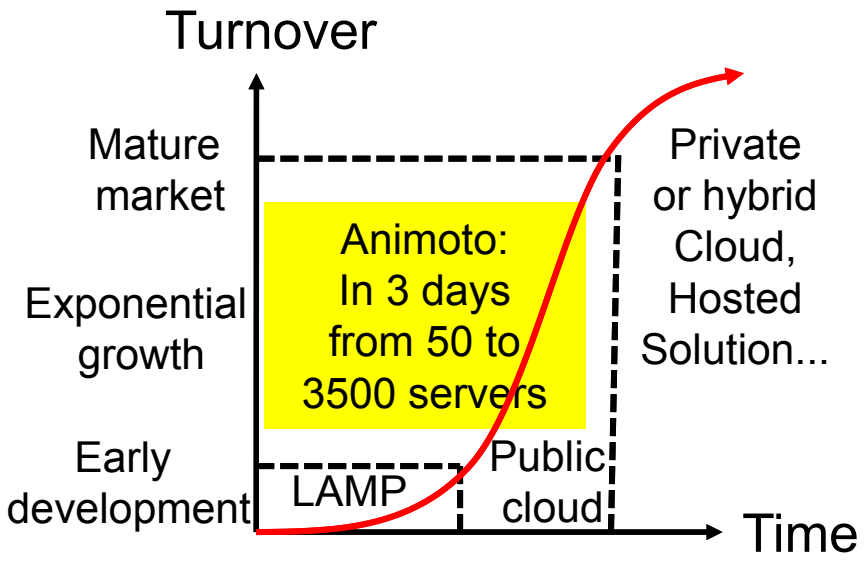

Source: Rolf Harms and Michael Yamartino: The Economics of the Cloud, Nov. 2010.

Key technology issues

- Mobile communication (19.9, 21.11)
- Virtualization, storage and programming models (10.10)
- Mobile offloading and HTML5 (17.10)
- Energy (31.10)
- Data Centres (7.11)
- Security (14.11)
- Resource provisioning (28.11)

Mobile offloading

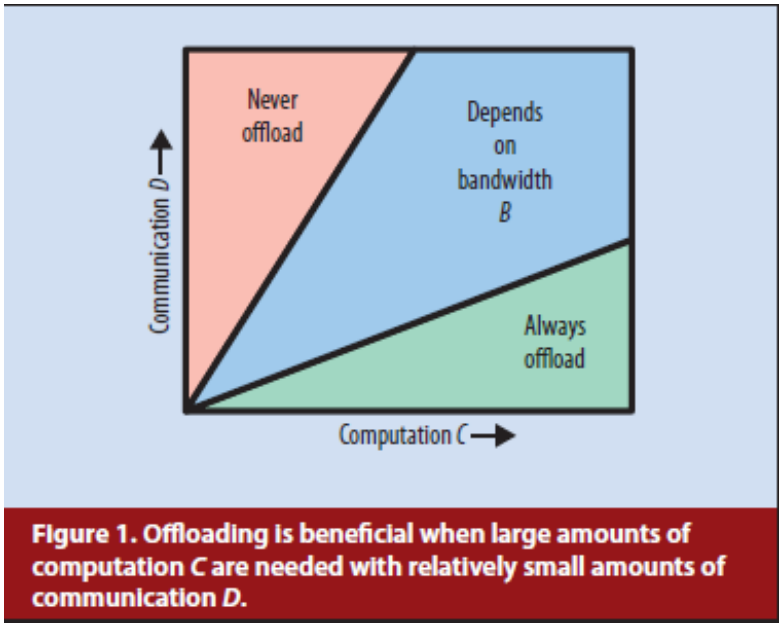

Figure 1. Offloading is beneficial when large amounts of computation *C* are needed with relatively small amounts of communication *D*.

Source: Kumar & Lu, ”Cloud Computing for Mobile Users: Can Offloading Computation Save Energy ”, 2010

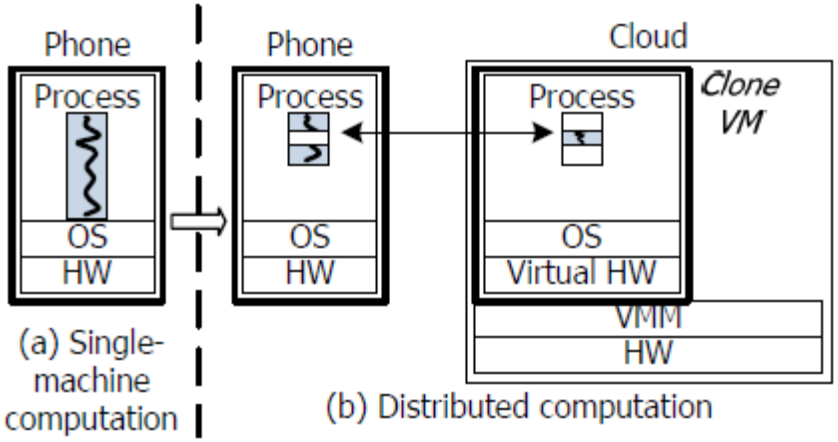

Source: Chun and Maniatis, ”Augmented Smartphone Applications Through Clone Cloud Execution”, 2009

Virtualization

1) **Cost Savings**
2) **Better Resource Utilization**
3) **Better Memory Management**
4) **Increased Availability**
5) **Better Resource Provisioning**
6) **Energy Saving**

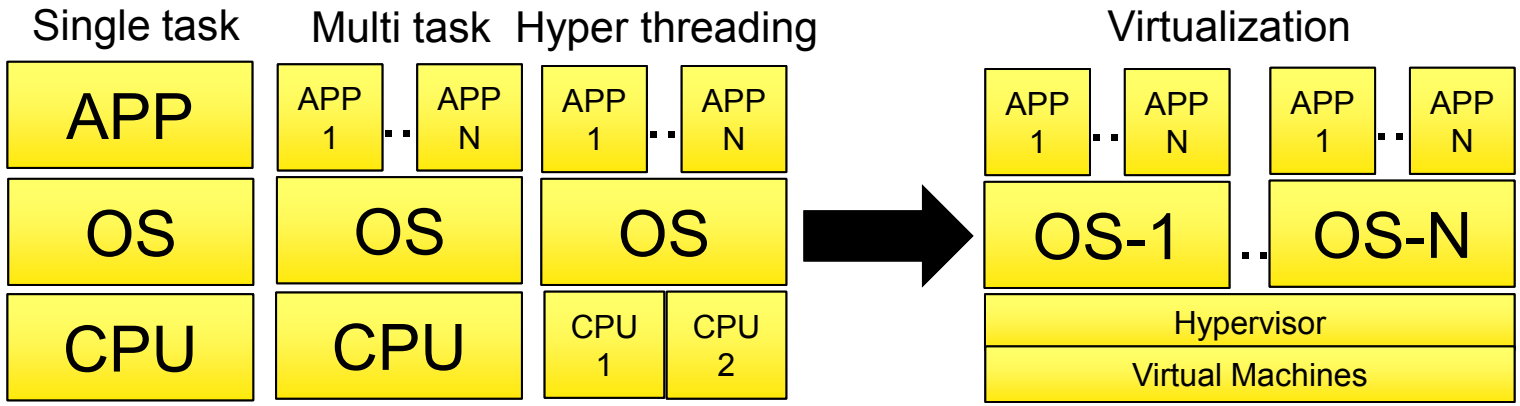

Source: Z. Ou, Virtualization Technology, T-110.7100, Autumn 2010.

Open Data and APIs

- **Freely available data from government, science, business and other organizations, citizens…**
  - free **access** and **absence** of technological **restrictions**
  - License allows **redistribute** and **reuse**
  - License may require **attribution** and **integrity**
- **Open data promotes**
  - Democracy and transparency
  - New and improved processes (and easier live)
  - New business opportunities
  - Maker – Do-It-Yourself– culture

More info from Open Knowledge Foundation: http://okfn.org/

Cloud computing technologies

- A collection of technologies aimed to provide elastic "pay as you go" computing
- Virtualization of computing resources: Amazon EC2, Eucalyptus, OpenNebula, Open Stack Compute, . . .
- Scalable file storage: Amazon S3, GFS, HDFS, . . .
- Scalable batch processing: Google MapReduce, Apache Hadoop, PACT, Microsoft Dryad, Google Pregel, Spark,..
- Scalable datastore: Amazon Dynamo, Apache Cassandra, Google Bigtable, HBase,. . .
- Distributed Coordination: Google Chubby, Apache Zookeeper, . . .
- Scalable Web applications hosting: Google App Engine, Microsoft Azure, Heroku, . . .

Brewer's CAP theorem

- **In a PODC 2000 conference invited talk Eric Brewer made a conjecture that it is impossible to create a distributed asynchronous system that is at the same time satisfies all three CAP properties:**
  - Consistency
  - Availability
  - Partition tolerance
- **This conjecture was proved to be a Theorem in the paper: "Seth Gilbert and Nancy A. Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. SIGACT News, 33(2):51-59, 2002."**
- **However, Google published Spanner concept based on atom clock 10/2012…**

Power management

- **Power Usage Efficiency (PUE) defined as:** $PUE = \frac{TotalFacilityPower}{ITEquipmentPower}$
- **Typical PUE 2-3, state of the art 1.1**
- **15% of all costs**
  - 59% IT equipment
  - 33% cooling
  - 8% power loss
- **Power off idle machines**
- **Raise temperature**
- **Cold location, cheap energy (for example Finland)**
- **More fine-grain accounting**
- **Better algorithms**
- **E2E model required including public clouds**

Source: A. Greenberg, J. Hamilton, D.A. Maltz abd P. Patel, The Cost of a Cloud: Research Problems in Data Center Networks, ACM SIGCOMM Computer Comm. Review, Jan 2009.

Security

- **Top threats**
  - Abuse and Nefarious Use of Cloud Computing
  - Insecure Application Programming Interfaces
  - Malicious Insiders
  - Shared Technology Vulnerabilities
  - Data Loss/Leakage
  - Account, Service & Traffic Hijacking
  - Unknown Risk Profile

Source: Cloud Security Alliance, “Top Threats to Cloud Computing V1.0”, March 2010

Pros and cons

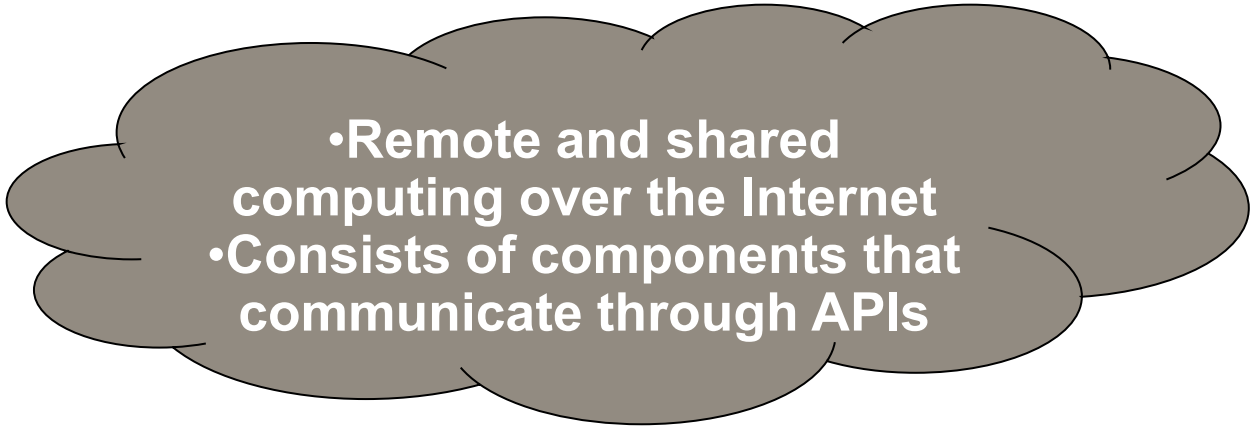

!
- Simple architecture
- Efficient usage of CPU (>50%)
- Scalability
- Load balancing
- Low capex
- High availability

?
- Security & Privacy
- High usage of certain CPUs
- Interoperability
- Vendor lock-in
- High opex
- SLA critical

Conclusions

- **Cloud computing is a new business model**
- **Great tool to startups**
- **Biggest challenge: (lack of) trust**
- **Next target: utility computing (similarly to water, electricity, gas and telephony)**
- **Future research topics**
  - Energy efficiency in data centers
  - Cloud interoperability
  - Security
  - HTML5
  - Dynamic resource provisioning algorithms
  - Game theory

Questions?

Contacts:
Teacher: Yrjo Raivio, A122
Assistant: Eetu Jalonen A121

Course staff: t-110.5121(at)tkk.fi