

# Advanced Services Engineering- Introduction

Hong-Linh Truong

Distributed Systems Group, TU Wien

[truong@dsg.tuwien.ac.at](mailto:truong@dsg.tuwien.ac.at)  
<http://dsg.tuwien.ac.at/staff/truong>  
Twitter: @linhsolar

- Why do we need a course on advanced services engineering?
- What is the course about?
- Course administrative information

# Current trends: emerging systems

- **Internet of Things (IoT)/cyber-physical systems**
  - Integration and virtualization of sensors/actuators and edge networks
- **IoT and cloud integration → IoT cloud systems**
  - Dealing with sensors/actuators and gateways integration with cloud data centers
- **Fog, Edge-centric and mobile-edge computing**
  - nano/micro data centers + cloud-based data centers
- **Social-cyber-physical clouds**
  - Core elements: software, people and things
  - Systems: human computation platforms+ IoT platforms + cloud systems

# Current trends: data, software, and services

- „Big“ and „small“ data
  - Data from Things (Internet of Things),
  - Human-sensing data, data marketplaces
- **Software**
  - High performance, scalable data analytics at data centers
  - Hybrid data analytics
  - Individuals, crowds, and collectives augmenting machine intelligence (cognitive computing)
- **Services**
  - Dynamic, flexible data, computation, and analytics provisioning and integration models
  - Human services for complex computation and analytics

# ASE – complex requirements (1)

- Big and near real-time data must be handled in a timely manner to extract insightful information
- Cross-boundary, Internet-scale computation, data and network services integration must be done
- Complex applications/systems executed atop multiple, diverse distributed computing environments
  - Data centers/cloud infrastructures, IoT systems, human computation environments, etc.
- Multiple concerns w.r.t trustworthiness, quality, regulation and cost/benefits must be assured.
- Flexible and dynamic management, e.g., virtualization, and software-defined and elastic capabilities

## ASE – complex requirements (2)

For complex functions offered atop distributed cloud and edge computing environments

- We want to have **a coherent, uniform view** of diverse types of resources and platforms
- We want to **coordinate** capabilities of these resources and platforms

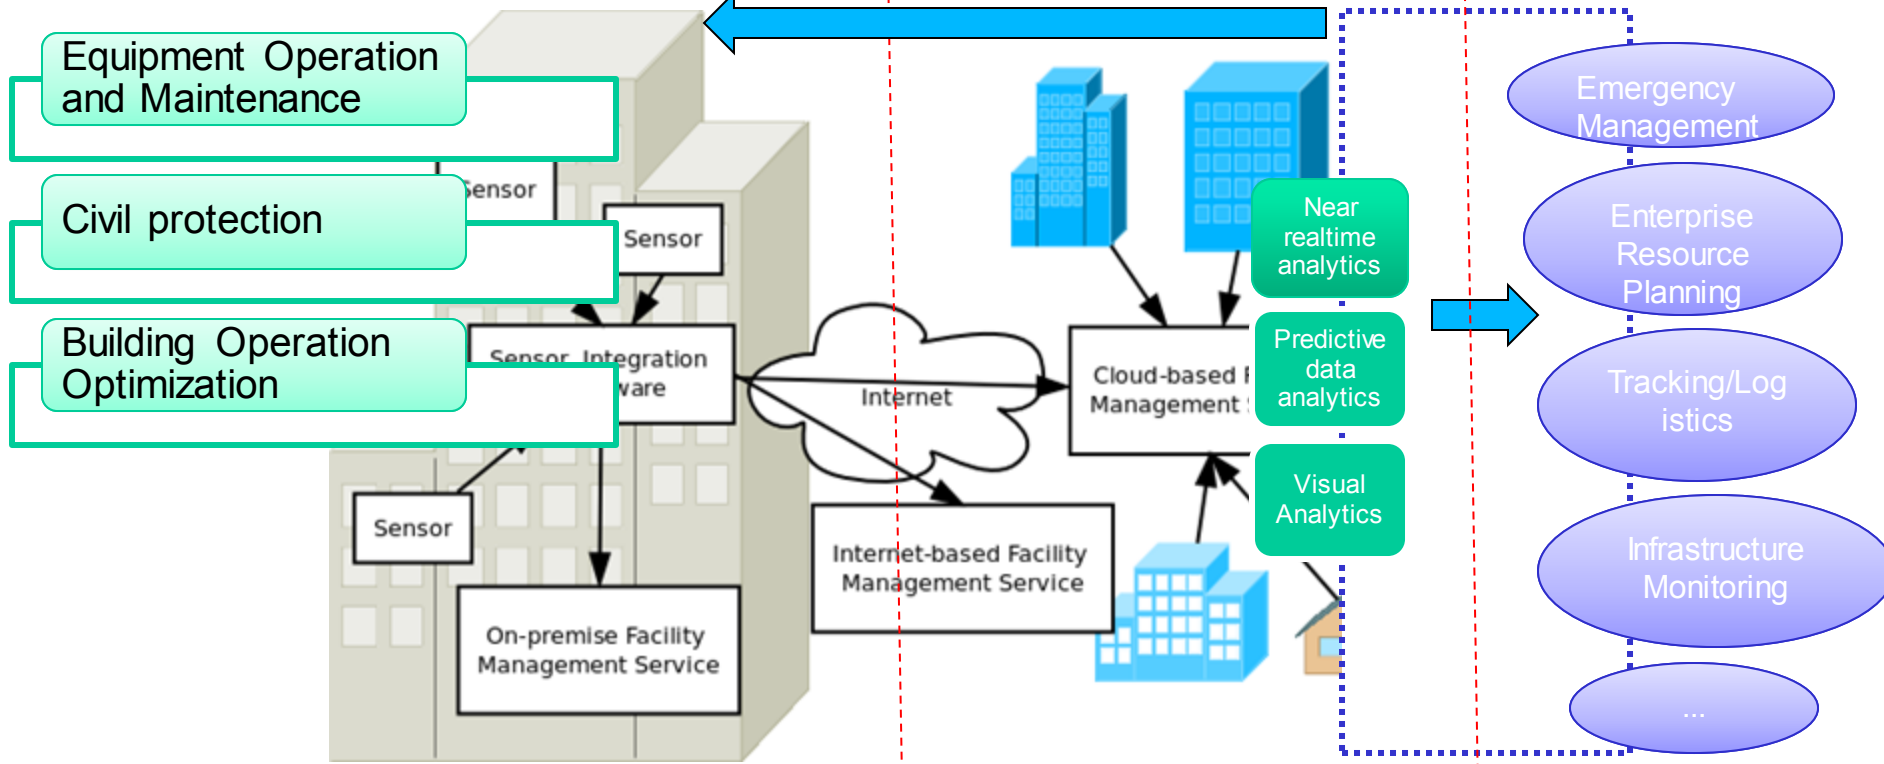
→ Engineering service-based systems for these requirements is very challenging

# ASE -- application examples (1)

Infrastructure/Internet of Things

Internet/public cloud boundary

Organization-specific boundary

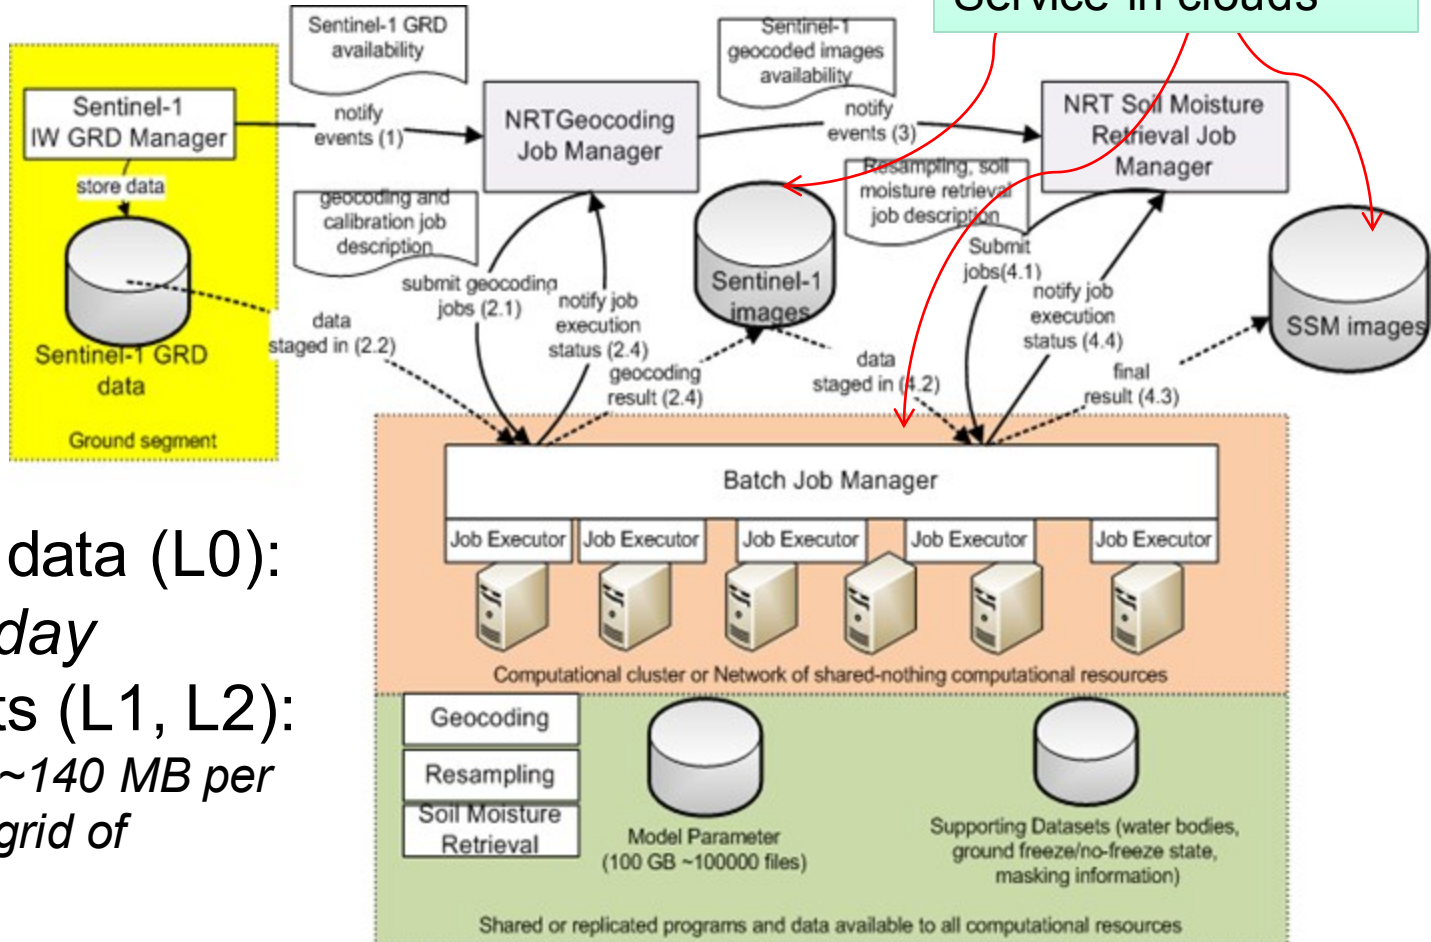


Cities, e.g. including:  
 10000+ buildings  
 1000000+ sensors

# ASE – application examples 2012 (2)

Data-as-a-Service and Platform-as-a-Service in clouds

## Soil moisture analysis for Sentinel-1



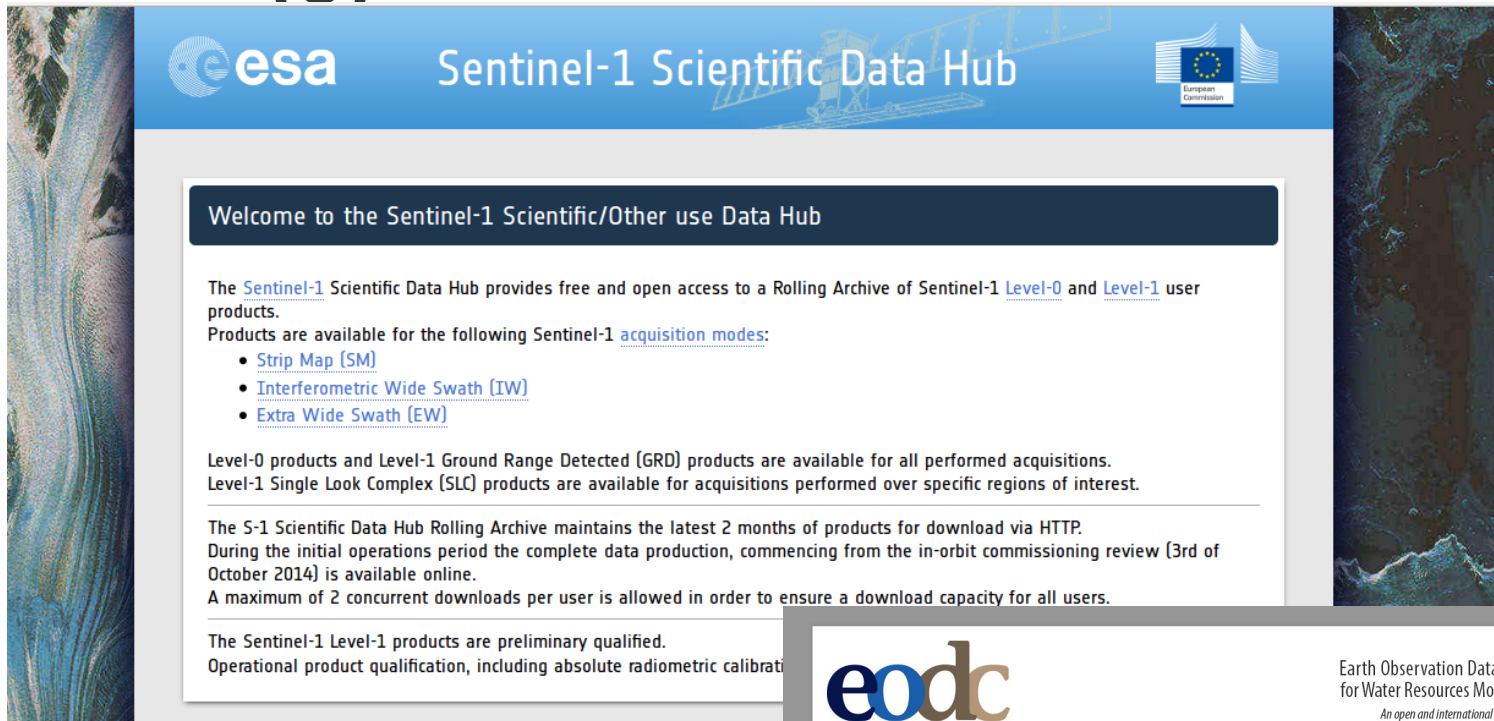
A lot of input data (L0):  
~2.7 TB per day

A lot of results (L1, L2):  
e.g., L1 has ~140 MB per day for a grid of 1kmx1km



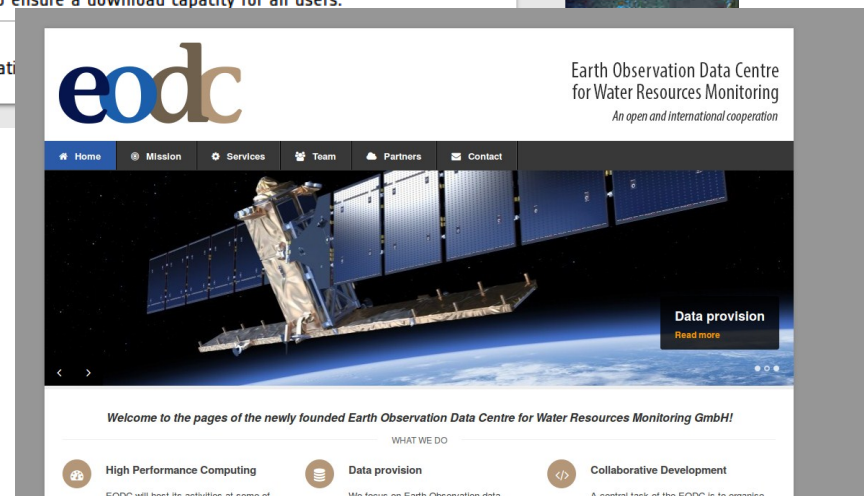
# ASE – application examples - 2015

## (3)



The screenshot shows the Sentinel-1 Scientific Data Hub website. The header includes the ESA logo, the text "Sentinel-1 Scientific Data Hub", and the European Commission logo. The main content area has a dark blue banner that says "Welcome to the Sentinel-1 Scientific/Other use Data Hub". Below this, it states: "The Sentinel-1 Scientific Data Hub provides free and open access to a Rolling Archive of Sentinel-1 Level-0 and Level-1 user products. Products are available for the following Sentinel-1 acquisition modes:" followed by a bulleted list: "Strip Map (SM)", "Interferometric Wide Swath (IW)", and "Extra Wide Swath (EW)". Further text explains that Level-0 products and Level-1 Ground Range Detected (GRD) products are available for all performed acquisitions, while Level-1 Single Look Complex (SLC) products are available for acquisitions performed over specific regions of interest. It also mentions that the S-1 Scientific Data Hub Rolling Archive maintains the latest 2 months of products for download via HTTP, and that a maximum of 2 concurrent downloads per user is allowed.

See: <https://www.eodc.eu/>



The screenshot shows the EODC website. The header features the "eodc" logo and the text "Earth Observation Data Centre for Water Resources Monitoring" with the tagline "An open and international cooperation". A navigation menu includes "Home", "Mission", "Services", "Team", "Partners", and "Contact". The main content area has a large image of a satellite in orbit with a "Data provision" button and a "Read more" link. Below the image, it says "Welcome to the pages of the newly founded Earth Observation Data Centre for Water Resources Monitoring GmbH!". A section titled "WHAT WE DO" lists three key areas: "High Performance Computing" (EODC will host its activities at some of...), "Data provision" (We focus on Earth Observation data), and "Collaborative Development" (A central task of the EODC is to organise...).

# Mobile-edge computing

Source: Mobile-Edge Computing: Introductory Technical White Paper, ETSI. September 2014  
[https://portal.etsi.org/portals/0/tbpages/mec/docs/mobile-edge\\_computing\\_-\\_introductory\\_technical\\_white\\_paper\\_v1%2018-09-14.pdf](https://portal.etsi.org/portals/0/tbpages/mec/docs/mobile-edge_computing_-_introductory_technical_white_paper_v1%2018-09-14.pdf)

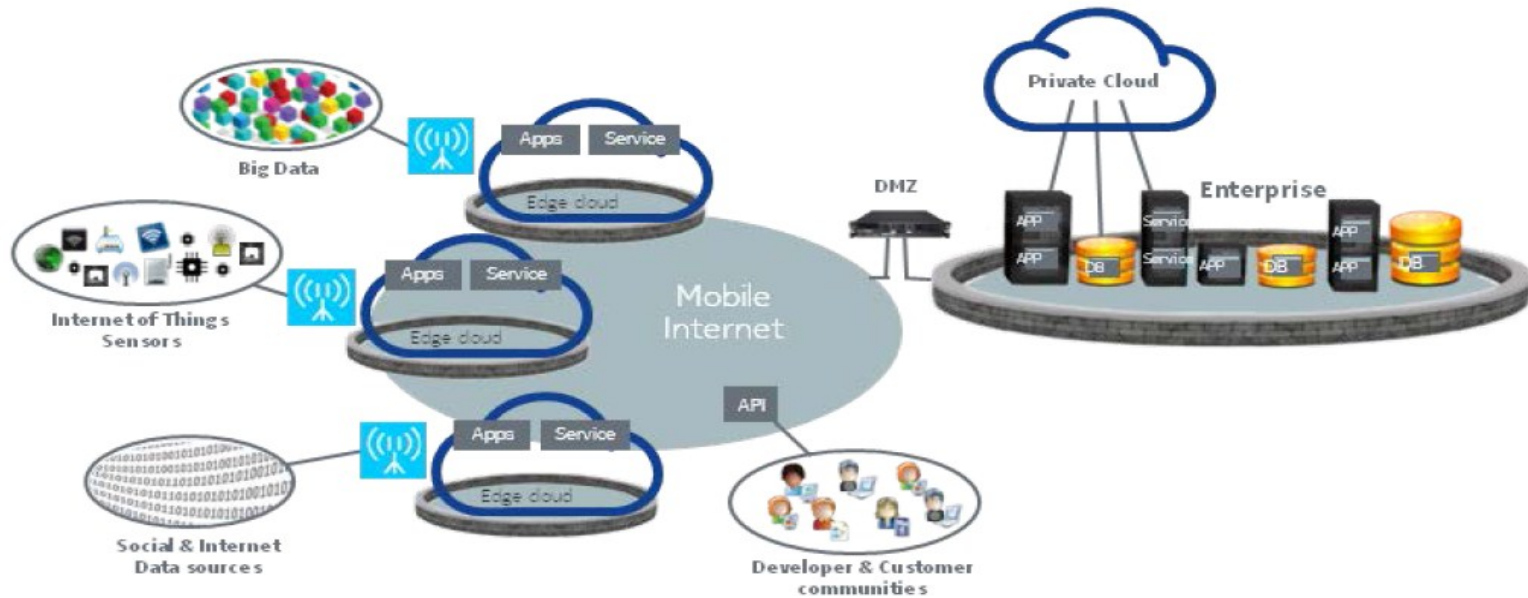


Figure 1: IT and Telecommunications networking convergence

Domains: Retail & M2M

Places: Station, Shopping Centers, and Airports

# ASE – application examples (4)

Source: <http://www.undata-api.org/>

```
<dataset organisation="IMF" category="International Financial statistics" name="BOP: capital account credit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: capital account debit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: current transfers credit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: current transfers debit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: exports of goods, f.o.b., US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: financial account, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: imports of goods, f.o.b., US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: income credit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: income debit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: net errors and omissions, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: services credit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: services debit, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: trade balance, goods and services, US$"/>
<dataset organisation="IMF" category="International Financial statistics" name="BOP: trade balance, goods, US$"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Adolescent fertility rate"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Adult literacy rate (percent)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Gross national income per capita (PPP international $)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Total enrollment ratio female (percent)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Total enrollment ratio male (percent)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Total population (thousands)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Urban growth rate (percent)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Urban areas (percent)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Population below the poverty line (percent living on less than US$1 a day)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Life expectancy at birth (years)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Infant mortality rate (per 1,000 live births)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Maternal mortality ratio (per 100,000 live births)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Crude death rate (per 1,000)"/>
<dataset organisation="WHO" category="Demographic and socioeconomic statistics" name="Crude birth rate (per 1,000)"/>
```

INTERNATIONAL OPEN GOVERNMENT DATA CONFERENCE

Sponsored by: U.S. General Services Administration

Hosted by: U.S. Department of Commerce

November 11 - 17, 2010 Washington, DC

Most Popular Datasets

- U.S. Overseas Loans and Grants (Greenbook)
- Worldwide M1+ Equivalents, Fall 7 Days
- Fuel and Drug Administration-Recalls
- FDIC Failed Bank List
- Clear Air Markets Clear Air Interstate Rule...

SEARCH OUR CATALOGS

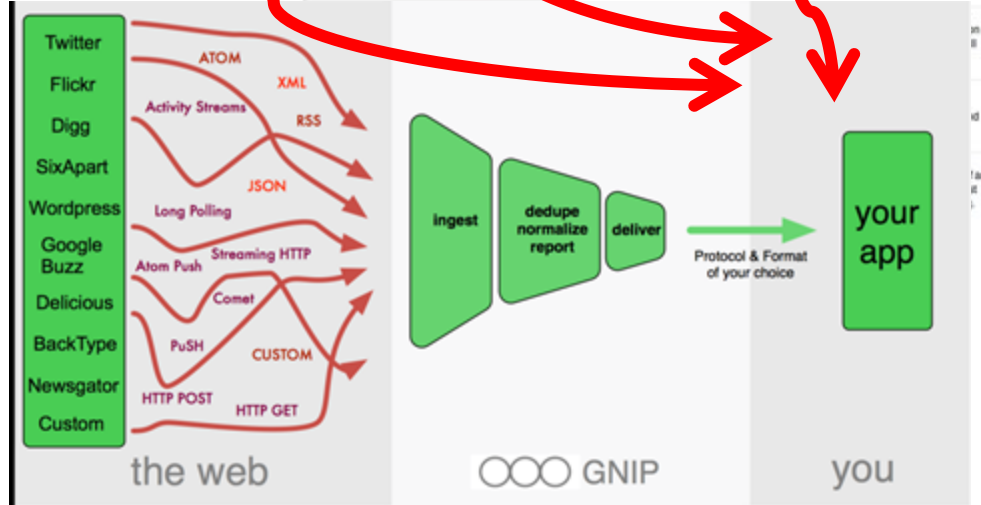
SEARCH our catalogs... SEARCH

APPS COMMUNITY SEMANTIC WEB

Read more buttons for each section.

Source: <http://www.strikeiron.com/Catalog/StrikeIronServices.aspx>

<b>ZIP and Postal Code Information</b> <a href="#">Try It Now</a>	Turns out you can learn a lot from a US Zip or Canadian Postal Code. This service instantly retrieves the city, state, county, time zone, latitude and longitude, ZIP and Postal Codes within a radius, along with US census information for a given ZIP or Postal Code.
<b>Cortera Business Vitals</b> <a href="#">Try It Now</a>	Think valuable insights about companies are tough to come by? Think again. Cortera Business Vitals delivers the key elements of a business location including a unique ID, complete address, phone number, URL, SIC, NAIC, employees, sales and year started.
<b>Zacks Company</b> <a href="#">Try It Now</a>	Provides live financial information on public companies from Zacks, including descriptions, pricing information, market caps, peer information, EPS, dividend information and other key financial ratios.
<b>Zacks Consensus</b> <a href="#">Try It Now</a>	Provides live financial information on public companies from Zacks, including current consensus on EPS estimates, estimates added, removed, changed, or revised, EPS perspectives on 30-day, quarter, 6-month, this year, and coming year, long term growth consensus, and more.
<b>Zacks Summary</b> <a href="#">Try It Now</a>	Provides live financial information on public companies from Zacks, including broker buy/hold/sell recommendations, earnings estimates, consensus EPS trends, EPS surprises both current and historical, and earning comparisons to the industry.
<b>NCOA Verification</b>	Did you know 40 million Americans change their address annually? Need help keeping your data clean and deliverable? StrikeIron can help. NCOA Verification can not only identify who has moved, but will provide the new address.



Source: <http://docs.gnip.com/w/page/23722723/Introduction-to-Gnip>

# ASE – complex, diverse and elastic properties

- Different platforms and multiple types of data and services from multiple providers for multiple stakeholders
- Complex service-based systems
  - Not just big data in a single organization which can be dealt by using, e.g., MapReduce/Hadoop
  - Not just take the data and do the computation: how to guarantee multitude of data/service concerns?
  - Not just things and software: when do we need human services?
  - Not just local actions: we need coordination-aware techniques for multiple resources

→ **Quality expectation (from the users)** are elastic: they are not fixed and dependent on specific contexts!

# ASE – relevant courses

- **Existing courses provide foundations**
  - Advanced Internet Computing
    - Give you some advanced technologies about SOC, Cloud Computing and (business) processes/workflows
  - (Advanced) Distributed Systems
    - Give you fundamental distributed system concepts and technologies
  - Distributed Systems Technologies:
    - Give you fundamental distributed technologies and how to use them
- **But they do not deal with engineering such large-scale, complex service-based systems**

Questions

**ARE YOU WORKING ON SUCH  
SYSTEMS? IS THIS COURSE  
SUITABLE FOR YOU?**

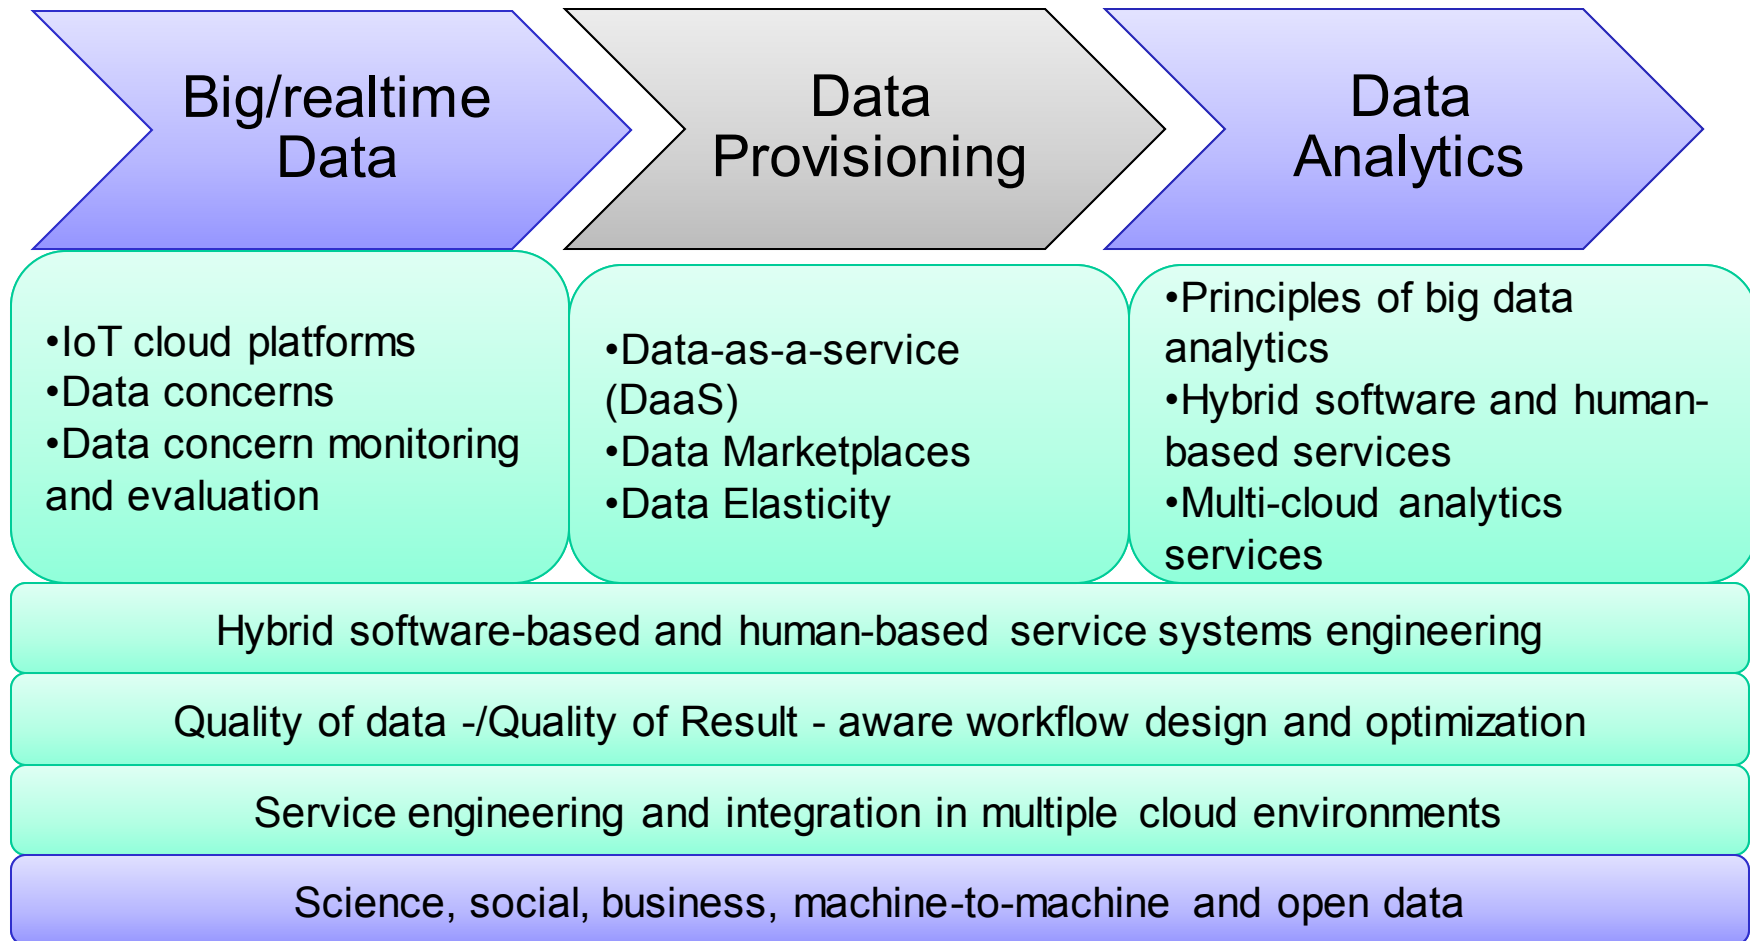
# What is the course about? (1)

- Discuss new concepts and techniques for engineering advanced, Internet-scale, elastic service-based systems
- Focus on service systems for complex data analytics, programming elasticity, and principles for engineering IoT cloud systems and for social-physical cloud systems
- Consider a wide range of applications for real-world problems in machine-to-machine (M2M), science and engineering, and social media

We **research and explore** emerging techniques for **interesting scenarios** by utilizing **existing, advanced technologies!**



# What is the course about? (2)



# EXAMPLES FROM PREVIOUS YEARS

# References for the course

- No text book designed for this course
- Some references from recent scientific papers
- Relevant research in big data
  - But not very much on data management or individual data processing framework (e.g., MapReduce/Hadoop)
- Relevant work in IoT, humans and software integration
- Distributed and Cloud computing
- Edge computing

# Course administration (1)

- Lectures are held through the whole semester
  - But not every week – check the course website!
  - Also backup dates
- Some tools from TU Wien
  - <http://tuwiendsg.github.io/>
  - <http://tuwiendsg.github.io/iCOMOT/>

## Course administration (2)

- Who could participate in the course?
  - Master students in advanced stages (e.g., seeking for master thesis) in informatics and business informatics
  - PhD students: normal PhD track, PhD School of Informatics, and Doctoral Colleges
  - Students should have knowledge about fundamental distributed systems, internet computing and distributed computing technologies

# Course administration (3)

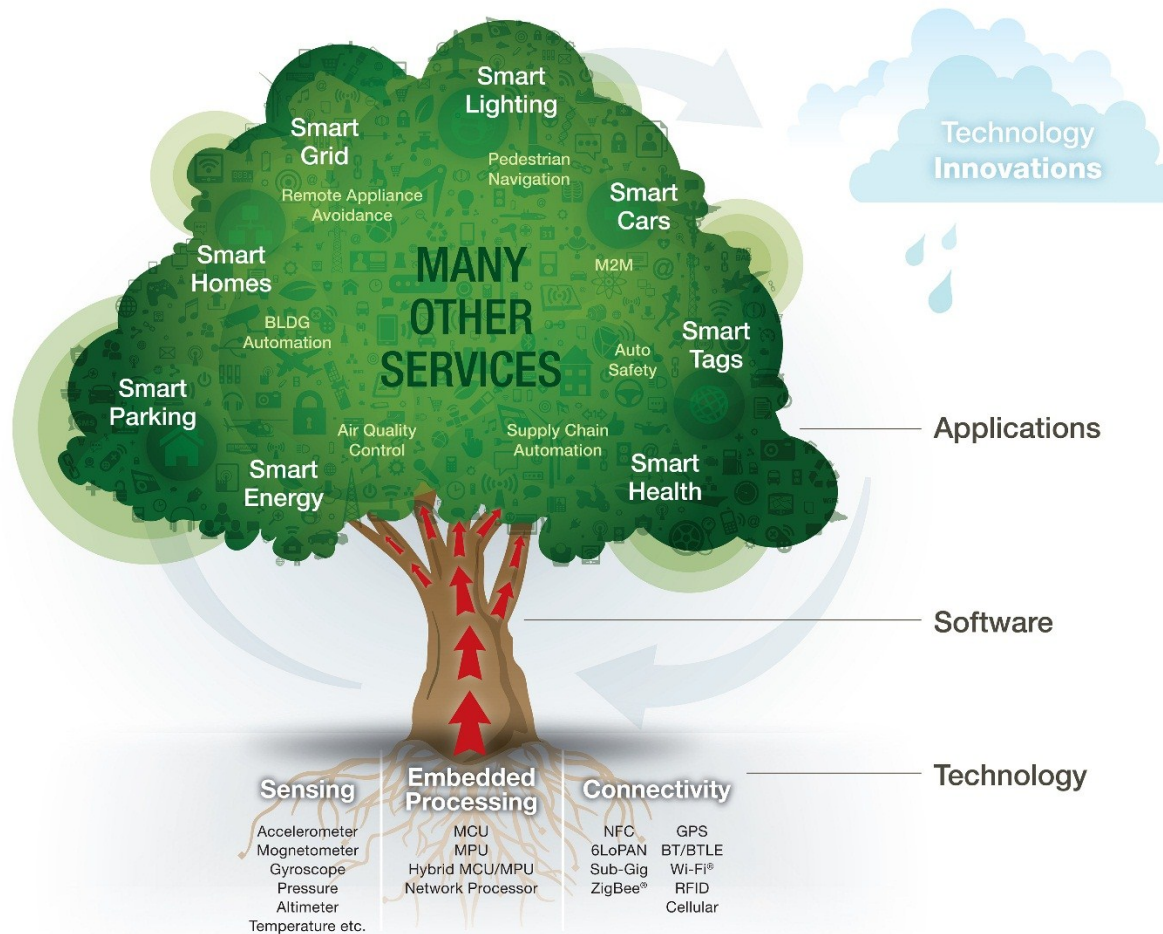
- Learning methods
  - Discussion, individual and team work, design, engineering and evaluation actions
- Evaluation methods
  - Assignments, a mini project and a final examination
- Assignments (hard deadline)
  - 4 home assignments resulting in some design/deployment and analysis summaries
- Mini project (hard deadline)
  - One mini project resulting in a small prototype/conceptual design
- Oral final exam (flexible, until Sep 2016)



# Assignments and Mini Project

- Define your interesting scenario
  - Look around, imagine and create your own scenario!
- Analyze and implement some concepts in the lectures
- Prototype and demonstrate your work
  - Code the prototype and present your work
  - We use github.com and we would like to make all code available (unless you have a reason to hide it)
  - Send your github account to me:  
<https://github.com/AdvancedServicesEngineeringTUWien2016>
- Some best results will be shown in the course website

# Running out of topics for your mini projects?



Source: <http://eecatalog.com/loT/files/2014/04/Freescale-Internet-of-Things-Tree.jpg?file=2014/04/Freescale-Internet-of-Things-Tree.jpg>



# Grades

- Participations + discussions: 10 points
- Assignments: 40 points
- Mini project: 20 points
- Final oral examination: 30 points

Point	Final mark
90-100	1 (sehr gut)
75-89	2 (gut)
56-74	3 (befriedigend)
40-55	4 (genügend)
0-39	5 (nicht genügend)

**Failed ? → retake the final oral examination part!**

**THANKS! ANY QUESTION?**

# Thanks for your attention

Hong-Linh Truong  
Distributed Systems Group  
Vienna University of Technology  
[truong@dsg.tuwien.ac.at](mailto:truong@dsg.tuwien.ac.at)  
<http://dsg.tuwien.ac.at/staff/truong>