From Cloud to Fog Computing

10 years of Deployment Experience
of Internet of Things Applications

Ioannis Chatzigiannakis

Sapienza University of Rome
Department of Computer, Control, and Management Engineering (DIAG)

Friday, March 19th, 2021

From Vacuum Tubes . . . millimeter-scale SoC

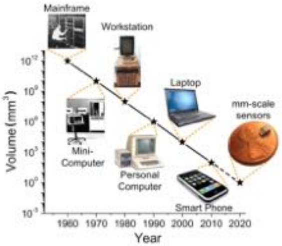

The rise of the IoT semiconductor

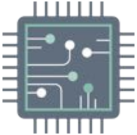

Semiconductor components that either individually or collaboratively contribute to the functionality of an IoT device.

- Microcontrollers (MCU),
- Microelectromechanical Systems (MEMS),
  - Sensors,
  - Actuators,
  - Energy harvesting.
- Connectivity chipsets,
- Embedded AI chipsets,
- Security chipsets.

Capabilities of Smart, Connected Products

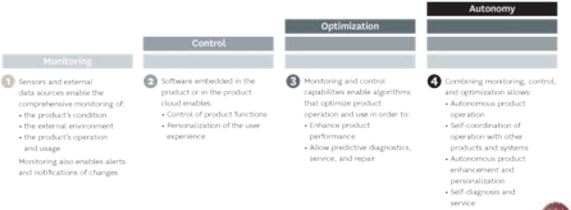

Redefining Industry Boundaries

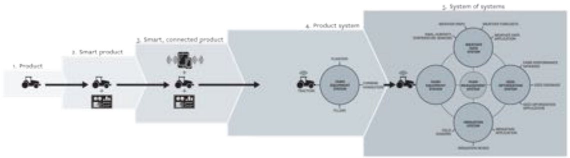

The Value of Sensor Data

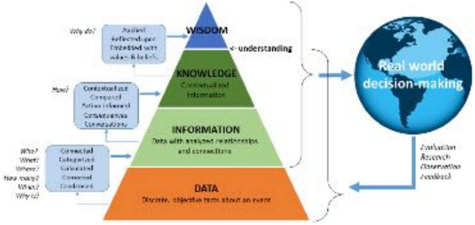

The New Technology Stack

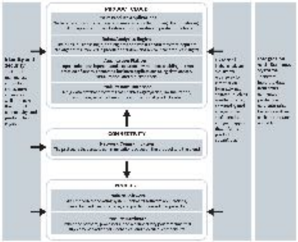

SmartSantander: A City-Scale IoT Infrastructure

- City-scale IoT deployment.
- 18,000 sensing points.
- Outdoor deployments, Mobile nodes, Human interaction.
- Integrated City Services – Continuous operation.
- Large variety of sensors – Large data.
- Business models and sustainable exploitation combining research & service support.

2010 . . . 2013 – FP7-ICT-2009.1.6 - Future Internet experimental facility and experimentally-driven research.

Monitoring Off-Street and Loading Spaces

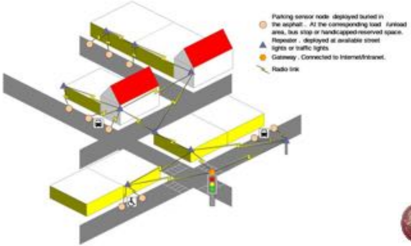

Park Irrigation Monitoring

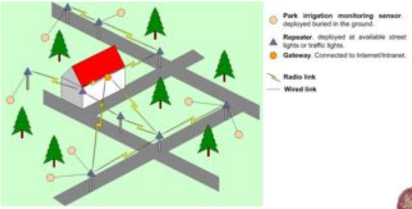

Public Transportation Monitoring

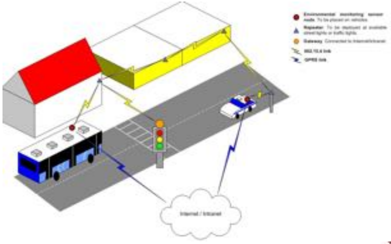

Platform Architecture & Services

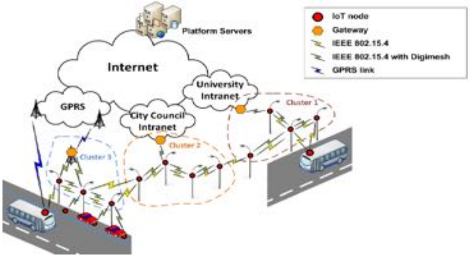

Platform Architecture & Services

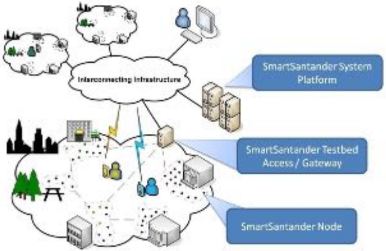

Platform Architecture & Services

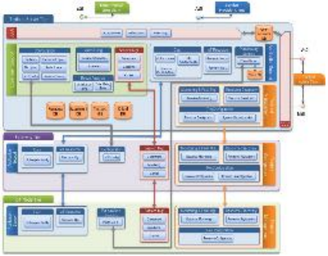

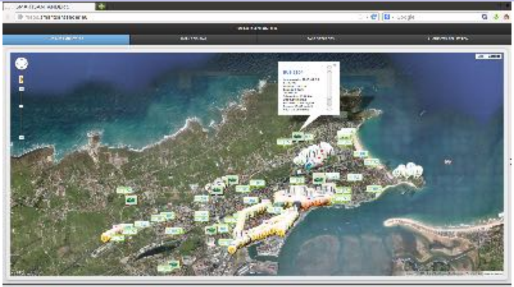

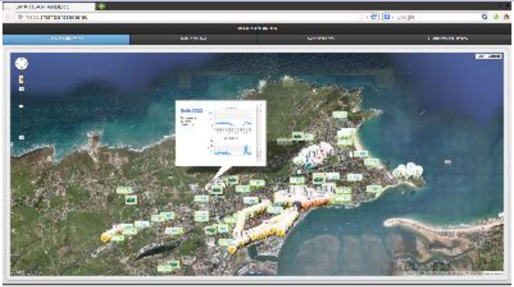

Data gathering: a general problem

- It is impossible to reach every corner of every neighborhood in the city.
- Even if you choose to do it . . .
- We need a secure, reliable infrastructure that enables interconnectivity and scale.
- Building this infrastructure is difficult, expensive and time-consuming.
- Maintenance of such an infrastructure is expensive and requires continuous investements.

Data gathering: a general problem

- Water: "U.N. has a limited success to get accurate information on water infrastructure and treatment systems".
  - Poor data, weak agencies hamstring U.N. environmental oversight, NY Times, 2009.
- Food: "Agricultural statistics has deteriorated over time" - weak estimation of global rice/wheat productions - fisheries data outdated.
  - Food and Agriculture Organization, Audit 2009.
- Health: "Exposure measures are sometimes completely lacking, frequently incomplete or otherwise uncertain".
  - Uncertainty and Data Quality in Exposure Assessment, World Health Organization, 2008.

An Example: Monitoring Pollution

- Important environmental issues in cities
- (long term) health, social and economic impacts
- An increasing problem, especially in developing countries
- Growing public concern & effort (European Directive -2002)
- but limited success of environmental policies.
- Complexity of monitoring the real exposure of the population.


Los-Angeles


Mumbai

OrganiCity: Enabling Co-Creation of Smart City Services

Expand IoT facilities by involving People: Participatory Sensing

- Access to powerful, rich-sensor smartphones.
- Use sensors for gathering quantitative information.
- Use people as sensors for gathering qualitative information.

2015 . . . 2018 – H2020-EU.2.1.1.3 - Future Internet: Software, hardware, Infrastructures, technologies and services.

Platform Architecture & Services

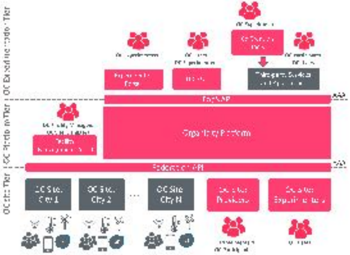

Platform Architecture & Services

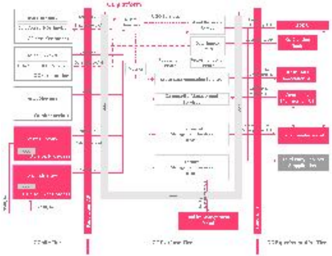

OrganiCity Web Platform

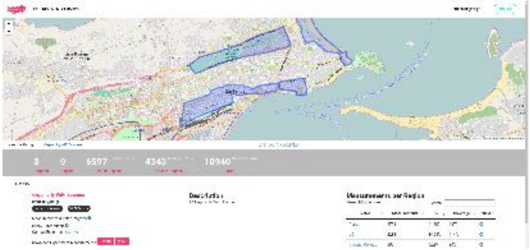

OrganiCity Web Platform

- Detect ambient noise in city centers throughout the day.
- Citizens participate using their smartphones to extend IoT infrastructure.
- Data collected are forwarded to the Cloud-based Organicity platform for analysis.

Stationary IoT nodes Readings 12:00 – 18:00 Readings 18:00 – 24:00

OrganiCity Web Platform

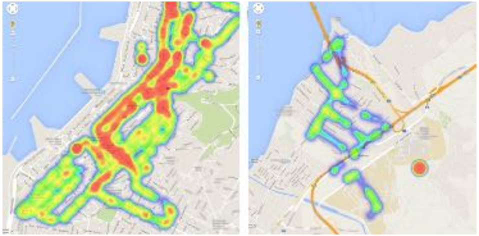

City center (left) and suburb/campus (right) average noise levels between 18:00-24:00

GAIA: Green Awareness in Action

- Educational buildings constitute 17% of the non-residential EU building stock
- Wide variety of buildings of different ages (some built around 1950s)
- Expensive to renovate existing buildings
- Affect the behavioral characteristics of the buildings users.

2016 . . . 2019 – H2020-EU.3.3.1. - Reducing energy consumption and carbon foorpint by smart and sustainable use.

Platform Architecture & Services

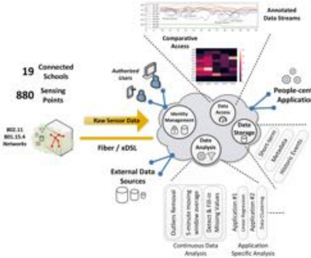

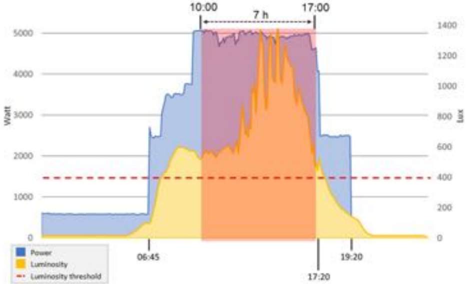

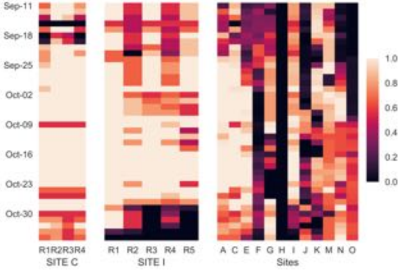

It's all in the Cloud

Today's Cloud and its Characteristics

Common Characteristics:

- Massive Scale
- Homogeneity
- Virtualization
- Low Cost Software
- Resilient Computing
- Geographic Distribution
- Service Orientation
- Advanced Security

Essential Characteristics:

- On Demand Self-Service
- Broad Network Access
- Resource Pooling
- Rapid Elasticity
- Measured Service

Batch processing of IoT Data

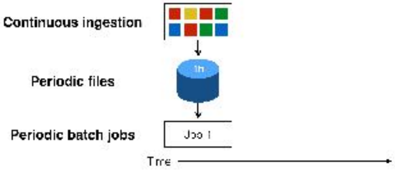

- Data cleaning – erroneous values
- Data enrichment – missing values
- Produce statistics – data mining

Batch processing of IoT Data

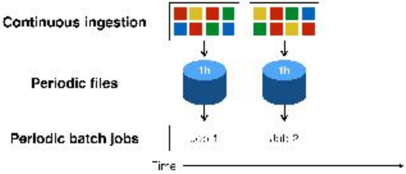

- Data cleaning – erroneous values
- Data enrichment – missing values
- Produce statistics – data mining

Batch processing of IoT Data

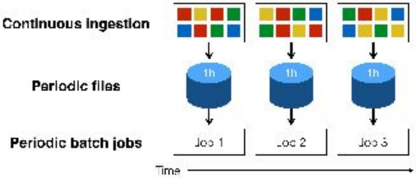

- Data cleaning – erroneous values
- Data enrichment – missing values
- Produce statistics – data mining

Batch processing of IoT Data

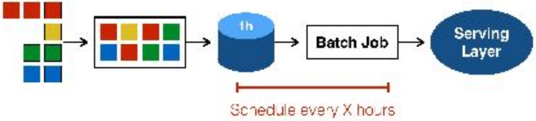

- High Latency from event to serving layer usually in the range of hours.
- Historical fact finding.
- Find and analyze information stored – pull model.
- Query-driven: bring data to the analytics
- High Latency from event to actionable insights.

Cloud-only Computing Models Inadequate

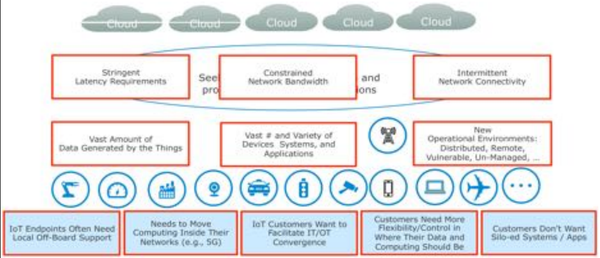

Cloud-to-Device Continuum

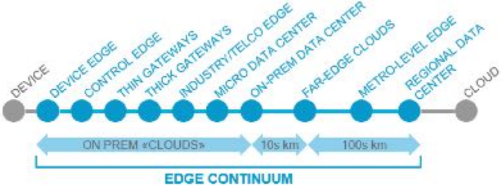

- Enable computing anywhere along the continuum.
- Software deployment closer to the devices.
- Distributed cloud computing.

Deliver Data to the Analytics

Traditional approach so far...

Analyze Data in the Right place

Traditional approach so far...

Fog computing: distributed analytics from Cloud to Source

Stream-based Data Processing

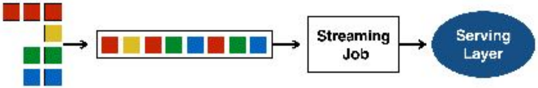

- Stream-based processing is enabling the obvious: continuous processing on data that is continuously produced.
- Current fact finding – Monitor data and react in real time.
- Analyze data in motion – push model
- Data-driven: bring analytics to the data
- Low Latency from reception of event to actionable insights.

Heterogeneity Challenges in Cloud-to-Device Continuum

CPU

Intel Ice Lake (10nm)
8 cores HT, AVX(512 SIMD)
~1TFlops* (including the iGPU)
~ TDP 28W

GPU

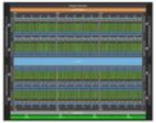

NVIDIA GP 100 – Pascal - 16nm
60 SMs, 64 cores each
3584 FP32 cores
10.6 TFlops (FP32)
TDP ~300 Watts

FPGA

Intel FPGA Stratix 10 (14nm)
Reconfigurable Hardware
~ 10 TFlops
TDP ~225Watts

Heterogeneity Challenges in Cloud-to-Device Continuum

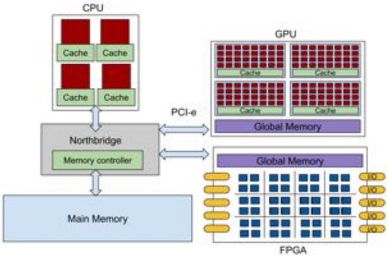

Heterogeneity Challenges in Cloud-to-Device Continuum

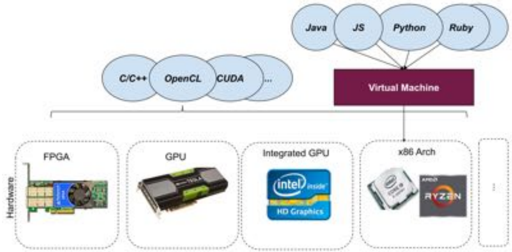

Heterogeneity Challenges in Cloud-to-Device Continuum

- Today a variety of mature stream processors are available: Flink, Spark . . . .
  - Data analytics stacks cannot exploit heterogeneous hardware dynamically.
  - Rely on existing, pre-complied operators or kernels for specific HW resources.
- Enabling dynamic compilation of arbitrary code to any* heterogeneous hardware device.
- Designing an intelligent elastic system where applications are executed on the best hardware resource

2018 . . . 2020 – H2020-ICT-16-2017 - Big data PPP: addressing technology challenges of the data economy.

Heterogeneous execution of Big Data Tasks

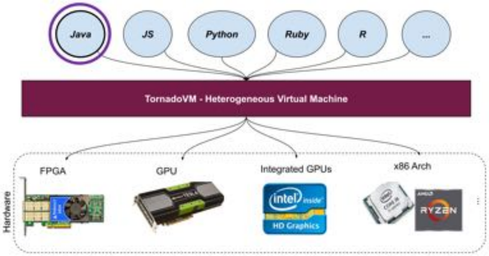

Heterogeneous execution of Big Data Tasks

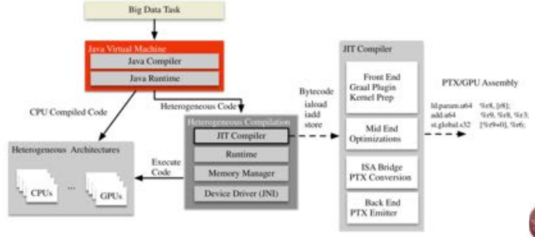

Heterogeneous execution of Big Data Tasks

```
public static void matrixMultiplication(final float[] A, final float[] B,
                                        final float[] C, final int size) {
    for (@Parallel int i = 0; i < size; i++) {
        for (@Parallel int j = 0; j < size; j++) {
            float sum = 0.0f;
            for (int k = 0; k < size; k++) {
                sum += A[(i * size) + k] * B[(k * size) + j];
            }
            C[(i * size) + j] = sum;
        }
    }
}
```


Fog-Computing based Smart Metering

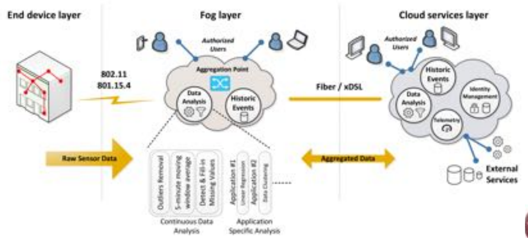

Fog-Computing based Smart Metering

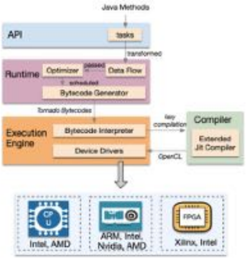

- Data Analytics stack executed directly over TornadoVM.
- Suitable for applications that require hardware acceleration but do not benefit from or fit a Big Data framework.
- Enables hot swapping of underlying runtimes with TornadoVM.

Performance Evaluation

- On average 15 sensing devices are available on each school.
- Each of them is equipped with 1 to 5 sensors.
- Sensing rate is 30*sec* thus on average each edge device is receiving 2.5 sensor updates every second.

	ARM	**Intel**	**Intel Xeon**
Processor	Cortex-A53	i3-3120M	E5-2630V4
Frequency	1.2Ghz	2.2GHz	2.2GHz
Cores	4	2	8
Memory	1 GB	8 GB	24 GB
Disk	64GB	120GB	600GB
Type	SD Card class10	SSD	SSD

Performance Evaluation

- Data engineering over windows of data
  - Min/Max/Average/Sum.
  - Outliers Detection with a Hampel Filter.
- Maximum number of messages processed per second

Device	Processing Rate (Msg/Sec)
Cortex-A53	15,36
i3-3120M	2,692.31
E5-2630V4	5,833.33

- Processing time per message (Sec)

Service	Cortex A53	i3-3120M	E5-2630V4
Window Setup	1.620	0.016	0.004
Data Engineering	63.817	0.186	0.145

Performance Evaluation

- AArch64 Mali G71 MP8 GPU available on the ARM Cortex-A.
- Data engineering over windows of data
- Speed-up from 3.5x to 9.3x

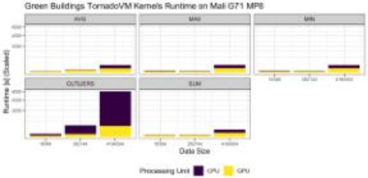

Performance Evaluation

- Tesla V100-SXM2-32GB GPU available on the Intel Xeon.
- Data engineering over windows of data
- Speed-up from 75x to 449x

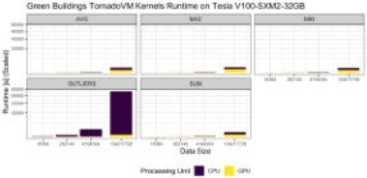

Hardware and Software Diversity in Big Data and IoT

Hardware and Software Diversity in Big Data and IoT

- Big Data and IoT have separate development routes
  - Diversification of the programming models
    - Big Data – Map/Reduce paradigm
    - IoT – Event-driven
  - Lack of interoperability
    - Big Data – Java/Python
    - IoT – C/C++
- Big Data and IoT have separate orchestration approaches
  - Cannot dynamically select which parts of the code should run on which side
  - Big Data – Rich orchestration tool-chain
  - IoT – Limited to Over-The-Air-Programming

Cloud-centric Internet of Things
From Cloud to Fog Computing
ELEGANT – Secure and Seamless Edge-to-Cloud Analytics

Hardware and Software Diversity in Big Data and IoT

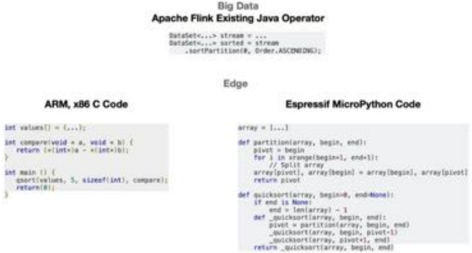

Cloud-centric Internet of Things
From Cloud to Fog Computing
ELEGANT – Secure and Seamless Edge-to-Cloud Analytics

Secure and Seamless Edge-to-cloud Analytics

- Developers will be able to view an IoT/Big Data deployment as a single system.
- Provide a unified API for data operators
  - Can be executed on both the IoT devices and the data analytics side.
  - Using a uniform orchestration process.
- Intelligent orchestration of code motion between IoT and cloud analytics
  - Gather profiling information of all devices
  - Gather profiling information of all data operators
  - Scheduling decisions regarding execution
  - Instruct code motion

2021 . . . 2023 – H2020-ICT-50-2020 - Software Technologies.

Cloud-centric Internet of Things
From Cloud to Fog Computing
ELEGANT – Secure and Seamless Edge-to-Cloud Analytics

Secure and Seamless Edge-to-cloud Analytics

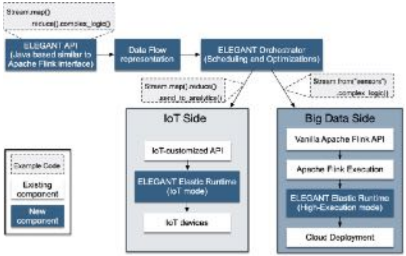

Our Vision: Unify IoT and Big Data fragmentation concerns

- Utilize such software and hardware diverse deployments as a single system
  - Heterogeneous hardware capabilities
  - Operate on demand depending on their specific use case requirements
- Enable computing anywhere along the continuum
  - Not just at specific edge devices.
  - Using a uniform execution environment.
- Orchestrate resources in clouds, fogs, and devices
  - Not just at isolated edge devices, systems, or apps
  - Using a uniform orchestration process.

. . . could be analogous to previous Internet revolutions

Thank you!

Questions ?