Real-Time Cloud & Big-Data Processing Infrastructures

Tommaso Cucinotta
Real-Time Systems Laboratory
Scuola Superiore Sant'Anna

Data Science Colloquium

Overview

- Presentations
- Scuola Superiore Sant'Anna
- RETIS group
- Research Overview
- Real-Time Cloud Computing
- Big-Data Streaming & Analytics
- Q&A

ISTITUTO
DI TECNOLOGIE DELLA
COMUNICAZIONE,
DELL'INFORMAZIONE
E DELLA
PERCEZIONE

Scuola Superiore
Sant'Anna

Presentations first...

Let me introduce myself…

- **2000**: MSc in Computer Engineering (**Computer Security**)
  - → Thesis: PKCS#11 module for Netscape
- **2004**: PhD in Computer Engineering (**Computer Security**)
  - → Interoperability in open-source **smart-card solutions**
  - → Open-source **MuscleCard** framework → RedHat CoolKey
- **2004-2012**: Researcher et al. at the **ReTiS**
  - → Adaptive scheduling for soft real-time systems
  - → **Deadline-based scheduler for the Linux kernel** for improved responsiveness of soft real-time, multimedia & virtualized services
- **2012-2014**: MTS in **Bell Labs**: research on **security** and real-time performance of cloud applications (NFV/IMS)
- **2014-2015**
  - → SDE in **AWS DynamoDB**: real-time performance and scalability of DynamoDB (NoSQL Real-Time DataBase)
- **2016**
  - → **Associate Professor** at the **ReTiS**
  - → Course on "**Component-Based Software Design**" within joint UNIPI/SSSA MSc degree on Embedded Computing Systems
  - → Board member for the **PhD in Emerging Digital Technologies** at SSSA
- **2017**
  - → Board member for the **PhD in Data Science** jointly run by SSSA, SNS, UNIPI, CNR, IMT Lucca

IN SUPREMÆ DIGNITATIS 1343
UNIVERSITÀ DI PISA

Scuola Superiore Sant'Anna
di Studi Universitari e di Perfezionamento

Bell Labs

amazon DynamoDB

Scuola Superiore Sant'Anna
di Studi Universitari e di Perfezionamento

Scuola Superiore Sant'Anna

An overview

ISTITUTO
DI TECNOLOGIE DELLA
COMUNICAZIONE,
DELL'INFORMAZIONE
E DELLA
PERCEZIONE

Scuola Superiore
Sant'Anna

Scuola Superiore Sant'Anna at a glance

Scuola Superiore Sant'Anna is a public university institute working in the field of applied sciences.

Rector: Prof. Pierdomenico Perata

- **Class of Social Sciences**
  - Economic Sciences
  - Legal Sciences
  - Political Sciences
- **Class of Experimental Sciences**
  - Agricultural Sciences
  - Medical Sciences
  - Industrial and Information Engineering

Scuola Superiore Sant'Anna at a glance

Scuola Superiore Sant'Anna became an autonomous university in February 1987.

- It includes Masters and PhD programs.
- It works as the **Honors College** of the University of Pisa for undergraduates offering extra advanced courses to its students.
- Rigorous public examination and selection (e.g., only 10 out of 400 are selected every year in Engineering).
- **Honors College Students** live at the Scuola for free, but have to get top marks in all their exams and show proficiency in two foreign languages.

Structure: 6 Research Institutes

Experimental Sciences

- Biorobotics Institute
- Institute of Communication, Information and Perception Technologies — Pisa
- Institute of Life Sciences — Pontedera

Social Sciences

- Institute of Economics — Livorno
- Institute DIRPOLIS (Law, Policy and Development)
- Institute of Management

SANT'ANNA RESEARCH PARK

Institute of Communication, Information and Perception Technologies (TeCIP)

TeCIP

Director: prof. Massimo Bergamasco

Research Units

- COMMUNICATIONS & NETWORKS (InRete)
- REAL-TIME & EMBEDDED SYSTEMS (ReTiS)
- PERCEPTUAL ROBOTICS (PERCRO)

Communications & Networks (InReTe Lab)

Focus: Optical communication systems

CNIT

Director:
Prof. F. Di Pasquale

Coordinator:
Prof. P. Castoldi

LNRF | **InReTe *Lab***

IRCPhoNeT

Optical Networks & Services	*Optical Communication Systems*	*Optical Communication Theory & Techniques*	*Optical Amplification & Sensing*	*Ultrafast Communication & Photonic Switching*	*Digital Photonics*
Prof. P. Castoldi	**Prof. E. Ciaramella**	**Prof. E. Forestieri**	**Prof. F. Di Pasquale**	**Ing. L. Potì**	**Ing. A. Bogoni**

Perceptual Robotics (PERCRO Lab)

Focus: Human-robot interfaces

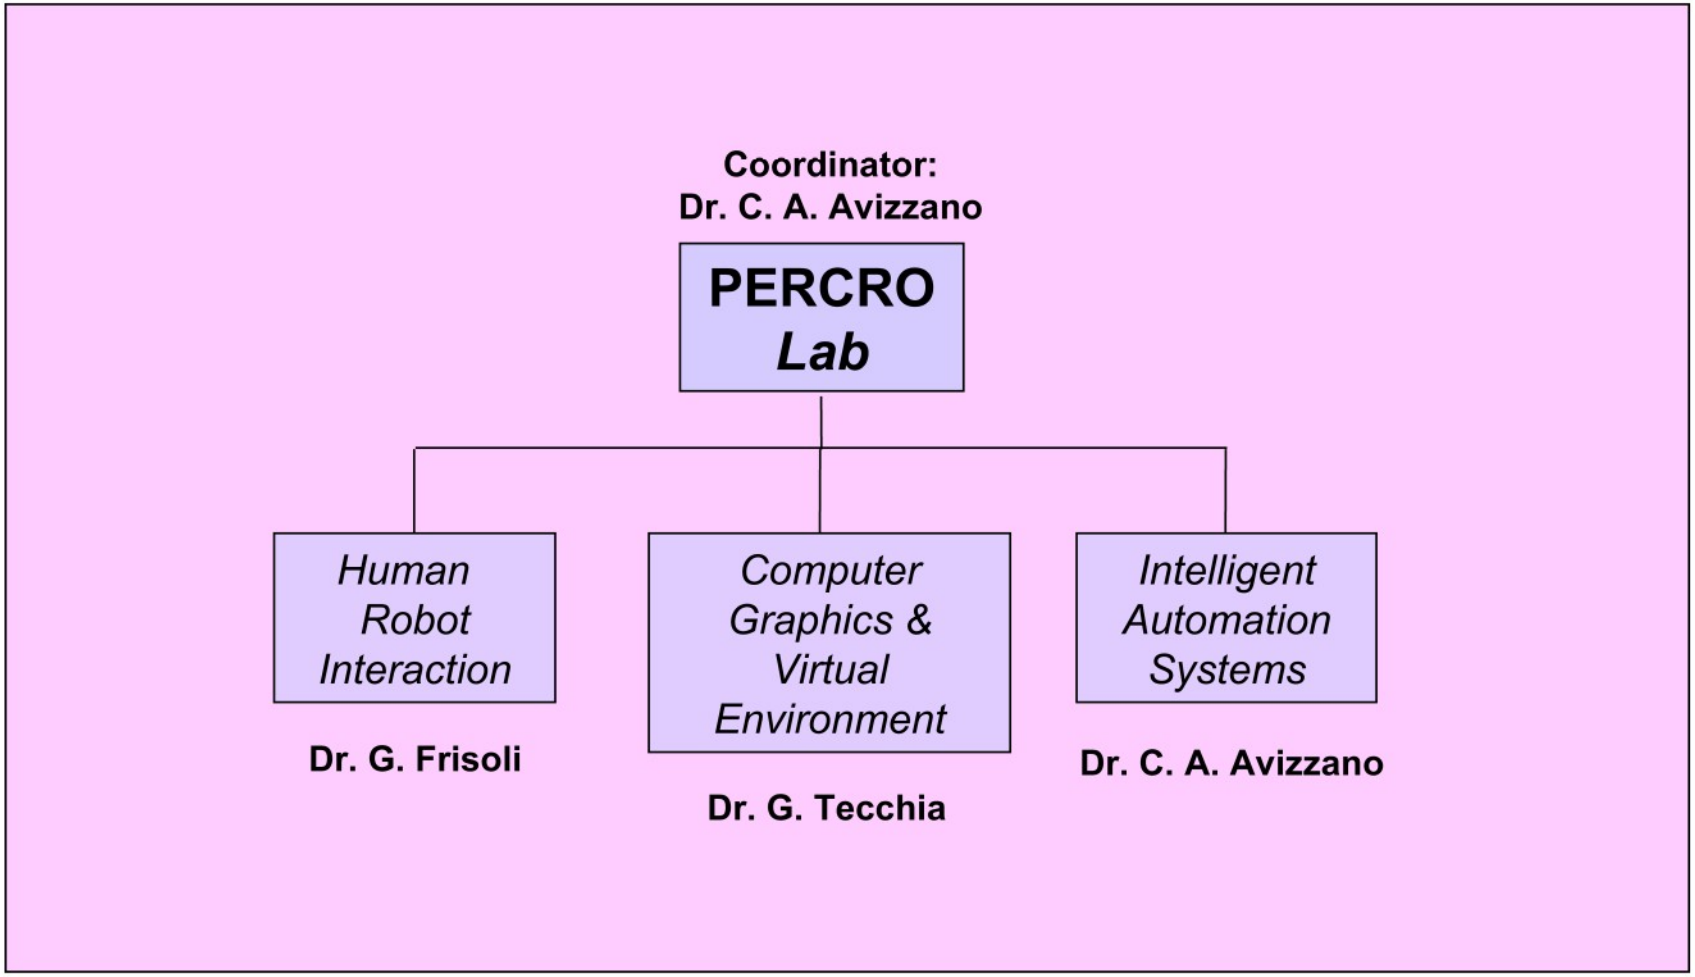

Real-Time Embedded Systems (RETIS Lab)

Focus: Predictable computing systems

Coordinator:
Prof. G. Buttazzo

RETIS *Lab*

Real-Time Scheduling & Resource Manag.	*Real-Time Operating Systems*	*Design of Embedded Systems*	*Cloud ComputingSystems*	*Real-Time Multimedia*	*Energy Aware Computing*
Prof. G. Buttazzo	**Prof. G. Lipari**	**Prof. M. Di Natale**	**Prof. T. Cucinotta**	**Prof. L. Abeni**	**Dr. M. Marinoni**

Real-Time Systems Laboratory

An overview

ISTITUTO
DI TECNOLOGIE DELLA
COMUNICAZIONE,
DELL'INFORMAZIONE
E DELLA
PERCEZIONE

Scuola Superiore
Sant'Anna

The RETIS Group (established in 1993)

It Includes 26 people:

- 2 Full professors
- 2 Associate professor
- 2 Assistant professors
- 5 Post Docs
- 5 Research Associates
- 10 PhD students

Mission of the RETIS Lab

- Increase software predictability through suitable
  - Operating systems mechanisms
  - Design methodologies and tools
  - Timing and performance analysis
- Provide real-time support for new computing platforms
  - multi-core, distributed, cloud
  - heterogeneous, FPGA-accelerated, massively parallel
- Make embedded systems resource efficient
  - w.r.t. time, memory, bandwidth, energy, …
- Prevent and manage overload conditions

→ through adaptation mechanisms

Research Topics

Resource management

Efficiency

Real-Time scheduling

Timing Analysis

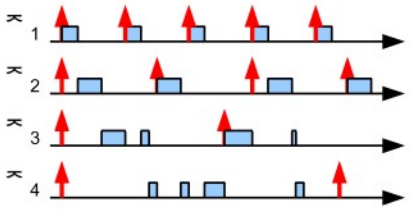

Real-Time Communication

Protocols

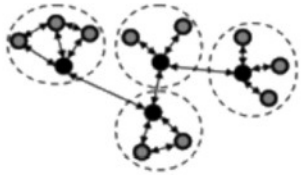

Retis

Real-Time Systems Laboratory

Design methods and Tools

Optimization

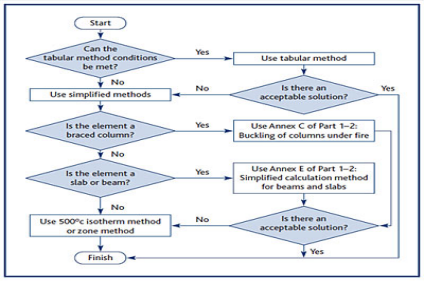

Adaptive behavior

Real-time Control

Real-Time Operating Systems

Multico platforn

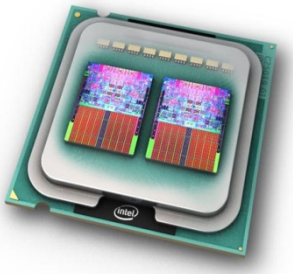

A research perspective

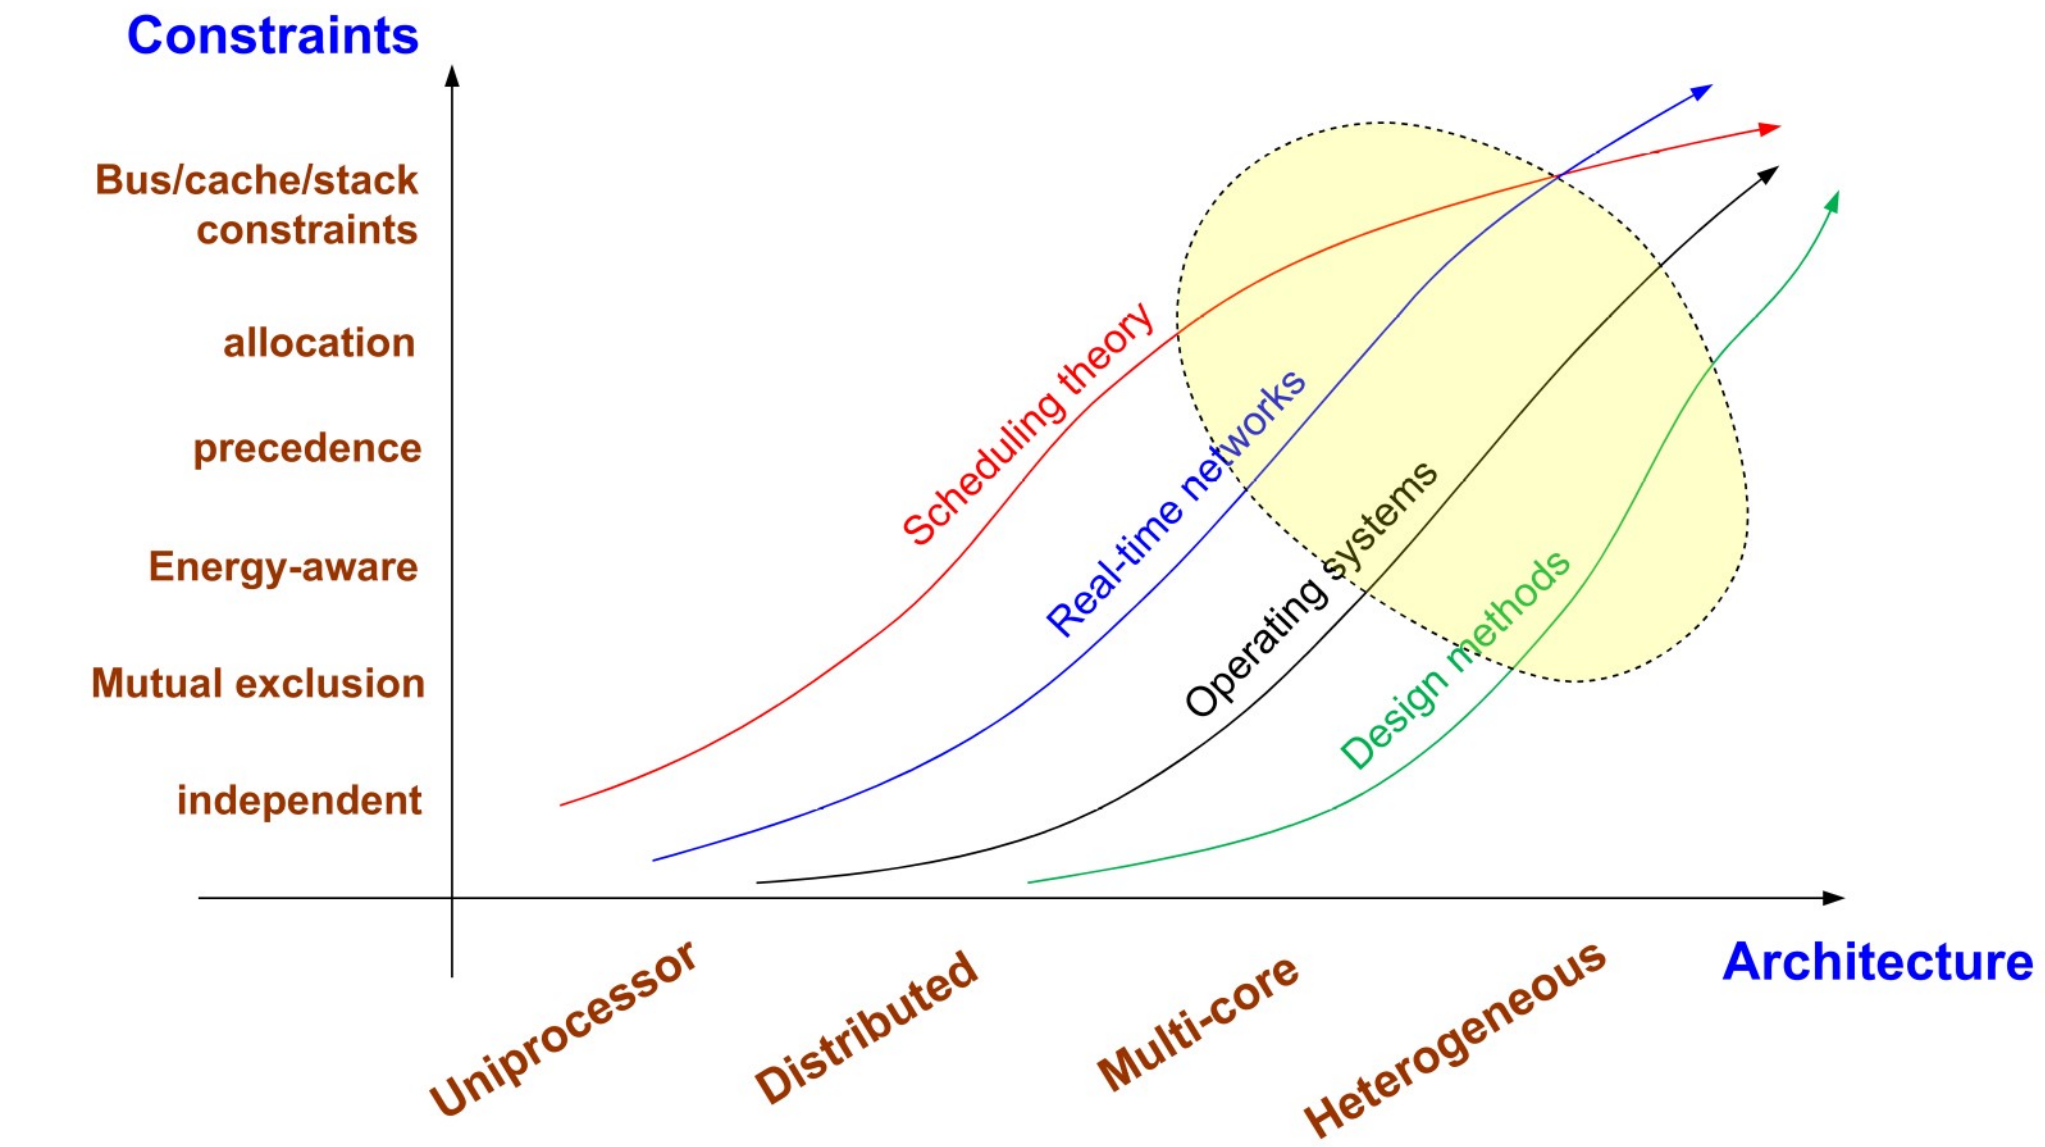

Applications

- Avionic systems
- Robotics
- Automotive systems
- Multimedia systems
- Sport & Healthcare
- Web & Cloud Services
- Agriculture
- Education

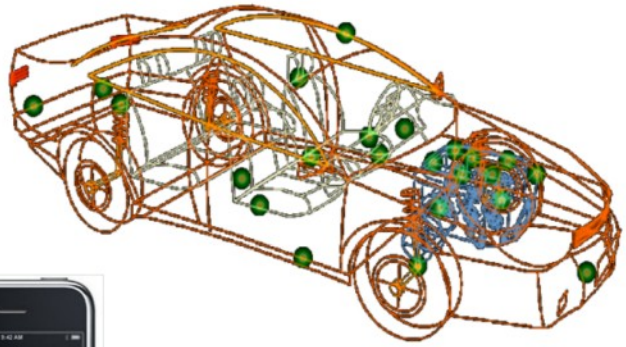

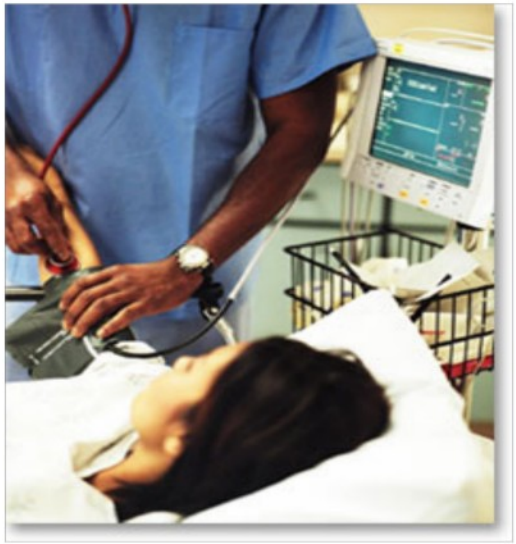

Real-Time Operating Systems

S.Ha.R.K.

for PC architectures

http://shark.sssup.it

Key features

- High predictability
- Posix interface
- Modularity - portability
- Advanced scheduling
- Adaptive resource mgm.
- Programming flexibility

E.R.I.K.A

for small microcontrollers

http://erika.tuxfamily.org

Key features

- Small footprint (1 Kb)
- OSEK certified!
- Support for multi-core
- and many ⌙ controllers
- High time resolution
- Advanced scheduling

Linux

In Linux **mainline** v. 3.14 since March 30, 2014

https://en.wikipedia.org/wiki/SCHED_DEADLINE

Key features

- Efficient scheduling
- Support for deadlines
- EDF scheduling
- Resource Reservation
- Temporal Isolation
- Support for multi-core

Industrial collaborations

Industrial impact of RETIS results

- **Sched_deadline**: Real-Time scheduler in the **Linux** mainline kernel
- **Erika:** Open source **OSEK certified kernel** for automotive systems (with Evidence)
- Predictable support for safety-critical systems (**Airbus**, **Bosch**, **Magneti Marelli**)
- Tele-rehabilitation and tele-monitoring (**Telecom**)

EU collaborations

Austria: Vienna

Czech Rep: Prague

France: Paris, Grenoble, Renne, Nantes, Nancy, Toulouse, Lille

Germany: Munich, Kaiserslautern, Dresden, Karlsrue, Saarland, Stuttgart

Greece: Athens

Ireland: Dublin, Cork

Italy: Pisa, Pavia, Catania, Siena, Florence, Bologna, Trento, Padova, Modena, Ancona, Rome

Portugal: Porto, Aveiro, Lisbon

Spain: Madrid, Cantabria,Valencia, Barcelona,, Palma de Mallorca

Sweden: Stockholm, Lund, Vasteras, Uppsala, Halmstad, Linkoping

Switzerland: Zurich, Lausanne,

UK: York, Cambridge

US & other collaborations

US

- **University of California at Berkeley** (Alberto Sangiovanni Vincentelli)
- **Carnegie Mellon University** (Raj Rajkumar, John Lehozcky)
- **University of Illinois at Urbana-Champaign** (Lui Sha, Tarek Abdelzaher)
- **University of North Carolina at Chapel Hill** (Sanjoy Baruah, James Anderson)
- **University of Texas at Austin** (Aloysius Mok)
- **University of Pennsylvania** (Insup Lee)
- **Florida State University** (Ted Baker)
- **University of Pittsburgh** (Daniel Mossé)
- **University of Virginia** (John Stankovic)
- **University of Nebraska-Lincoln** (Steve Goddard)
- **George Mason University** (Hakan Aydin)
- **Indiana-Purdue University Indianapolis** (Yao Liang)

Other

- **Universidad del Sur, Bahìa Blanca** (Rodrigo Santos)
- **Federal University of Santa Catarina, Brazil** (Romulo S. de Oliveira)

Research projects

RETIS developed timing analysis methods and tools for increasing software predictability in safety-critical systems.

In collaboration with

AIRBUS

BOSCH

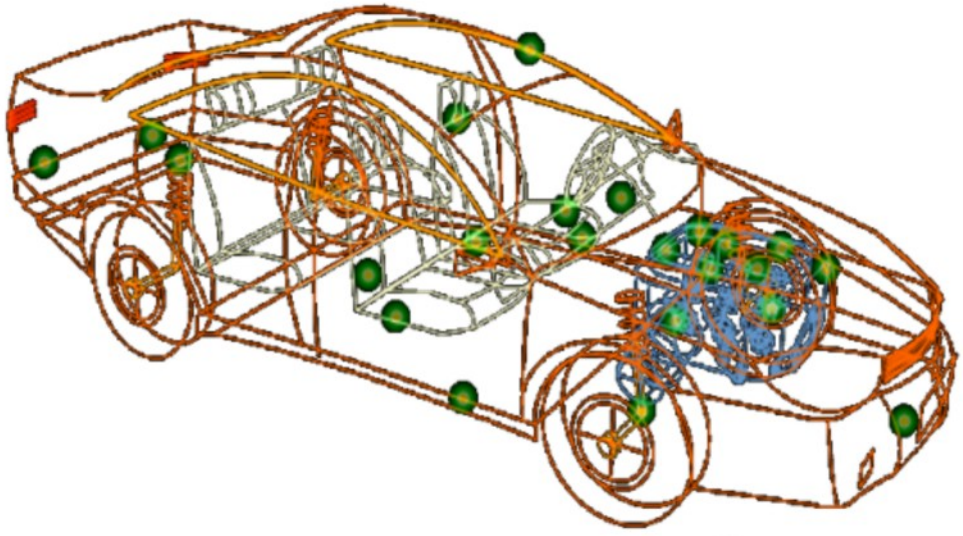

Framework for Real-time Embedded Systems based on COntRacts

frescor

- Soft and hard real-time
- Support for **CPU**, **disk**, **network**
- Distributed real-time systems
- Portability across RTOSs
- **Adaptive** real-time systems
- Application-level QoS control
- QoS+**power-aware** optimization
- Atomic negotiations

- http://www.frescor.org

THALES ENEA

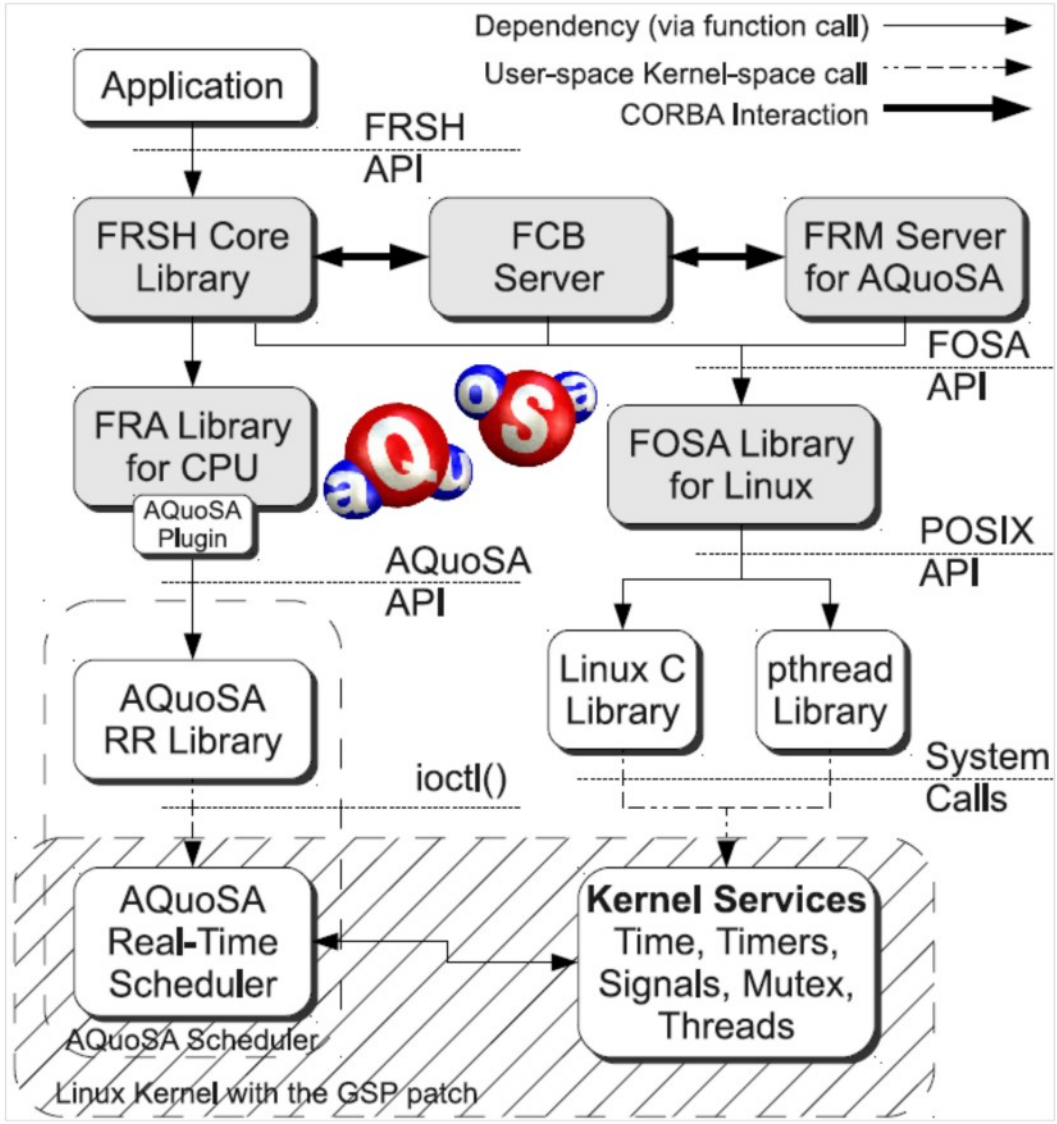

Framework for Real-time Embedde Systems based on COntRacts

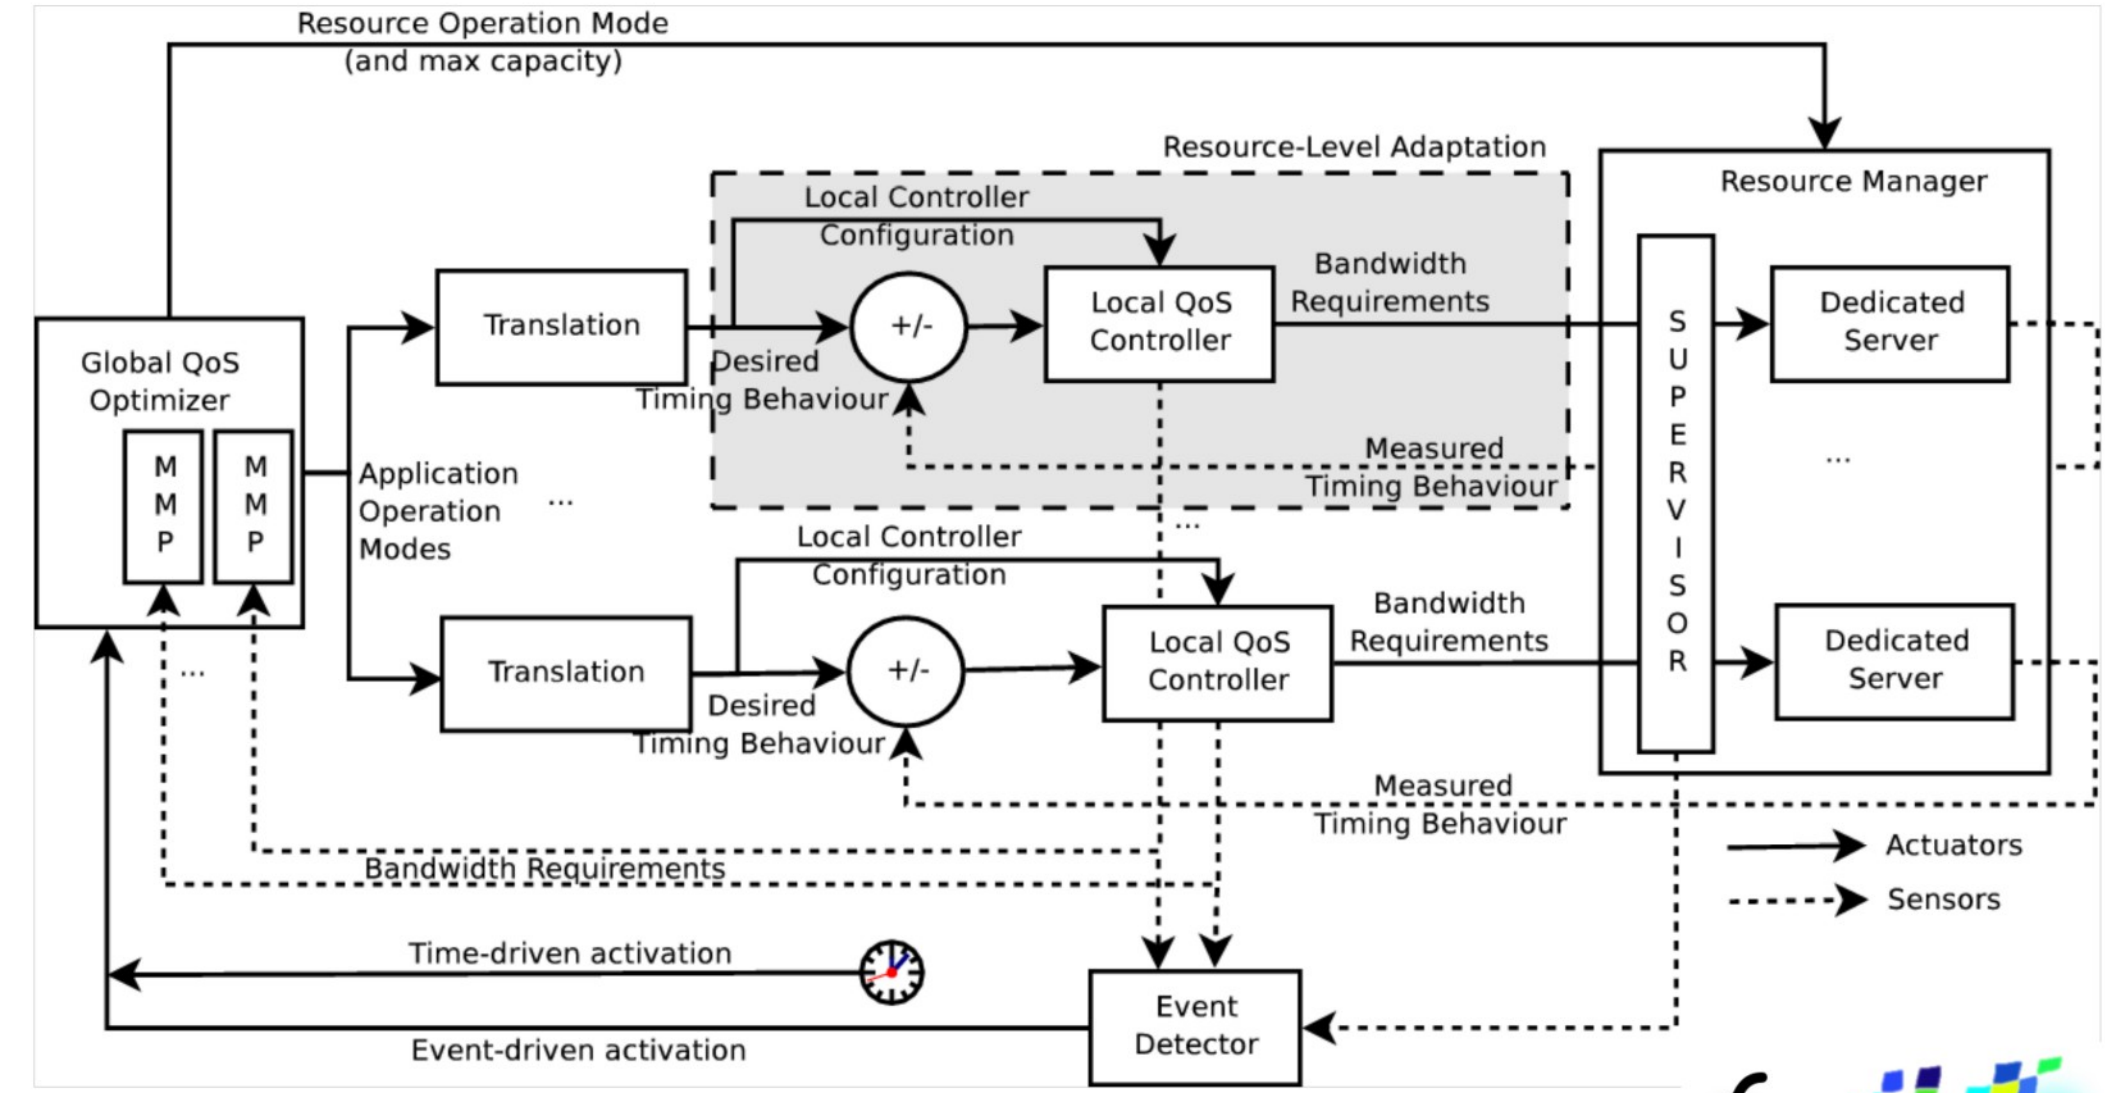

ACTORS

RETIS provided new kernel algorithms for adaptive resource management in next-generation cell phones.

Linux was modified to include a new scheduler able to reduce interference among concurrent software activities

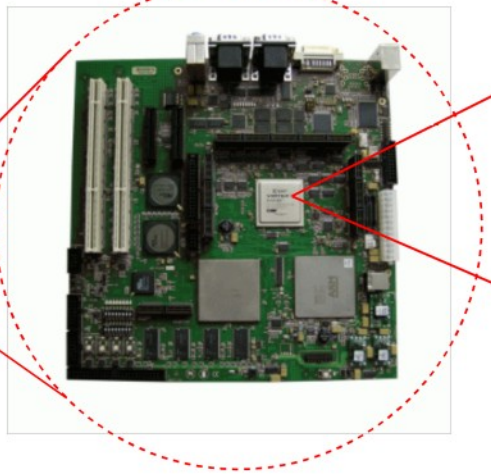

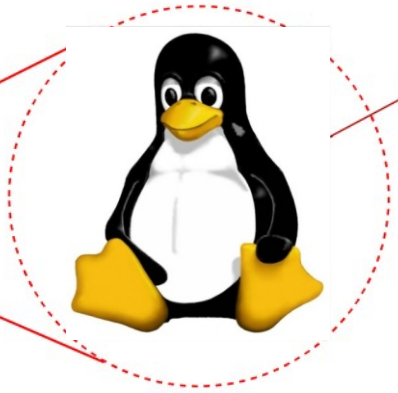

SCHED_DEADLINE

Included in Linux main line v3.14 since March 30, 2014

Linux 2.6.x

ARM11 (4 cores)

IRMOS

IRMOS
Interactive Realtime Multimedia Applications on Service Oriented Infrastructures

IaaS infrastructure for Real-Time cloud services

- Film post-production (on-line editing of 4K contents)
- Virtual & augmented reality
- Advanced e-Learning scenarios

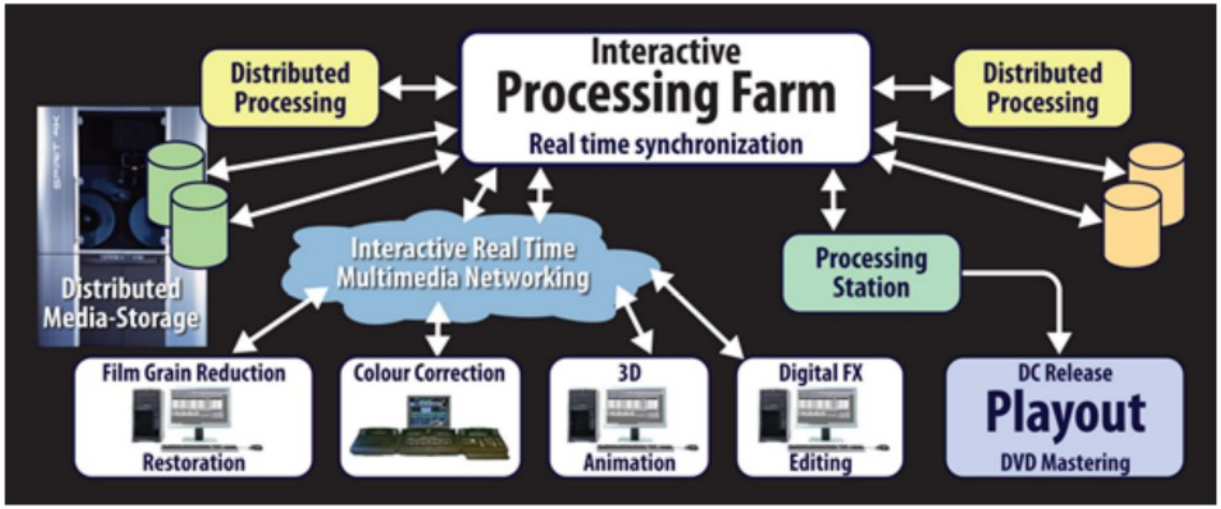

xyratex
Telefonica
Telefónica Investigación y Desarrollo
Alcatel·Lucent
GIUNTI labs
THOMSON images & beyond
THOMSON
santAnna school of advanced studies
iT innovation
SINTEF

Service-oriented Operating Systems

S(o)OS

Service-oriented Operating Systems

http://www.soos-project.eu

Future Massively Parallel Many-Core Systems

- Cloud and High-Performance Computing
- Manage heterogeneous hardware
- Workload distribution, replication, load balancing and scheduling
- Minimize maximum latency
- Minimize energy consumption while keeping timing constraints
- OS with **distributed kernel architecture**
- Experimentation with **Intel Single-chip Cloud Computer (SCC) NoC**

- Performance, guarantees and scalability of large streaming and storage data systems.
- Create a ***Java Platform*** to support H-P Information Management systems with real-time constraints.

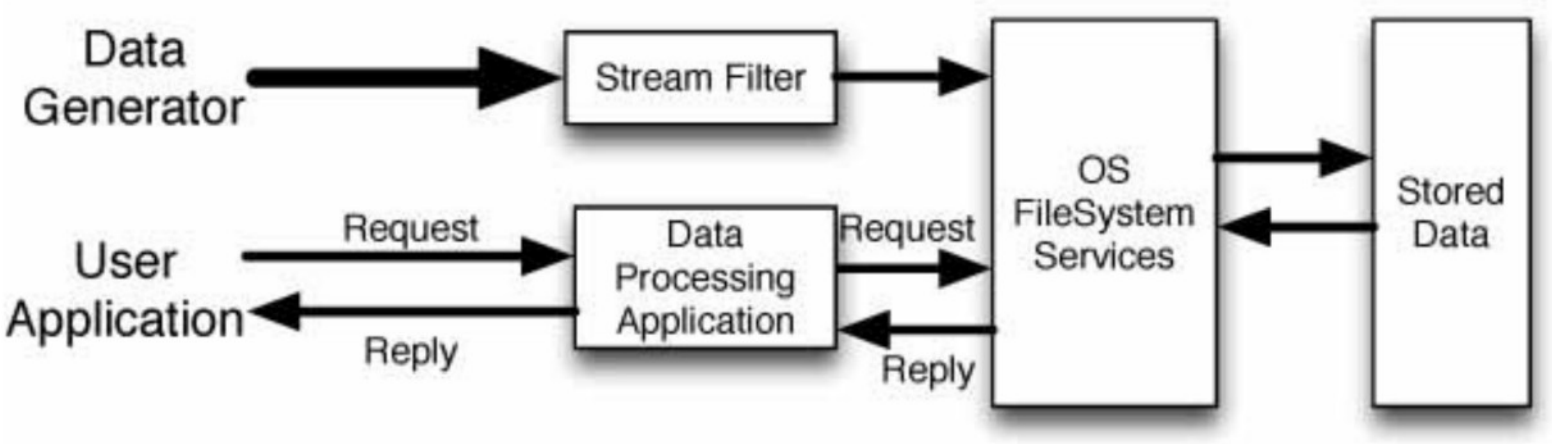

- **Financial** and **web streaming** case studies from industrial partners.

- Requirements on the OS and system software:
  - control the ***scheduling*** and the ***allocation*** of resources
  - isolate the timing behaviour of one component from others (temporal isolation)
- Enhanced kernel features:
  - CPU reservation for independent threads (*SCHED_DEADLINE*)
  - CPU reservation with mutexes (*M-BWI*)
  - Hierarchical CPU reservation (*BDM*)
  - Disk I/O proportional share scheduling (*BFQ*)
  - Network proportional share scheduling (*QFQ+*)
  - PREEMPT_RT compatibility

Safety and Security by design for interconnected mixed-critical CPS

- Engineering methods for cyber-physical systems for **safety** and **security by design**.
- Tools and capabilities to **prevent attacks in real-time**.
- **Guidelines** to assist designers and developers during the whole engineering process.

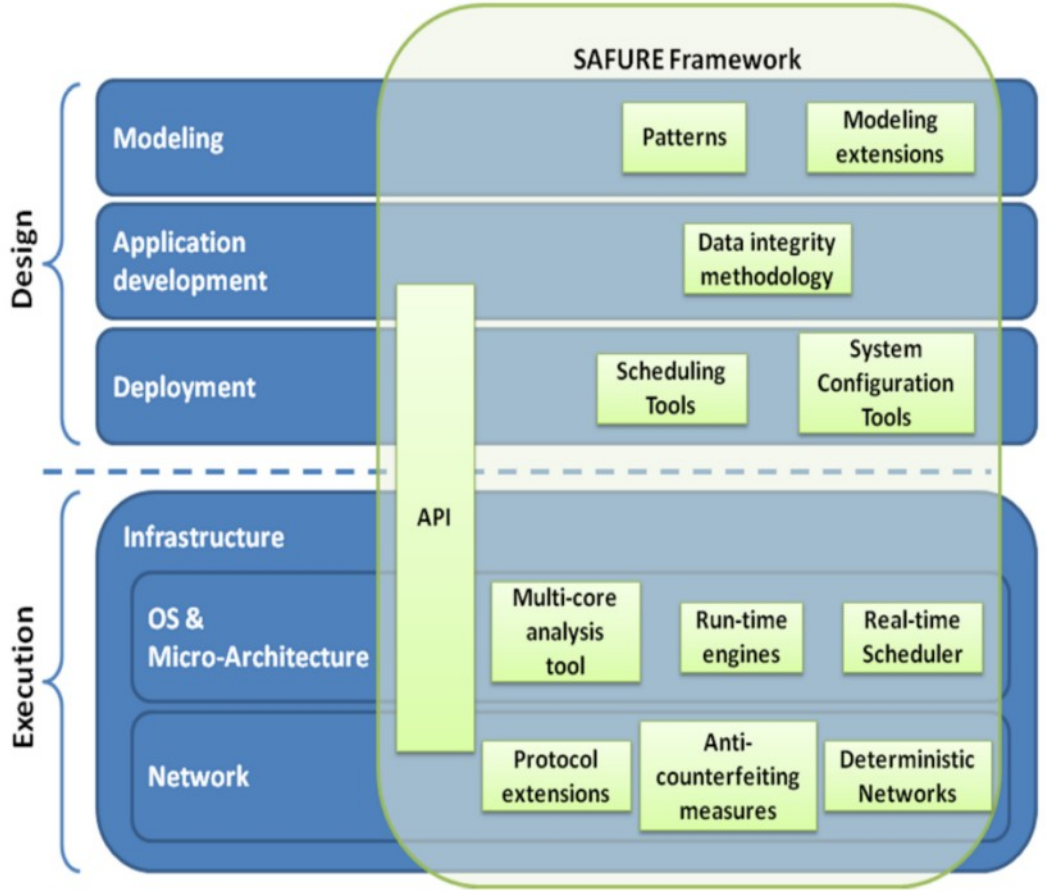

Other research topics

Engine control

Timing analysis of periodic and angular tasks activated at specific angles crankshaft.

The activation rate of these tasks is proportional to the angular velocity of the engine:

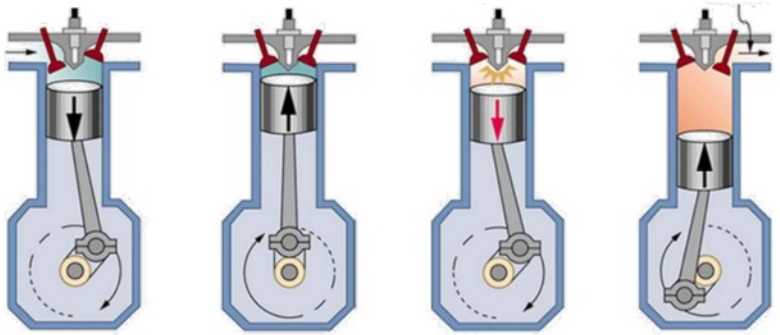

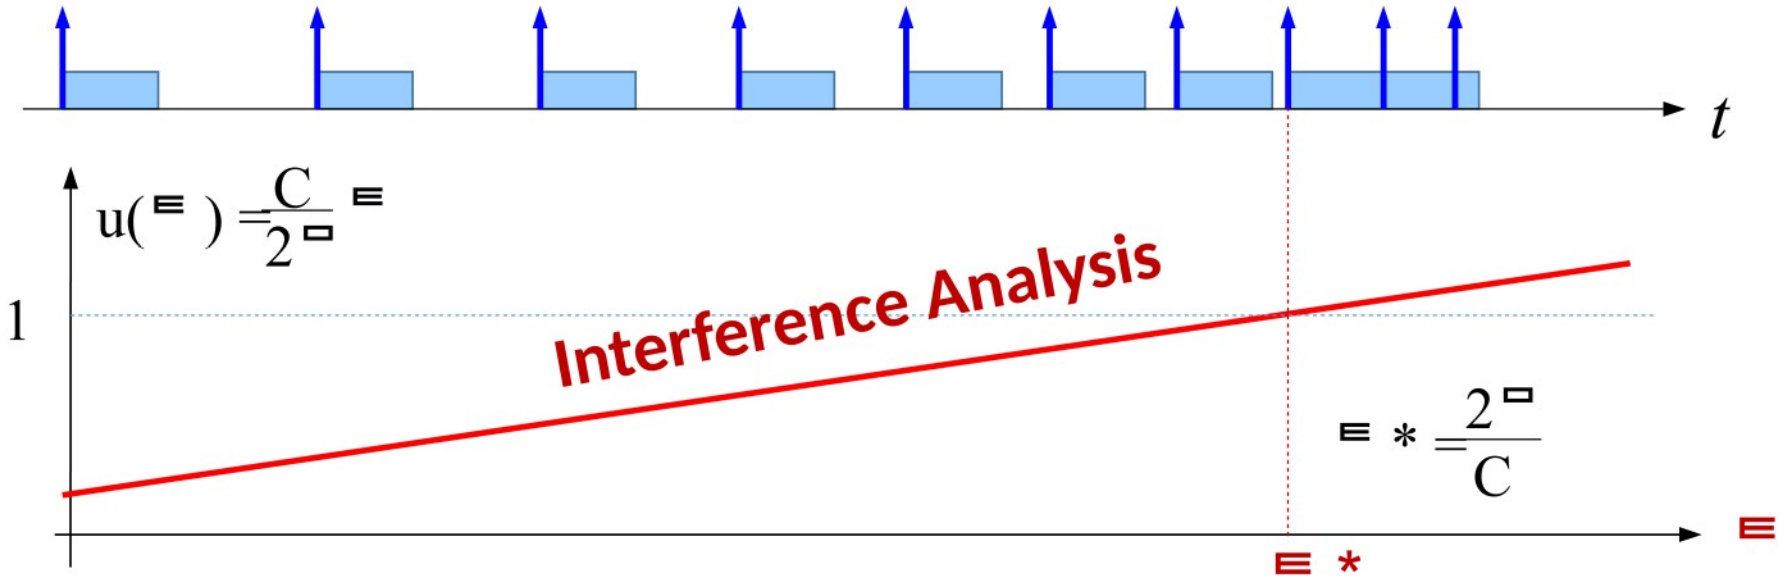

Real-Time support for Arduinc

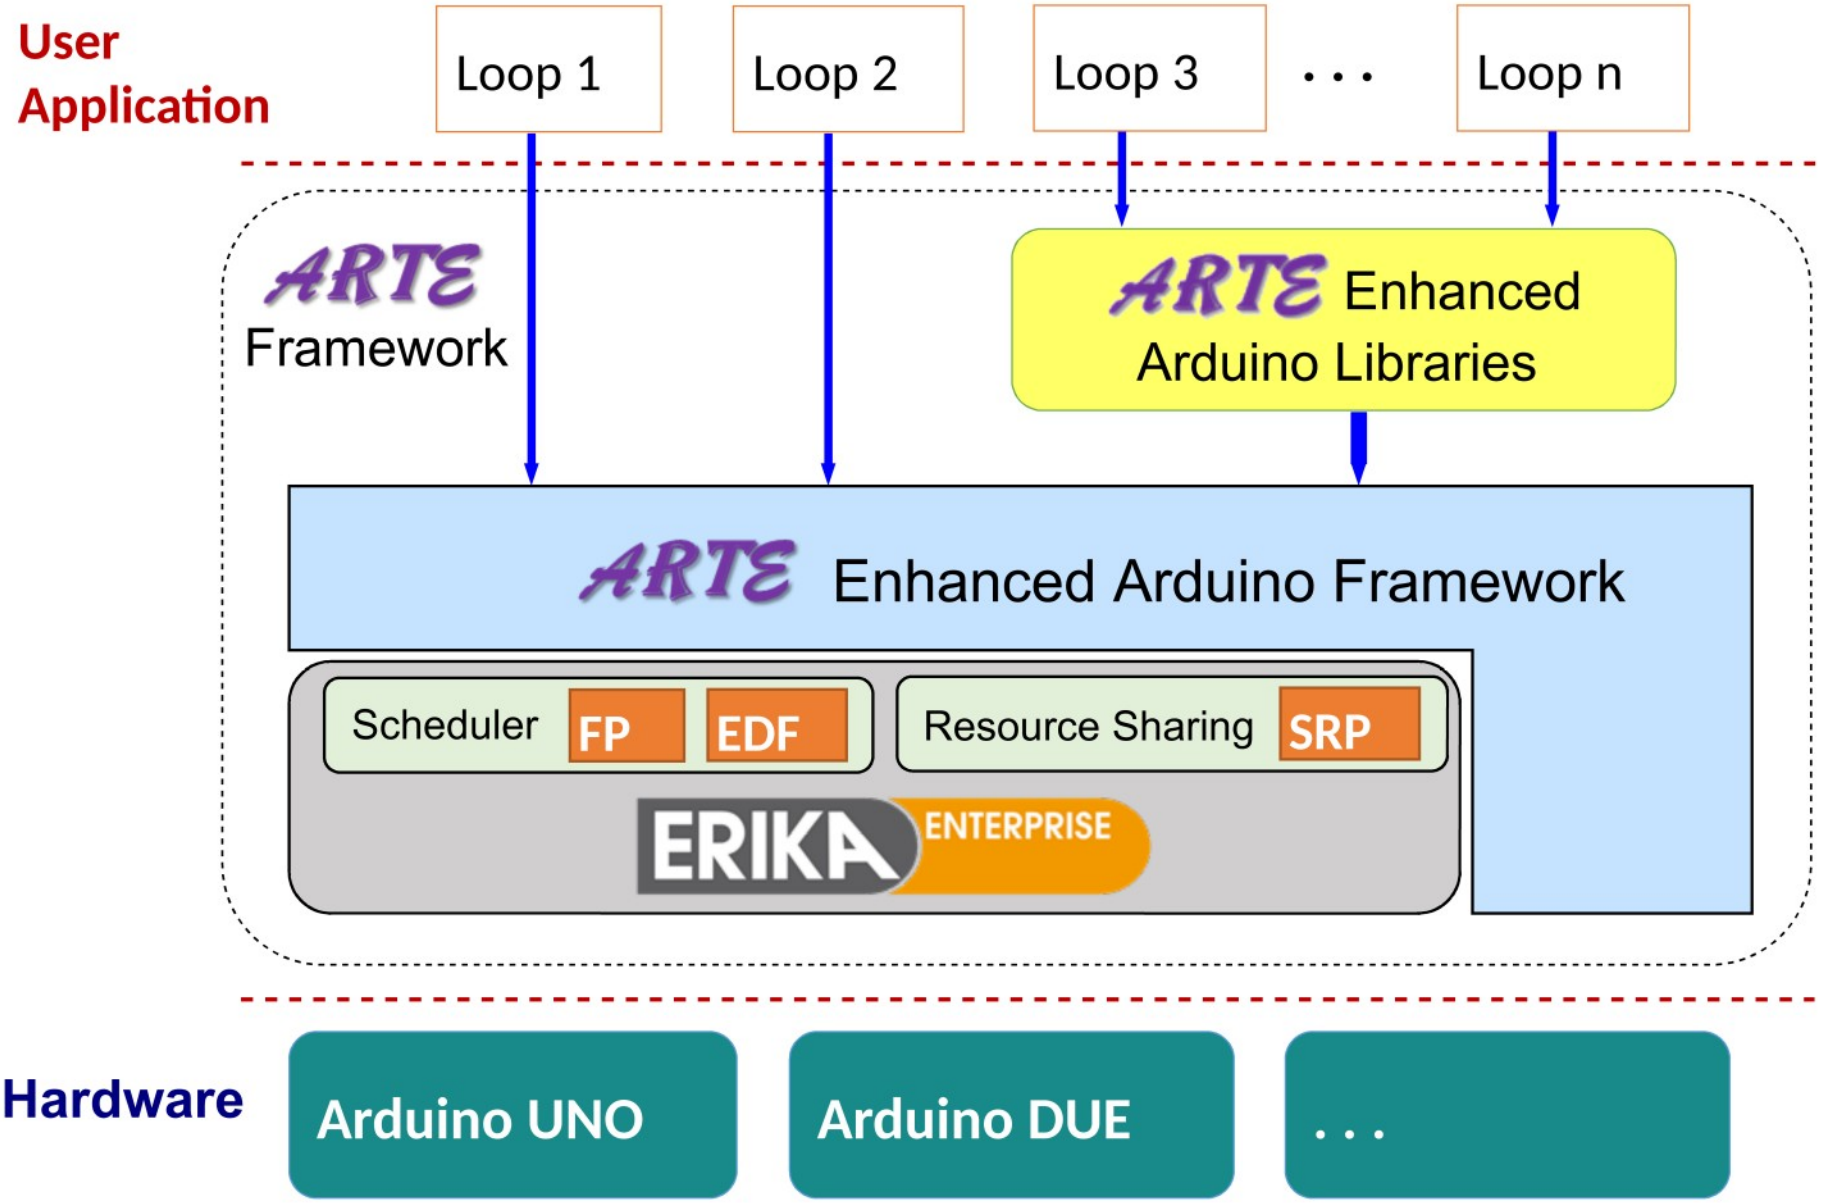

e-Health

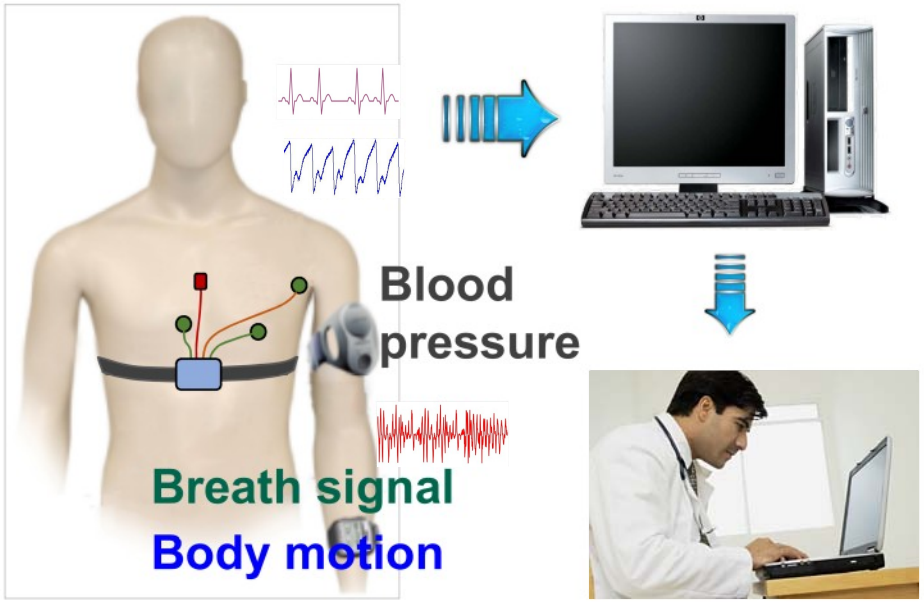

Remote monitoring

- Multisensor node (ECG, SPO2, IMU)
- **Cloud storage**
- Online **big-data analysis**

Ascolta regional project

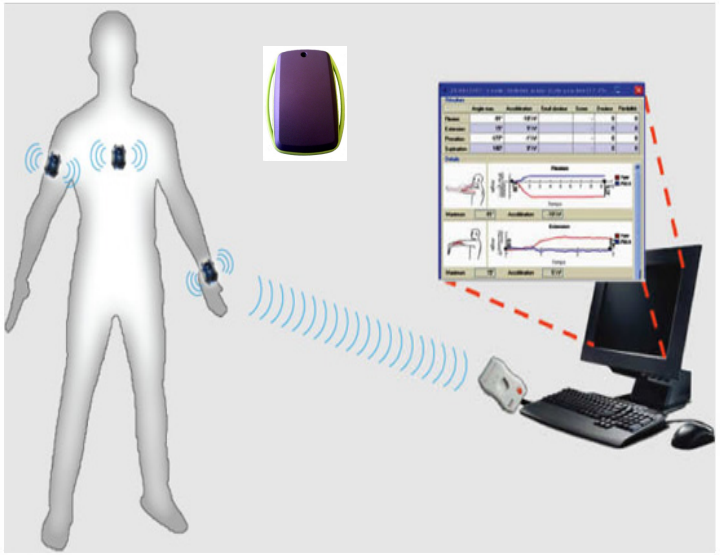

Tele-rehabilitation

- IMU-based body sensor network
- Data fusion
- Limb tracking
- Posture/action recognition

TELECOM ITALIA

Drones in Agriculture

Efficient monitoring of agriculture areas for automatic **vegetation control** and map reconstruction of geographical areas

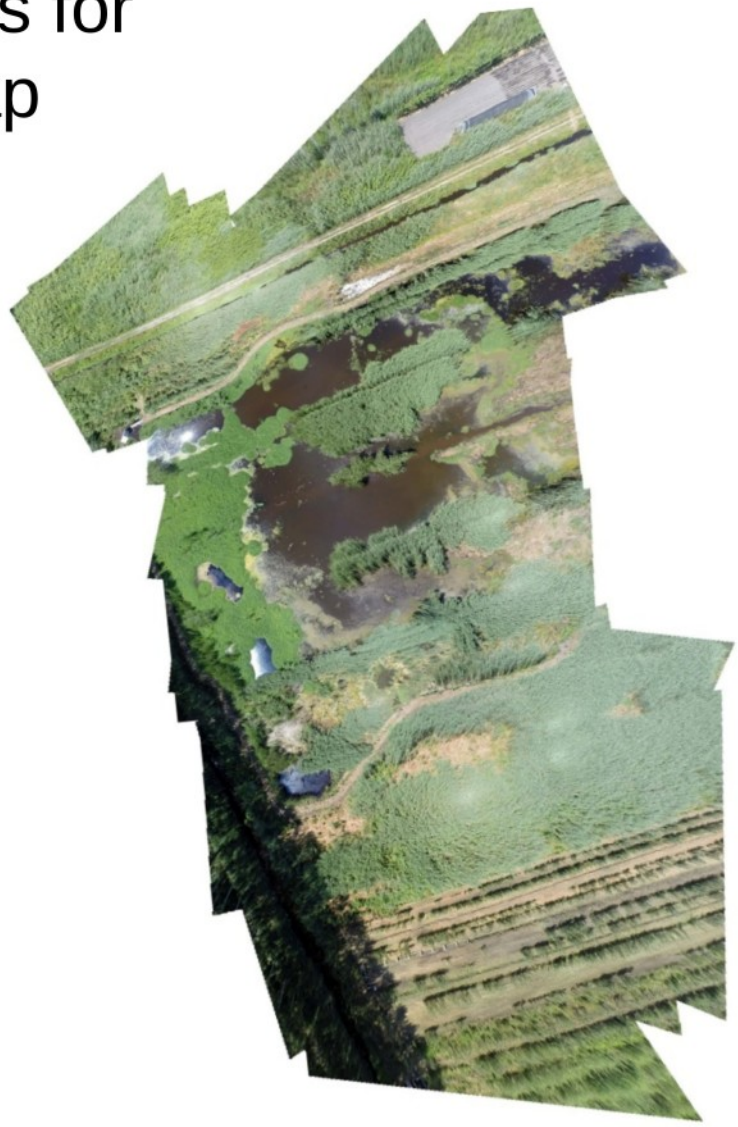

Energy-Aware path planning

energy model for main maneuvers

- constant speed, acceleration, deceleration
- hovering, climbing, …

Sharing FPGA by DPR

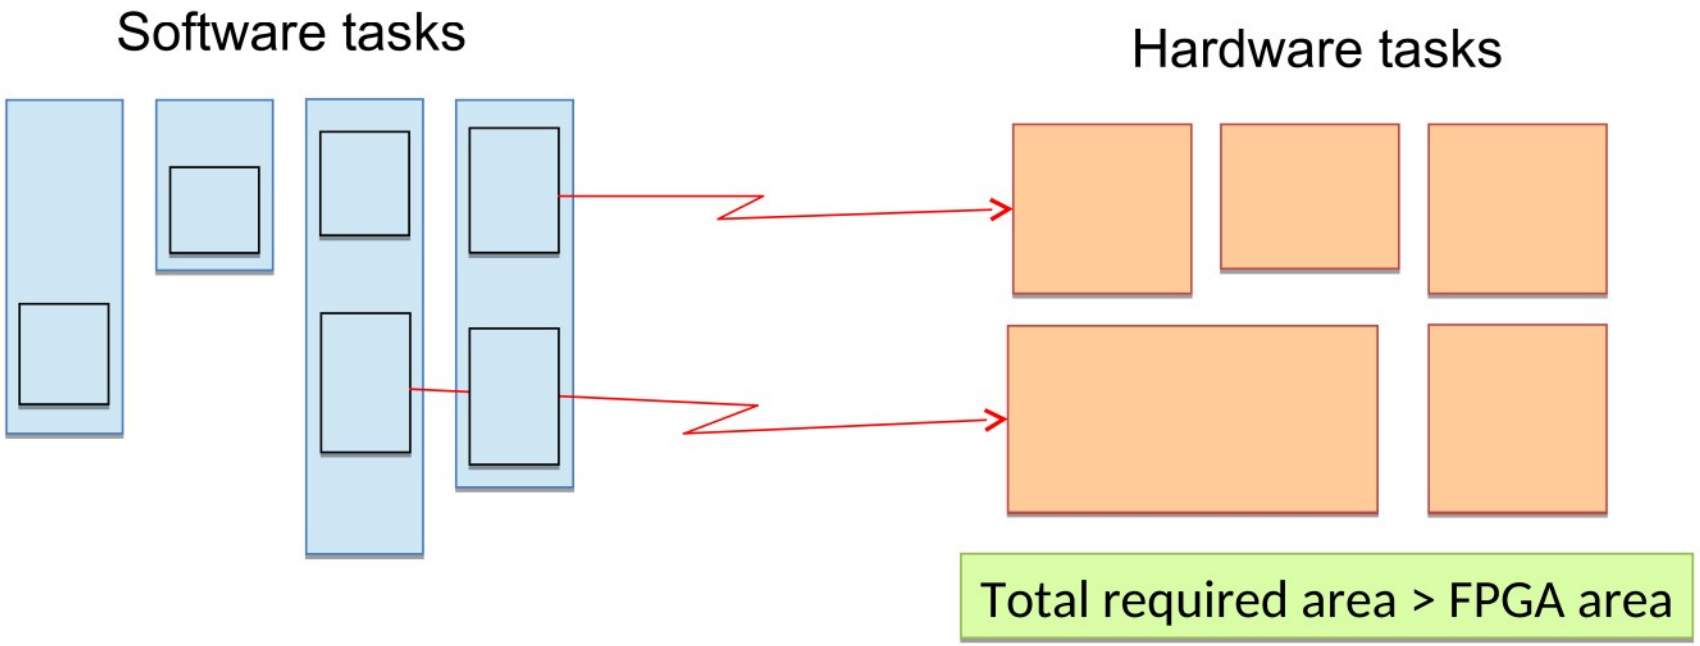

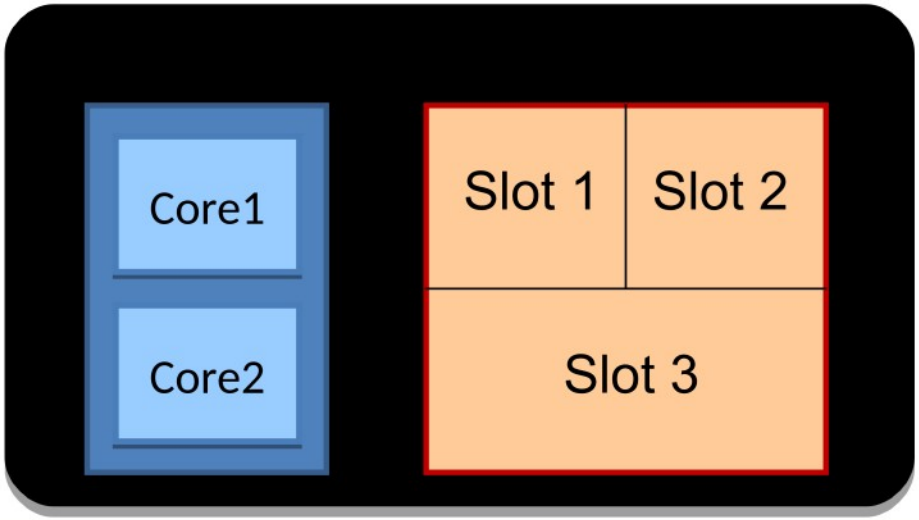

Sharing FPGA by DPR

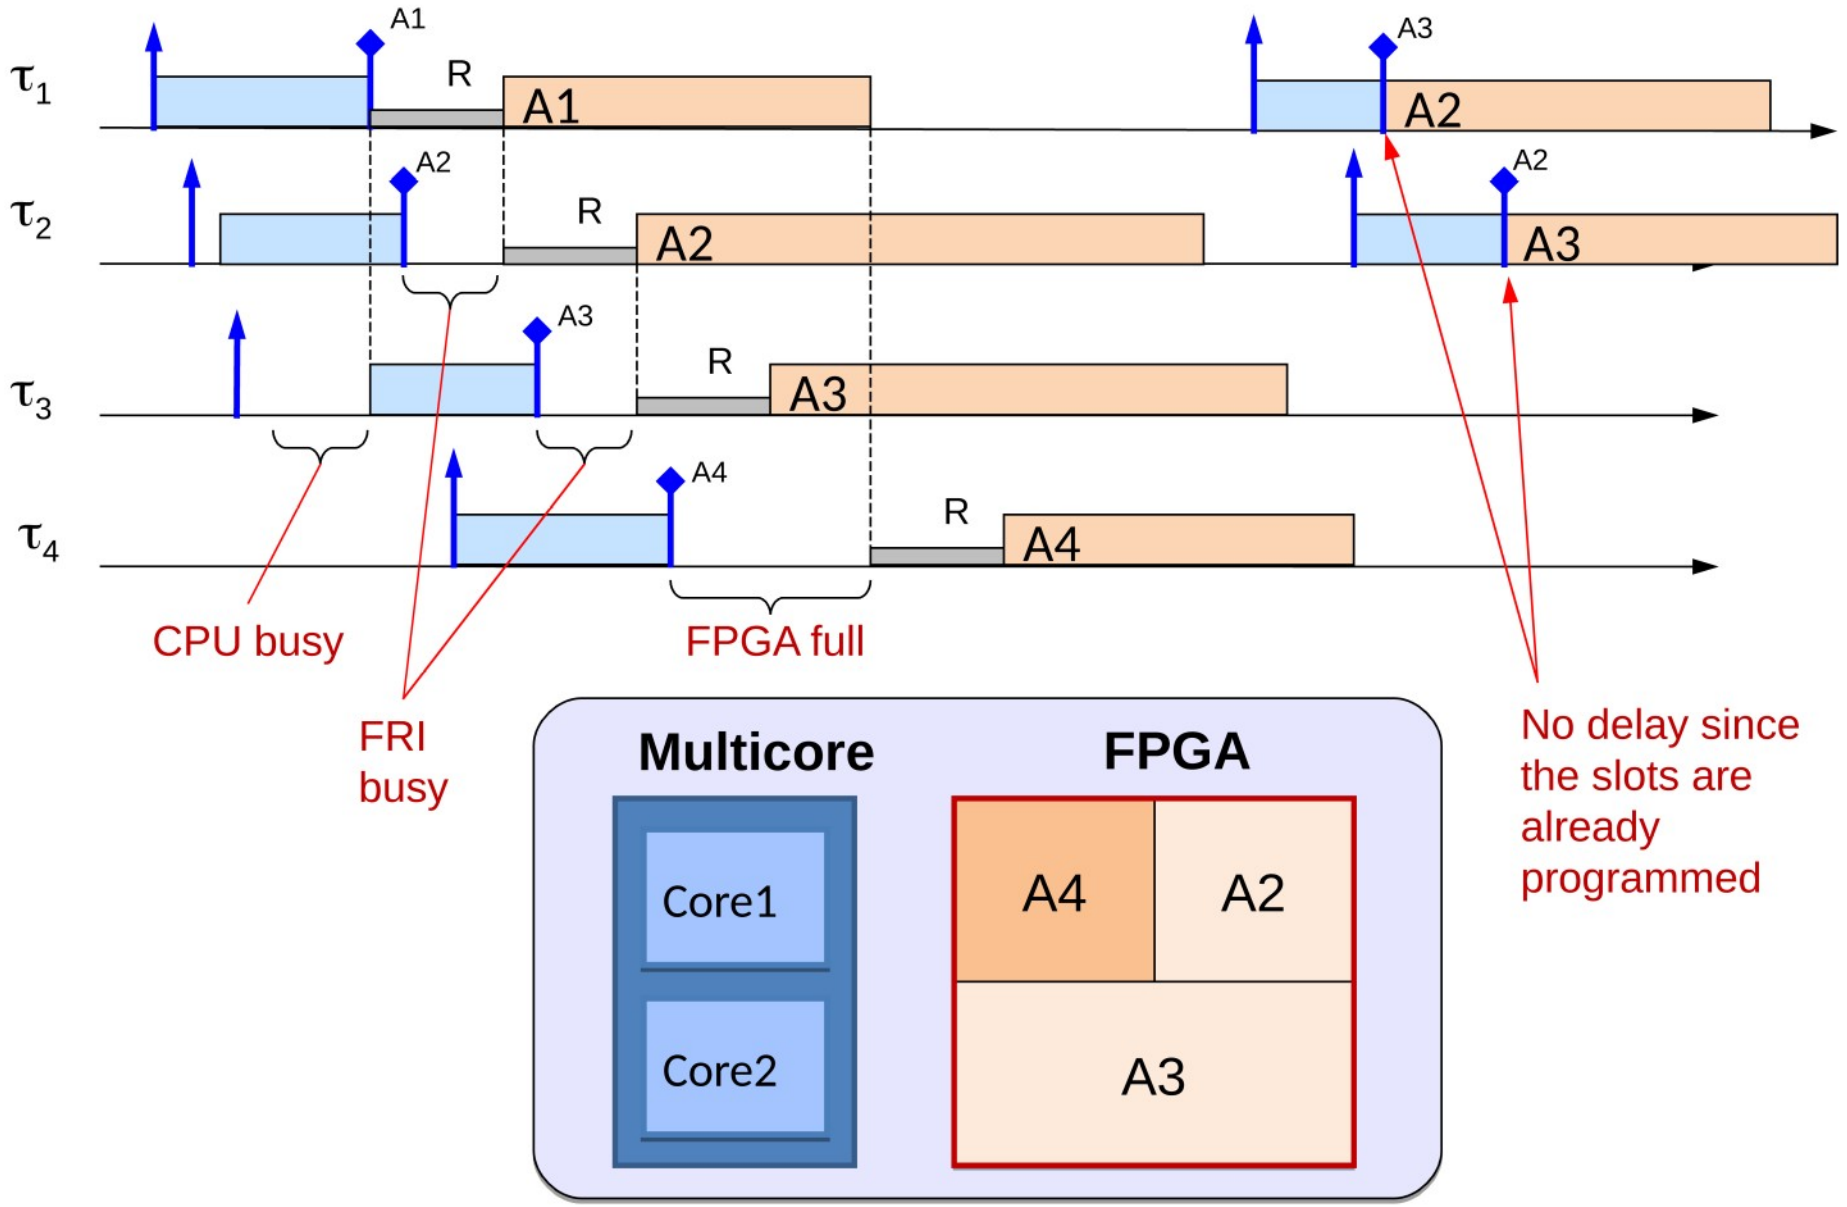

Cloud Computing Research

IRMOS Two-Phases Approach

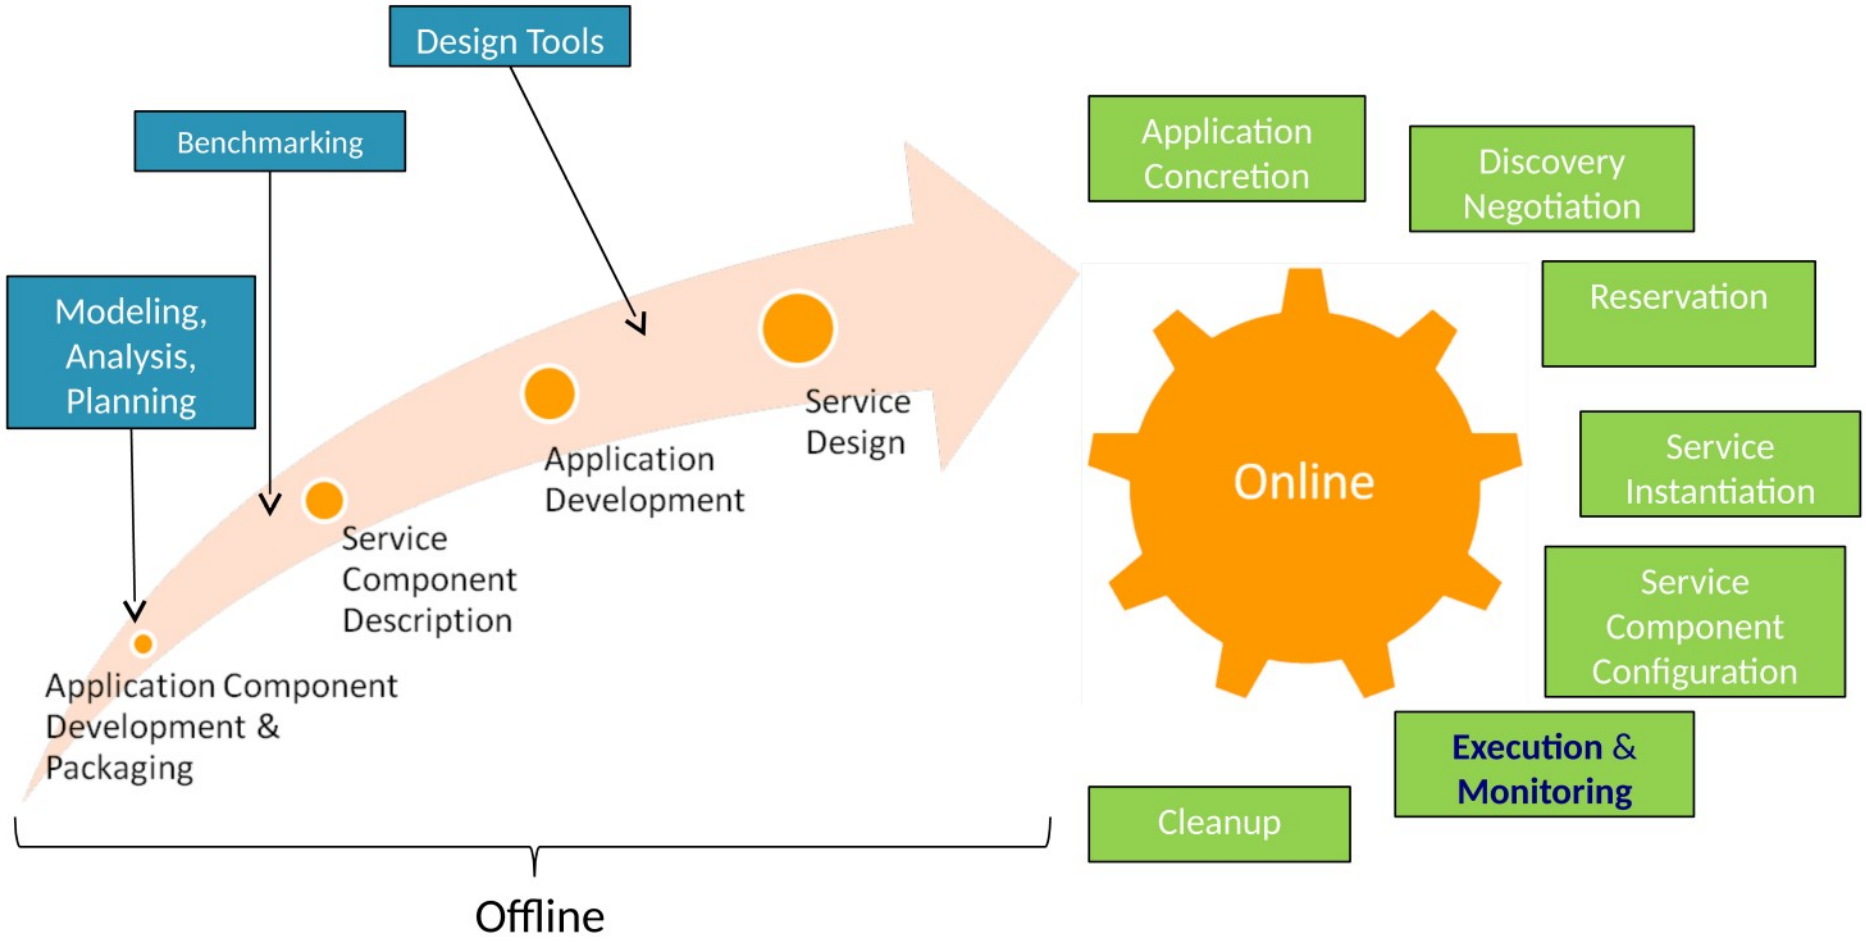

VM Placement – Distributed Cloud Services with e2e Constraints

Deployment of VSNs on PNs

- Given computing/network/memory requirements
- Respecting end-to-end timing constraints

IRMOS
Interactive Realtime Multimedia Applications on Service Oriented Infrastructures

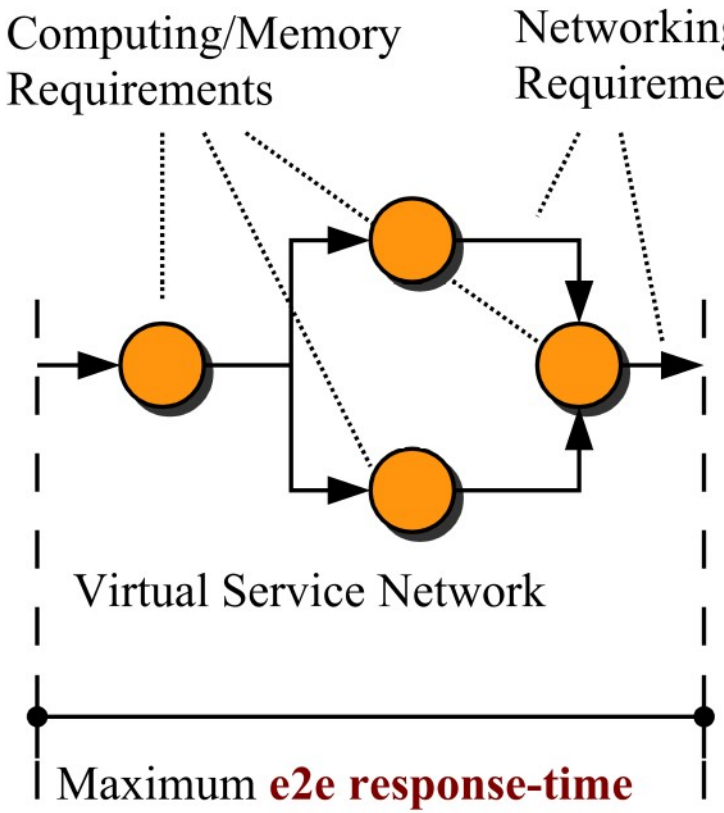

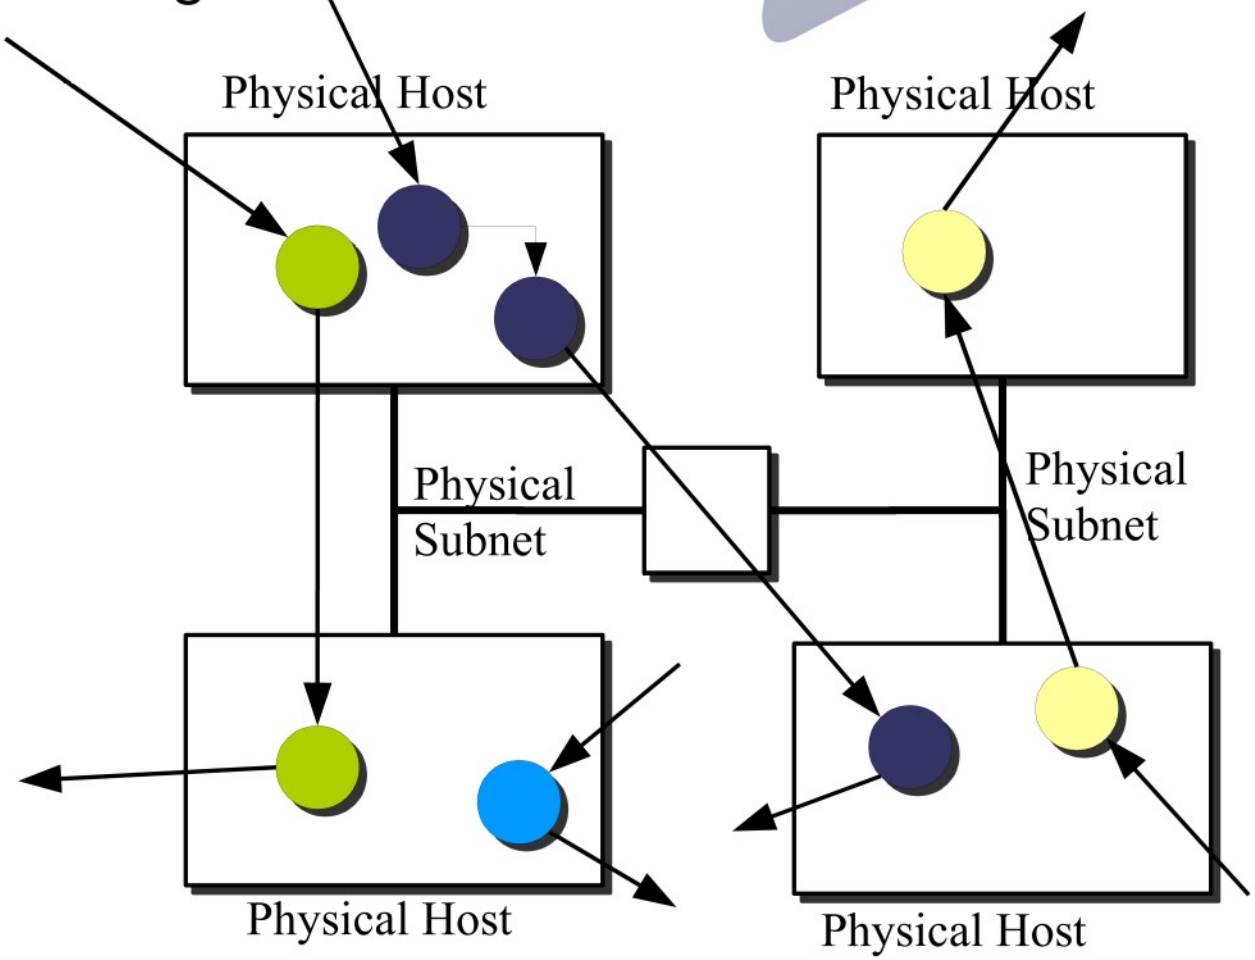

Probabilistic response-time guarantees

Tune allocation on computation-time **percentiles** (instead of WCET)

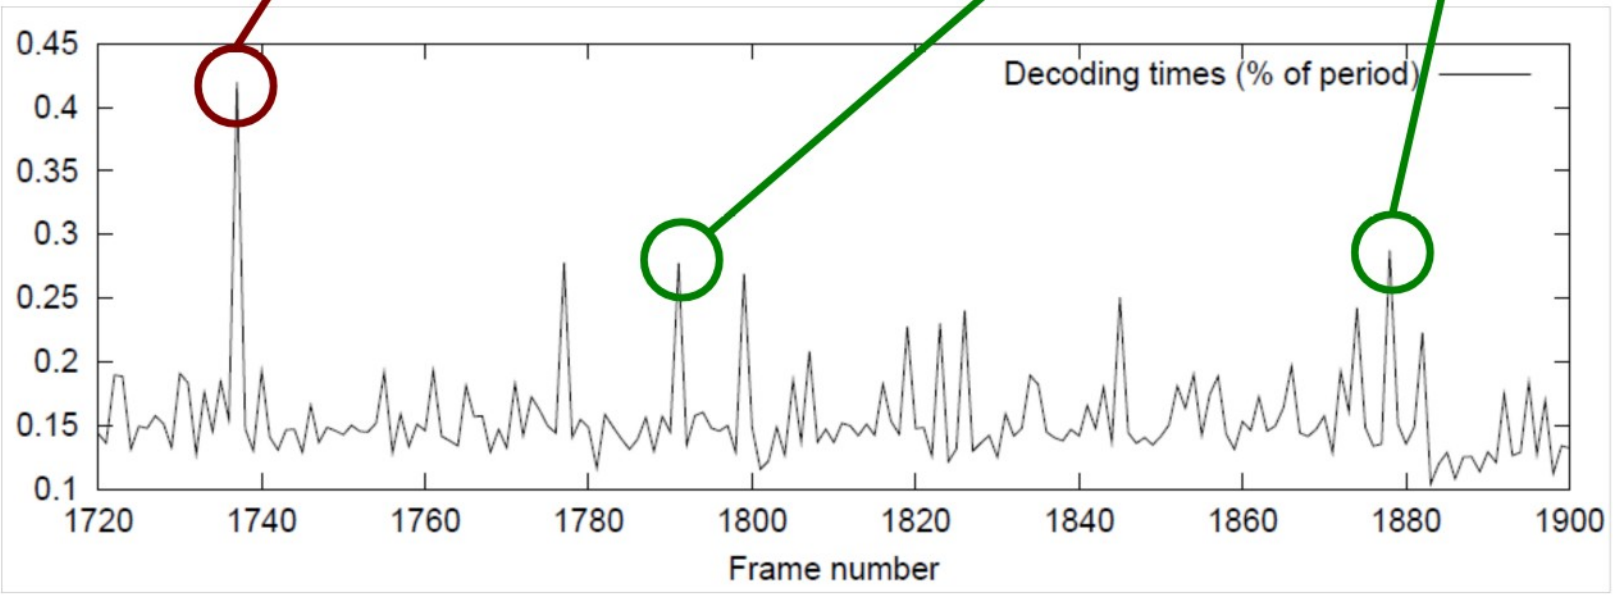

Probabilistic availability guarantees

Applications sharing the same PH

- may be independently activated

Over-subscription policy where provider relies on

- **probabilities of activation** for admitted & new service

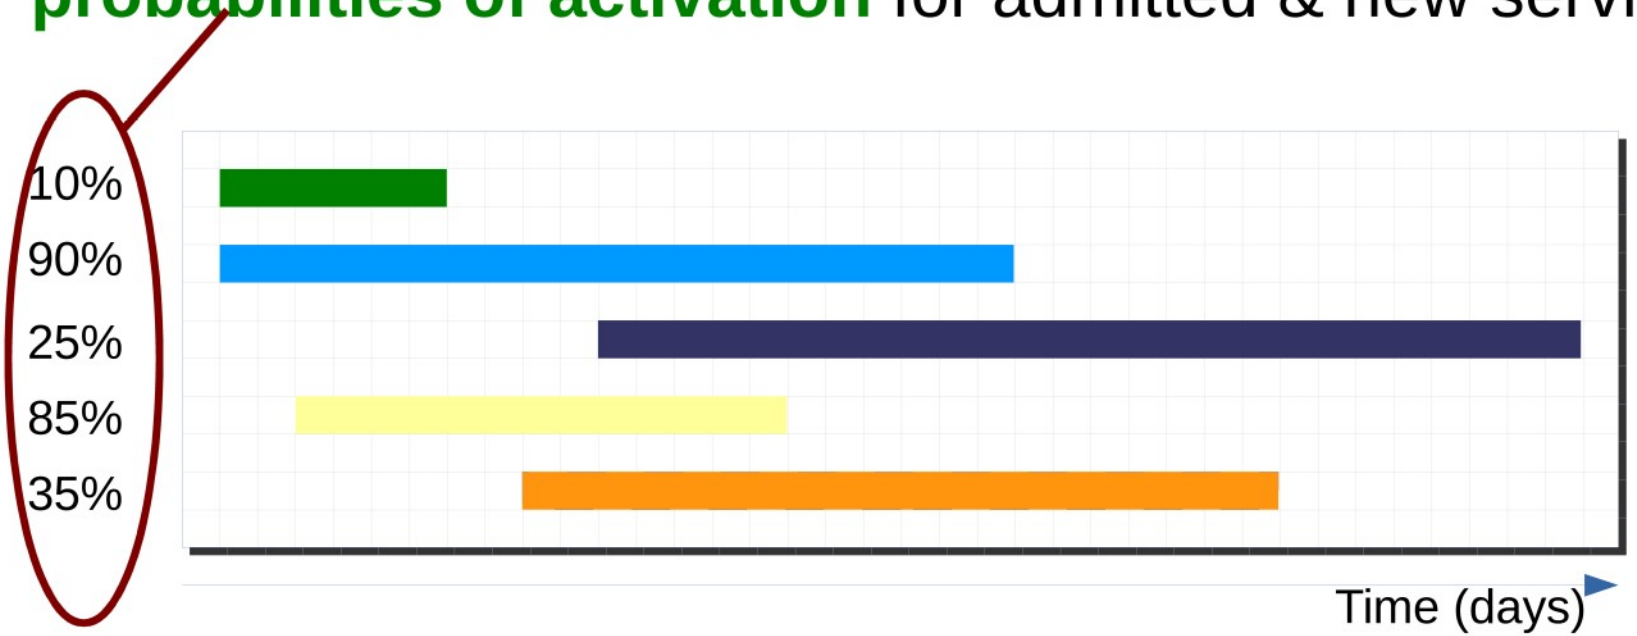

Real-Time Scheduling of VMs

Use of real-time CPU scheduling

- As **basic mechanism** for isolating VMs concurrently running on the same CPU and core

IRMOS Real-Time Scheduler

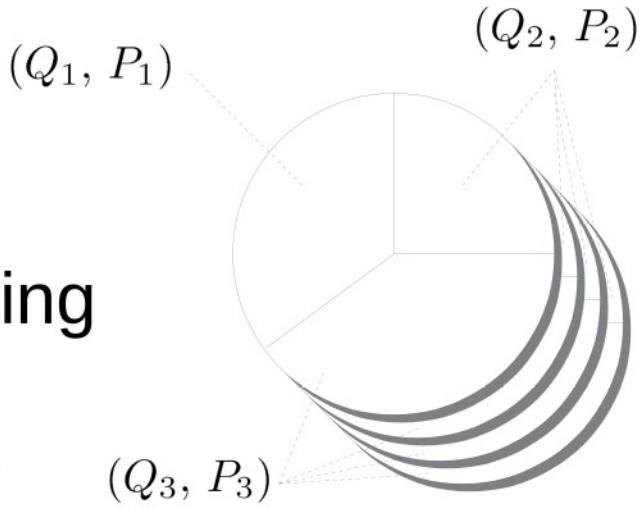

- For the Linux kernel
- Provides hierarchical **EDF/FP** scheduling
  - EDF-based **resource reservations**
    - (Q, P): a **budget** Q is granted every **period** P
    - Both a **guarantee** and a **limitation**
  - FP scheduling within each EDF reservation
- **CBS**-based: provides **temporal isolation**

Hierarchical CPU reservation

- One reservation for each thread
  - Too many parameters
  - No statistical multiplexing
- One reservation for a group of threads
  - Easier to set-up for the designer
  - Provide isolation **where is needed**
  - Take advantage of **intra-application variability**
  - Enables component-based analysis

End-to-end QoS brokering

T. Cucinotta, D. Lugones, D. Cherubini, K. Oberle

Bell Labs

Introduction

Cloud customers' requirements are evolving

- Not merely storage and best-effort services
- Towards enterprise services with stringent real-time and interactivity requirements
  - → Virtual desktop
  - → Professional on-line multimedia authoring
  - → On-line Gaming
  - → Virtualized Network Functions (NFV)
    - e.g., IMS, CDN, ...

Increasing interest in hybrid deployments

- Private cloud accompanied by federated public clouds for possible peak workload hand-over

Problem Presentation #1/3

Potential benefits impaired by

- best-effort Internet
- lack of end-to-end QoS

Requests cross various NSPs

- access network
- metro network
- core network

and different CSPs

- SaaS
- IaaS

still requiring tight interactivity

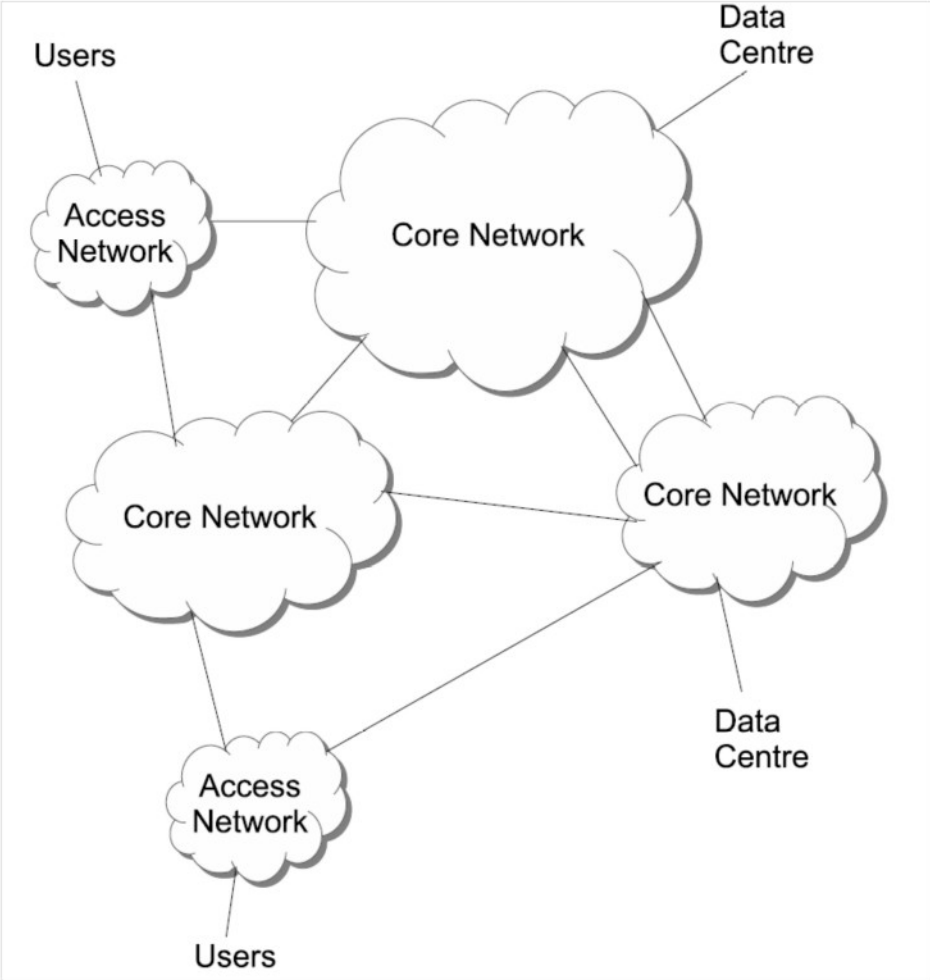

Problem Presentation #2/3

If multiple NSPs and CSPs provide various offerings with various QoS/cost trade-offs

- NSPs: from A to B with latency l1 costs c1, with latency l2 costs c2
- CSPs: VMs with different access latency and unit capacity cost

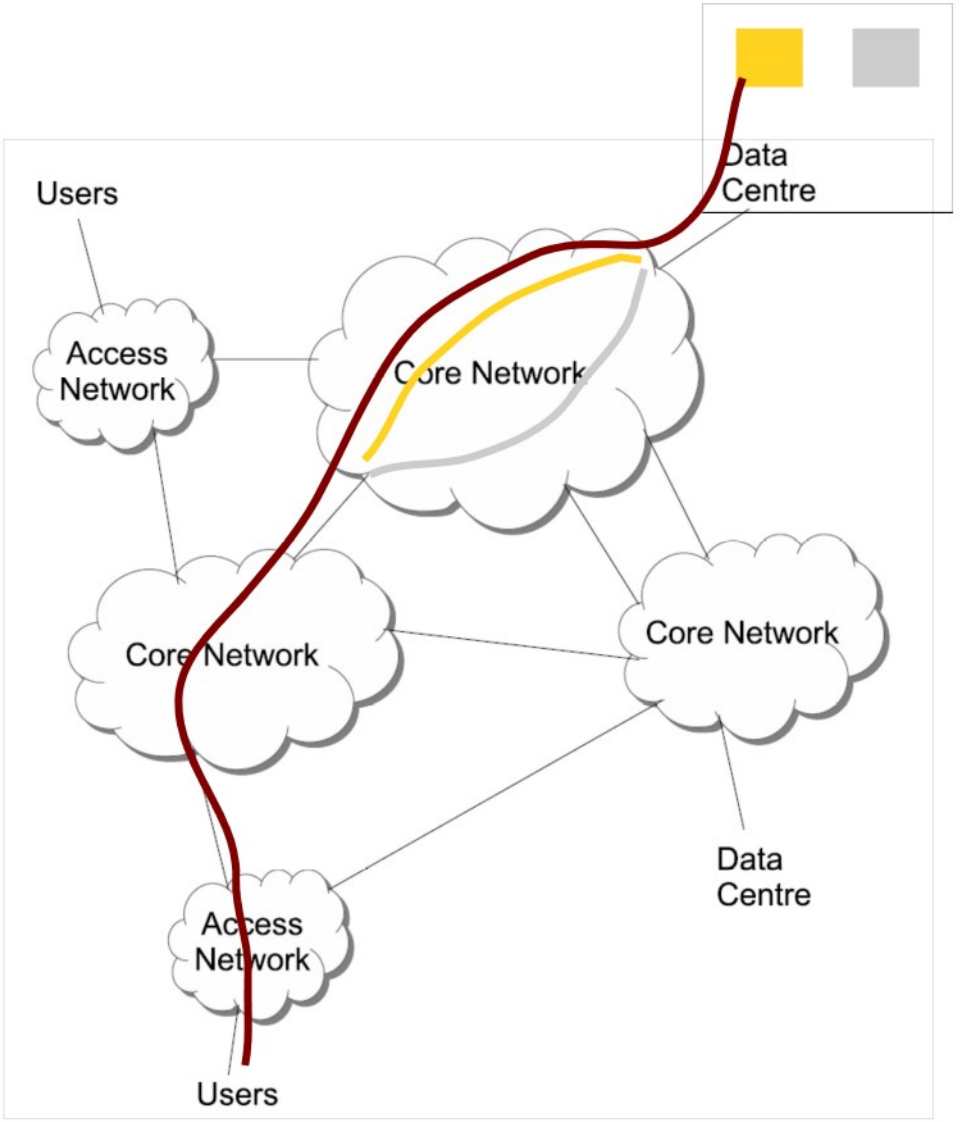

Problem Presentation #3/3

If multiple NSPs and CSPs provide various offerings with various QoS/cost trade-offs

- NSPs: from A to B with latency l1 costs c1, with latency l2 costs c2
- CSPs: VMs with different access latency and unit capacity cost

How to combine these in a brokering logic that meets end-to-end QoS constraints minimizing cost ?

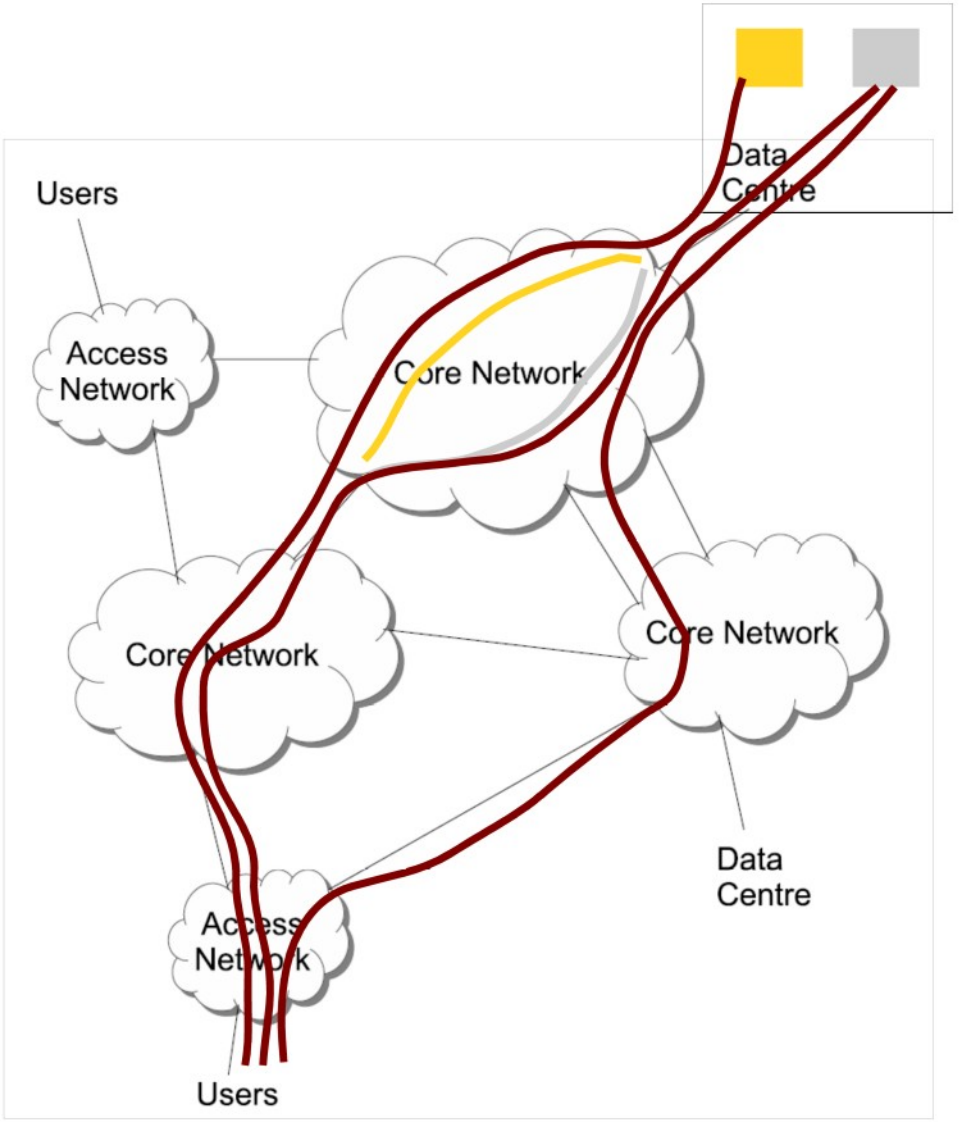

Data Centre Optimization Enhanced by SDN

T. Cucinotta, D. Lugones, D. Cherubini, E. Jul

Bell Labs

Problem Presentation

Increasing Computing/Communication demand of Cloud Applications

- Need for transferring massive amounts of data
  - → e.g., "big-data" workloads
- Predictable performance, real-time responsiveness, ultra low-latency

More network capacity and versatility needed in DCs

- Multi-path challenges traditional network equipment
- VLANs used for mitigating issues and isolating customers
  - → Limited availability => insufficient in multi-tenant environments

Need for flexible, dynamic and scalable resource management solutions

- Intelligent placement decisions considering computing/networking/storage requirements
- Flexible & elastic network management solutions => SDN, NFV.

Proposed Approach

- Max gain
- Min occupation (maximise saturation & sharing)
- Minimise max saturation (minimise temporal interference)

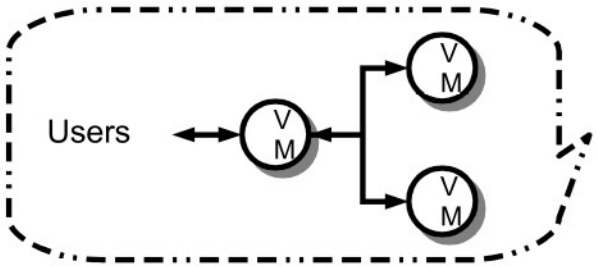

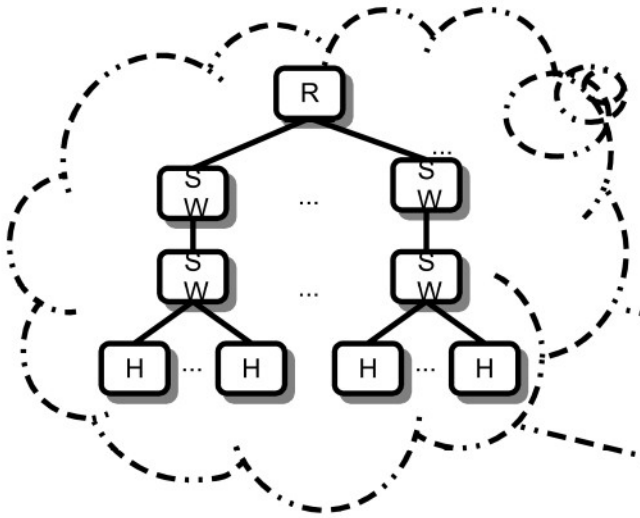

Application Models and Requirements

Physical Topology Technology Constraints

Optimization Logic (ILP Formulation)

Flow Placement

VM Placement

SDN Controller

Cloud Management VM instantiation

Evaluation Framework

CloudNetSim

- Co-simulation of network/CPU scheduling
- Implemented in OMNeT++
  - → Reuse of existing models (INET for TCP/IP)
- We added models for
  - → Hierarchical CPU scheduling
  - → Linux CFS, focusing on Linux + KVM
- Non-functional simulation
  - → We run apps models to mimic realistic behavior
  - → We gather end-to-end performance metrics

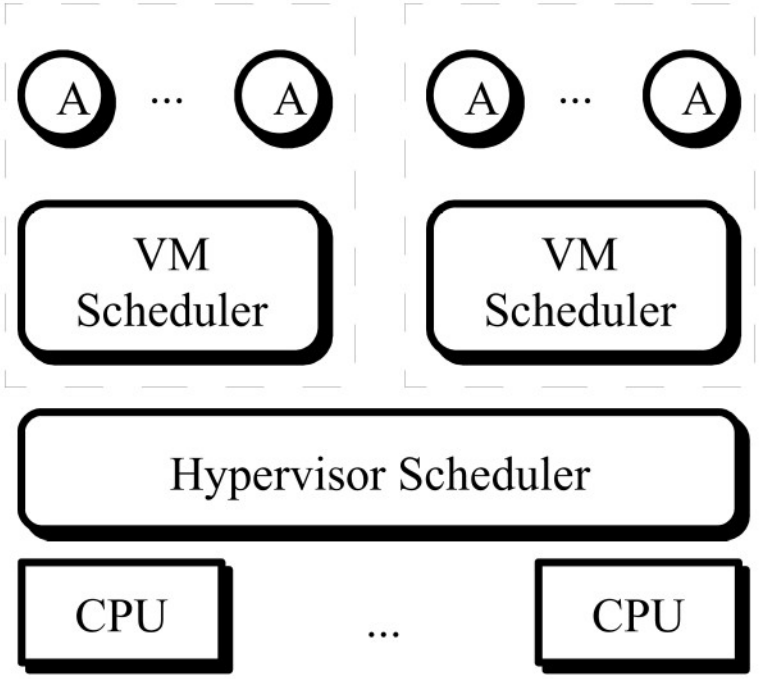

Simple "show-case" scenario

- Simple application models
  - → Randomly generated parameters
- Simple physical topology

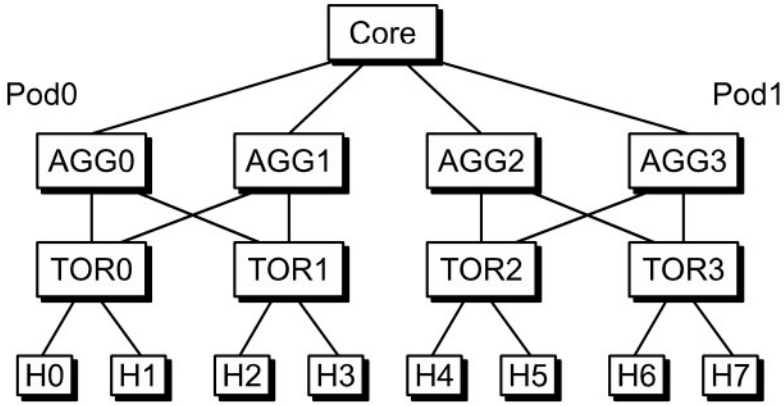

Simulation Results

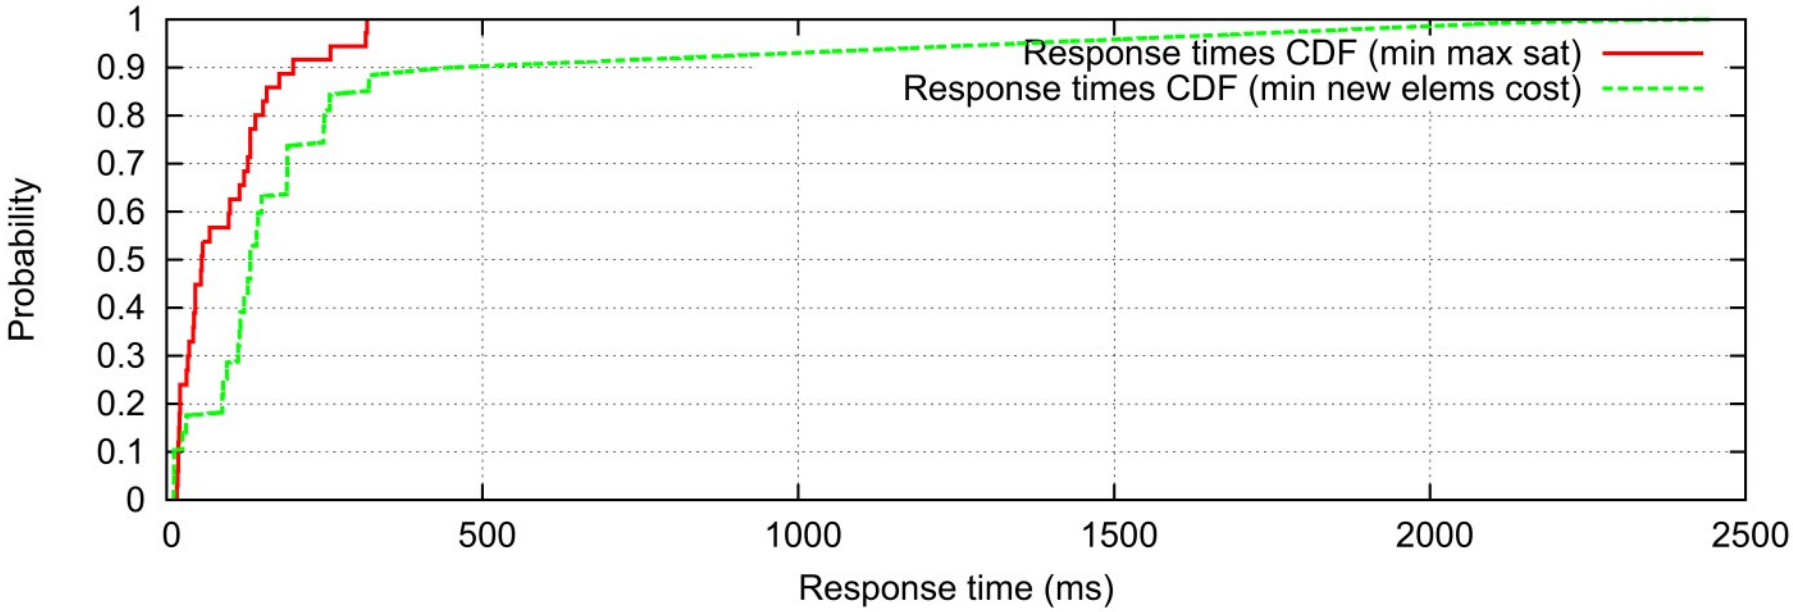

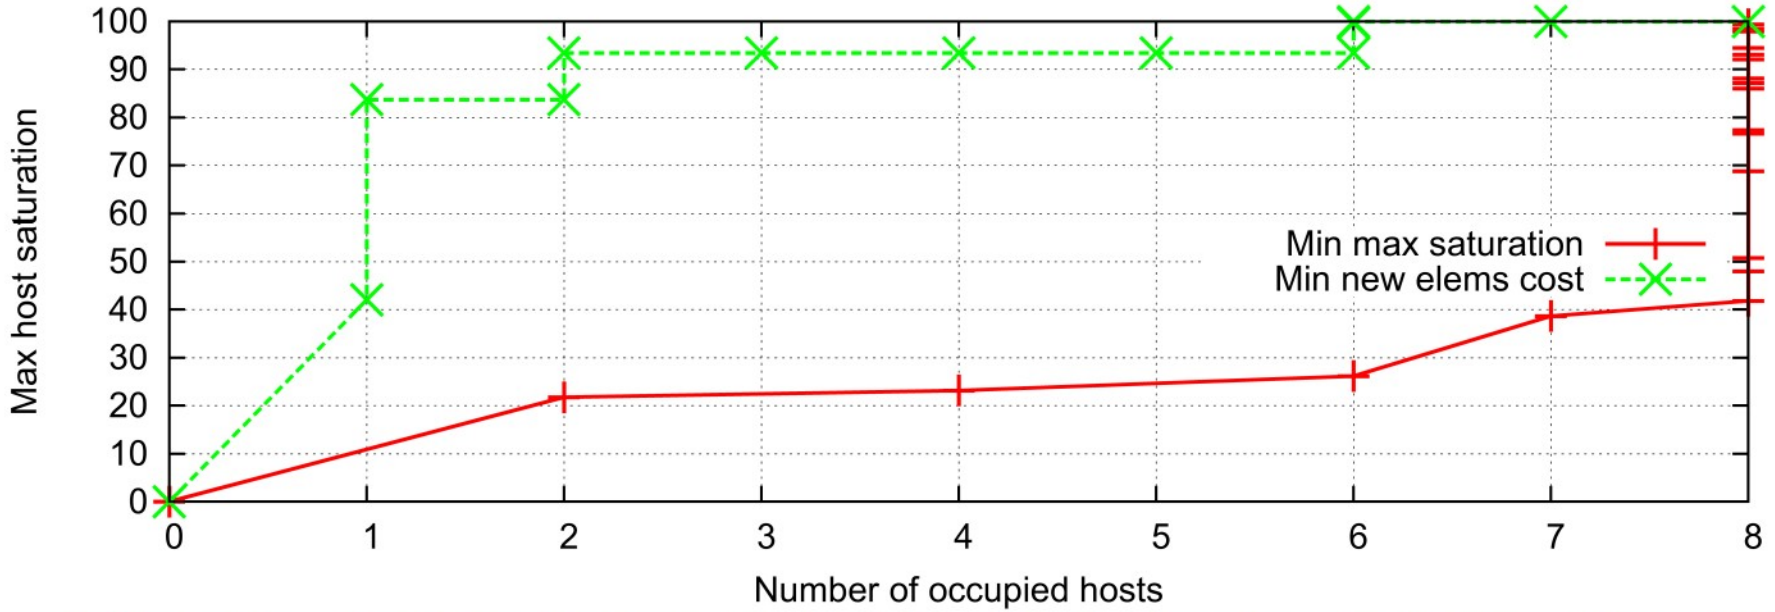

Controlling Elastic Virtualized Applications

T. Cucinotta, K. Oberle, M. Stein, P. Domschitz, S. Mullender

Bell Labs

Plethora of Cloud Providers Tools and Frameworks

Cloud IaaS

- Amazon EC2, Rackspace, Google Compute, ...
- OpenNebula, OpenStack, CloudStack
- CloudBand, ...

Configuration Management

Monitoring and Orchestration

- AWS AutoScaling, Heat+Ceilometer, Cloudify, CloudFoundry, Chef Recipes, ...

Elasticity Loop

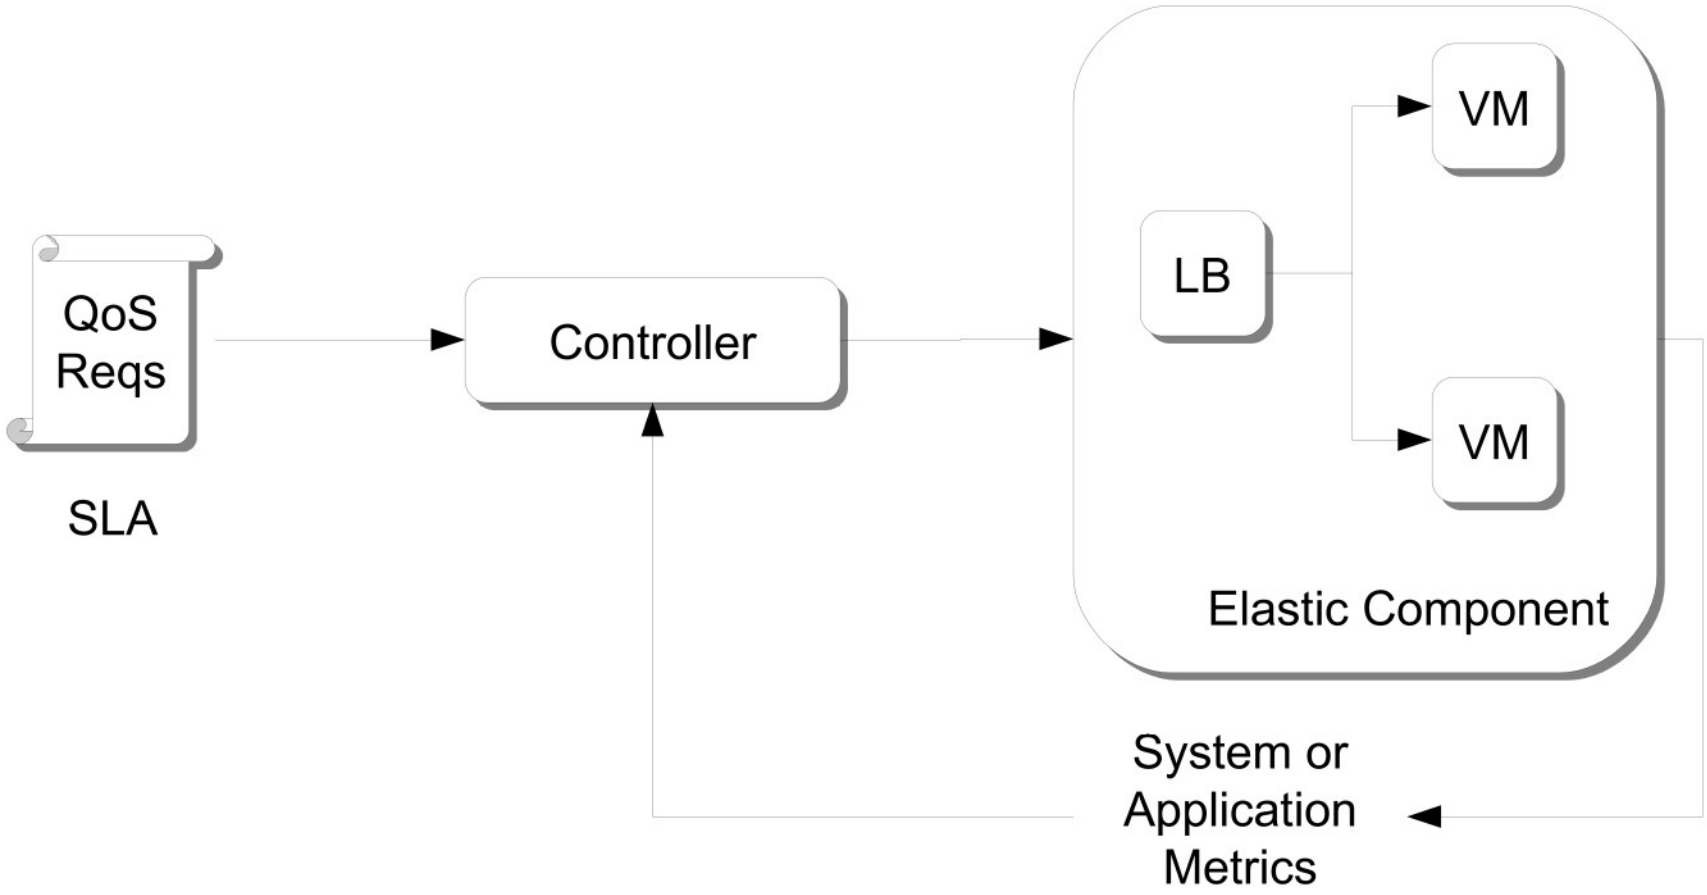

Elastic IMS
experimental set-up

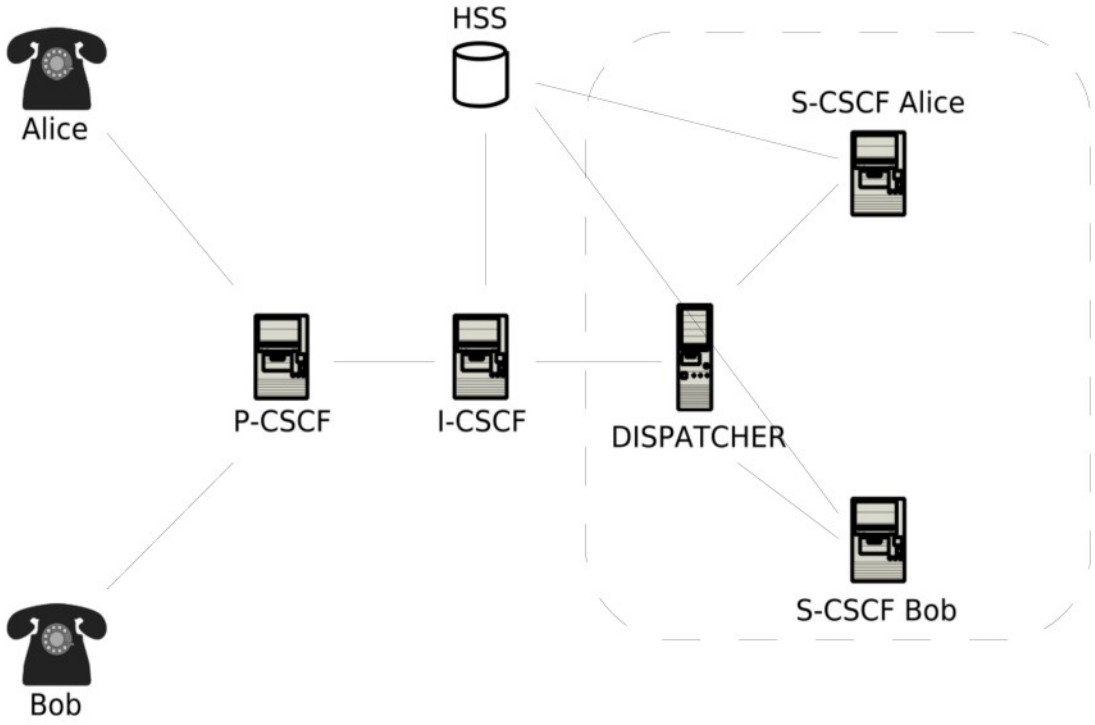

Elastic IMS
experimental set-up

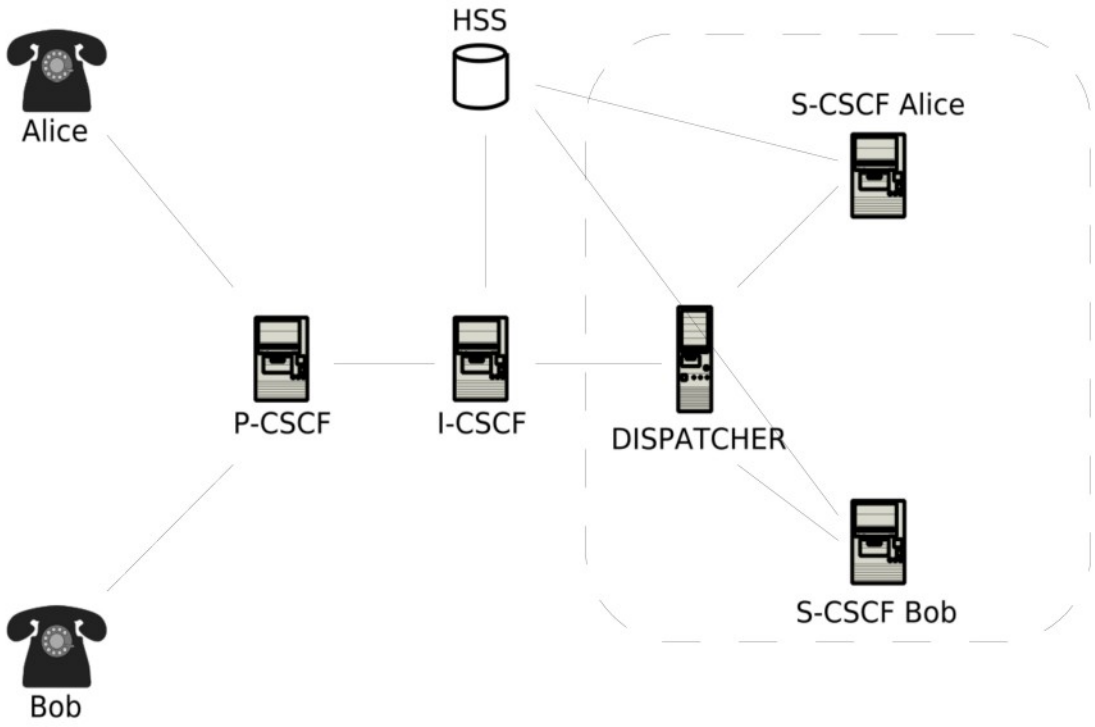

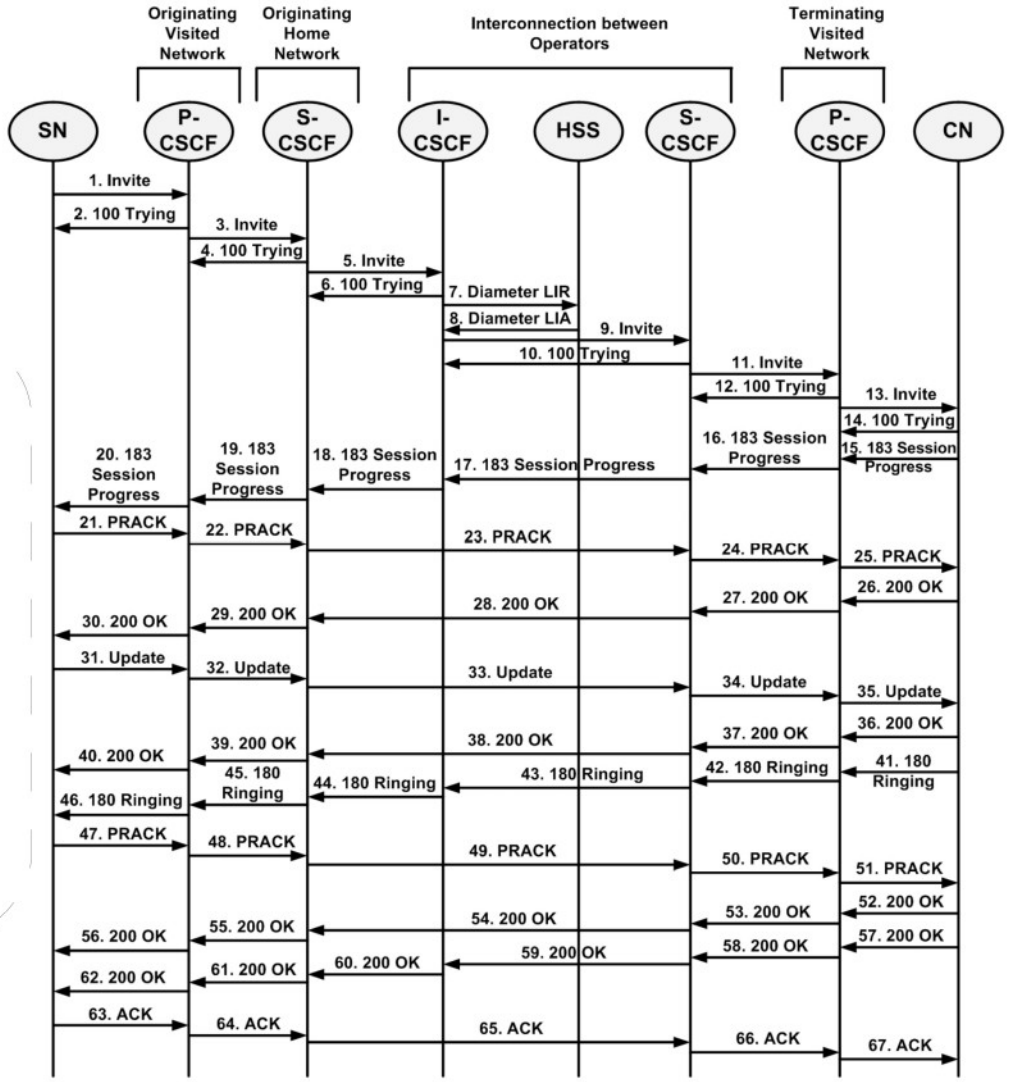

SN: Source Node
CN: Correspondent Node
HSS: Home Subscriber Server
Diameter LIR: Diameter Location-Information-Request
Diameter LIA: Diameter Location-Information-Answer
PRACK: Provisional Response Acknowledgment
P/I/S-CSCF: Proxy/Interrogating/Serving-Call Session Control Function

But...

Adaptation logic built on unstable terrain!

But...

Adaptation logic built on unstable terrain!

Can we build anything better?

Osprey

Osprey is an OS intended for Cloud Computing, HPC, servers and embedded systems

- **Reliable** and available
  - → Network interruptions do not cause sessions to fail
- **Real time support**
  - → Useful for multimedia processing
  - → Heavily based on **asynchronous operations** among components
- Multi-core performance and **energy aware**
  - → Heterogeneous cores supported
  - → Switch off unused cores for greener footprint
- **Per-process kernel state is minimized**
  - → Library OS approach: OS state attached to user-space applications
  - → Smaller kernel, simple process migration

MediaCloud

Context / background

- Increase of media traffic: ~40% per year
- Not only media consumption, also production and editing

MediaCloud is a dataflow-oriented multimedia processing framework

- **Lightweight container** architecture
- **Decentralized processing** of media
- Highly flexible and dynamic:
  - → stateless relocatable data flows
  - → **Instantiation and relocation in 2-3 ms**
- Enhanced performance and predictability w.r.t. standard virtualised media processing

Optimization of a Parallel Streaming Engine for Telco Applications

B. Theeten, I. Bedini, P. Cogan, A. Sala, **T. Cucinotta**

Bell Labs

Performance Model of a Telco Storm Application

Telelcom Application

- CDR analysis domain

  Spout → Bolt A → Bolt B

Model encompassing

- Storm topology configuration
  - → no. nodes, spouts, bolts, workers
- Resource consumption
  - → CPU%, network bandwidth
- Obtained throughput
- Obtained latency

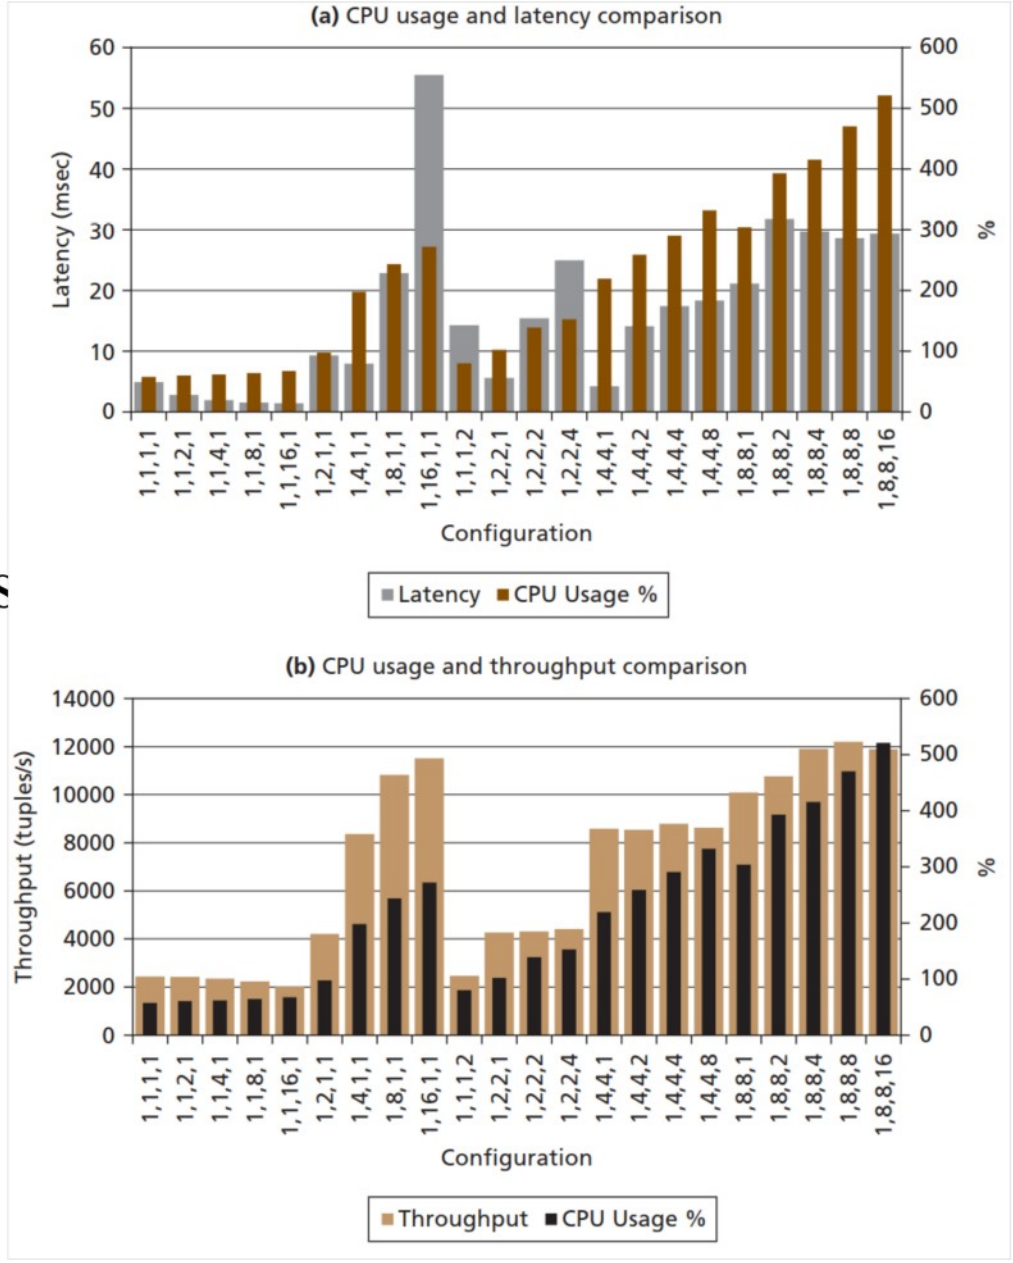

Temporal Isolation Among LTE/5G Network Functions by Real-time Scheduling

T. Cucinotta, M. Marinoni, A. Melani, A. Parri, C. Vitucci

ERICSSON

RT Packet Processing in Mobile Networks

Need for controlling temporal interference among VNFs

- use of kernel/hypervisor scheduler → (Q, P)
- MoC under specific assumptions reduced to standard Markov Chains

Modelled behavior

- Input traffic
- Scheduling parameters (Q, P)
- Obtained throughput and latency percentiles

Example results

- Min. Q/P for SLO: $\frac{Q}{P} \geq \frac{1}{\mu}\left[\lambda_i - \frac{ln(1-\phi)}{R^*}\right]$
- Max. input rate: $\lambda_G = n\lambda_i \leq n\left[\frac{Q}{P}\mu + \frac{\ln(1-\phi)}{R^*}\right]$

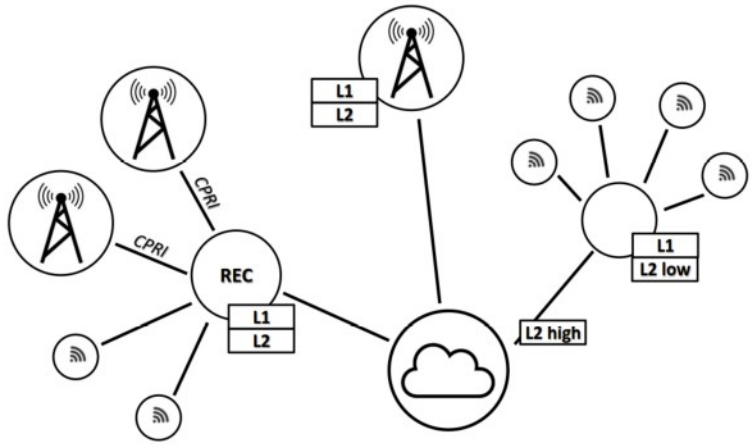

Figure 1: Possible RAN/VRAN deployment models.

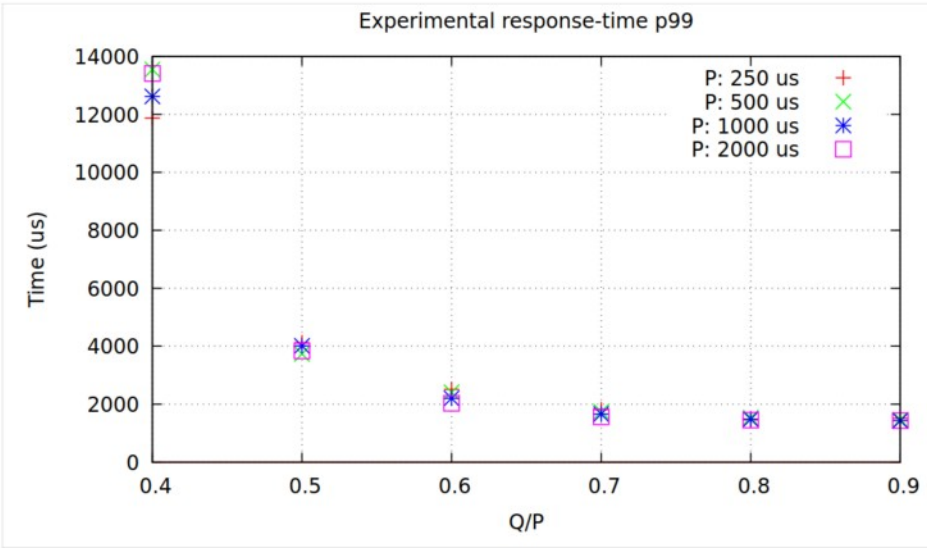

SDN-Controlled Energy-Efficient Mobile Fronthaul: An Experimental Evaluation in Federated Testbeds

L. Valcarenghi(a), K. Kondepu(a), A. Sgambelluri(a), F. Cugini(b), P. Castoldi(a), R. Aparicio Morenilla(c), D. Larrabeiti(c), B. Vermeulen(d)

(a)Scuola Superiore Sant'Anna, Pisa, Italy
(b)CNIT, Pisa, Italy
(c)Universidad Carlos III, Madrid, Spaim
(d)iMinds, Ghent, Belgium

EuCNC 2016
June 27-30, 2016. Athens, Greece

Experiment: setup before reconfiguration

- Two User Equipments (UEs) are connected to antenna sites (C1 a
- The antenna sites are connected to two servers (NF1 and NF2) hosting network functions (NF) (e.g., Radio Resource Control) through VLANs

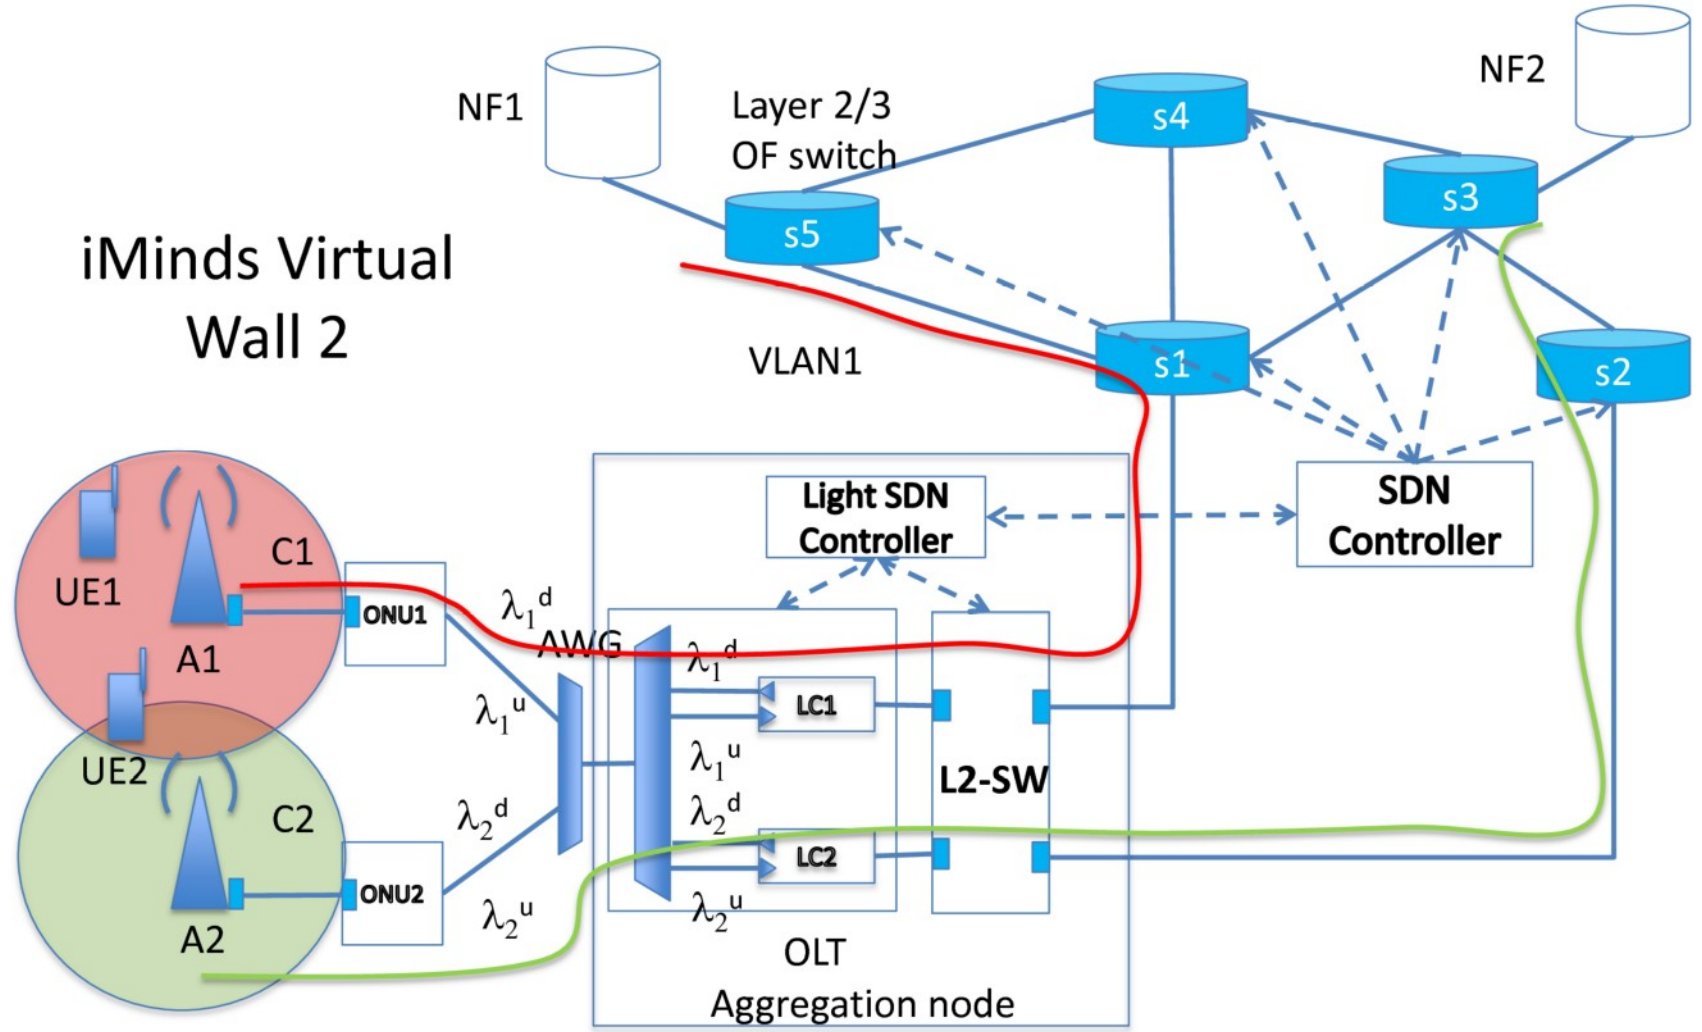

Experiment: setup after reconfiguration

- C2 is turned off for energy efficiency purposes
- VLAN between A2 and NF2 is reconfigured to reach NF2

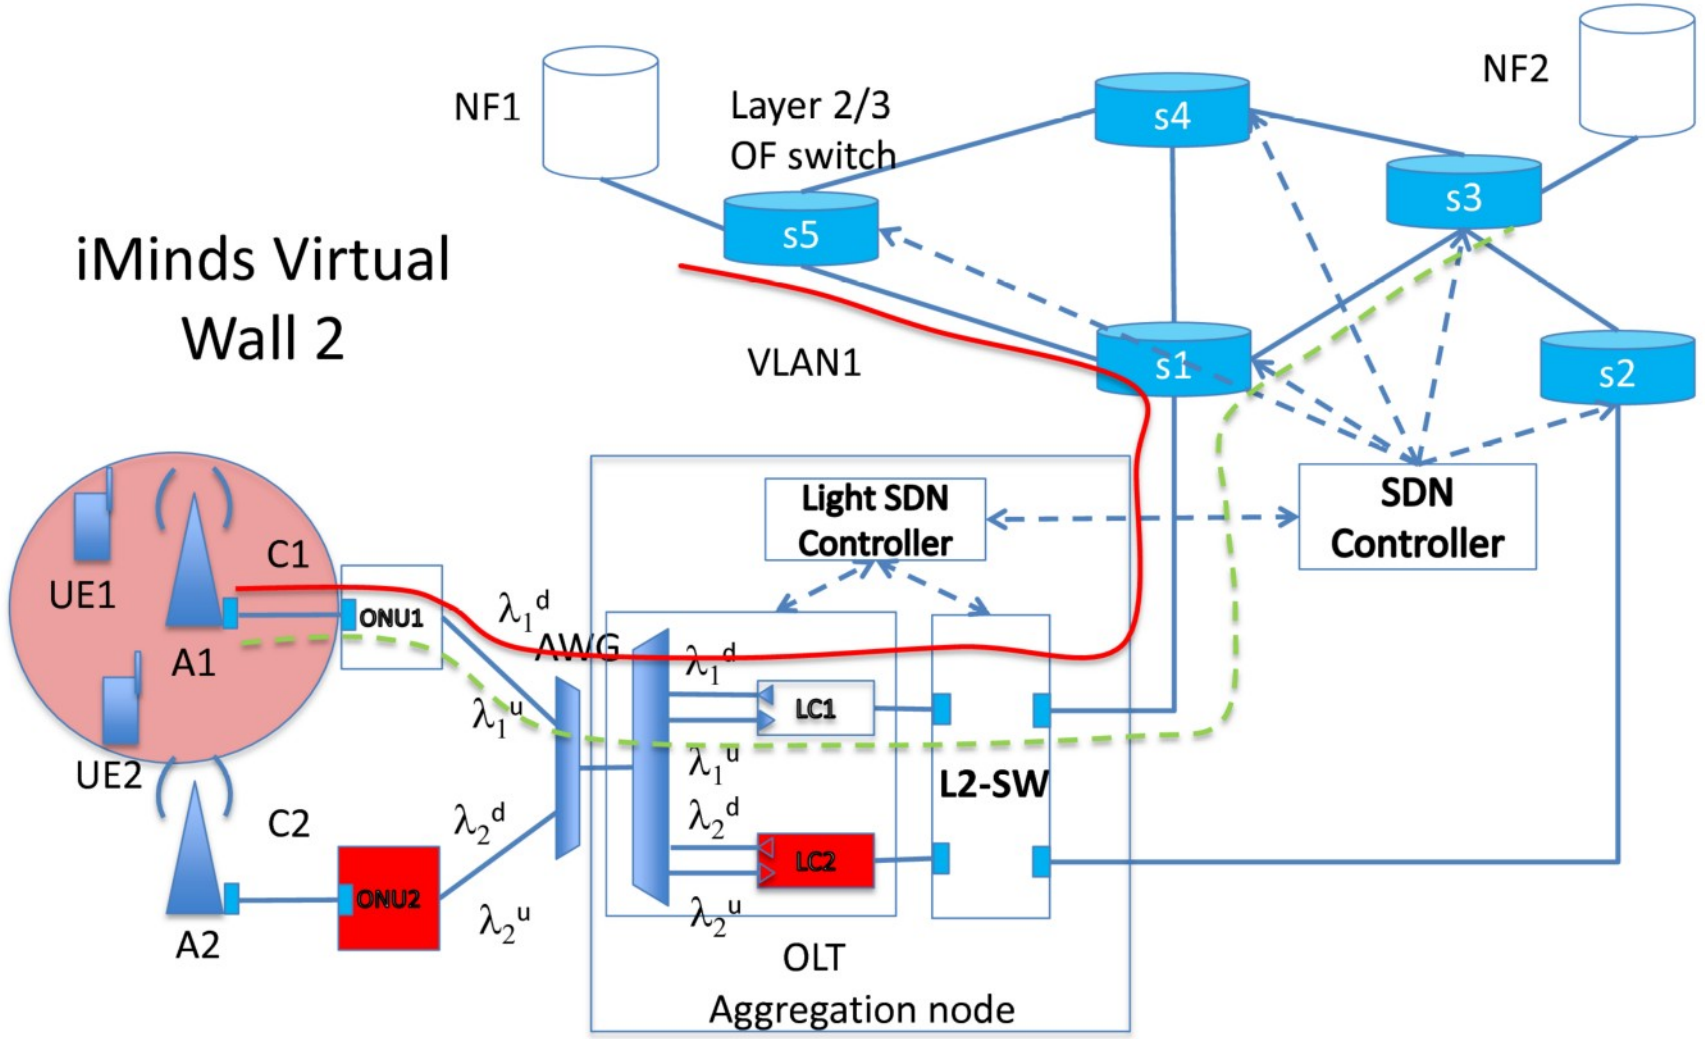

End-to-End Service Orchestration in 5G

B. Martini[(a)], M. Gharbaoui[(b)], S. Fichera[(b)], P. Castoldi[(b)]

(a)Scuola Superiore Sant'Anna, Pisa, Italy - (b)CNIT, Pisa, Italy

End-to-End Service Orchestration in 5G

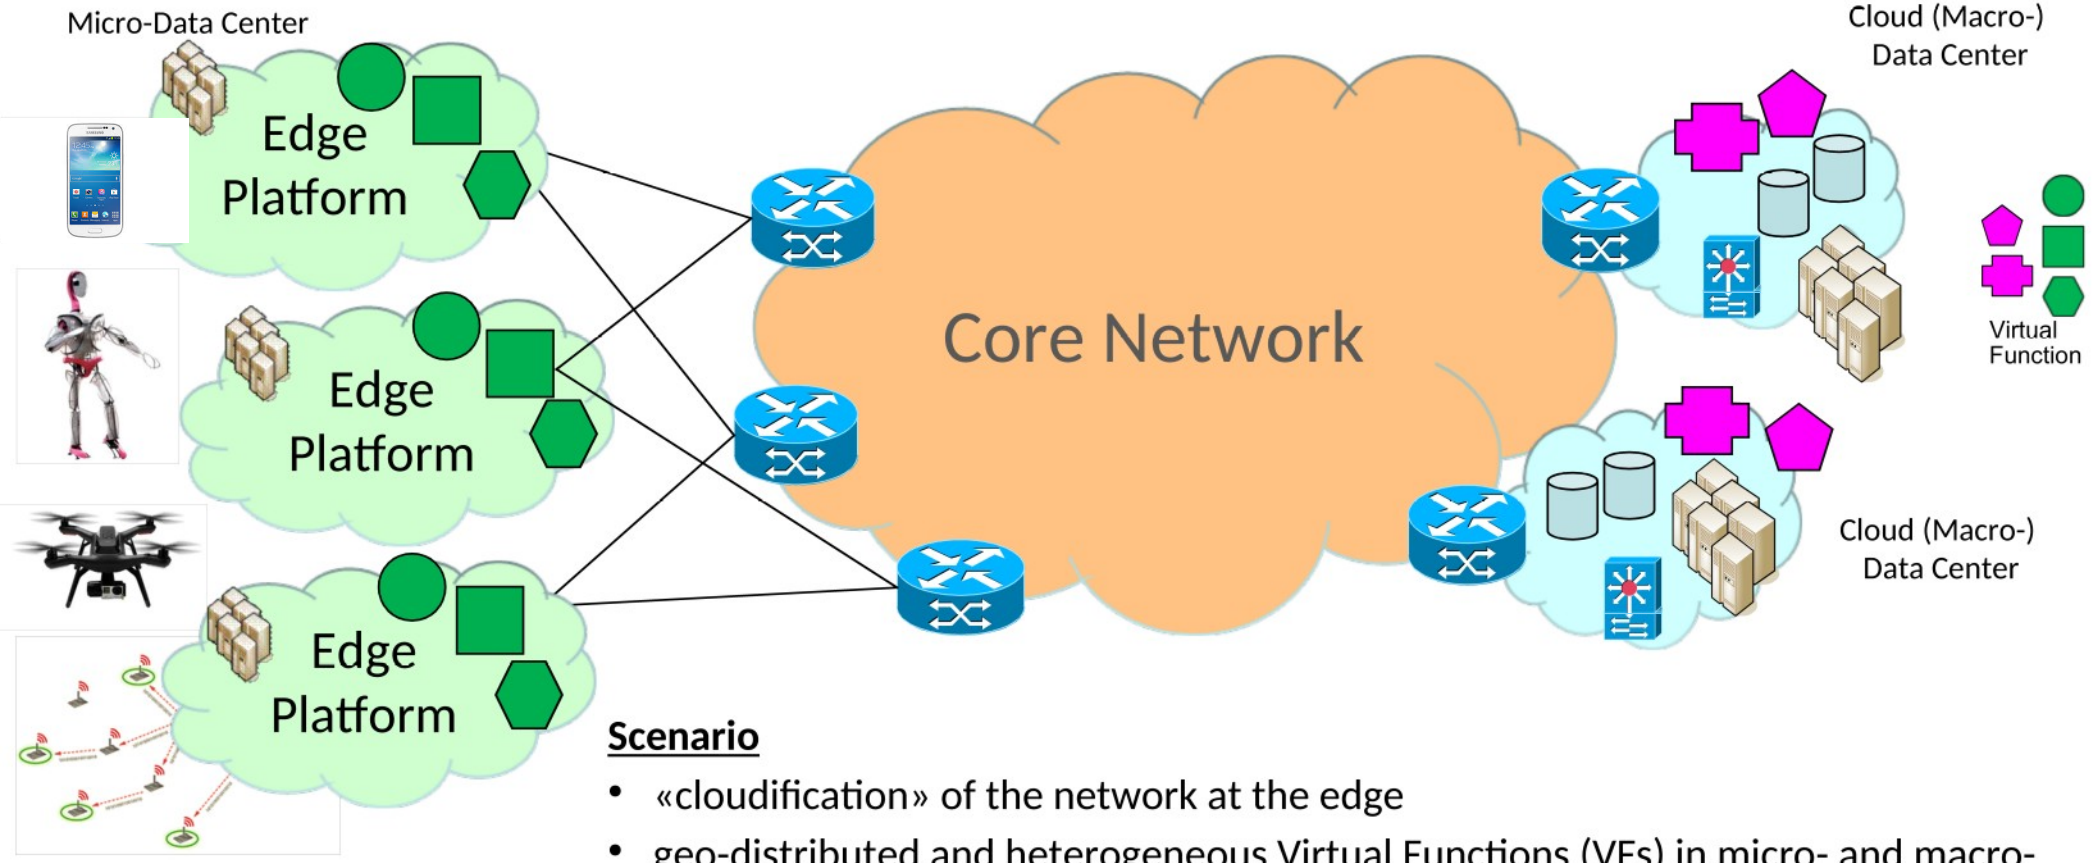

Scenario

- «cloudification» of the network at the edge
- geo-distributed and heterogeneous Virtual Functions (VFs) in micro- and macro-clouds controlled by different managers/orchestrators

Challenges

- **latency-aware VF placement in geo-distributed clouds**
- optimize infrastructure usage (scarce-vs-powerful resources at the edge-vs-macroDC)
- generalized **service chains** including both computing and network VFs for high availability

Approach

- integration of **cloud, NFV and SDN** paradigms
- orchestrated deployment of VFs at local, edge and cloud level (dynamic service chaining)
- strategies to minimize (or don't exceed) overall processing+networking **e2e delays**
- Evaluation by simulation, EU FIRE infrastructure, OpenStack & POX/Floodlight/ONOS SDN ctrls

Selected Publications
Journals

- **Improving Resp. of Time-Sensitive Applications by Exploiting Dynamic Task Dependencies**
  (to appear on) Software: Practice & Experience, Wiley, 2018
  T. Cucinotta, L. Abeni, J. Lelli, G. Lipari
- **Elastic Admission Control for Federated Cloud Services**, IEEE TCC 2014
  K. Konstanteli, **T. Cucinotta**, K. Psychas, T. Varvarigou
- **Challenges in Real-Time Virtualization and Predictable Cloud Computing**, Elsevier JSA 2014
  M. Garcia-Valls, **T. Cucinotta**, C. Lu
- **Modular Software Architecture for Flexible Reservation...**, Elsevier JSA 2011
  M. Sojka, P. Pisa, D. Faggioli, **T. Cucinotta**, F. Checconi, Z. Hanzalek, G. Lipari
- **A robust mechanism for adaptive scheduling of multimedia applications**, ACM TECS 2011
  T. Cucinotta, L. Abeni, L. Palopoli, G. Lipari
- **Adaptive Real-Time Scheduling for Legacy Multimedia Applications**, ACM TECS 2012
  T. Cucinotta, F. Checconi, L. Abeni, L. Palopoli
- **Optimum Allocation of Distributed Service Workflows with Probabilistic Real-Time Guarantees**, Springer SOCA 2010
  K. Konstanteli, **T. Cucinotta**, T. Varvarigou
- **QoS Control for Pipelines of Tasks using Multiple Resources**, IEEE TC 2010
  T. Cucinotta, L. Palopoli
- **A Real-time Service-Oriented Architecture for Industrial Automation**, IEEE TII 2009
  T. Cucinotta, A. Mancina, G. F. Anastasi, G. Lipari, L. Mangeruca, R. Checcozzo

Selected Publications Conferences/Workshops

- **Temporal Isolation Among LTE/5G Network Functions by Real-time Scheduling**, CLOSER 2017
  T. Cucinotta, M. Marinoni, A. Melani, A. Parri, C. Vitucci
- **Constant Bandwidth Servers with Constrained Deadlines**, RTNS 2017
  D. Casini, L. Abeni, A. Biondi, **T. Cucinotta**, G. Buttazzo
- **Automata-Based Modeling of Interrupts in the Linux PREEMPT RT Kernel**, IEEE ETFA 2017
  D. B. de Oliveira, R. S. de Oliveira, **T. Cucinotta**, L. Abeni
- **Data Centre Optimisation Enhanced by SDN**, IEEE CLOUD 2014
  T. Cucinotta, D. Lugones, D. Cherubini, E. Jul
- **Brokering SLAs for end-to-end QoS in Cloud Computing**, CLOSER 2014
  T. Cucinotta, D. Lugones, D. Cherubini, K. Oberle
- **Virtualised e-Learning with RT Guarantees on the IRMOS Platform**, IEEE SOCA 2010 (best paper award)
  T. Cucinotta, F. Checconi, …
- **Providing Performance Guarantees to VMs using Real-Time Scheduling**, VHPC 2010
  T. Cucinotta, D. Giani, F. Checconi, D. Faggioli
- **An Exception Based Approach to Timing Constraints Violations...**, IEEE SIES 2010
  T. Cucinotta, D. Faggioli
- **The Multiprocessor BandWidth Inheritance Protocol**, ECRTS 2010
  D. Faggioli, G. Lipari, **T. Cucinotta**
- **Advance Reservations for Distr. RT Workflows with Probabilistic...**, IEEE SOCA 2009
  T. Cucinotta, K. Konstanteli, T. Varvarigou
- **Self-tuning Schedulers for Legacy Real-Time Applications**, EuroSys 2010
  T. Cucinotta, F. Checconi, L. Abeni, L. Palopoli
- **Hierarchical Multiprocessor CPU Reservations for the Linux Kernel**, OSPERT 2009
  F. Checconi, **T. Cucinotta**, D. Faggioli, G. Lipari
- **Respecting temporal constraints in virtualised services**, IEEE RTSOAA 2009
  T. Cucinotta, G. Anastasi, L. Abeni

Selected Publications InRete & CNIT

Journal Papers

- B. Martini and F. Paganelli, “A Service-Oriented Approach for Dynamic Chaining of Virtual Network Functions over Multi-Provider Software-Defined Networks”, Future Internet, Vol. 8, Issue 2, June 2016
- M. Gharbaoui, B. Martini, D. Adami, S. Giordano, P. Castoldi, “Cloud and network orchestration in SDN data centers: Design principles and performance evaluation, “ Computer Networks, Volume 108, 24 October 2016
- C.J. Bernardos, B.P. Gerö, M. Di Girolamo, A. Kern, B. Martini, I. Vaishnavi, “5GEx: realising a Europe-wide multi-domain framework for software-defined infrastructures”, Trans. Emerging Tel. Tech., July 2016
- R. Guerzoni, I. Vaishnavi, D. Perez, A. Galis, F. Tusa, P. Monti, A. Sgambelluri, G. Biczók, B. Sonkoly, L. Toka, A. Ramos, J. Melián, O. Dugeon, F. Cugini, B. Martini, P. Iovanna, G. Giuliani, R. Figueiredo, L. M. Contreras-Murillo, C. J. Bernardos, C. Santana, R. Szabo, "Towards End-to-End Management and Orchestration of Software Defined Infrastructures: an Architectural Survey", Trans. Emerging Tel. Tech., August 2016
- W. Cerroni, M. Gharbaoui, B. Martini, A. Campi, P. Castoldi, F. Callegati, “Cross-layer resource orchestration for cloud service delivery: A seamless SDN approach”, Computer Networks, Elsevier, 2015
- F. Paganelli, M. Ulema, B. Martini, "Context-aware service composition and delivery in NGSONs over SDN," in IEEE Communications Magazine, vol.52, no.8, pp.97-105, Aug. 2014

Conference Papers

- B. Martini, F. Paganelli, P. Cappanera, S. Turchi and P. Castoldi, "Latency-aware composition of Virtual Functions in 5G," NetSoft 2015, London, April 2015
- B. Martini, F. Paganelli, A.A. Mohammed, M. Gharbaoui, A. Sgambelluri, P. Castoldi, “SDN controller for context-Aware data delivery in dynamic service chaining “, NetSoft 2015, London, UK, April 2015
- A.A. Mohammed, M. Gharbaoui, B. Martini, F. Paganelli and P. Castoldi, "SDN controller for network-aware adaptive orchestration in dynamic service chaining," NetSoft 2016, Seul, Korea, June 2016, pp. 126-130
- M. Gharbaoui, B. Martini, D. Adami, P. Castoldi and S. Giordano, "Experiments on SDN-based network and cloud resource orchestration in FED4FIRE," NetSoft 2016, Seoul, 2016, pp. 131-135

thank you!

Q&A

t.cucinotta@sssup.it