

INPUT

In-Network Programmability for next-generation personal cloUd service support



MATILDA

A HOLISTIC, INNOVATIVE FRAMEWORK FOR THE DESIGN, DEVELOPMENT AND ORCHESTRATION OF 5G-READY APPLICATIONS AND NETWORK SERVICES OVER SLICED PROGRAMMABLE INFRASTRUCTURE



Modeling Performance and Energy Efficiency of Virtualized Flexible Networks

Franco Davoli

DITEN-University of Genoa, Italy

CNIT – University of Genoa Research Unit and National Laboratory of Smart, Sustainable and Secure Internet Technologies and Infrastructures (S3ITI)



<http://www.tnt-lab.unige.it>

<http://www.cnit.it/en/institutes/s3iti/>



SIMULTECH 2017

7th International Conference on Simulation and Modeling Methodologies, Technologies and Applications

Madrid, Spain . 26 - 28 July, 2017

Outline

- Today's bottlenecks: rigid, non-general-purpose IT infrastructure
- Keywords: Flexibility, Programmability, Energy Efficiency
- Possible ways to achieve the goals: SDN, NFV, Green capabilities
- Short account on SDN / NFV - Openflow
- Reasons for going green – The Carbon footprint

Outline

- Power management primitives
- Energy-Performance Tradeoff
- Conveying power management capabilities:
the GAL
- The importance of models
- Modeling energy consumption and
performance in NFV
- Some examples
- European Projects

Current bottlenecks in the networking infrastructure

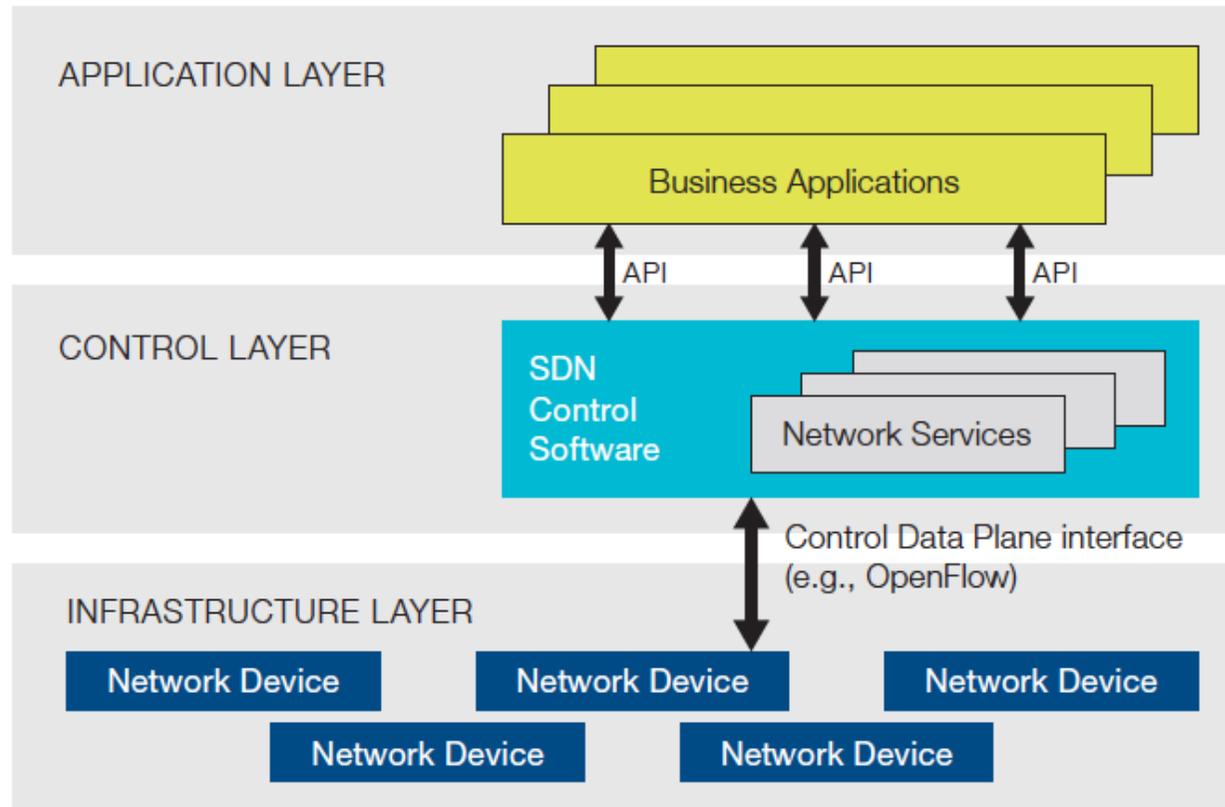
- Once it used to be bandwidth... (still to be administered carefully in some cases, though)
- However, with the increase of available bandwidth and processing speed, paralleled by an unprecedented increase in user-generated traffic, other factors that were previously concealed have become evident:
 - The networking infrastructure makes use of a large variety of hardware appliances, dedicated to specific tasks, which typically are *inflexible, energy-inefficient, unsuitable to sustain reduced Time to Market of new services.*

Keywords

- As one of the main tasks of the network is allocating resources, how to make it more dynamic, performance-optimized and cost-effective?
- Current keywords are
 - **Flexibility**
 - **Programmability**
 - **Energy-efficiency**

Flexibility/Programmability – Software Defined Networking (SDN)

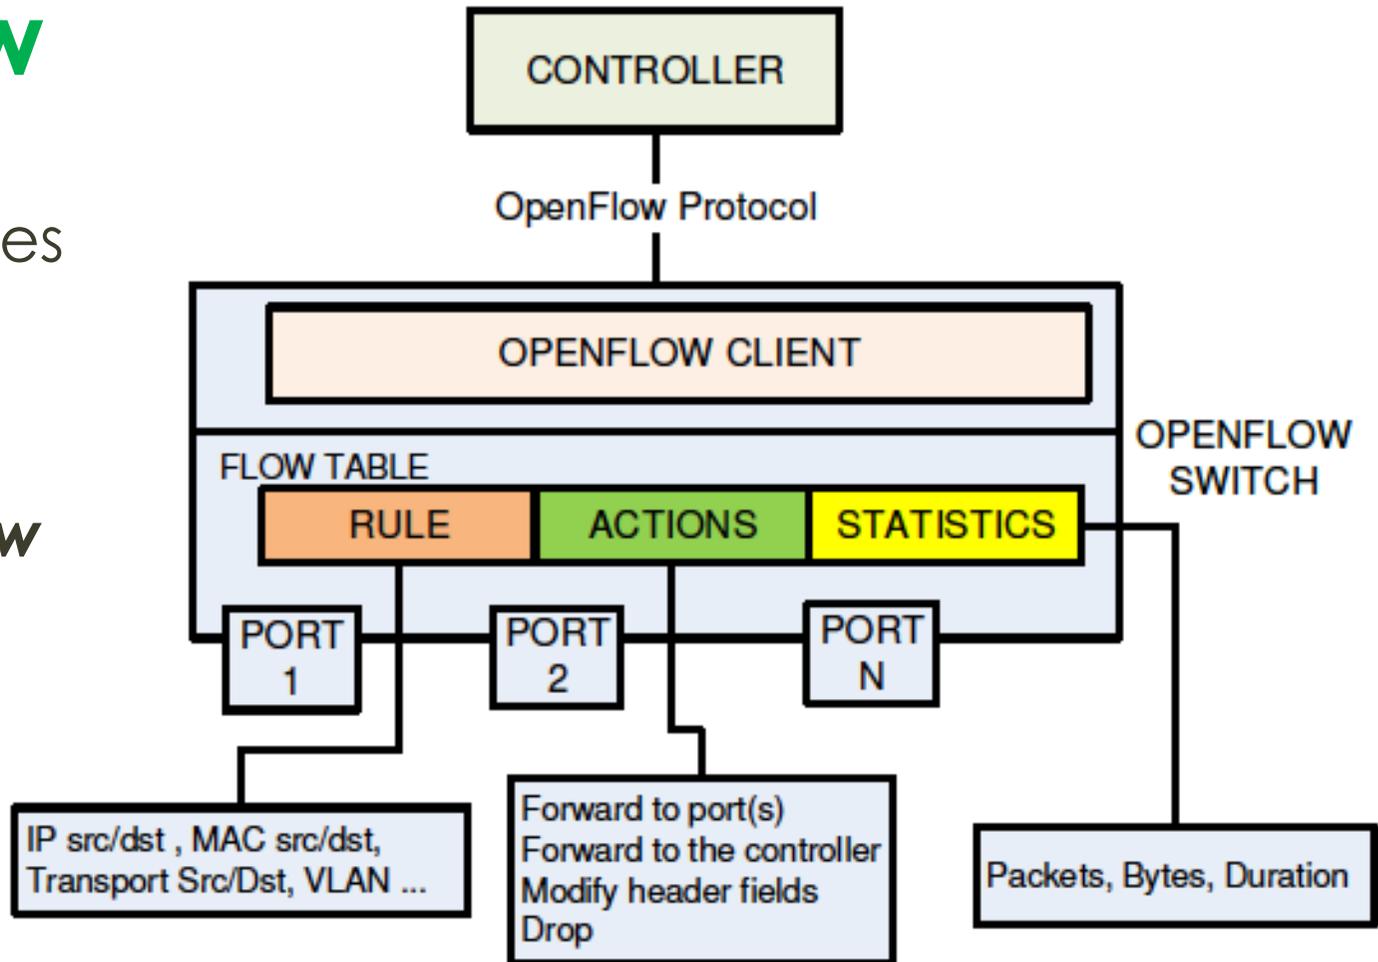
- SDN decouples the Control Plane and the Data Plane (Forwarding) Plane.



Source: *Software-Defined Networking: The New Norm for Networks*, Open Networking Foundation (ONF) White Paper, April 2012.

Flexibility/Programmability – OpenFlow

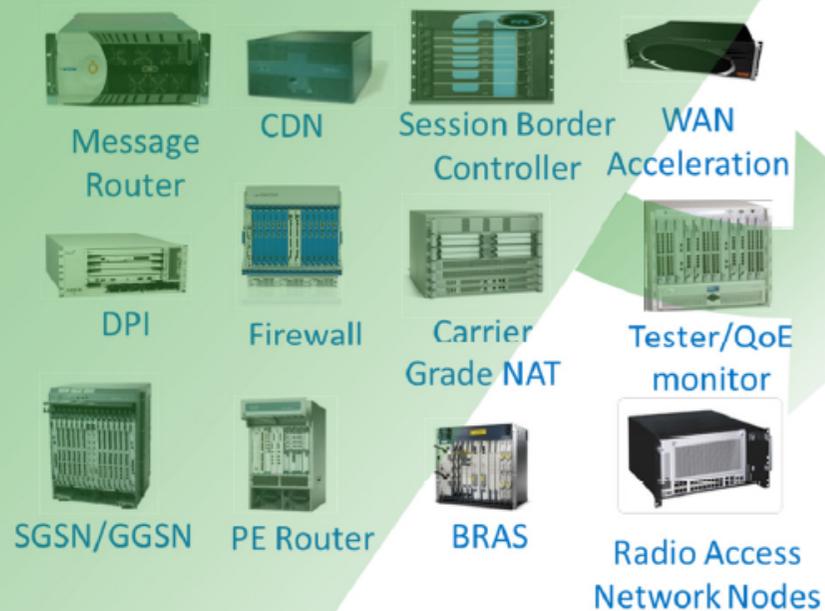
- Matching rules
- Actions
- Counters
- Acting at *flow* level



Source: B. A. A. Nunes, M. Mendonça, X.-N. Nguyen, K. Obraczka, T. Turletti, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks", *IEEE Communications Surveys & Tutorials*, vol. 16, no. 3, pp. 1617 – 1634, 3rd Qr. 2014.

Flexibility/Programmability – Network Functions Virtualization (NFV)

Classical Network Appliance Approach



- Fragmented non-commodity hardware.
- Physical install per appliance per site.
- Hardware development large barrier to entry for new vendors, constraining innovation & competition.



Leverages “...standard IT virtualisation technology to consolidate many network equipment types onto industry standard high volume servers, switches and storage, which could be located in Datacentres, Network Nodes and in the end user premises.”

Source: *Network Functions Virtualisation – Introductory White Paper*, SDN and OpenFlow World Congress, Darmstadt, Germany, Oct. 2012.

Flexibility/Programmability – Network Functions Virtualization (NFV)

- Improved equipment consolidation
- Reduced Time-to-Market
- Single platform, multiple applications, users, and tenants
- Improved scalability
- Multiple open eco-systems
- Exploits economy of scale of the IT industry
 - 2016 market for data-center servers ~ \$32 billion worldwide, growth rate ~ 6%; for routers & switches ~ \$27-billion worldwide, growth rate ~ 1%. But, in any case, *“The main disruption to the market is being provided by the growth of cloud and hosted solutions, which are redefining markets and enabling new competitors to emerge.”* *

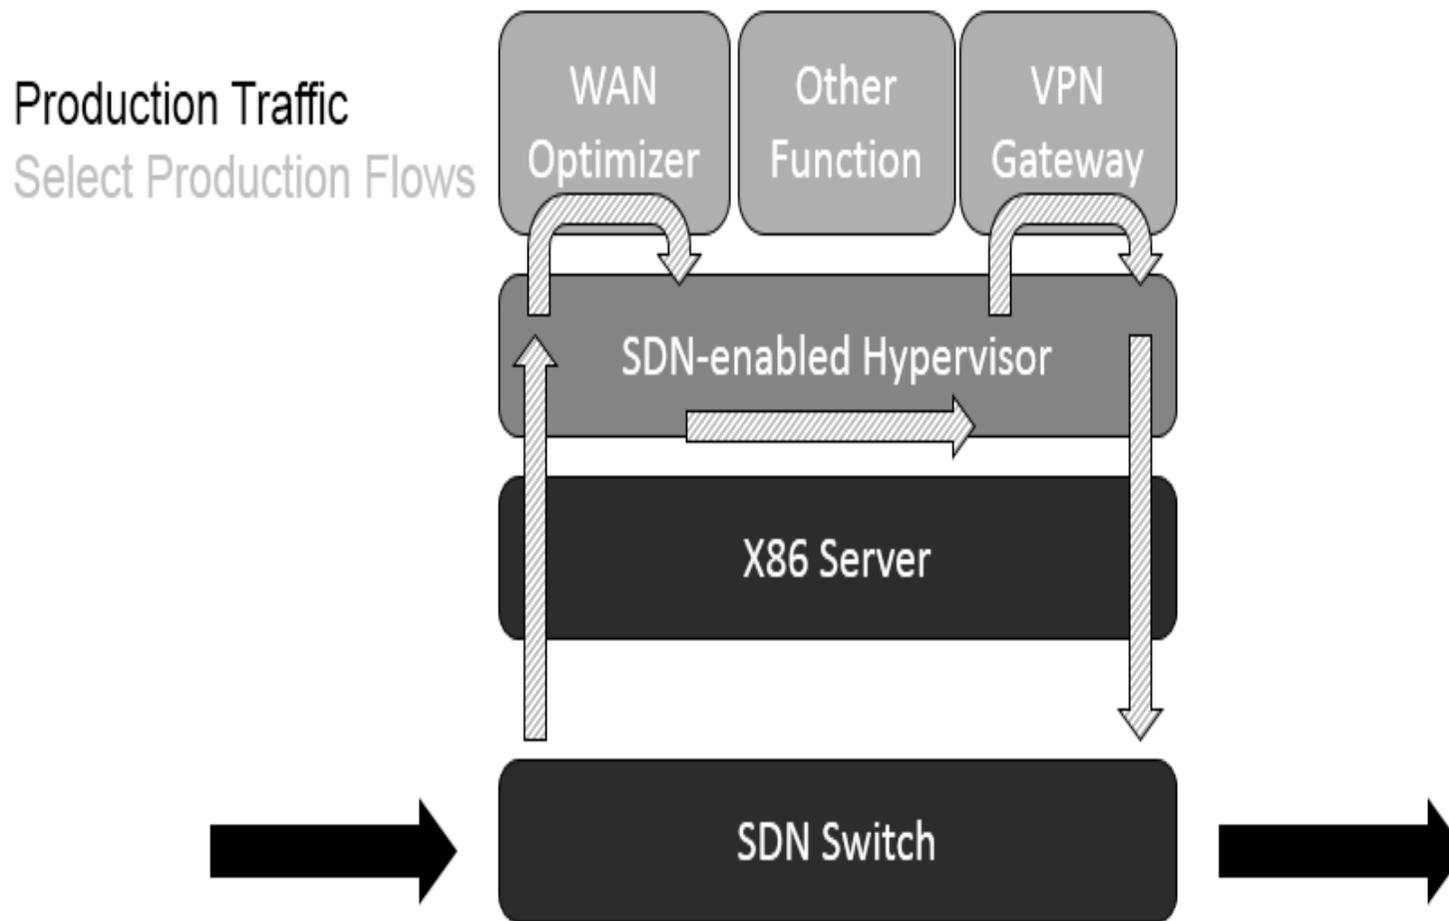
* Source: Synergy Research Group

<https://www.srgresearch.com/articles/enterprise-spending-nudged-downwards-2016-cisco-maintains-big-lead>

SDN *and* NFV

- NFV requires
 - swift I/O performance between the physical network interfaces of the hardware and the software user-plane in the virtual functions, to enable sufficiently fast processing
 - well-integrated network management and cloud orchestration system, to benefit from the advantages of dynamic resource allocation and to ensure a smooth operation of the NFV-enabled networks
- SDN is not a requirement for NFV, but NFV can benefit from being deployed in conjunction with SDN.

SDN and NFV – an example



Source: M. Jarschel, T. Hoßfeld, F. Davoli, R. Bolla, R. Bruschi, A. Carrega, "SDN-enabled energy-efficient network management," in K. Samdanis, P. Rost, A. Maeder, M. Meo, C. Verikoukis, Eds., *Green Communications: Principles, Concepts, and Practice*, Wiley, 2015, pp. 323-338.

Integrated management and control for Traffic Engineering

- The premises are there for a – **technically and operationally** – easier way to more sophisticated
 - Control

Quasi-centralized / hierarchical vs. completely distributed
 - Management

Tighter integration with control strategies, closer operational tools, perhaps only difference in time scales

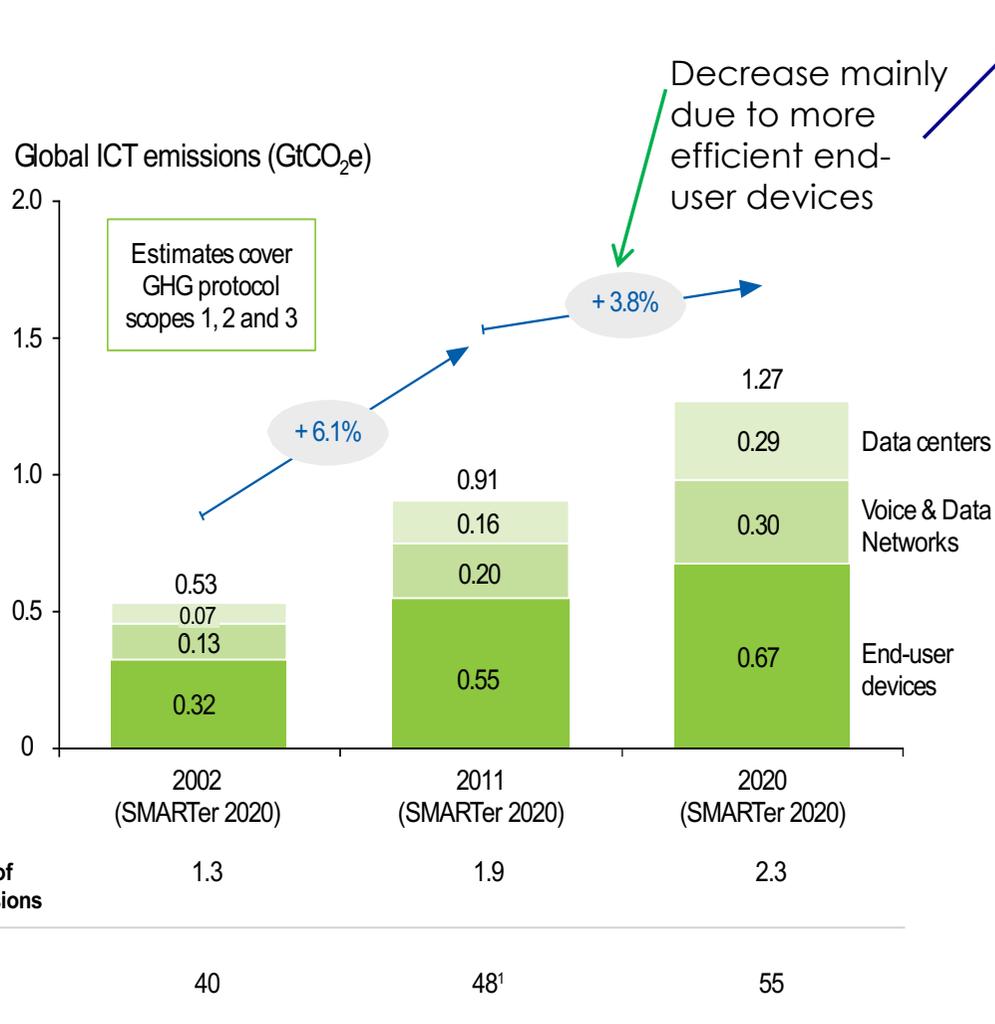
How does all this interact with network energy-efficiency?

- Making the network energy-efficient (“Green”) cannot ignore Quality of Service (QoS) / Quality of Experience (QoE) requirements.
- At the same time, much higher flexibility, as well as enhanced control and management capabilities, are required to effectively deal with the performance/power consumption tradeoff, once the new dimension of energy-awareness is taken into account in all phases of network design and operation.

Why “greening” the network?

- ICT has been historically and fairly considered as a key objective to reduce and monitor “third-party” energy wastes and achieve higher levels of efficiency.
 - Classical example: Video-Conferencing Services
 - Newer examples: ITS, Smart Electrical Grid
- However, until recently, ICT has not applied the same efficiency concepts to itself, not even in fast growing sectors like telecommunications and the Internet.
- There are **two main motivations** that drive the quest for “green” ICT:
 - the **environmental** one, which is related to the reduction of wastes, in order to impact on CO₂ emission;
 - the **economic** one, which stems from the reduction of operating costs (OPEX) of ICT services.

The Carbon Footprint of ICT



However, **ICT industry's footprint is projected to increase at a faster rate than the total global footprint between 2011 and 2020**. According to IEA data, global GHG emissions are expected to rise at 1.5 percent per year between 2011 and 2020.

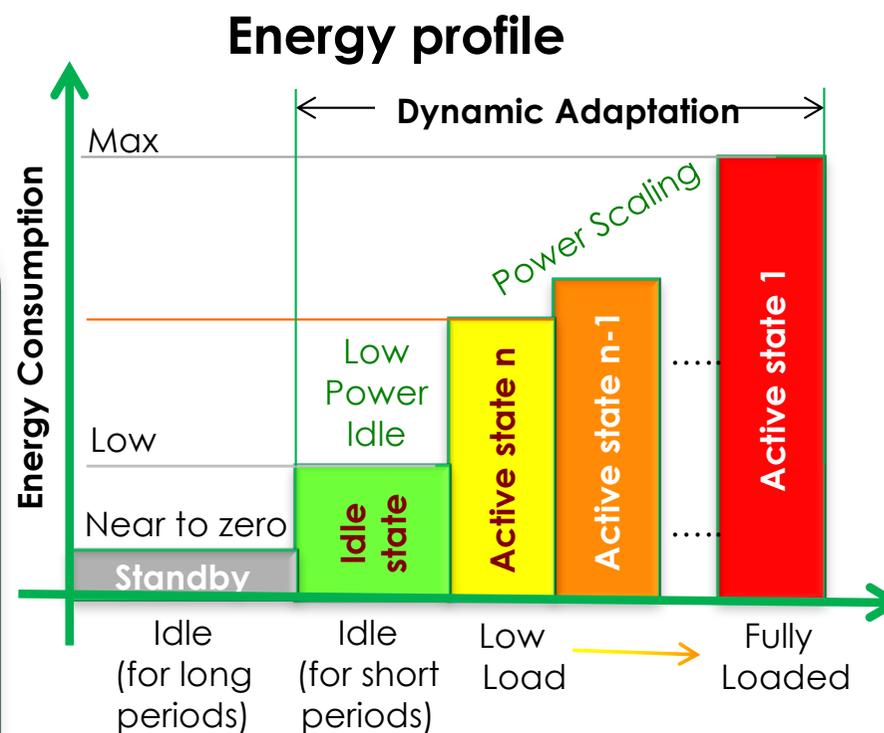
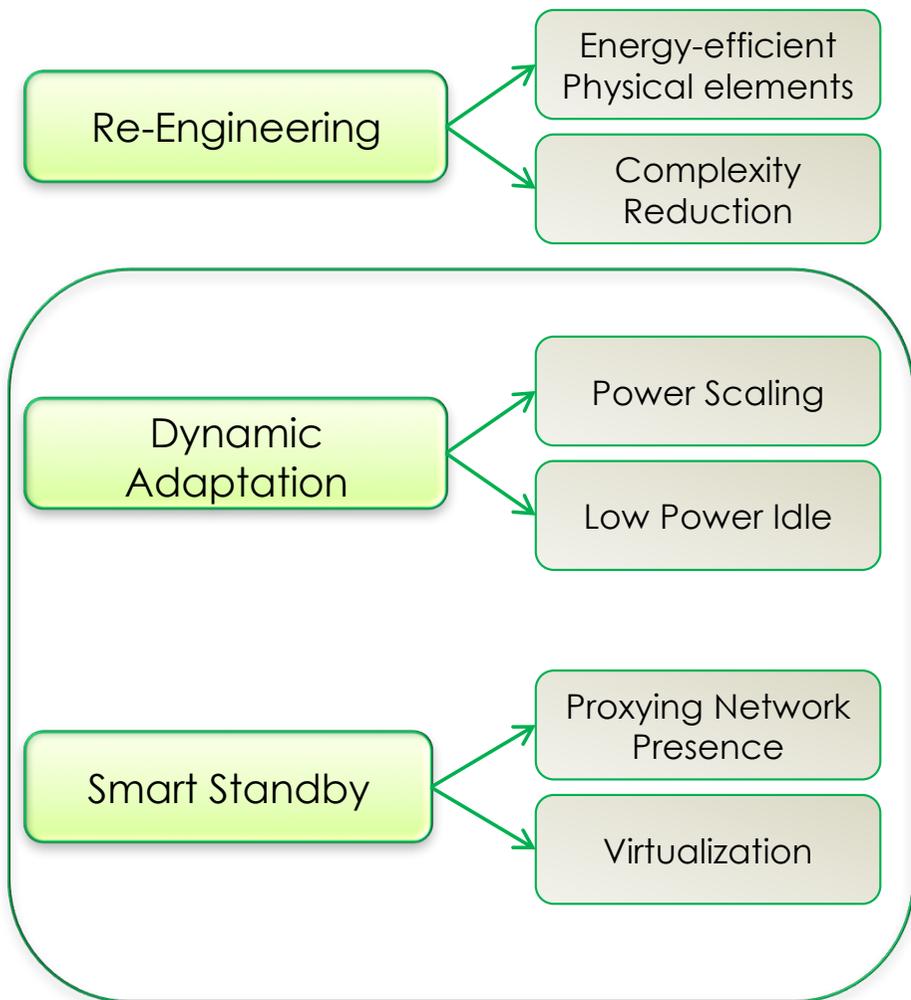
While projected ICT's own footprint is 1.3 GtCO₂e/y (2.3%), **ICT's abatement potential is 7 times higher (16.1%)**

Source: Global e-Sustainability Initiative (GeSI), "SMARTer2020: The Role of ICT in Driving a Sustainable Future," Report, URL: <http://gesi.org/SMARTer2020>.

How to manage this trend

- Today's (and future) network infrastructures characterized by:
 - **Design capable to deal with strong requests and constraints** in terms of resources and performance (large loads, very low delay, high availability,)
 - **Services characterized by high variability of load and resource requests** along time (burstiness, rush hours, ...)
- The current feasible solution:
 - **Smart power management**: energy consumption should follow the dynamics of the service requests.
 - **Flexibility in resource usage**: virtualization to obtain an aggressive sharing of physical resources

Taxonomy of Approaches



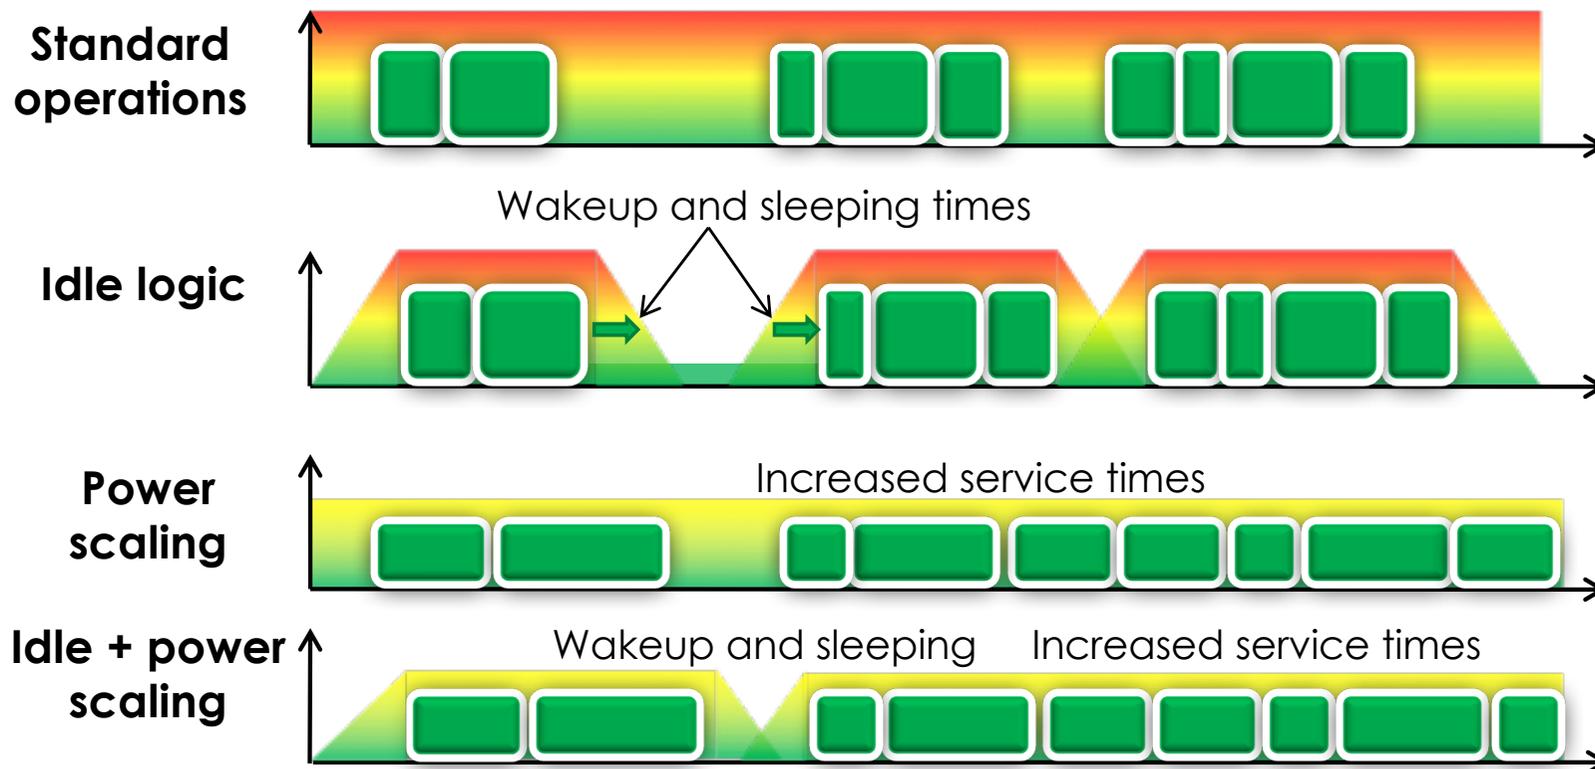
Source: R. Bolla, R. Bruschi, F. Davoli, F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 2, pp. 223-244, 2nd Qr. 2011.

Dynamic Adaptation

QoS vs Power Management

- The maximal power saving is obtained when equipment is actually turned off
- However, under such condition the performance is actually zero
- On the other extreme, it is also clear that the best performance equipment may provide is under no-power-limit mode. There is a whole range of intermediate possibilities between these two extremes.

QoS vs Power Management



Network processors - SW Routers & the ACPI

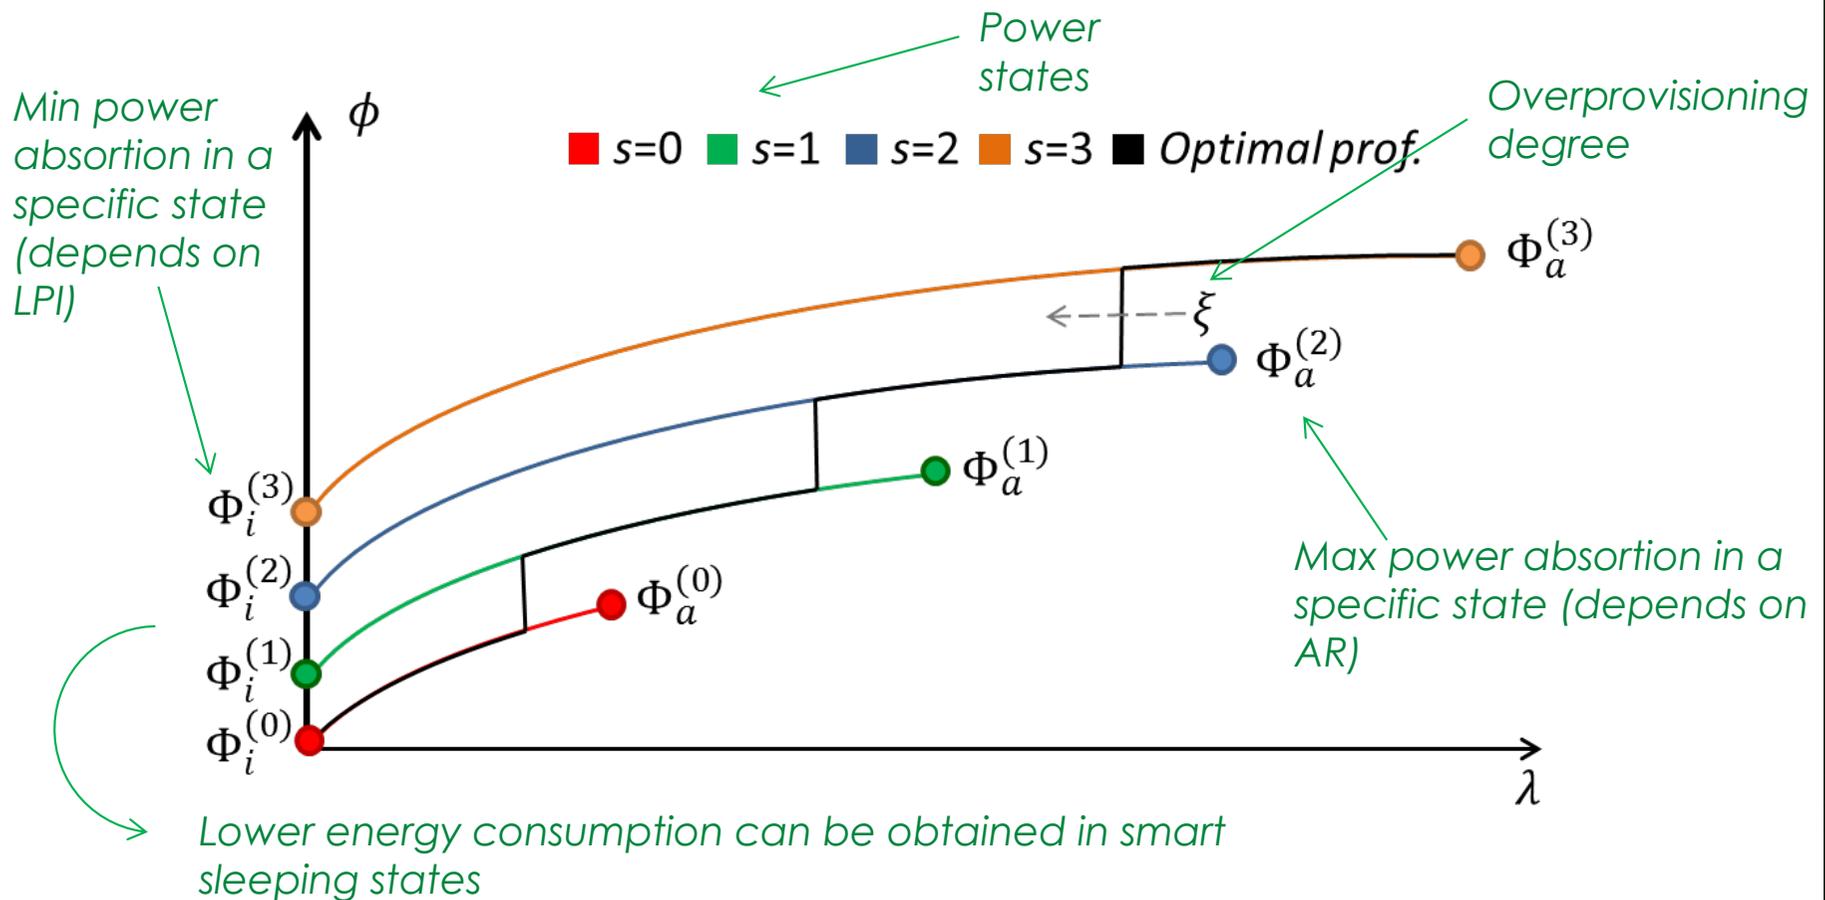
- In PC-based devices, the [Advanced Configuration and Power Interface \(ACPI\)](#) provides a standardized interface between the hardware and the software layers.
- ACPI introduces two power saving mechanisms, which can be individually employed and tuned for each core:
 - [Power States \(C-states\)](#)
 - ❑ C_0 is the active power state
 - ❑ C_1 through C_n are processor sleeping or idle states (where the processor consumes less power and dissipates less heat).
 - [Performance States \(P-states\)](#)

while in the C_0 state, ACPI allows the performance of the core to be tuned through P-state transitions.

P-states allow modify the operating energy point of a processor/core by altering the working frequency and/or voltage, or throttling the clock.

Beyond ACPI – Dynamic Adaptation (AR & LPI) and Smart Sleeping in All Network Segments

Energy profile and AR/LPI power states of a networking device



Source: R. Bolla, R. Bruschi, A. Carrega, F. Davoli, D. Suino, C. Vassilakis, A. Zafeiropoulos, "Cutting the energy bills of Internet Service Providers and Telecoms through power management: an impact analysis", Computer Networks, vol. 56, no. 10, pp. 2320–2342, July 2012.

Beyond ACPI – The Green Abstraction Layer (GAL)

ETSI ES 203 237 V1.1.1 (2014-03)

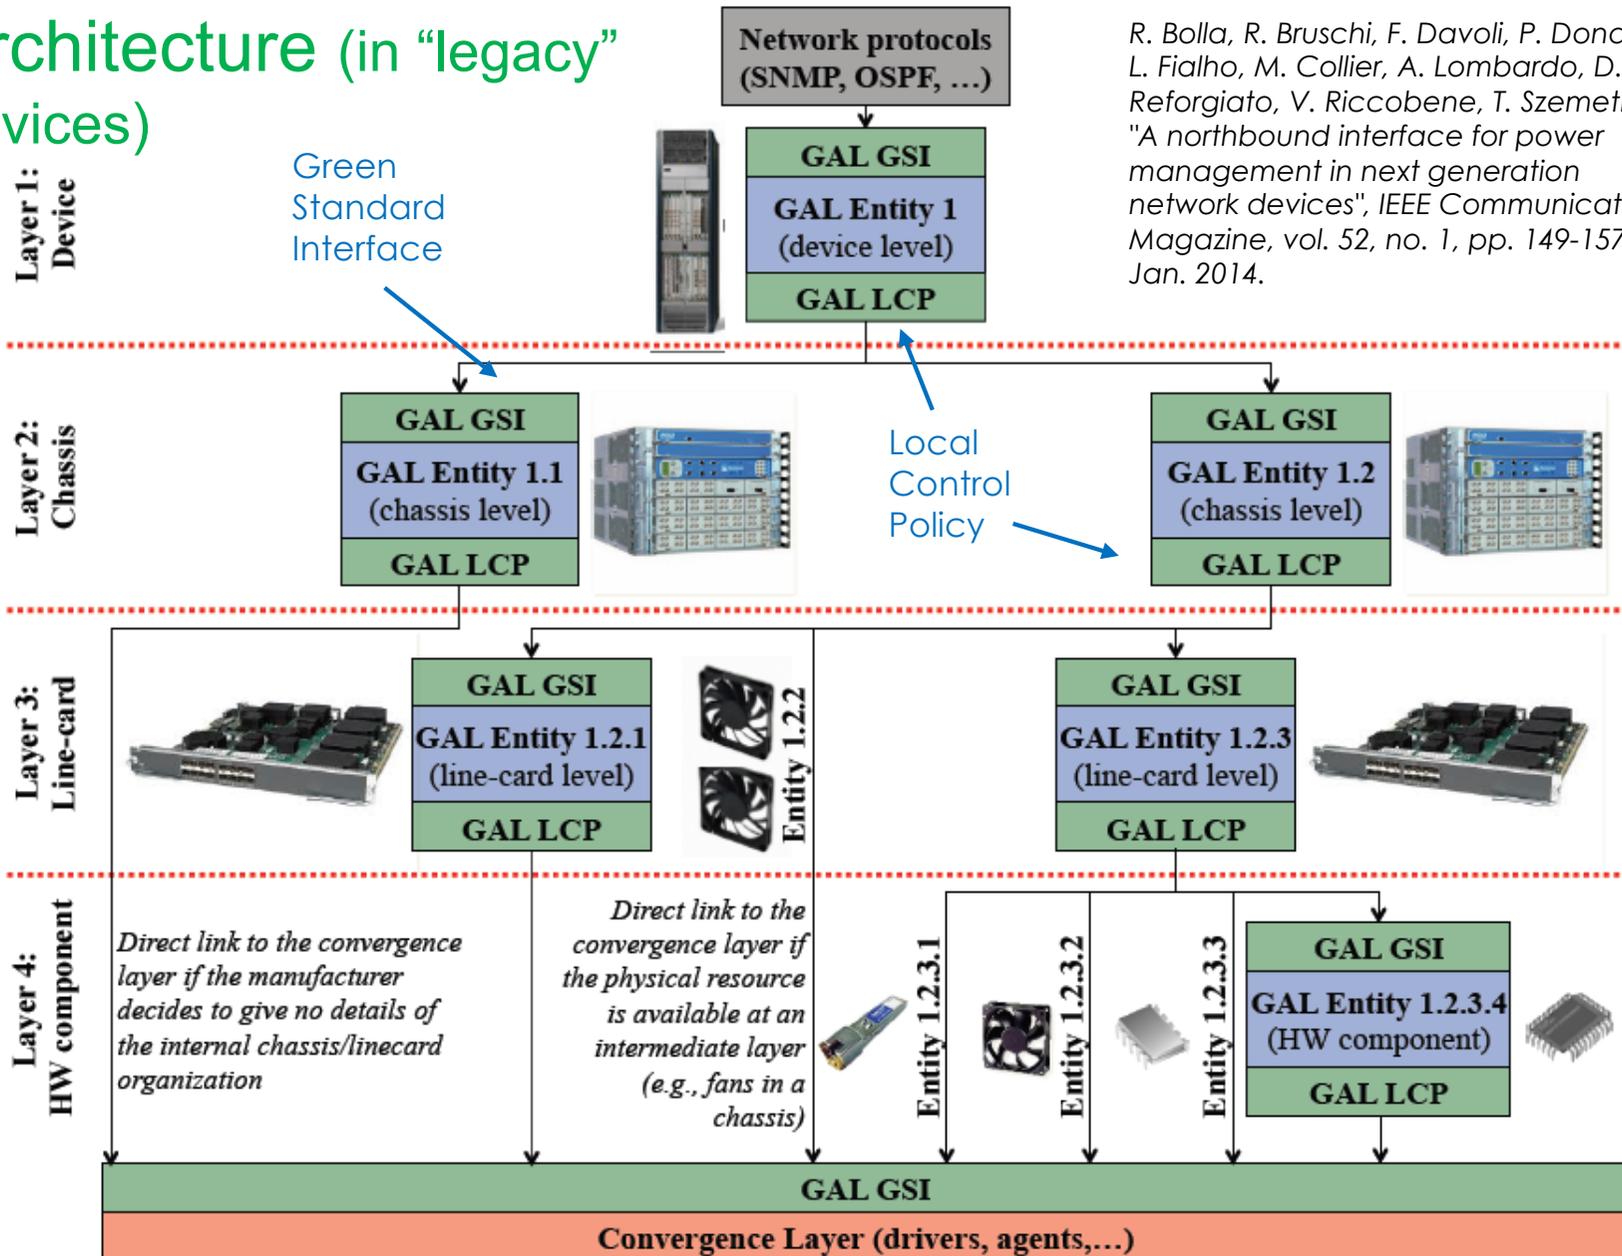


A simple method to expose energy-related parameters to the management entities inside the device architectures and to the management plane (back to this in more detail later on).

**Environmental Engineering (EE);
Green Abstraction Layer (GAL);
Power management capabilities of the future energy
telecommunication fixed network nodes**

R. Bolla, R. Bruschi, F. Davoli, L. Di Gregorio, P. Donadio, L. Fialho, M. Collier, A. Lombardo, D. Reforgiato Recupero, T. Szemethy, "The Green Abstraction Layer: A standard power management interface for next-generation network devices," IEEE Internet Computing, vol. 17, no. 2, pp. 82-86, 2013.

The GAL Hierarchical Architecture (in "legacy" devices)

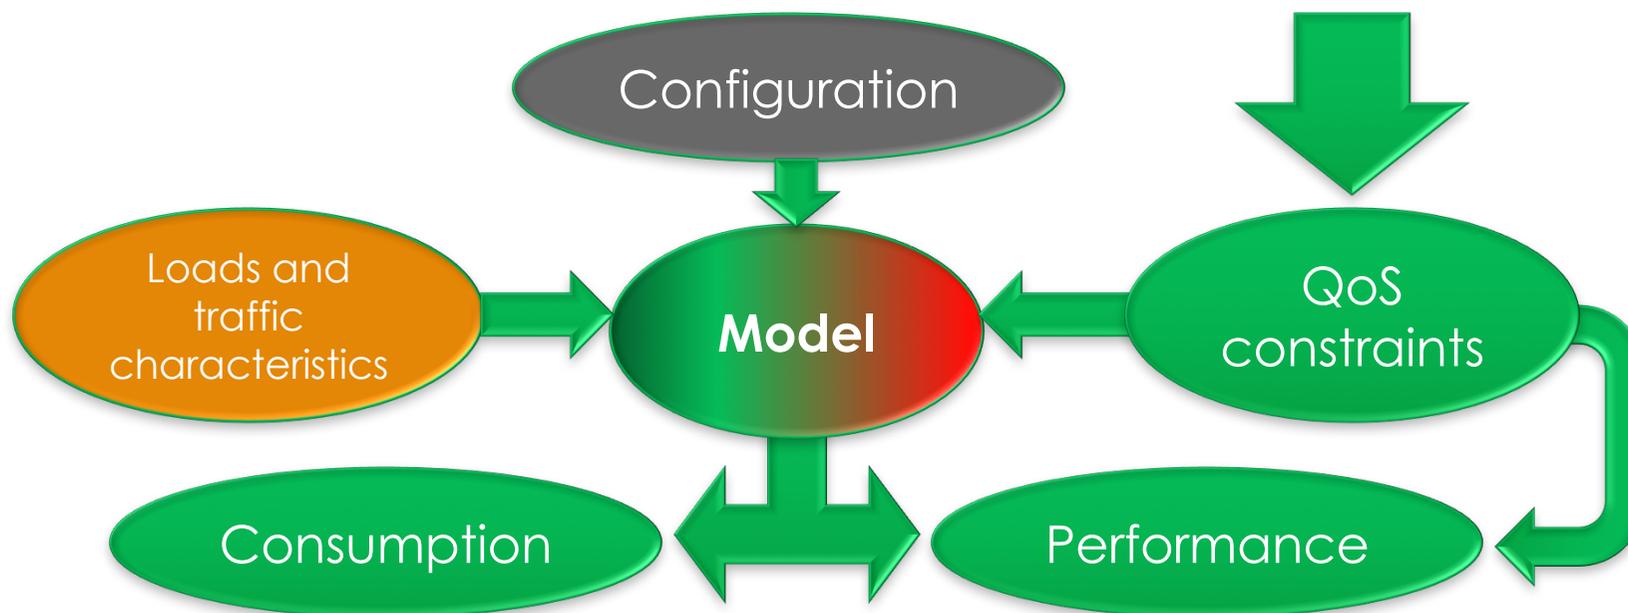


R. Bolla, R. Bruschi, F. Davoli, P. Donadio, L. Fialho, M. Collier, A. Lombardo, D. Reforgiato, V. Riccobene, T. Szemethy, "A northbound interface for power management in next generation network devices", IEEE Communications Magazine, vol. 52, no. 1, pp. 149-157, Jan. 2014.

Dynamic Adaptation

- If we change the consumption we change also the performance

We need to model a device in terms of consumption and performance versus loads and configurations



Queueing models for performance analysis and control

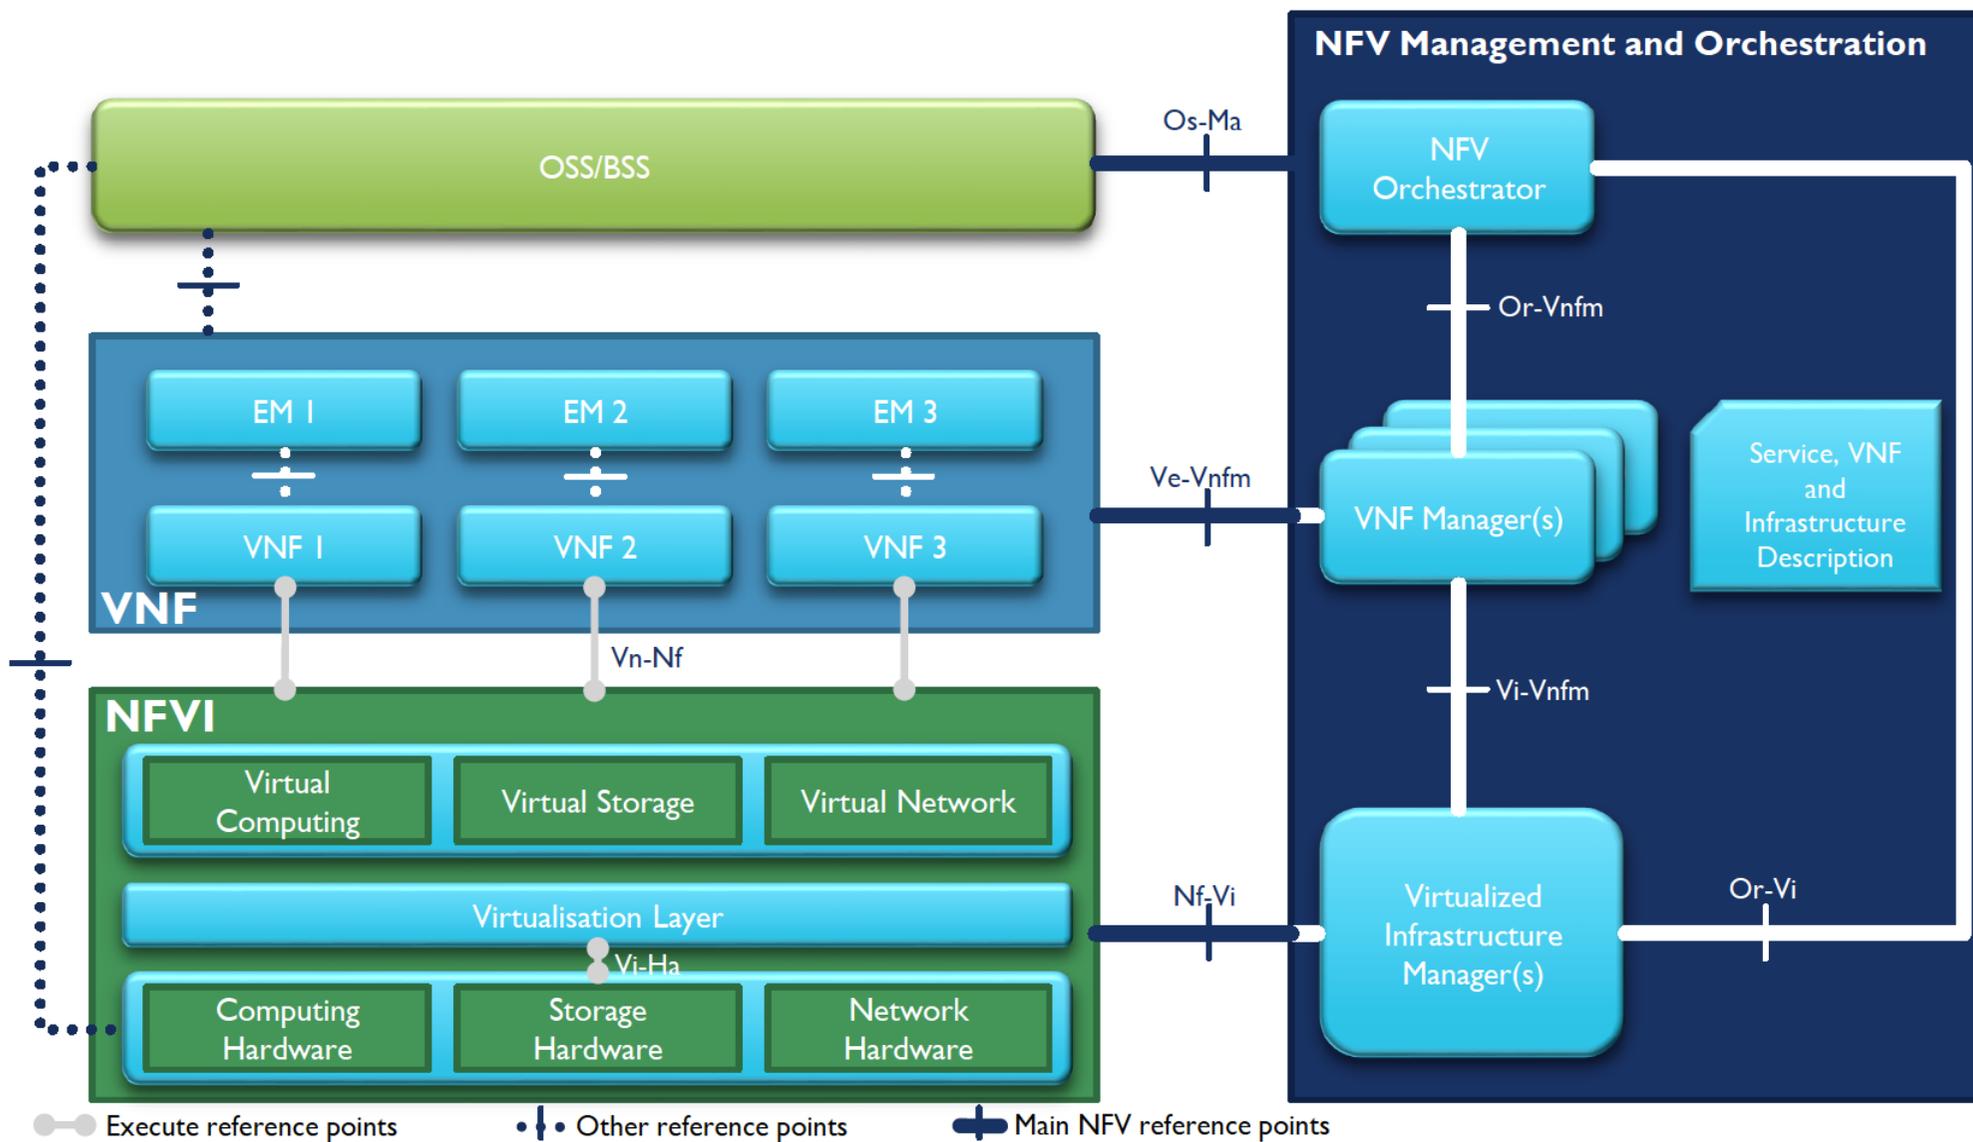
- Models based on classical queueing theory (e.g., the **$M^X/G/1/SET^{1,2}$**) lend themselves to performance analysis or parametric optimization for adaptive control and management policies over longer (with respect to queueing dynamics) time scales.
- Fluid models suitable for real-time control can stem from them ³, or even from simpler, measurement-based, stochastic continuous fluid approximations⁴.

1. R. Bolla, R. Bruschi, A. Carrega, F. Davoli, "Green networking with packet processing engines: Modeling and optimization", *IEEE/ACM Transactions on Networking*, vol. 22, no. 1, pp. 110-123, Feb. 2014 .
2. R. Bolla, R. Bruschi, A. Carrega, F. Davoli, P. Lago, "A closed-form model for the IEEE 802.3az network and power performance", *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 1, pp. 16-27, Jan. 2014.
3. J. Filipiak, *Modelling and Control of Dynamic Flows in Communication Networks*, Springer-Verlag, Berlin, Germany, 1988.
4. R. Bruschi, F. Davoli, M. Mongelli, "Adaptive frequency control of packet processing engines in telecommunication networks", *IEEE Communications Letters*, vol. 18, no. 7, pp. 1135-1138, July 2014.

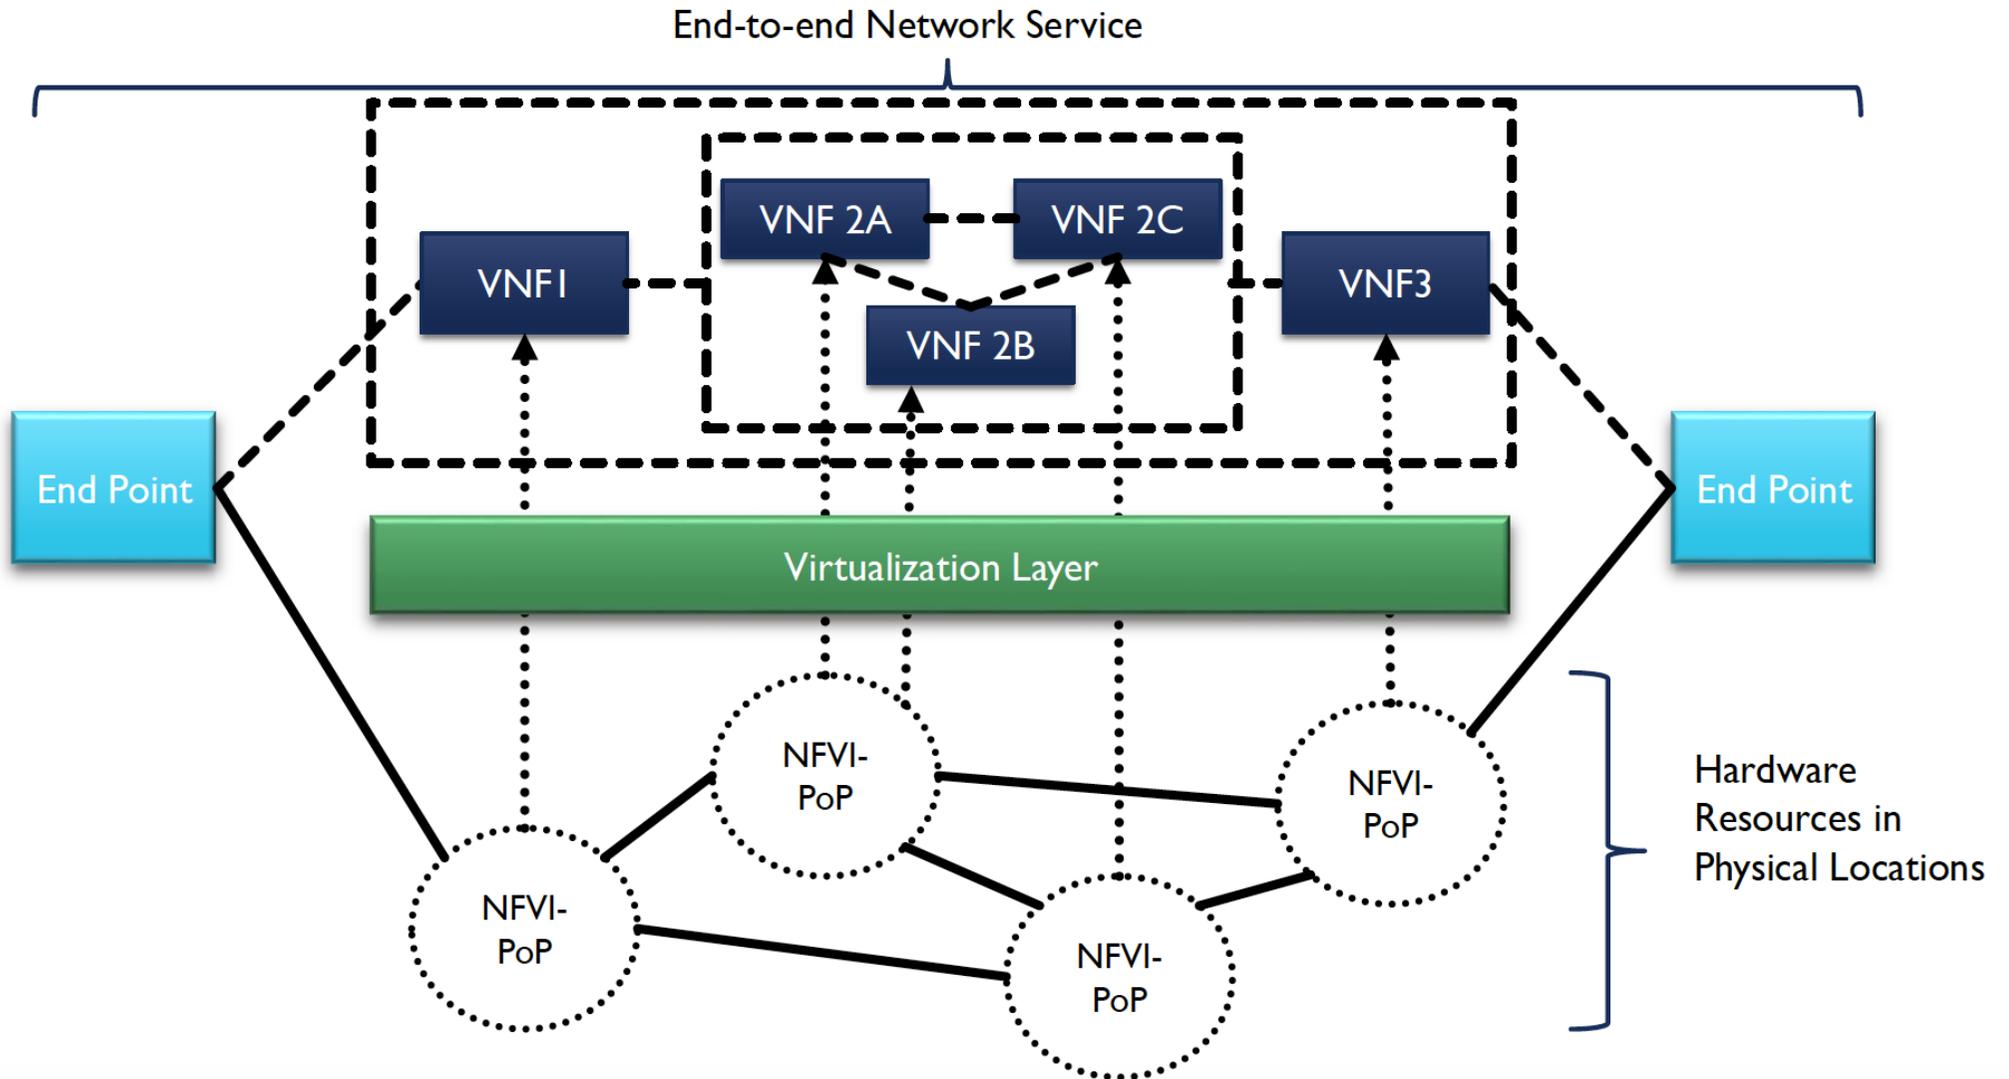
Queueing models for performance analysis and control

- These and other models have been applied for performance optimization and performance-energy tradeoff in traditional “legacy” networking architectures (based on L2-switches and routers).
- However, the introduction of NFV changes the perspective quite a bit:
 - The hardware (HW) that consumes energy belongs to the **Infrastructure Provider**, which in general may not coincide with the **Network Service Providers** (in a multi-tenant environment);
 - The HW is shared by multiple **Virtual Machines** (VMs) or by **Network Slices**, through a virtualization environment;
 - Queueing models can be identified and used to assess the performance of VMs as function of the virtual resources assigned to them (as well as to control their assignment), but the relation between the performance of the VMs and their energy consumption is not straightforward.

A look at ETSI NFV Architectural Framework

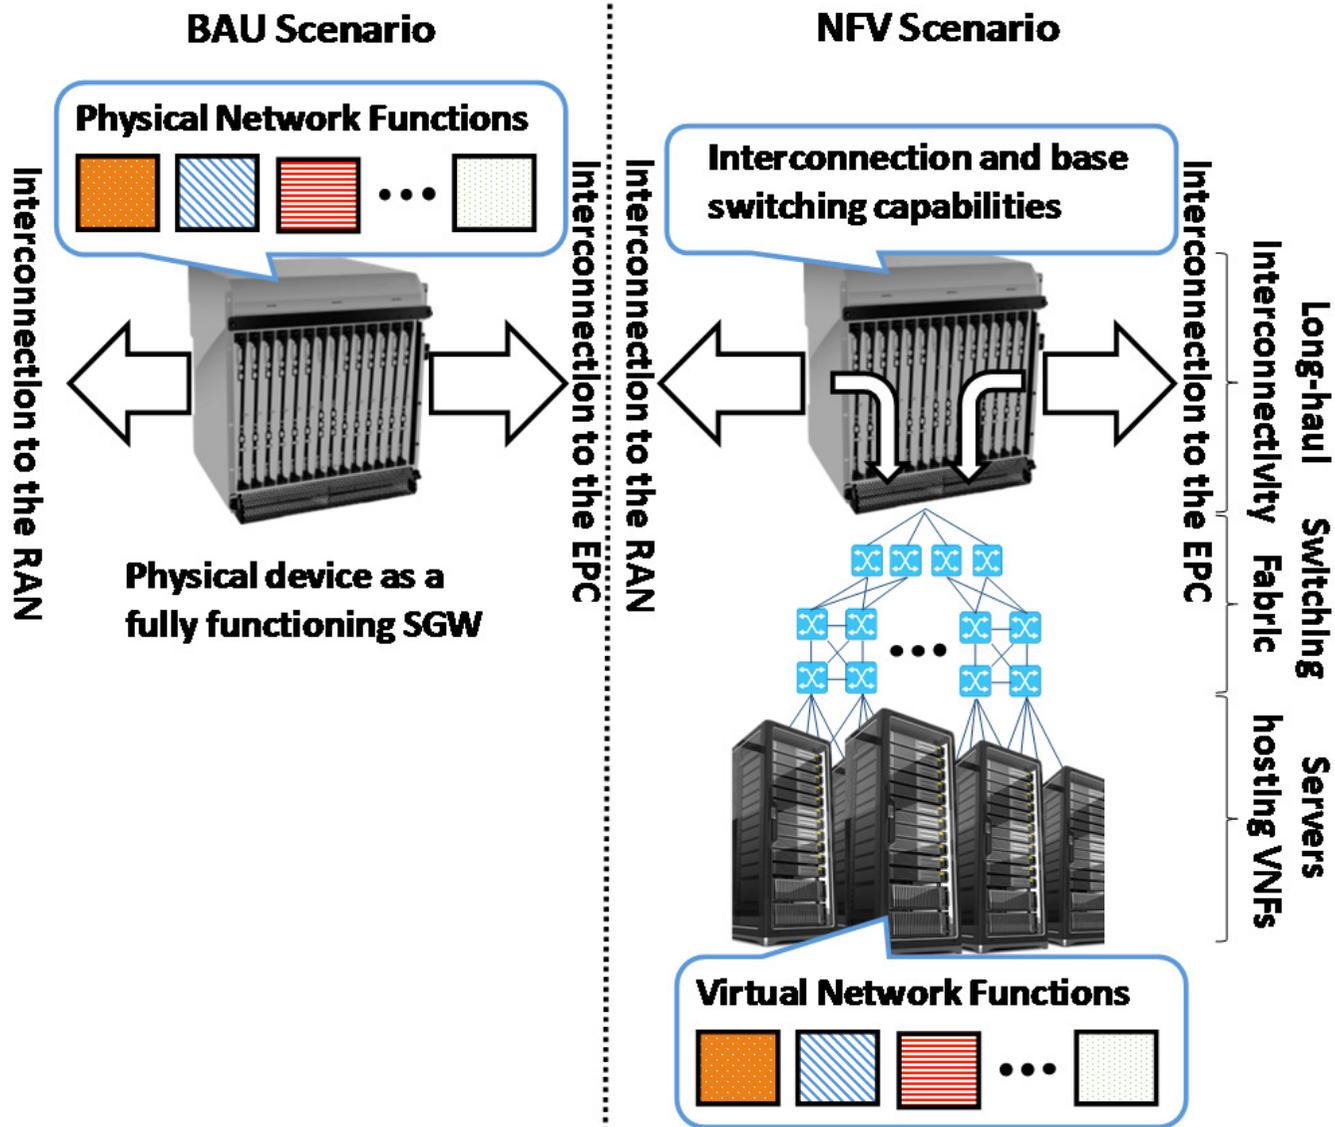


Network Services as NF Forwarding Graphs



An energy consumption

use case: virtualization of a Serving Gateway (SGW) in the Evolved Packet Core (EPC)



Source: R. Bolla, R. Bruschi, F. Davoli, C. Lombardo, J. F. Pajo, O. R. Sanchez, "The dark side of network functions virtualization: A perspective on the technological sustainability", Proc. IEEE Intern. Conf. Commun. (ICC 2017), Paris, France, May 2017.

An energy consumption

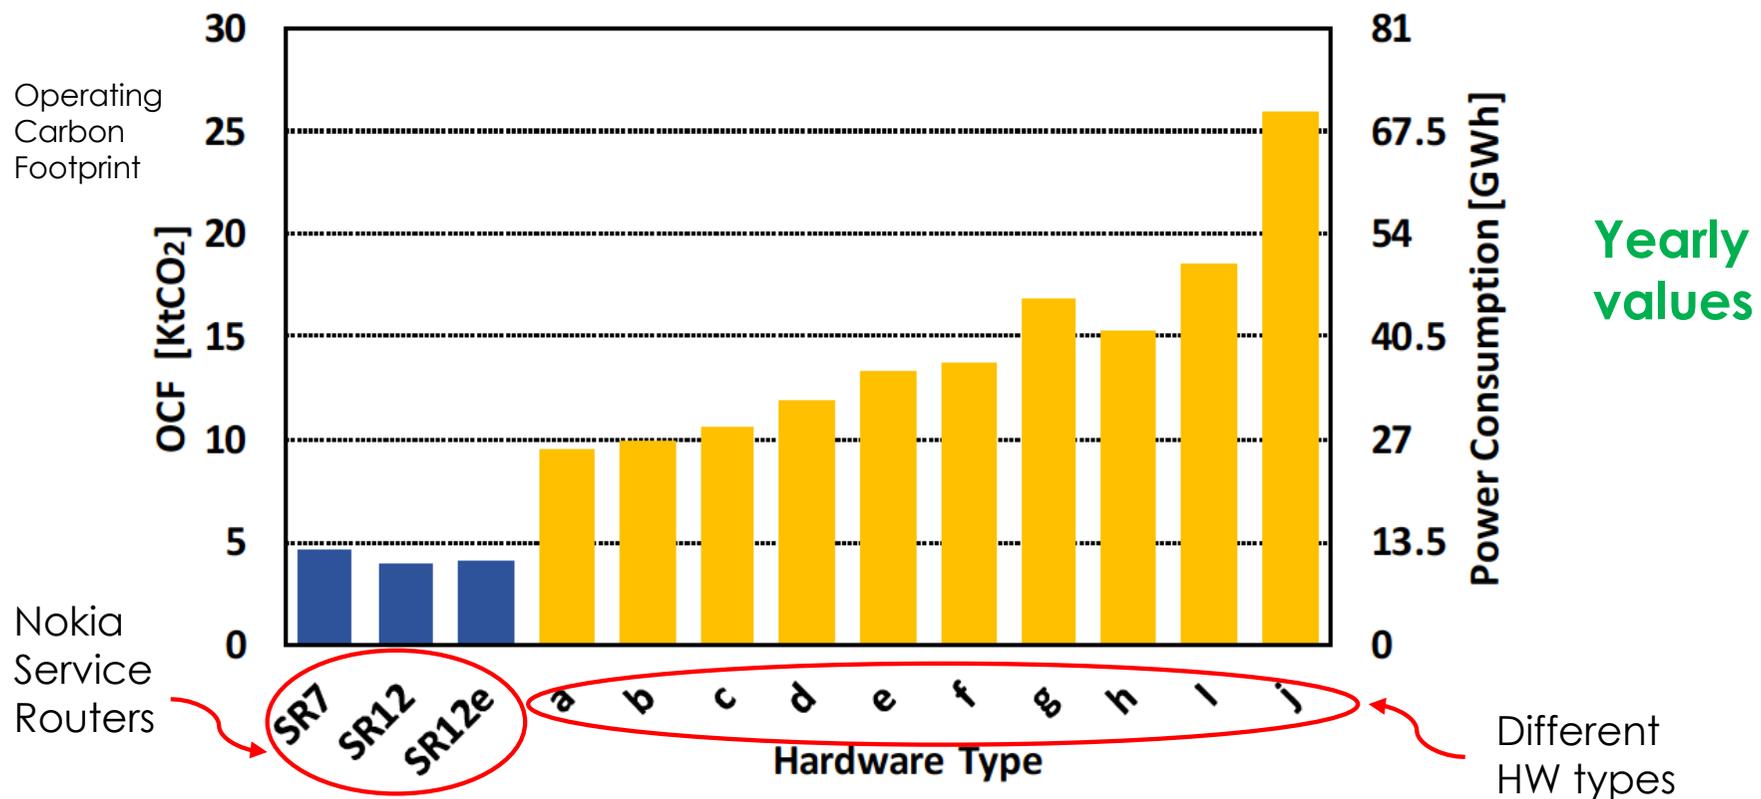
USE CASE: virtualization of a Serving Gateway (SGW)
in the Evolved Packet Core (EPC)

- Even though:
 - Appropriately downsizing the parts of SGW that cannot be completely virtualized;
 - Dynamically activating the minimum number of Virtual Machines (VMs) to support the current traffic and consolidating them to the minimum number of servers;
 - Optimizing the interconnection switches' topology and making them energy-aware;
 - Scaling the throughput of a server through Amdahl's law, to account for parallelization ...

An energy consumption

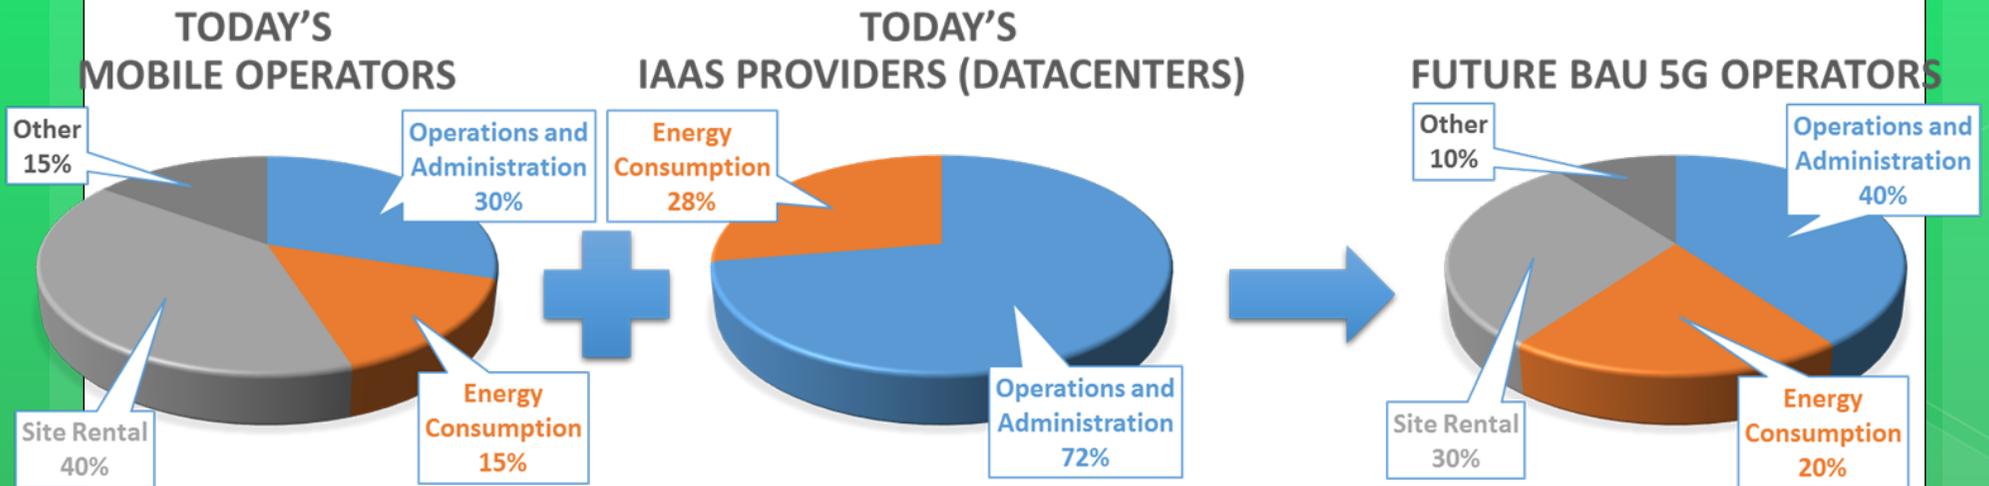
USE CASE: virtualization of a Serving Gateway (SGW) in the Evolved Packet Core (EPC)

... the power consumption of the virtualized Service Router (vSR), at the same target delay, is at least twice that of the “traditional” SR of the BAU (Business-As-Usual) solution.



Yearly values

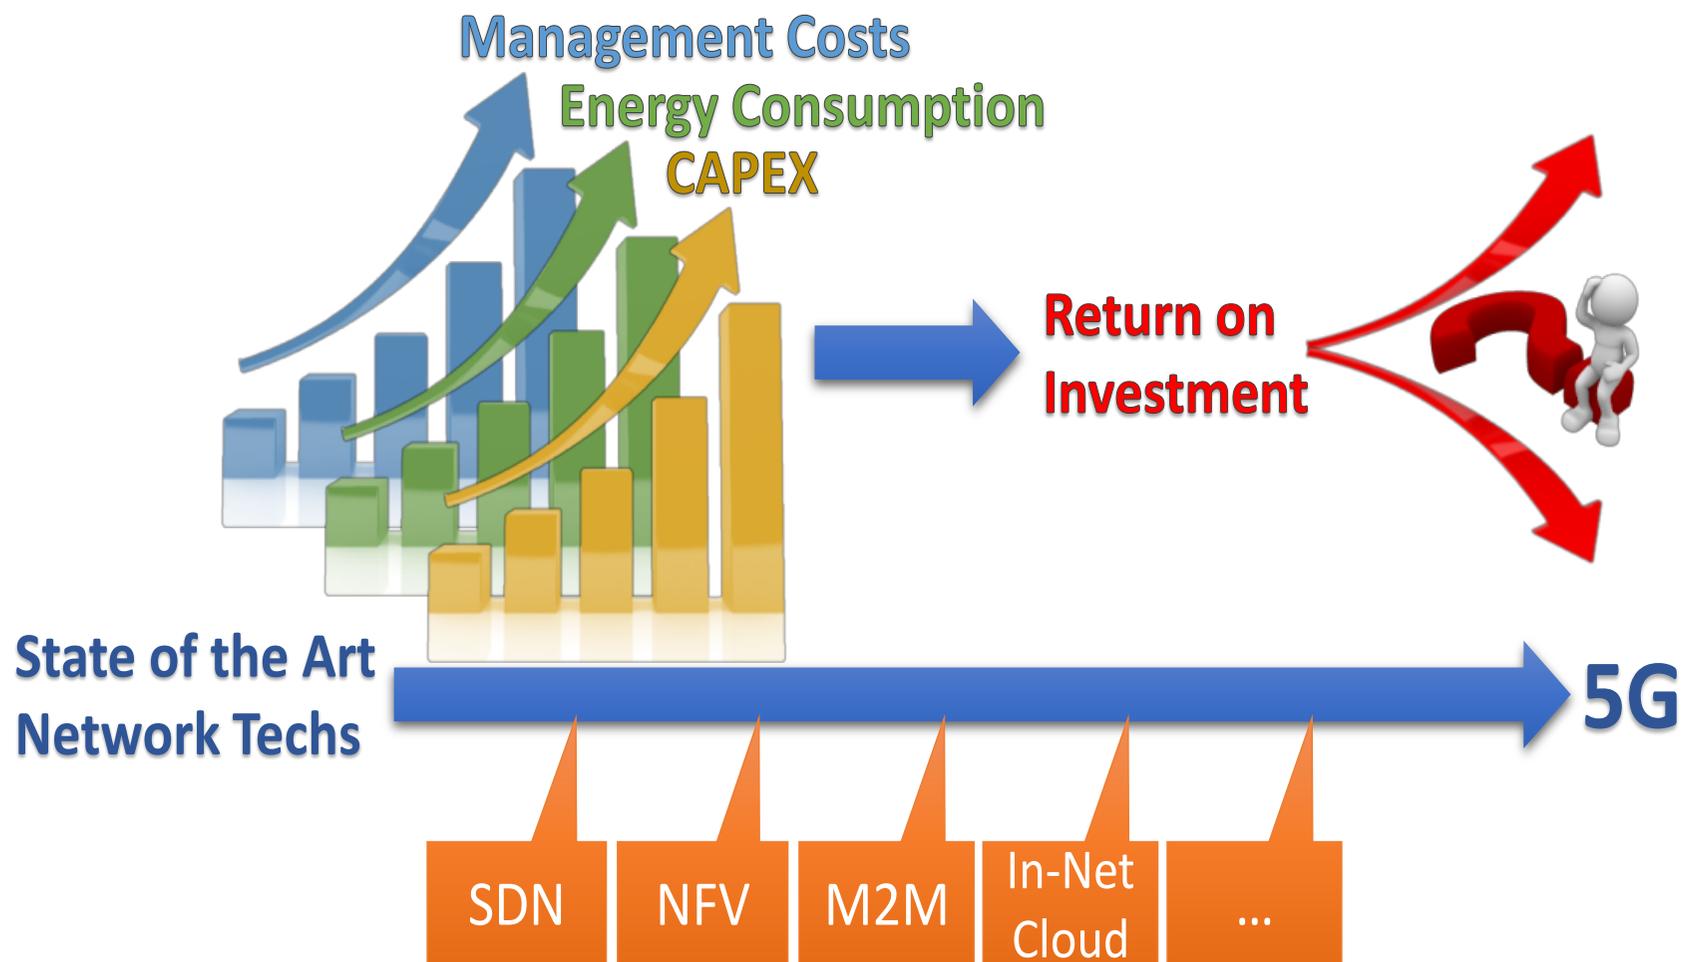
Trends



Main OPEX sources related to today's Mobile Operators (Source: Yankee Group), of IaaS (Datacenter) providers (Source: IDC, M. Bailey, K. Broderick, 2012) and projection to future Business-as-Usual 5G Operators

- The only significant OPEX sources that can be reduced by technological advancement appear to be the ones related to **energy consumption** and **network management**, which account for a figure equivalent to the entire infrastructure CAPEX.
- With the upcoming SDN/NFV technological revolution, while **CAPEX is envisaged to decrease**, owing to the adoption of commodity hardware cheaper than today's devices, the **OPEX spent in energy and in network operations is expected to significantly increase**, unless specific solutions are included in future 5G technologies.

Dilemmas



Energy Efficiency

- **Strong presence of programmable/general purpose HW inside networks and devices:**

it is the main element for introducing the flexibility necessary for Future Internet architectures development.

- **Costs & greenhouse gas emissions:**

sustainability is a must for the Future Internet deployment.

To introduce **effective but sustainable** technologies enabling “**network programmability**” (and realizing the complete integration with information technologies) is a cornerstone of the Future Internet architecture

Programmability vs Energy Efficiency

- Fixed the silicon technology, energy consumption largely depends on the number of gates in the network device/chip hardware.
- The number of gates is generally directly proportional to the flexibility and programmability levels of HW engines.
- If we fix a target number of gates by using
 - **General Purpose CPUs, we obtain:**
 - Maximum flexibility, 
 - Reduced performance/power (in the order of 100 Mbps/W) 
 - **Very specialized ASICs, we obtain:**
 - Minimum flexibility 
 - Greatly enhanced performance/power (in the order of 1 Gbps/W) 
 - *Other technologies (e.g., network/packet processors) provide performance between these boundaries*

Programmability today is energy consuming!

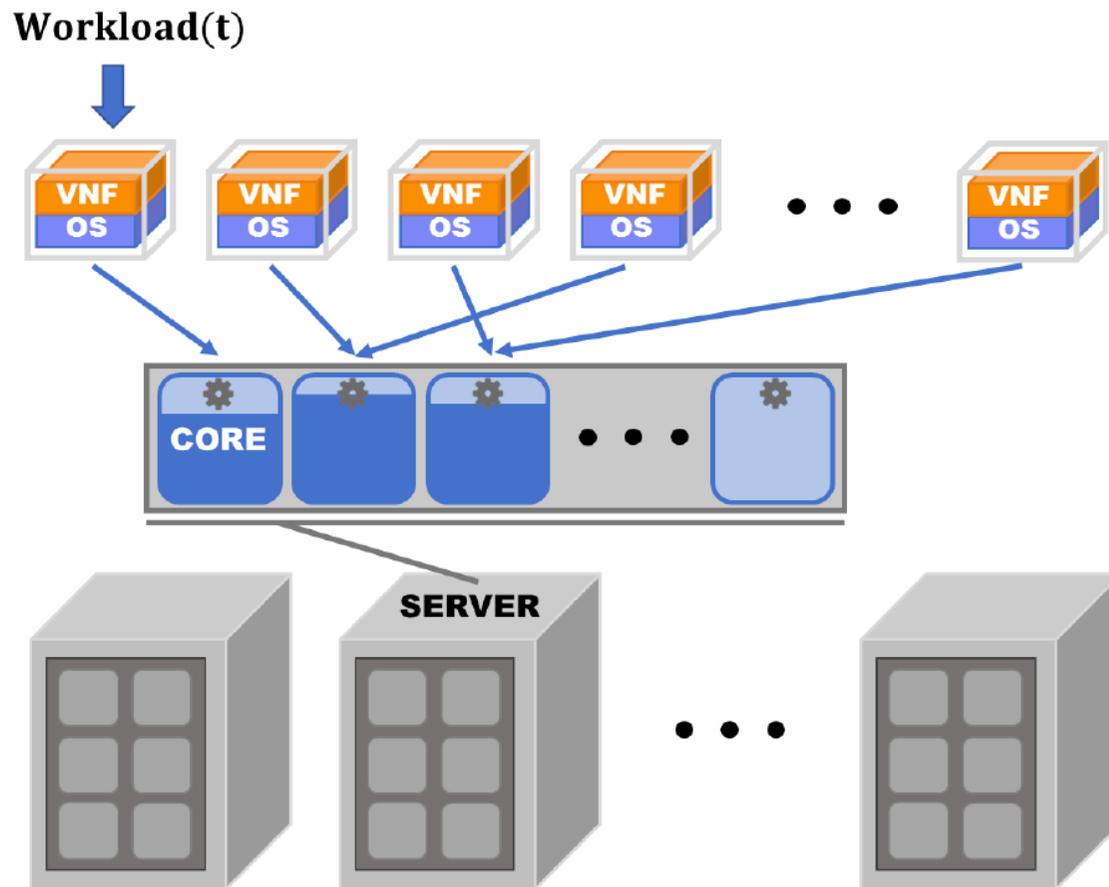
Programmability for Energy Efficiency

- Three basic enablers at chip/system level:
 - Dynamically programmable resources able to perform multi-purpose services;
 - Specialized HW for offloading to speed-up basic functionalities;
 - Standby capabilities to save energy if a resource is unused.
- The presence of general-purpose HW offers the possibility of moving services among the components of a node, or among nodes in a network.
- When the workload is low, many services can aggressively share single general-purpose HW resources.
- So, even if a general-purpose/programmable resource consumes more energy than ASIC-based solutions, a smaller number of HW elements can be left active, in order to effectively handle the current workload.

Programmability for Energy Efficiency

- Main issues:
 - which basic (sub-)functionalities need to be moved (and “frozen”) to the offloaded specialized engines (best performance in terms of bps/W)?
 - which ones have to remain in the programmability space (lower performance but stronger sharing and more evolution opportunities)?
- The solutions need to be identified by considering and effectively supporting the newest trends in Internet technology evolutions.

Back to modeling



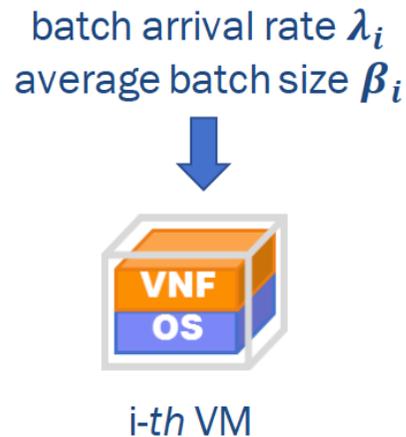
VNF Consolidation

- One-to-one correspondence between VNFs & VMs
- ACPI-enabled multi-core processors
- First-Fit-Decreasing (FFD) bin-packing
- Initial VM placement based on SLA-specified workflows
- Dynamic re-consolidation within each server

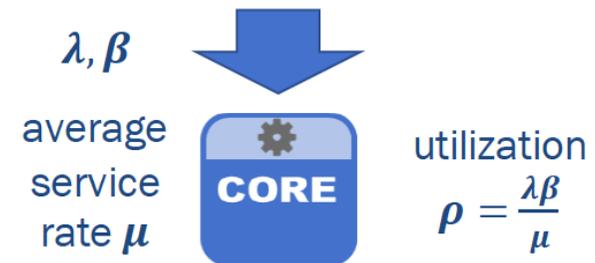
Source: R. Bruschi, F. Davoli, P. Lago, J. F. Pajo, "Joint power scaling of processing resources and consolidation of virtual network functions", Proc. 5th IEEE Int. Conf. on Cloud Networking (CloudNet), Pisa, Italy, Oct. 2016.

Back to modeling

An $M^X/G/1/SET$ queueing model can be effectively used to represent the energy consumption, when considering the aggregate workload produced by all VMs insisting on the core.

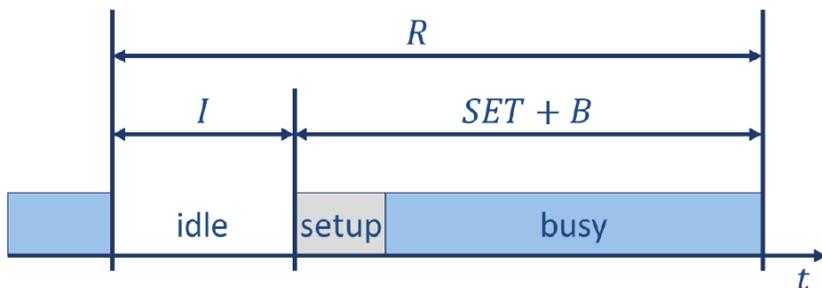


set J of VMs served by a core



Two ways to **enforce performance constraints**:

- *Roughly*, by imposing a limit on the maximum utilization for each core;
- *More precisely*, by computing the average system latency on the basis of the model. This requires



the knowledge of the second moments of the batch size (obtainable from measurements of the second moment of the busy period) and of the packet service time (directly measurable).

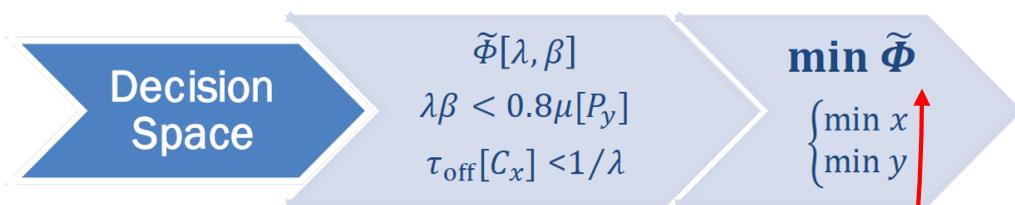
Back to modeling

Based on the model, decision rules in the $(\lambda, \lambda\beta)$ space can be defined

Aggregate batch arrival rate

Average batch size

to design an energy- and performance-aware consolidation policy.

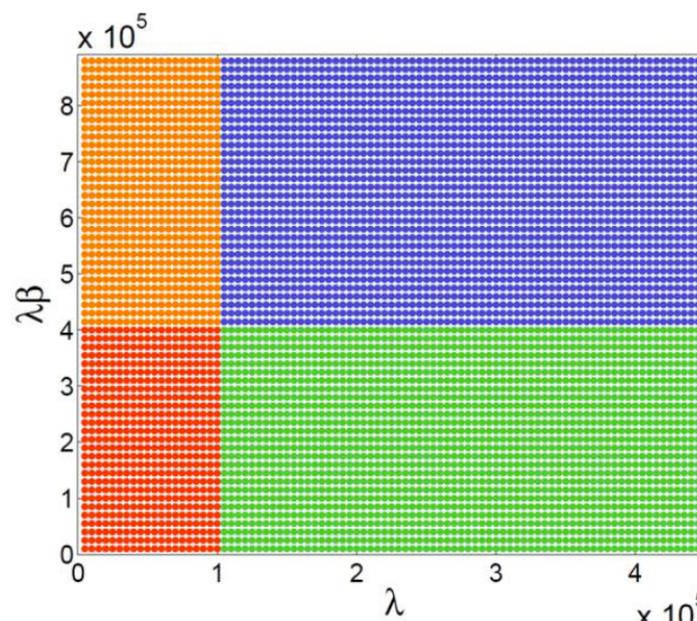


can be generated offline for a wide range of workloads and ACPI configurations

- Data center performs FFD bin-packing consolidation at every interval T
- In-server re-consolidation performed at every ΔT sub-interval within T

Average power consumption

- (C_1, P_0)
- (C_3, P_0)
- (C_1, P_8)
- (C_3, P_8)



Core workload classification

Back to modeling

A numerical example

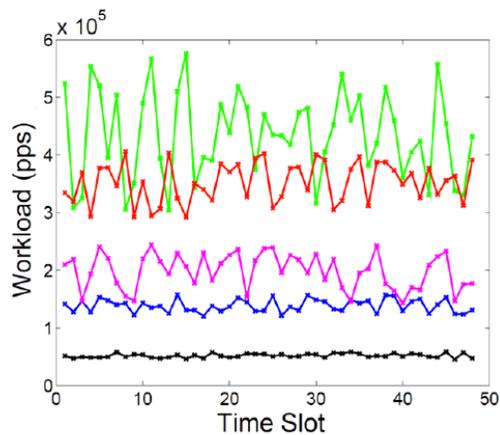
System:

- i. 500 servers, 1000 VMs
- ii. 2 octa-core processors per server
- iii. uniform distributions are used to vary VM workloads and burstiness

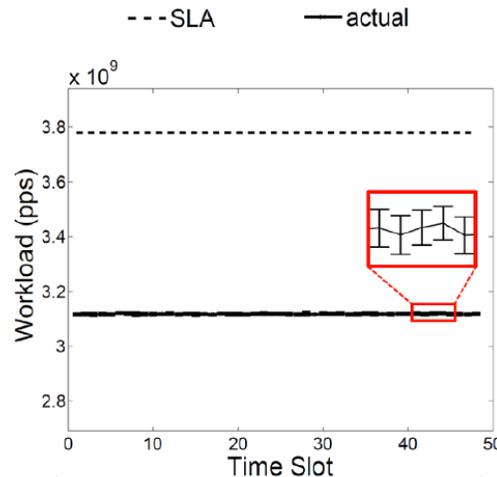
Evaluation:

1. despite VM workload variations, total system workload is relatively stable at **approx. 18% below SLA specs**
2. 2 baseline scenarios:
 - **xC+xPS** – no in-server consolidation, no power scaling
 - **C+xPS** – with in-server consolidation, no power scaling
3. potential annual savings of **C+PS: 19,000€ (0.12€/kWh)**

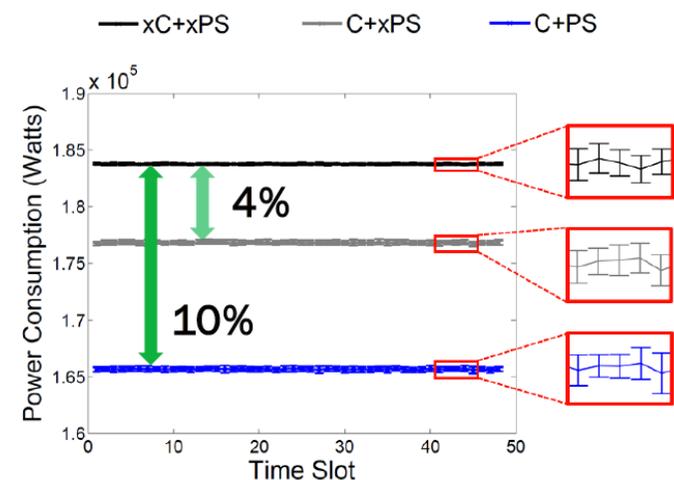
VM₁ VM₂ VM₃ VM₄ VM₅



VM workload variations

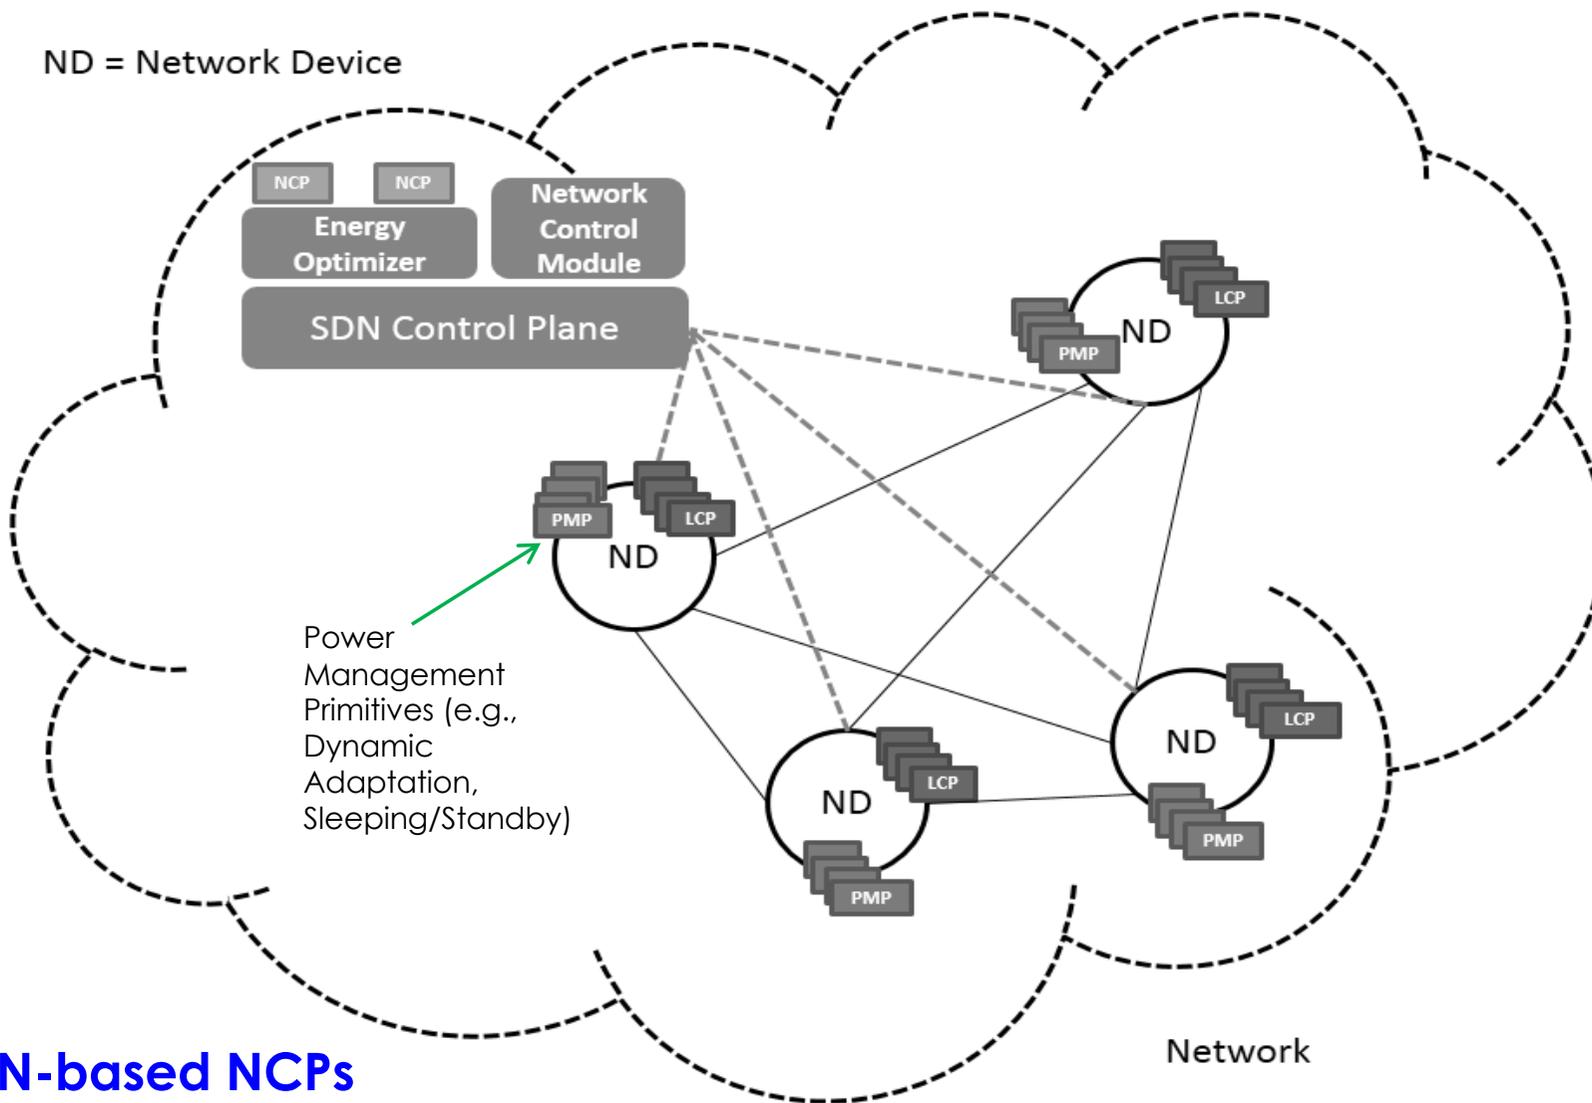


Datacenter workload



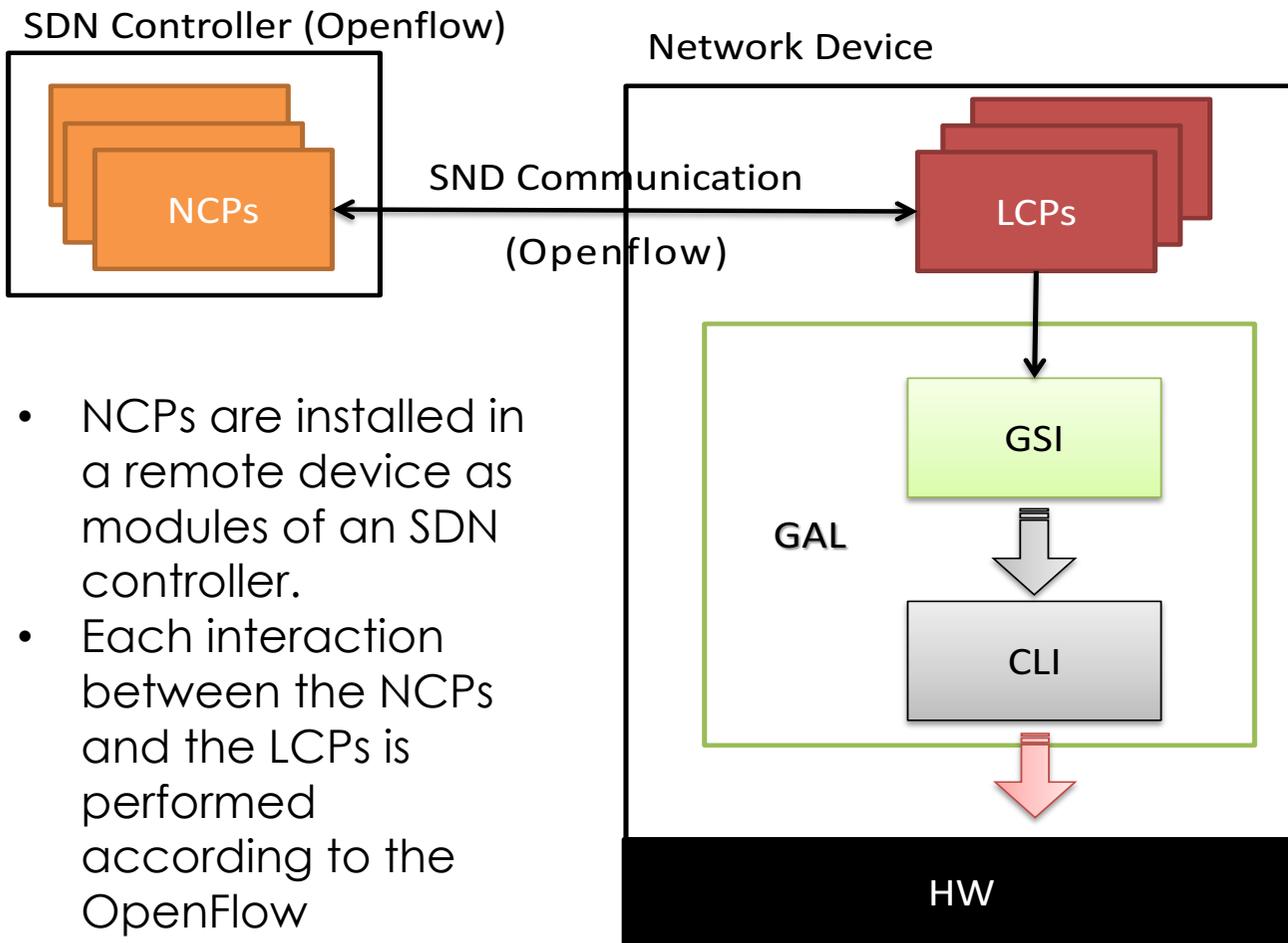
Datacenter power consumption

Back to SDN/NFV and the GAL



SDN-based NCPs

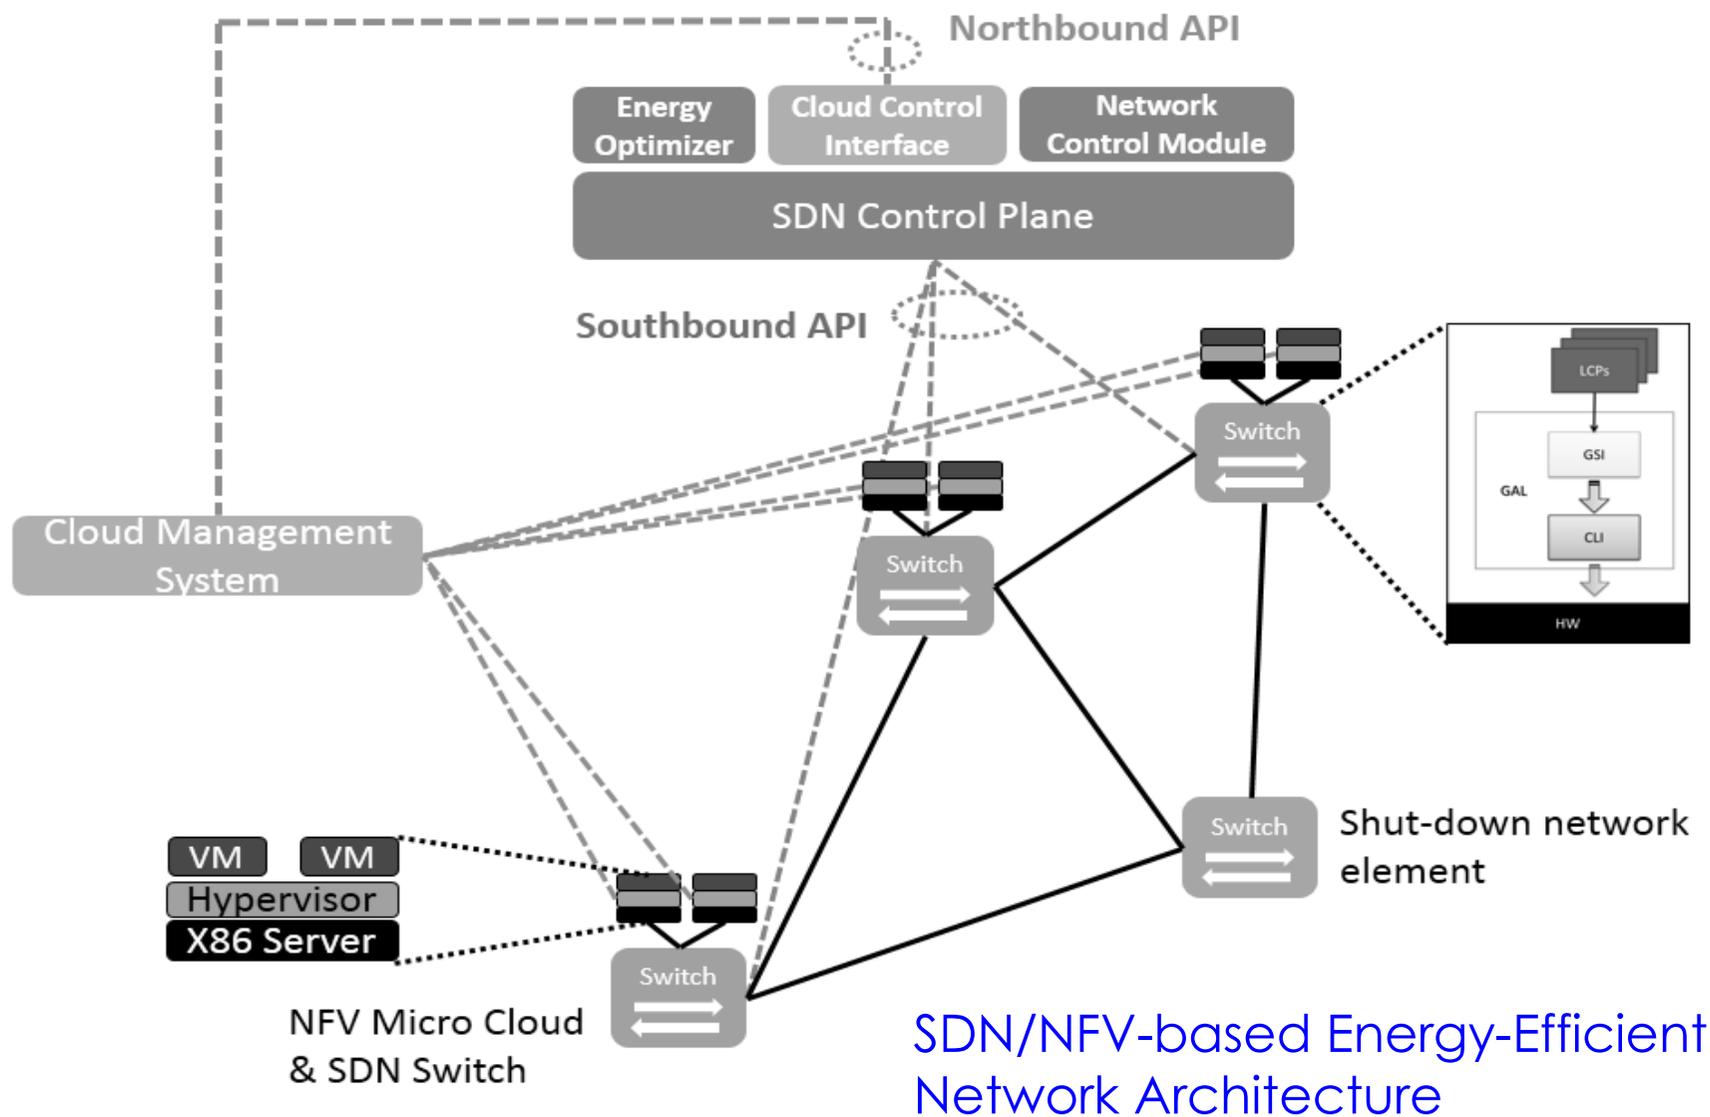
Back to SDN/NFV and the GAL



- NCPs are installed in a remote device as modules of an SDN controller.
- Each interaction between the NCPs and the LCPs is performed according to the OpenFlow Specification.

- LCPs set and get the energy-aware configuration by means of Energy Aware States (EASes) and by using the GSI.
- Inside the GAL framework, each GSI request is translated by the Convergence Layer Interface (CLI) into a specific command for the underlying HW components.

Back to SDN/NFV and the GAL



SDN/NFV-based Energy-Efficient Network Architecture

GALv2: how to extend the concept to NFV?

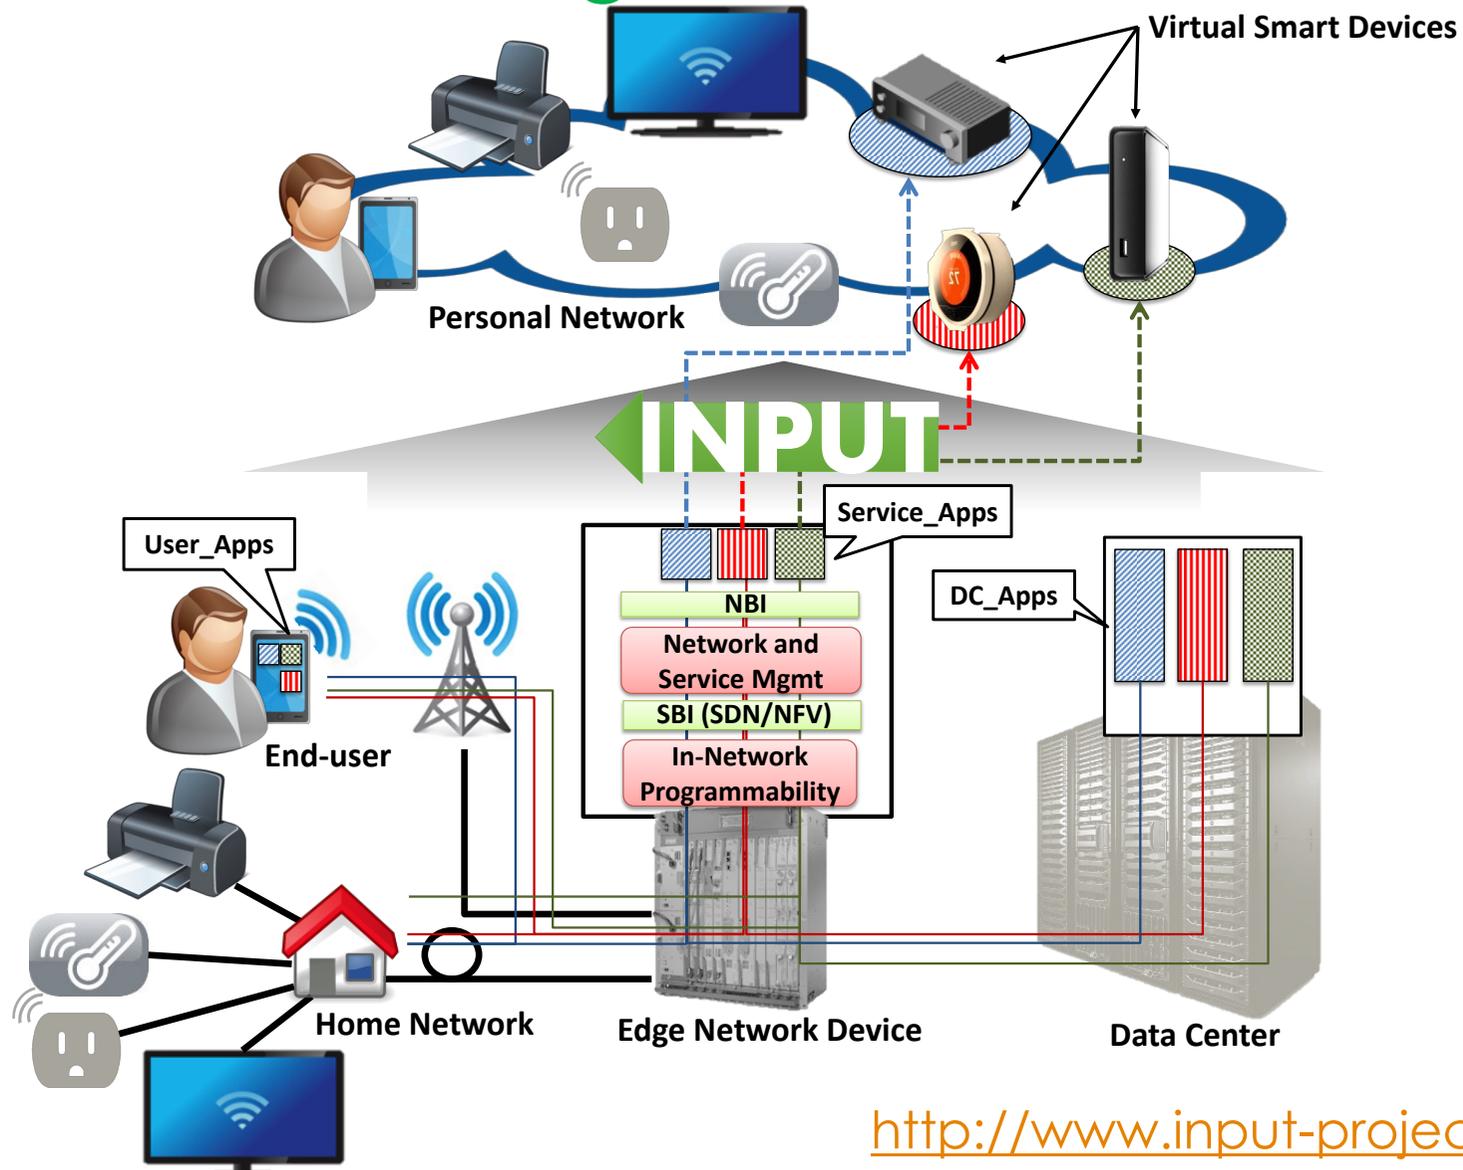
- Main challenge: the correspondence between the HW that consumes energy (and belongs to the **Infrastructure Provider**) and the virtualized objects (VMs, containers, ...) that execute Network Functions (and belong to the **Network Service Provider**) is not so straightforward as in the legacy environment:
 - Execution mediated by the hypervisor and its scheduling policies;
 - Multiple tenants;
 - No “direct” relation between virtual resources (e.g., vCPU) and *Energy Aware States* of the HW.

GALv2: how to extend the concept to NFV?

- Possible lines of action

- Queueing models for aggregated traffic only (per server/core);
- Simpler aggregated models for HW energy consumption (e.g., Generalized Amdhal's Law), more detailed queueing models for execution machines;
- Use of “virtualized” Energy Aware States as backpressure from the Infrastructure Provider to create incentives toward tenants to become energy-aware.

A glimpse on projects - INPUT: Offloading the Cloud



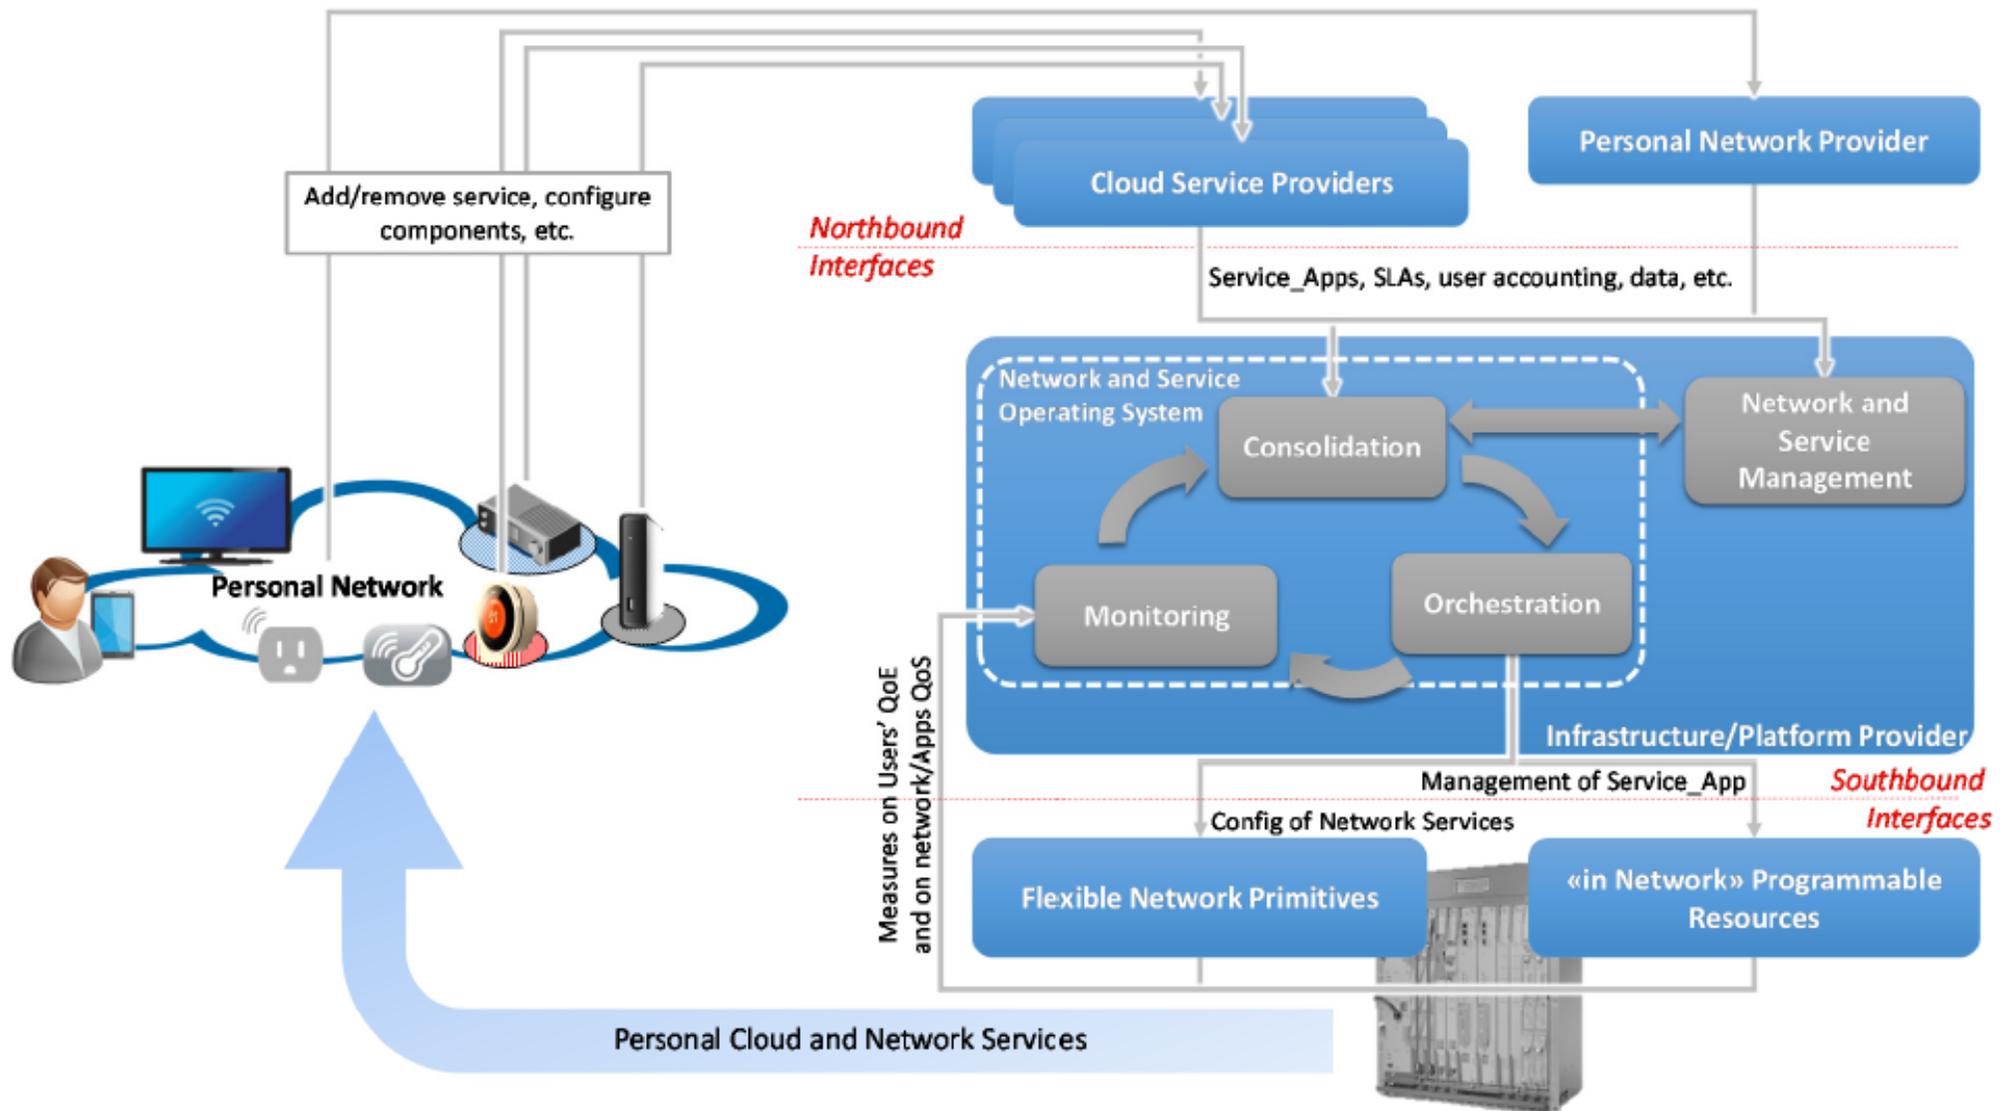
<http://www.input-project.eu>

A glimpse on projects -

INPUT: Offloading the Cloud

- Personal Cloud Services
- Replace physical Smart Devices (SDs) with their Virtual Images, providing them to users “as a Service” (SD as a Service – SDaaS)...
- ...by means of virtual cloud-powered Personal Networks
- ...introducing computing and storage capabilities to edge network devices (“**in-network programmability**”), to allow users/telecom operators create/manage private clouds “in the network”
- ...moving cloud services closer to end-users and smart devices, in order both to avoid pointless network infrastructure and datacenter overloading, and to provide lower latency reactivity to services.

INPUT: Management, Consolidation and Abstraction framework



MATILDA:

A Holistic, Innovative Framework for the Design, Development and Orchestration of 5G-Ready Applications and Network Services over Sliced Programmable Infrastructure



MATILDA

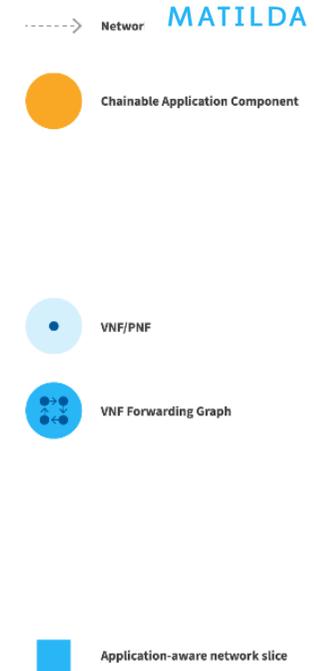
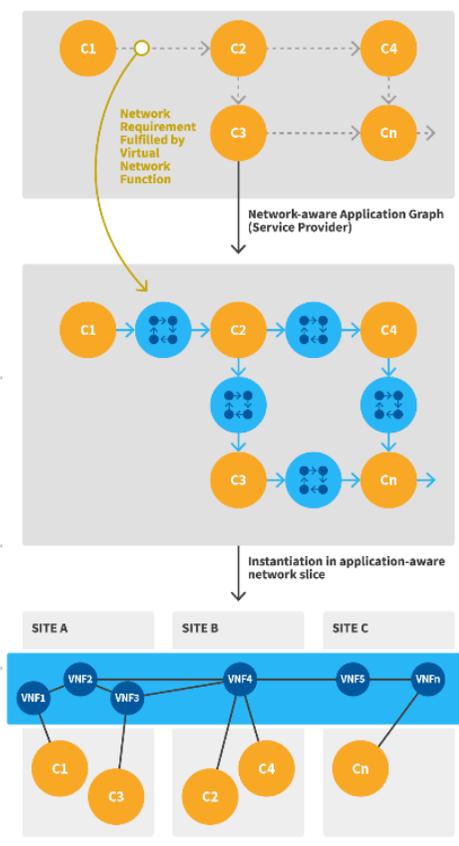
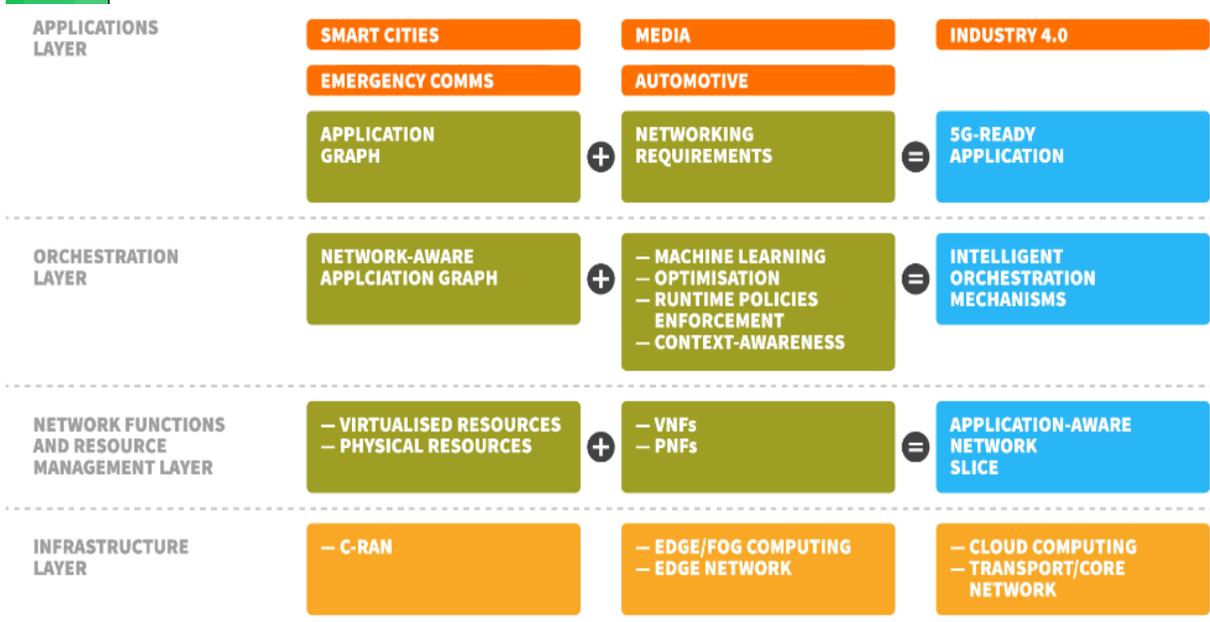
- Define appropriate **abstractions for the design of 5G-ready applications**;
- Develop an **agile programming and verification platform** for developing and verifying **industry vertical 5G-ready applications** and **network services**;
- Support mechanisms for **automated or semi-automated translation of application-specific** requirements to **programmable infrastructure** requirements;
- Support **intelligent orchestration mechanisms** for managing the entire lifecycle of 5G-ready applications and network services;
- Support mechanisms for **multi-site network, compute and storage** resource management;
- Involve **key actors of the value chain** in the operational model.



MATILDA

MATILDA 5G

High Level Architectural Approach

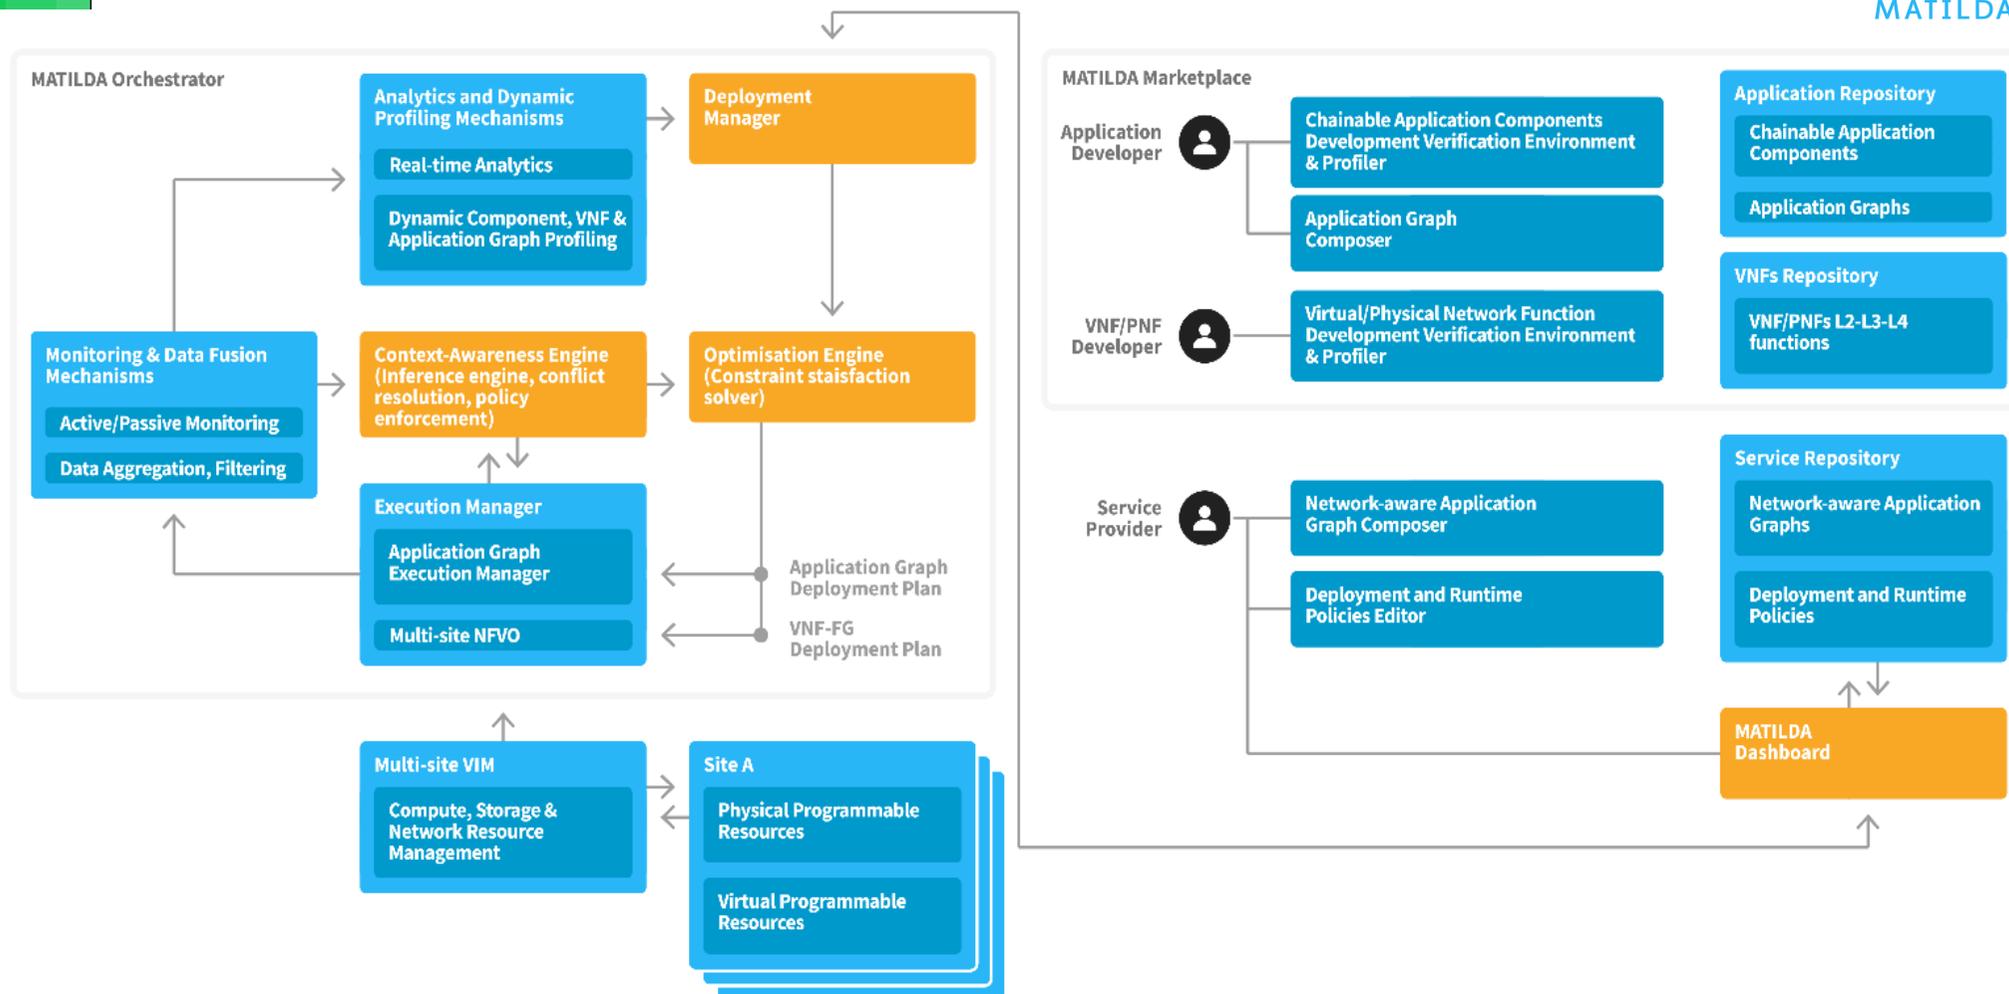




MATILDA

MATILDA

Conceptual Architecture



Conclusions

- Combining SDN, NFV and *energy-aware performance optimization* can shape the evolution of the Future Internet and contribute to CAPEX and OPEX reduction for network operators and ISPs.
- Many of the concepts behind this evolution are not new and ideas have been around in many different forms; however, current advances in technology make them feasible.
- Sophisticated control/management techniques can be realistically deployed – *both* at the network edge and inside the network – to dynamically shape the allocation of resources and relocate applications and network functionalities, trading off QoS/QoE and energy at multiple granularity levels.
- Modeling and analytics play a key role to enable this scenario.
- Network programmability and virtualization enable extensive use of emulation to test solutions in a controlled environment.