22nd ITC Specialist Seminar
on Energy Efficient and Green Networking (SSEEGN 2013)
November 20-22, 2013, Christchurch, New Zealand

2013 ATNAC
AUSTRALASIAN TELECOMMUNICATION NETWORKS & APPLICATIONS CONFERENCE

Green Networking and Network Programmability: A Paradigm for the Future Internet?

Franco Davoli
DITEN-University of Genoa, Italy
CNIT – University of Genoa Research Unit

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

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
- Taxonomy of Green Networking Approaches
  - Dynamic Adaptation
  - Smart Sleeping

Outline

- Dynamic Adaptation I – Link Protocols
  - Link Control: the Green Ethernet
    - IEEE 802.3az
- Dynamic Adaptation II – Packet Processing Engines
  - Idle Logic
  - Power Scaling
- Modeling and optimization
  - Modeling Line Card Queues
  - Modeling Green Ethernet
  - Power/Performance Trade-off
- Standby
  - Proxying the network presence
- Network-level optimization

Outline

- Implementing controls: The Green Abstraction Layer (GAL) approach
- SDN/NFV and the GAL
- Examples of virtualized functions: NCP to LCP to GAL, L2 virtualization, DROP router
- Conclusions

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 (Forwarding) Plane.

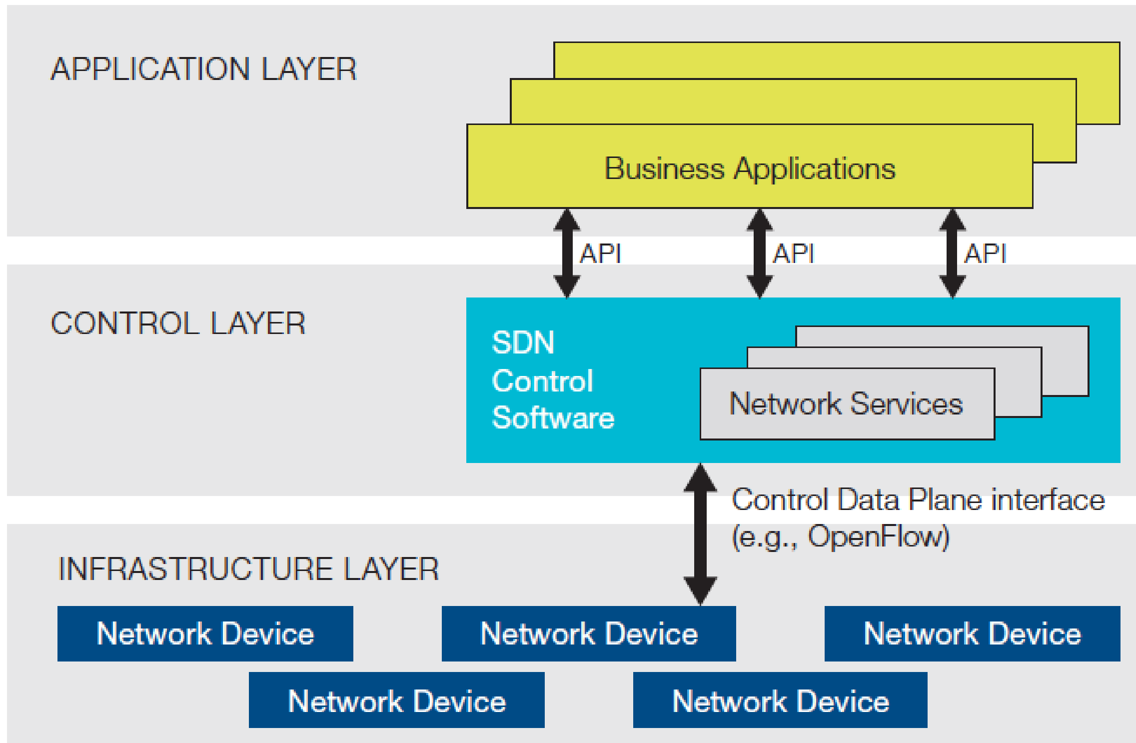

Source: *Software-Defined Networking: The New Norm for Networks, Open Networking Foundation (ONF) White Paper,* *April 2012*.

Flexibility/Programmability – **OpenFlow**

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

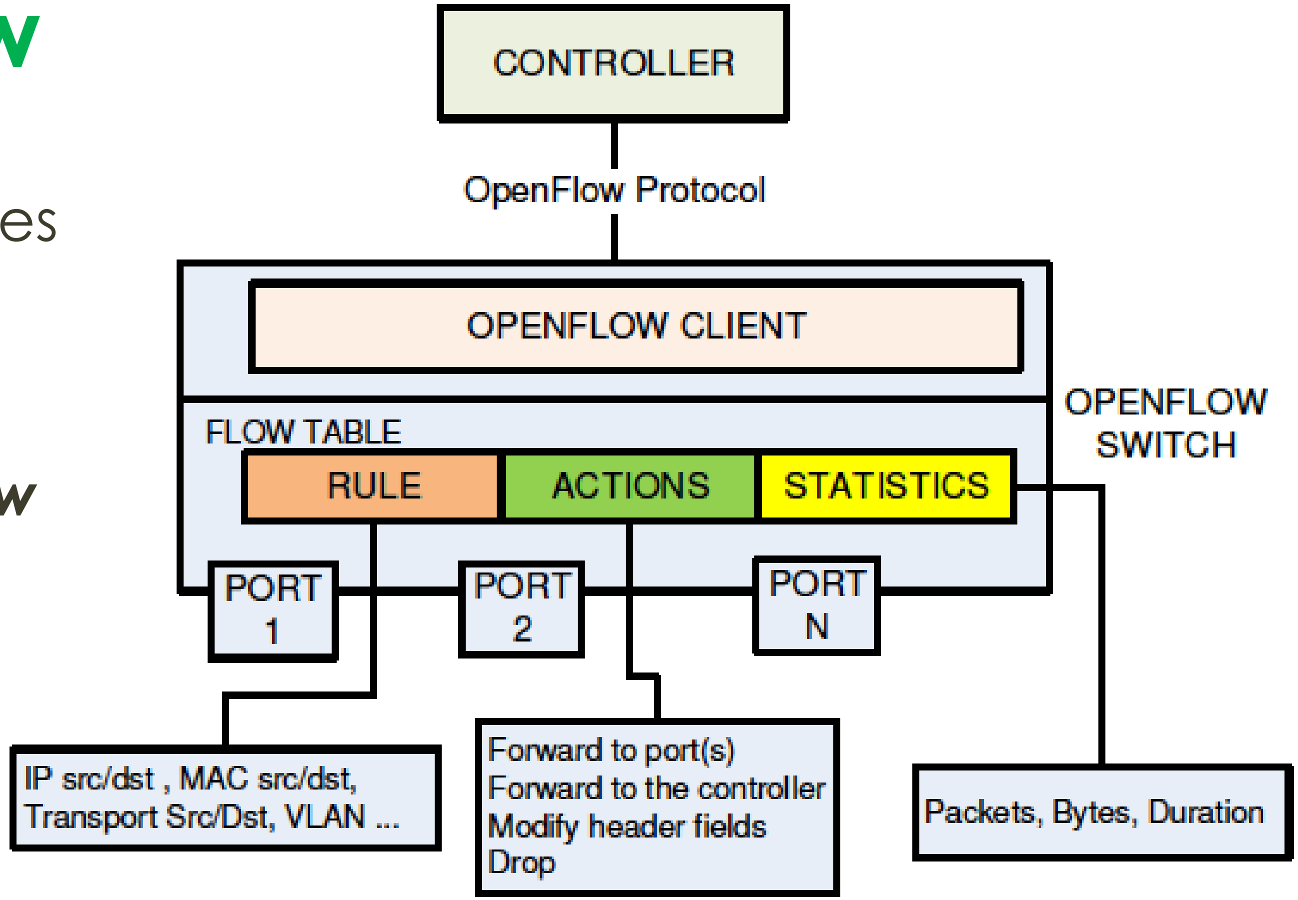

Source: B. A. A. Nunes, M. Mendonça, X.-N. Nguyen, K. Obraczka, T. Turletti, "A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks", Oct. 2013, in submission; http://hal.inria.fr/hal-00825087.

Flexibility/Programmability – Network Functions Virtualization (**NFV**)

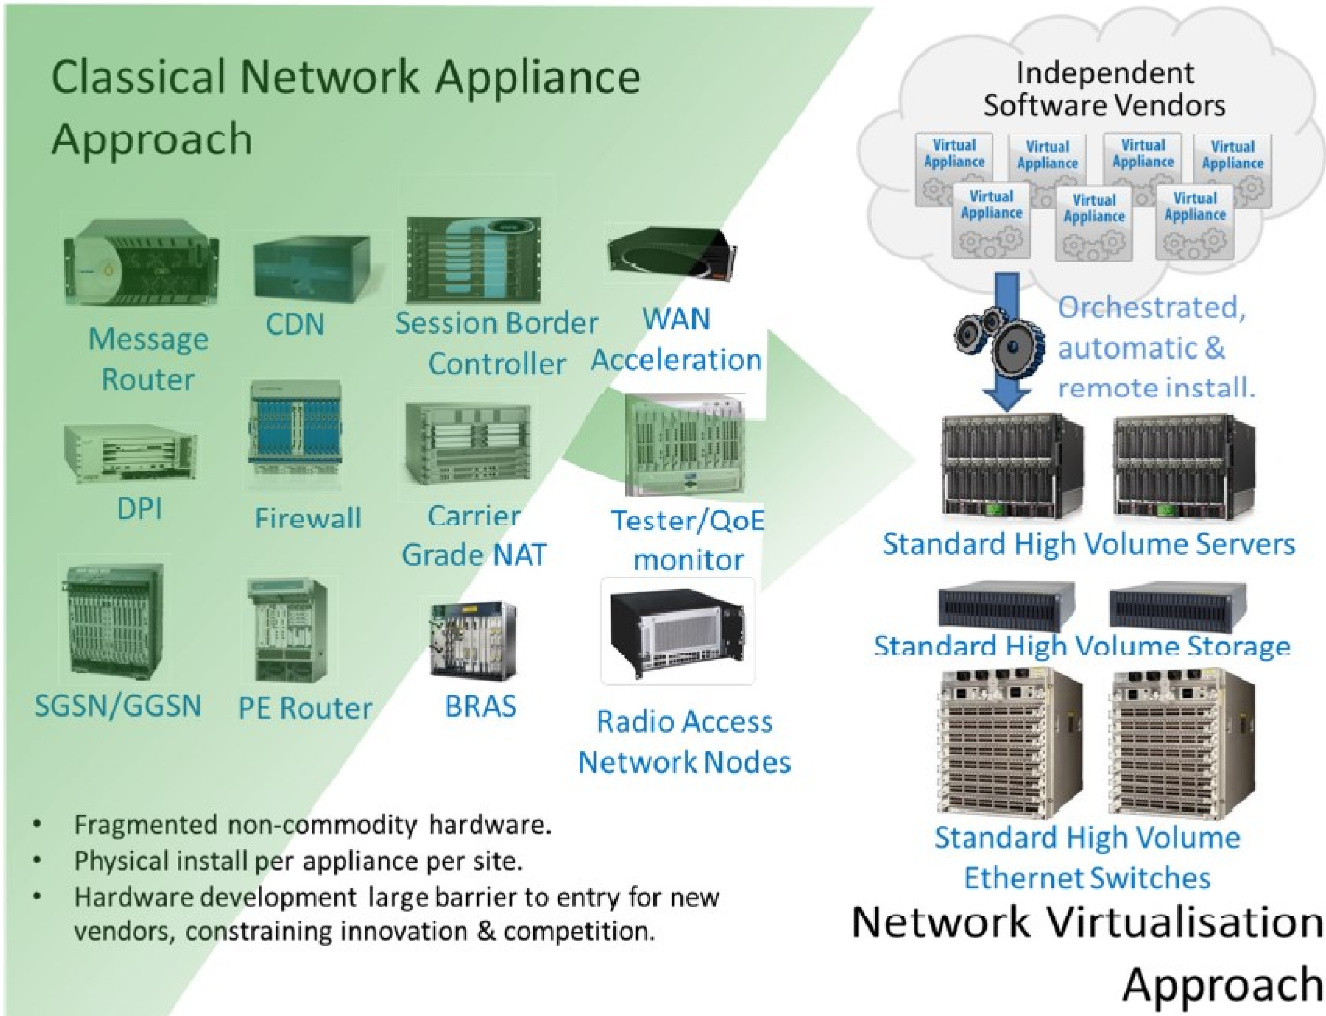

Leverages "…standard IT virtualisation technology to consolidate many network equipment types onto industry standard high volume servers, switches and storage, which could be located in Datacentres, Network Nodes and in the end user premises."

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 – *approx. 9.5 M servers shipped in 2011 against approx. 1.5 M routers*

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

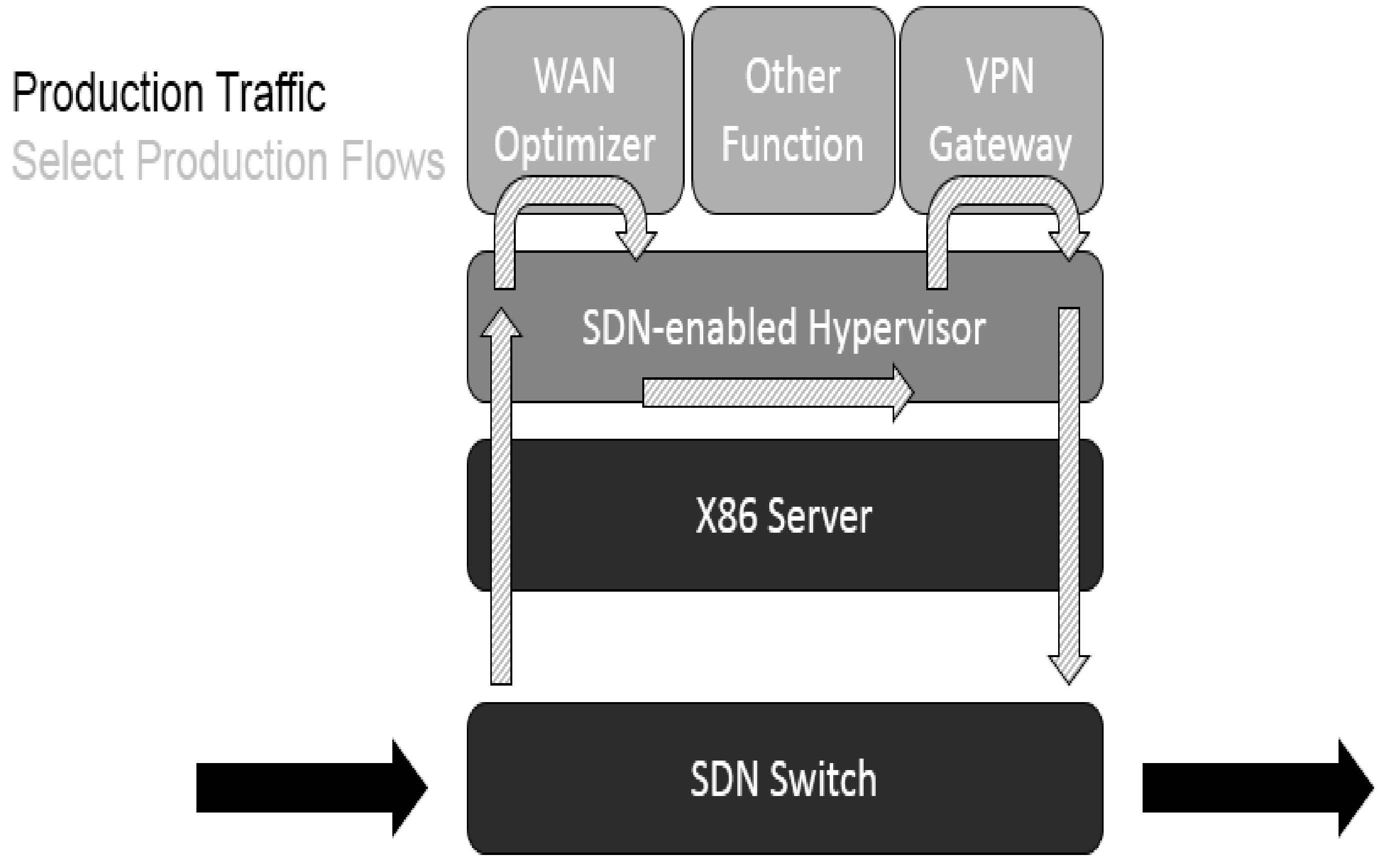

Source: *M. Jarschel, T. Hoßfeld, F. Davoli, R. Bolla, R. Bruschi, A. Carrega, "SDN-Enabled Energy-Efficient Network Management", to appear in K. Samdanis, P. Rost, A. Maeder, M. Meo, C. Verikoukis, Eds., Green Communications Book, Wiley, 2014.*

Integrated managament and control for Traffic Engineering

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

    Quasi-centralized / hierarchical vs. 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 CO_2 emission;
  - the **economic** one, which stems from the reduction of operating costs (OPEX) of ICT services.

The Carbon Footprint of ICT

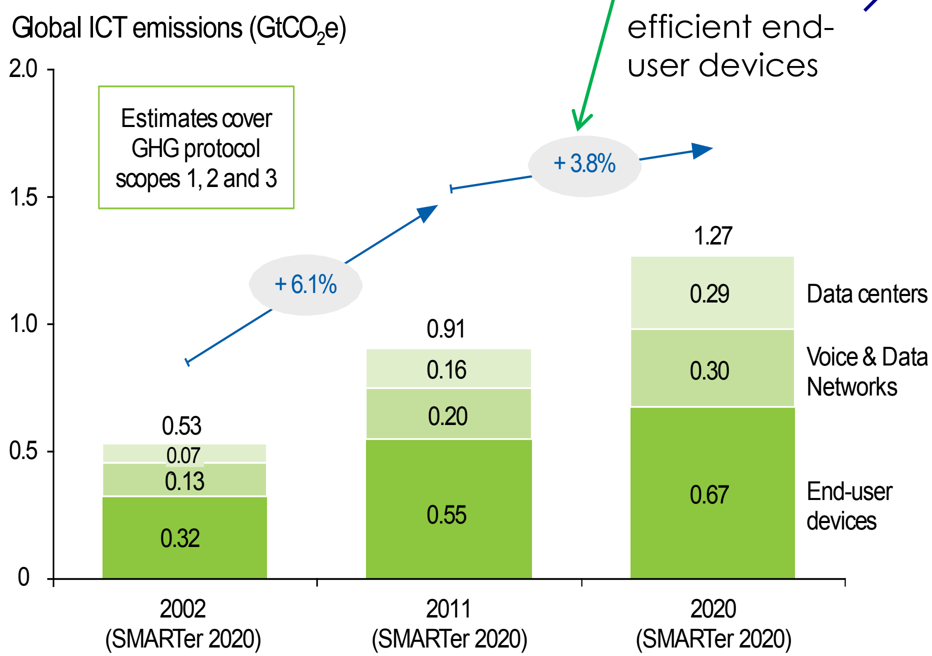

Decrease mainly due to more efficient end-user devices

	2002 (SMARTer 2020)	2011 (SMARTer 2020)	2020 (SMARTer 2020)
ICT as percentage of global emissions	1.3	1.9	2.3
Global GHG emissions ($GtCO_2e$)	40	48[1]	55

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.

	CAGR % 2002-2011	CAGR % 2011-2020
Data centers	8.6	7.1
Voice & Data Networks	4.7	4.6
End-user devices	6.1	2.3

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.*

Long-Term Sustainability

- The sole introduction of novel low consumption HW technologies cannot clearly cope with increasing traffic and router capacity trends, and be enough for drawing ahead current network equipment towards a **greener** and **sustainable Future Internet**.

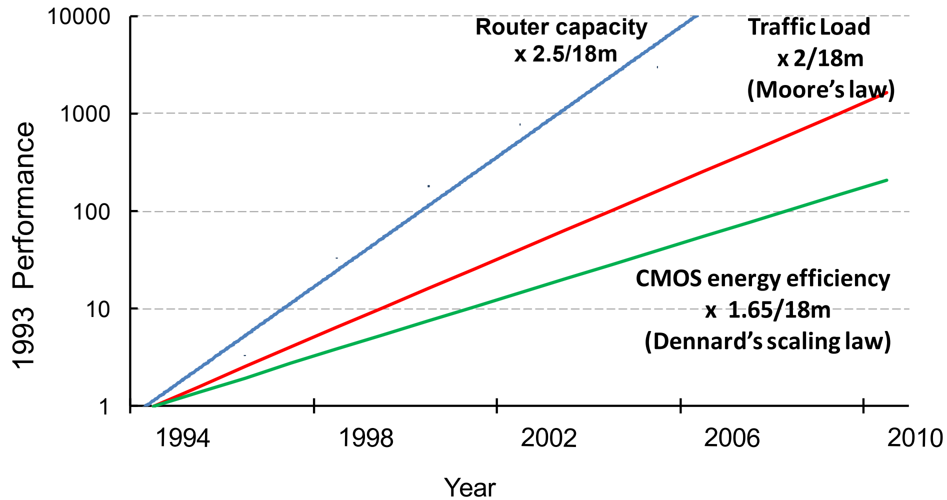

Evolution from 1993 to 2010 of high-end IP routers' capacity (per rack) vs. traffic volumes (Moore's law) and energy efficiency in silicon technologies.

Source: Neilson, D.T., "Photonics for switching and routing," *IEEE Journal of Selected Topics in Quantum Electronics (JSTQE)*, vol. 12, no. 4, pp. 669-678, July-Aug. 2006.

Reasons for energy inefficiencies…

- The origin of these trends can be certainly found in current **Internet infrastructures, technologies and protocols**, which **are designed to be extremely over-dimensioned and available 24/7**.
- Links and devices are provisioned for rush hour load.
- The overall power consumption in today's networks remains more or less constant even in the presence of fluctuating traffic loads.

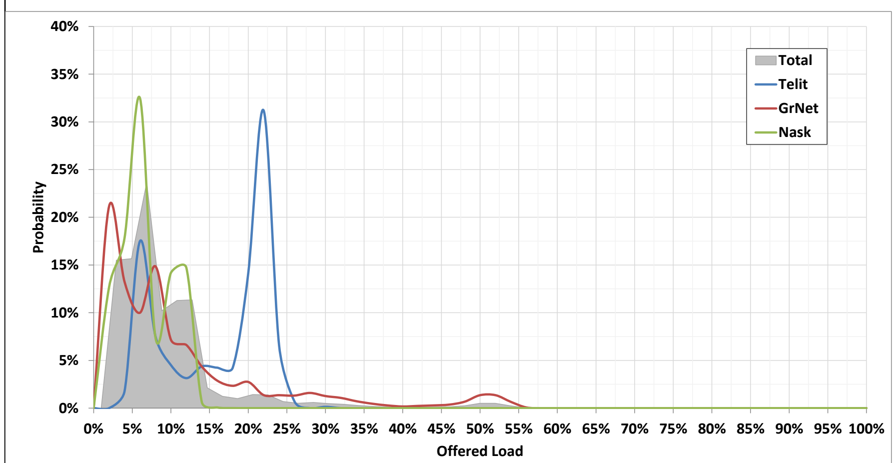

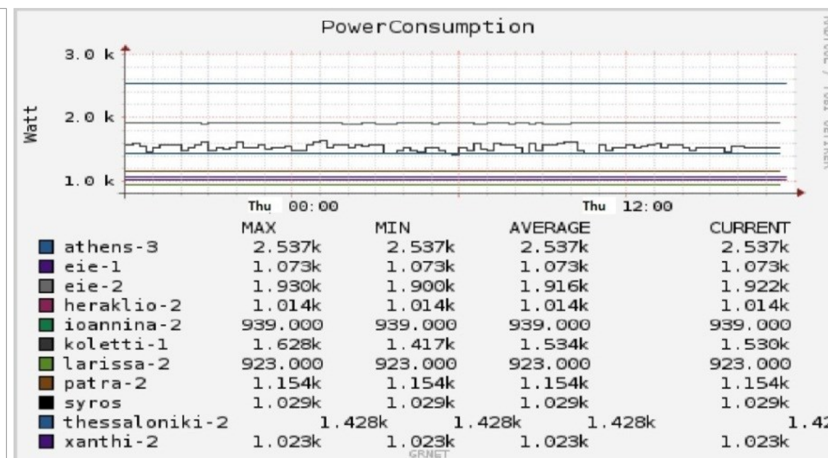

…despite wide traffic variations

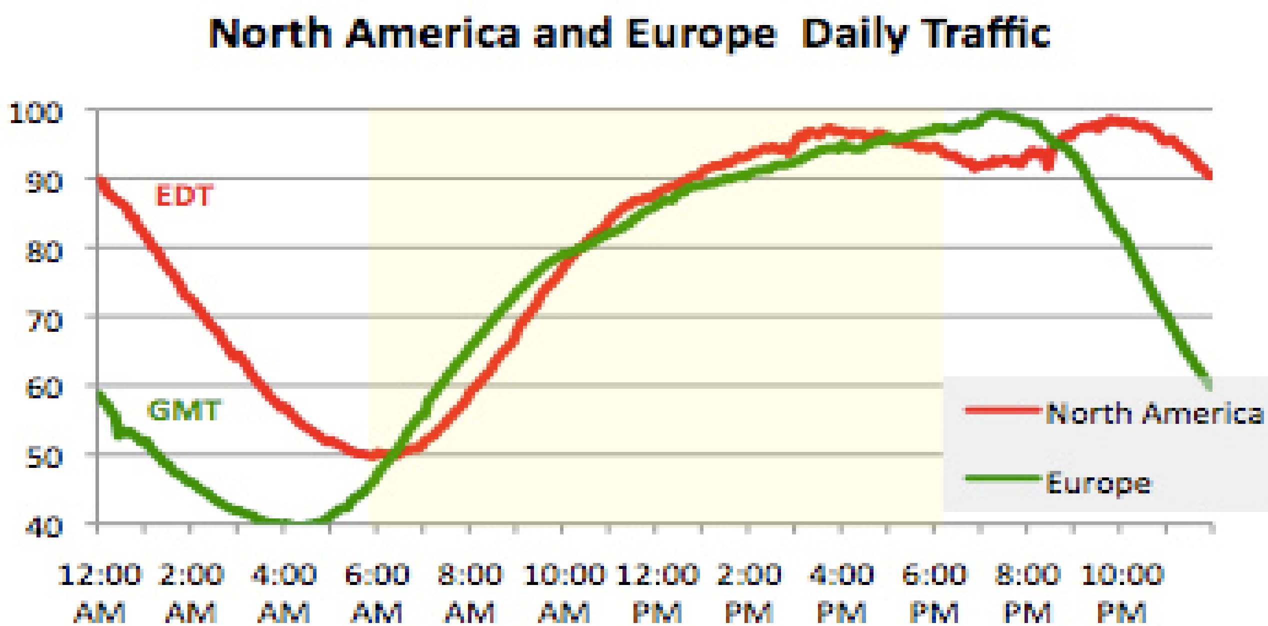

Percentage w.r.t. peak level.
The profiles exhibit regular, daily cyclical traffic patterns with Internet traffic dropping at night and growing during the day.

Traffic load fluctuation at peering links for about 40 ISPs from USA and Europe

Source: http://asert.arbornetworks.com/2009/08/what-europeans-do-at-night/

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

Decomposing the Energy Consumption

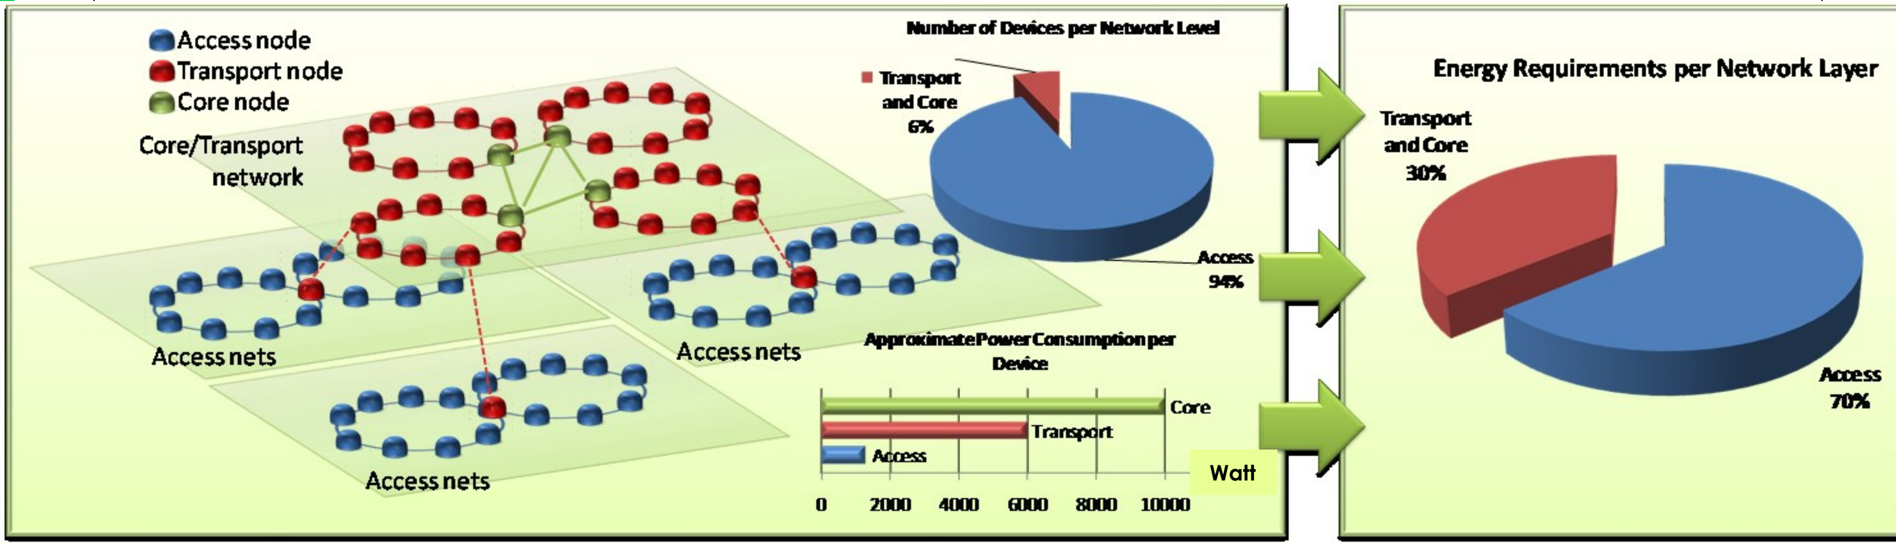

Typical access, metro and core device density and energy requirements in today's typical networks deployed by telcos, and ensuing overall energy requirements of access and metro/core networks.

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.

Taxonomy of Approaches

Re-Engineering
- Energy-efficient Physical elements
- Complexity Reduction

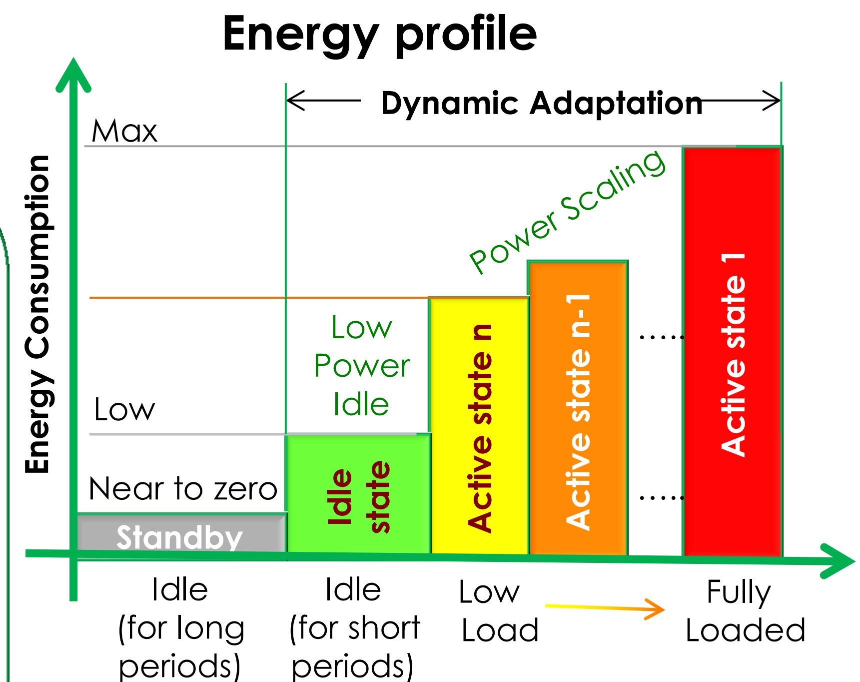

Dynamic Adaptation
QoS vs Power Management

- 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.

Dynamic Adaptation

QoS vs Power Management

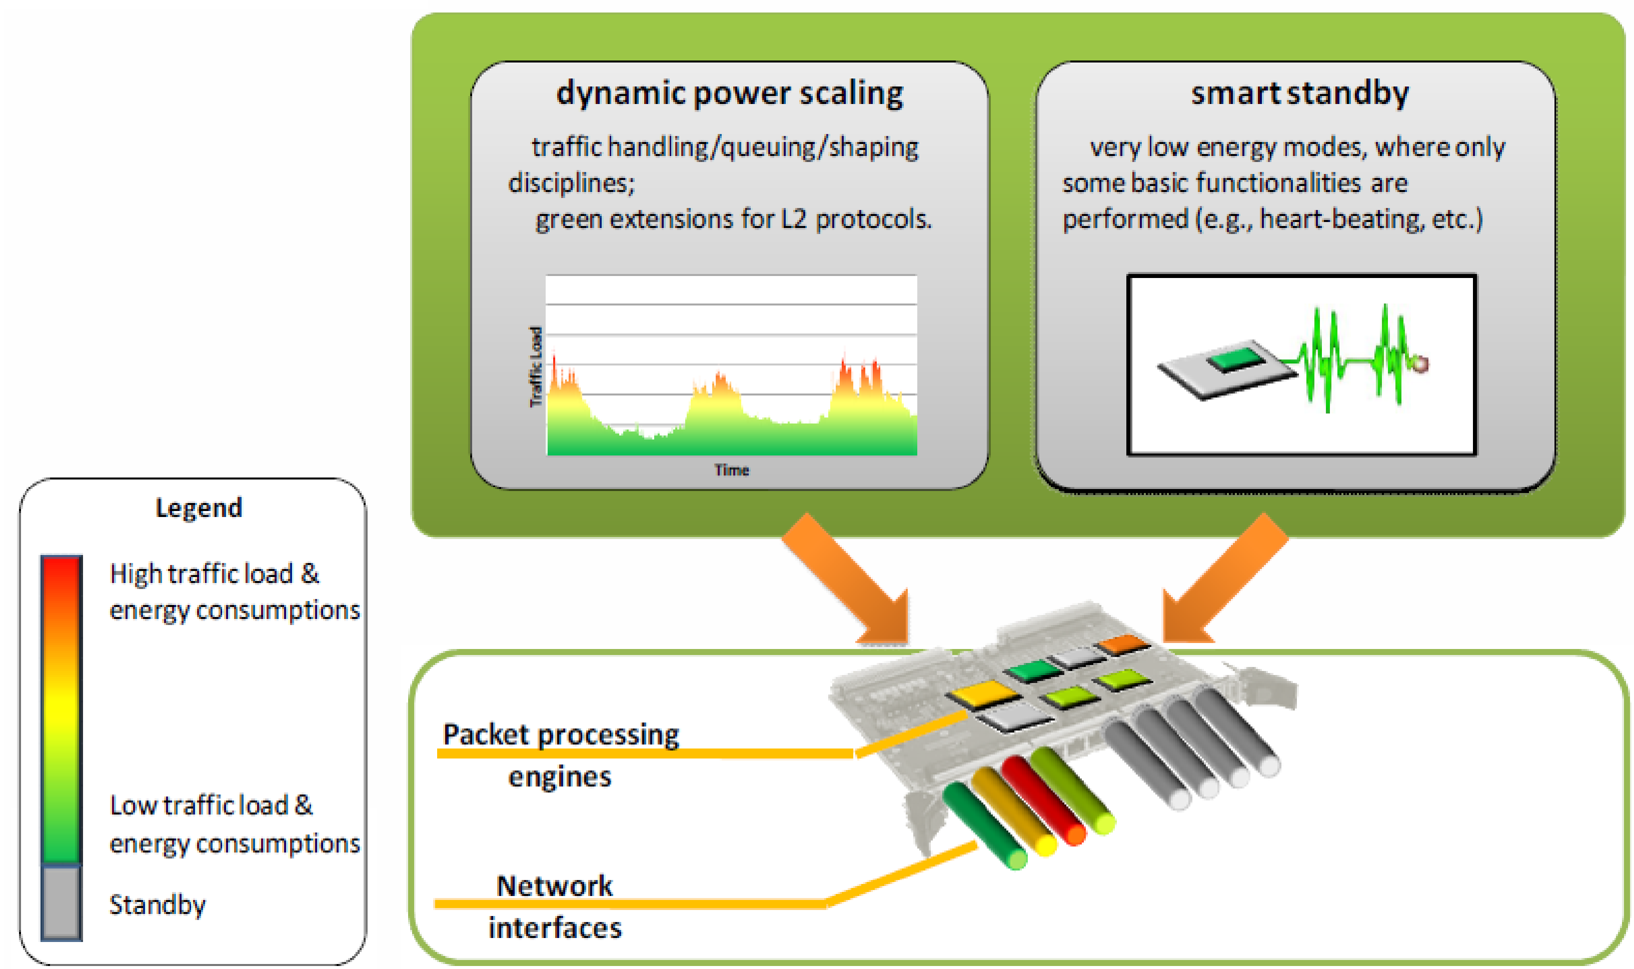

Dynamic Adaptation
QoS vs Power Management

Technology mapping

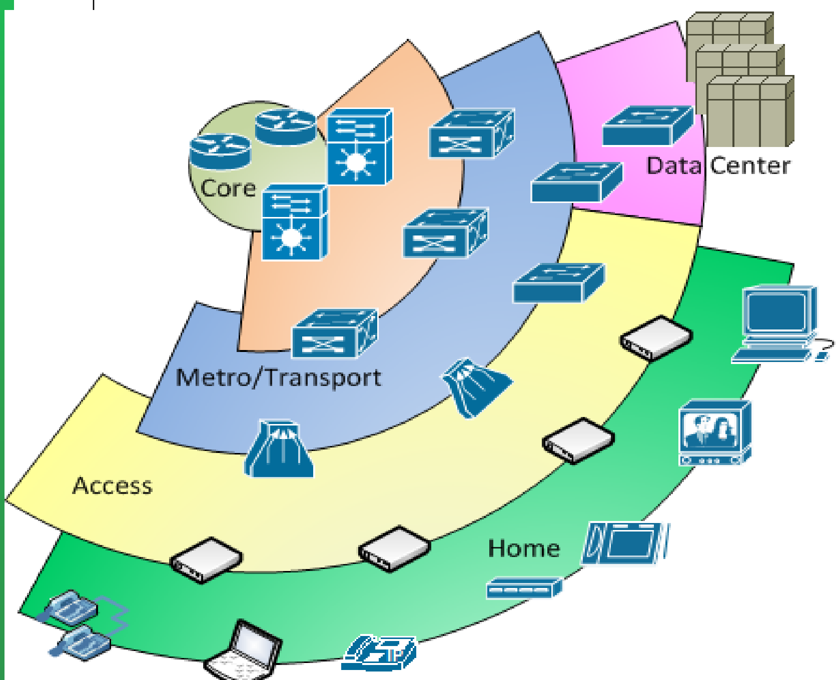

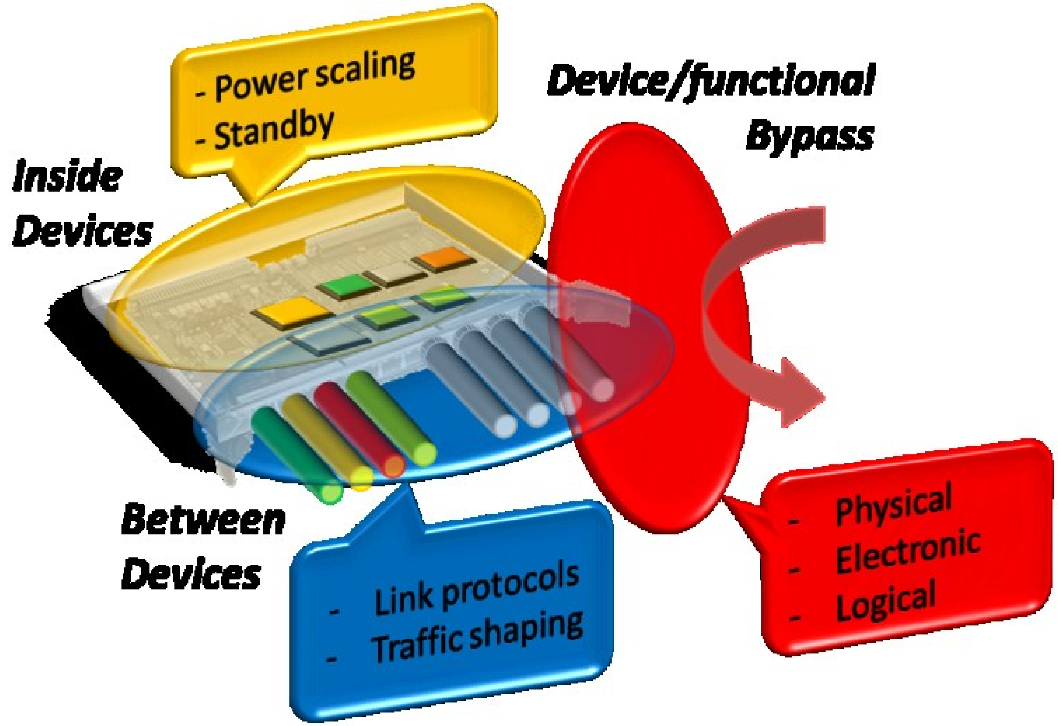

QoS vs Power Management

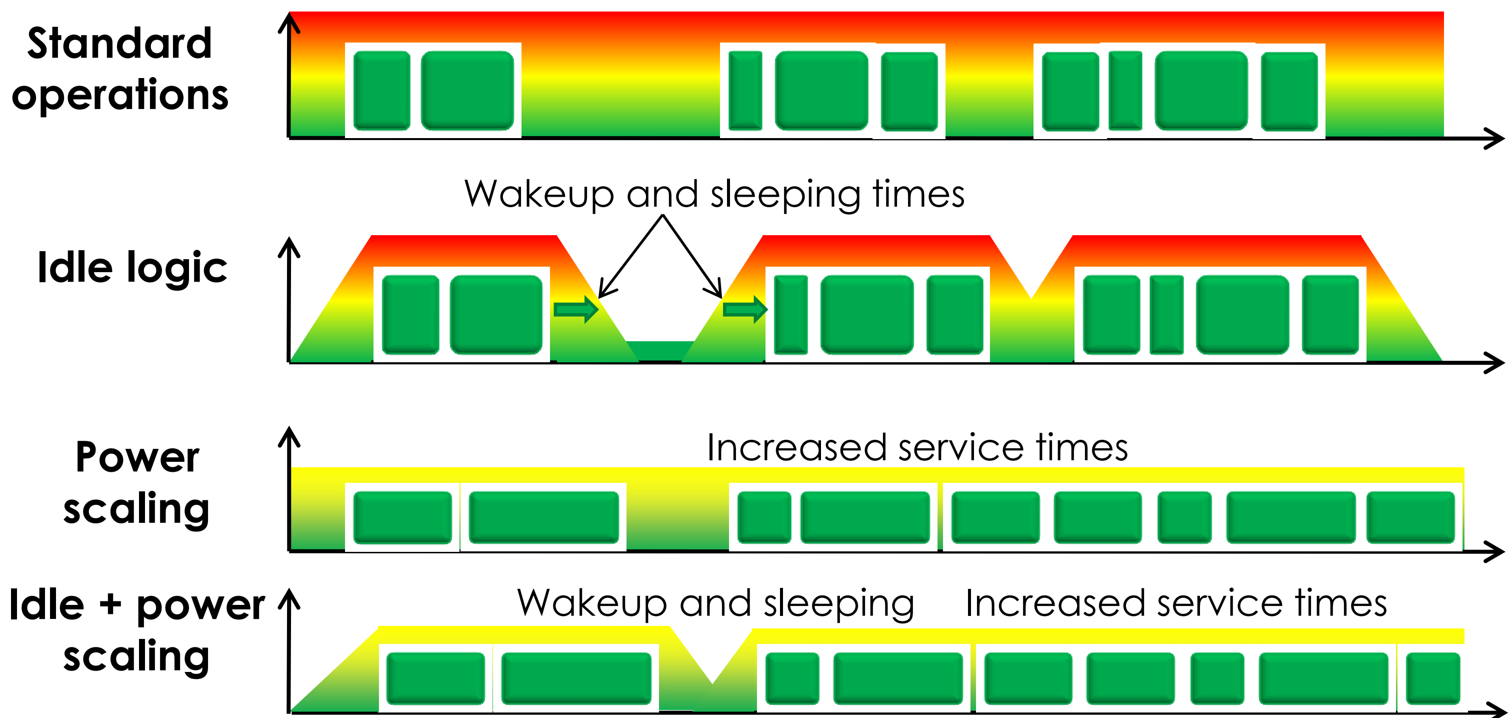

Link-level example: Green Ethernet (IEEE 802.3 az

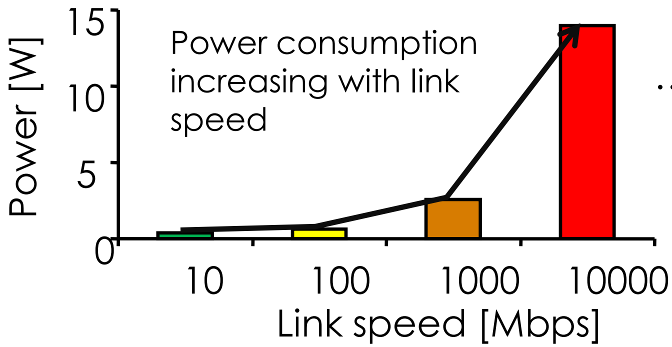

...hence...

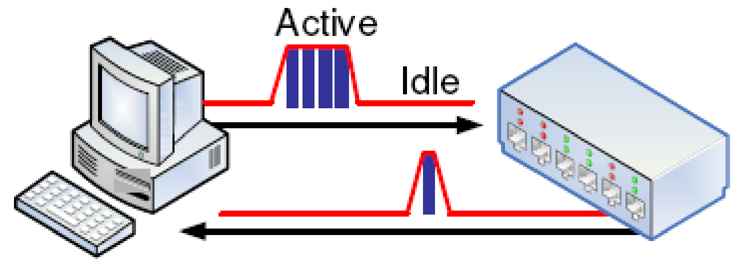

- Based on the “low power idle” concept.
- **Idea:** transmit data at the maximum speed, and put the link to sleep when it is idle.

- **Effect:** LPI has two transitions for each packet (or block of packets) : Link wake-up and sleep
  - LPI can possibly be asynchronous (one direction awake, the other asleep)
  - Retraining can be done via periodic on intervals (if no packets are being sent)
  - LPI requires no complicated handshaking

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

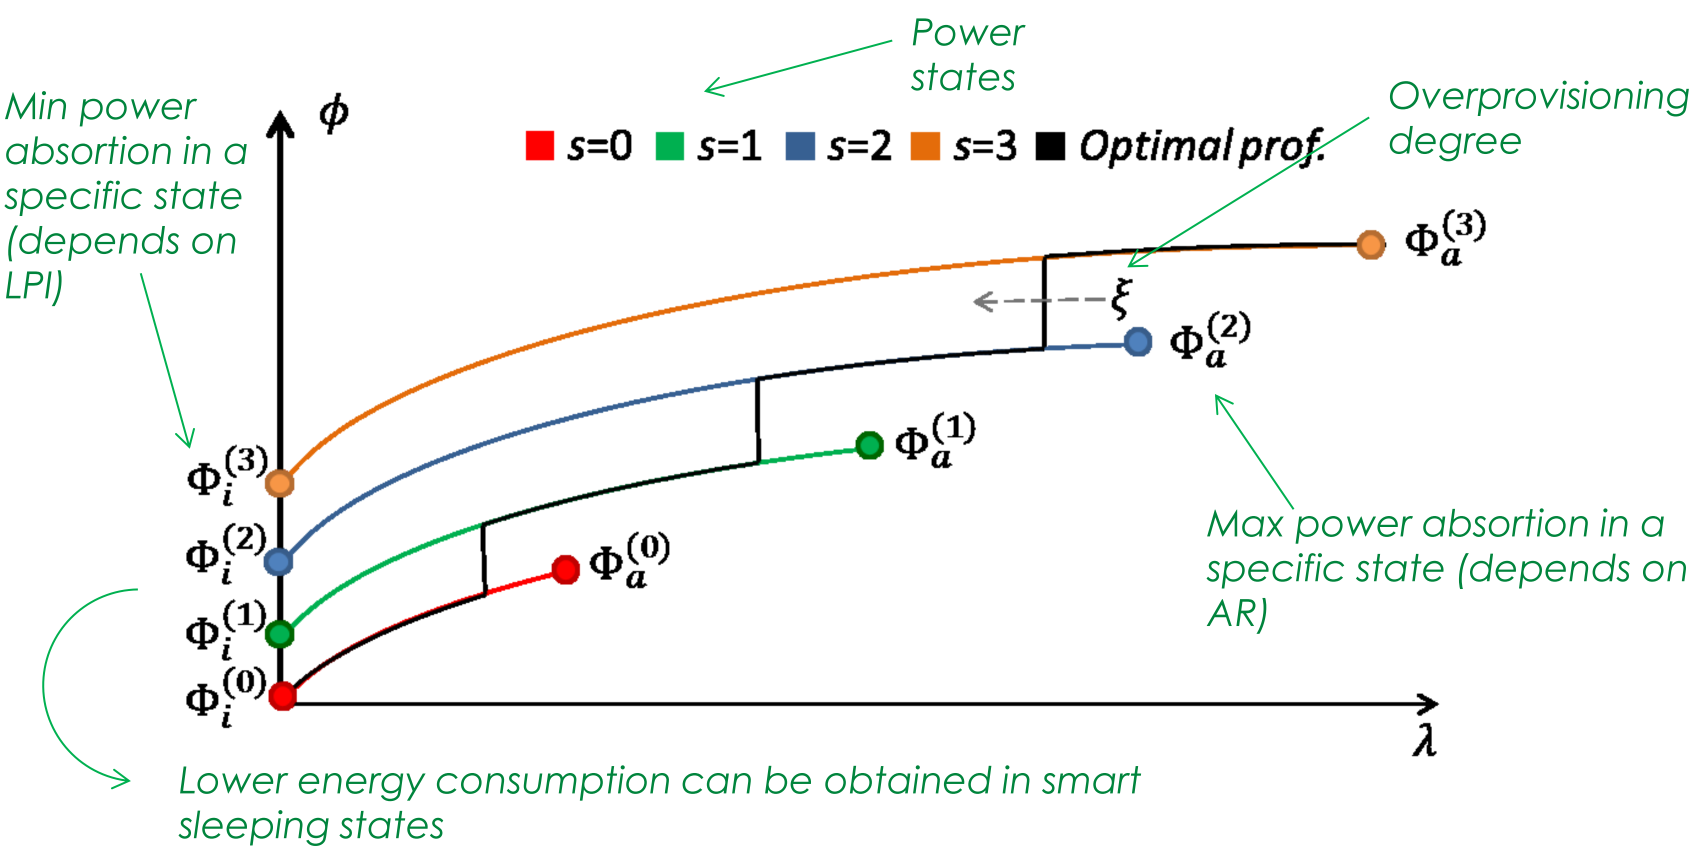

SW Routers & the ACPI

- The multi-core/cpu SW router architecture:

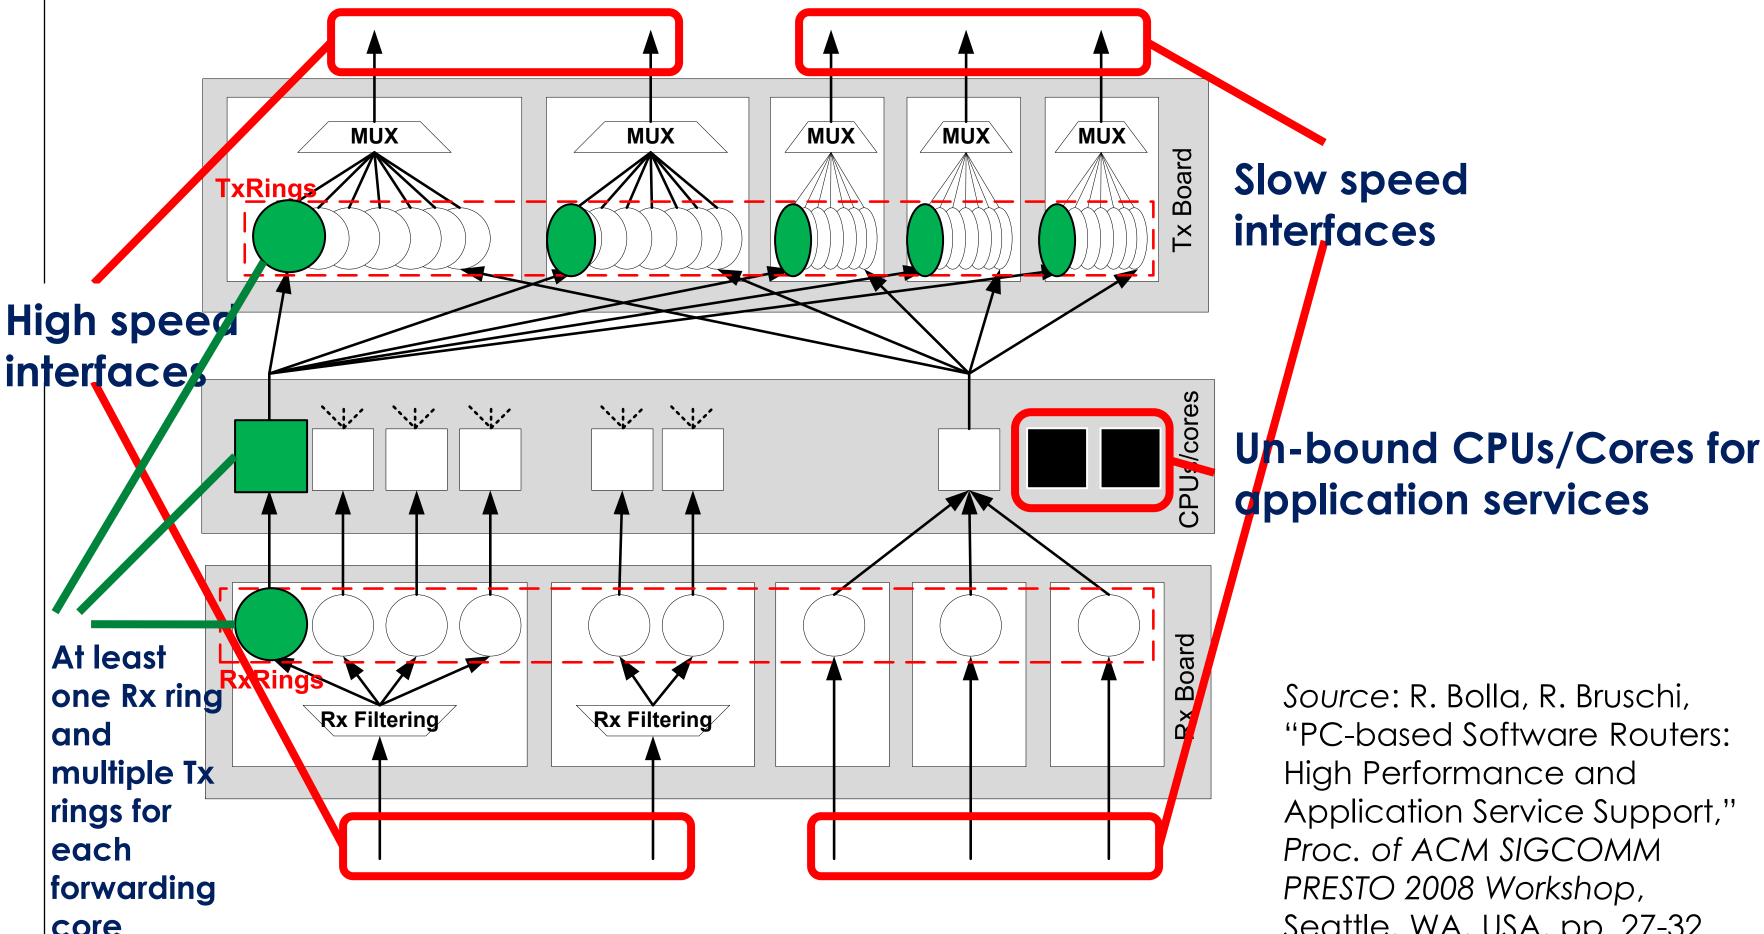

Source: R. Bolla, R. Bruschi, "PC-based Software Routers: High Performance and Application Service Support," *Proc. of ACM SIGCOMM PRESTO 2008 Workshop*, Seattle, WA, USA, pp. 27-32.

Measurements
An Experimental Test-Bed (1/2)

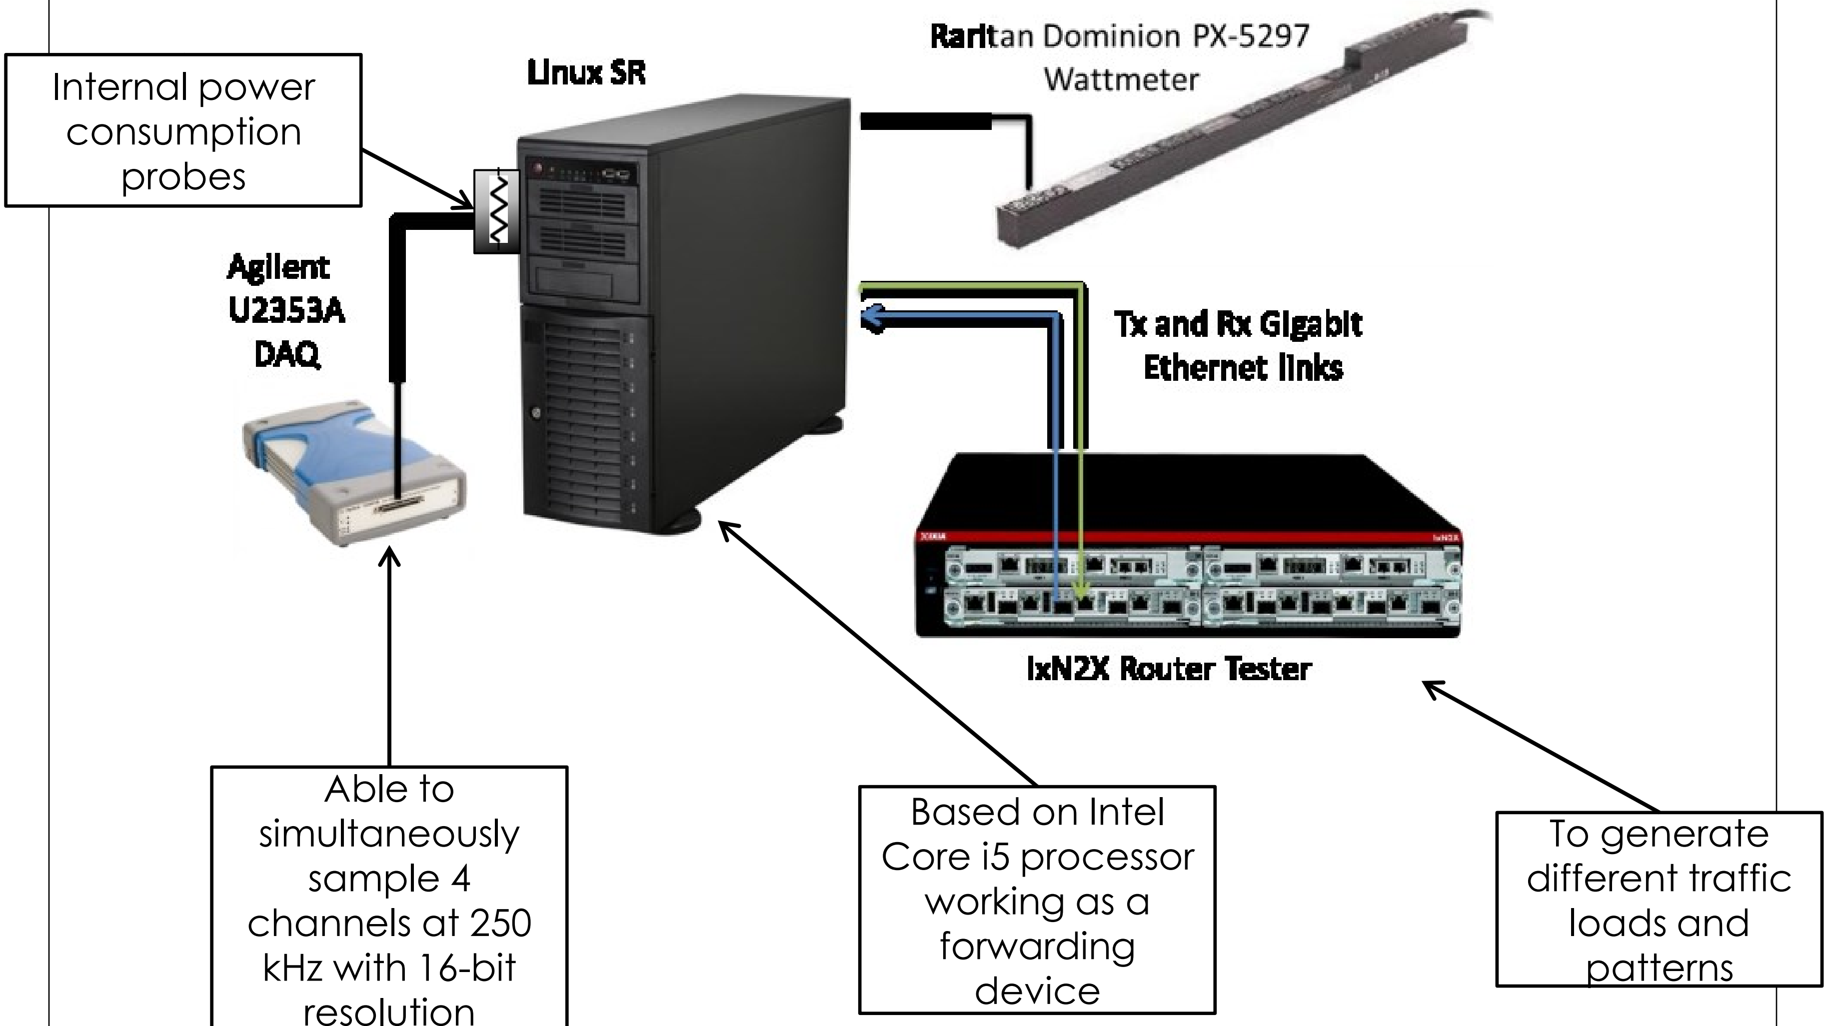

Experimental Test-Bed (2/2)

- To measure CPU power consumption we designed a riser board for ATX power connectors.
  - The 8-pin 12 V ATX rails provide power to the CPU cores only.
  - The 24-pin 12 V ATX rails supply other CPU sub-components (e. g., cache, and DDR controllers), RAM and Network Interface Card (NIC).
- By using commercial bus risers we isolated power absorption of NIC and RAM.

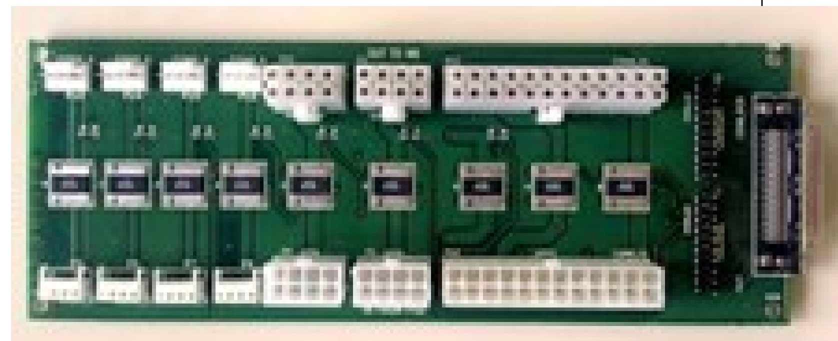

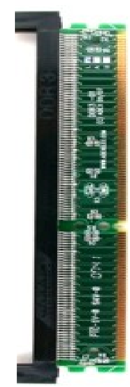

Energy profile of SR Data plane (1/2)

- C1 state: the Intel Core i5 controller turns off the clocks of all clock domains pertaining to the Core pipeline
- C3 state: the Core Phase-Locked-Loops (PLLs) are turned off, and the cores flush the contents of their Level 1 (L1) instruction cache, L1 data cache, and Level 2 (L2) cache to the shared Level 3 (L3) cache.
- Moving from C1 to C3 provides a significant power saving of about 10 W.
- Only the 12 V 24-pin rail changes the power consumption with the C-state level.
- This is because 12 V 24-pin rails supply, among other components, the CPU cache,

Idle power consumption for C1 and C3 LPI states

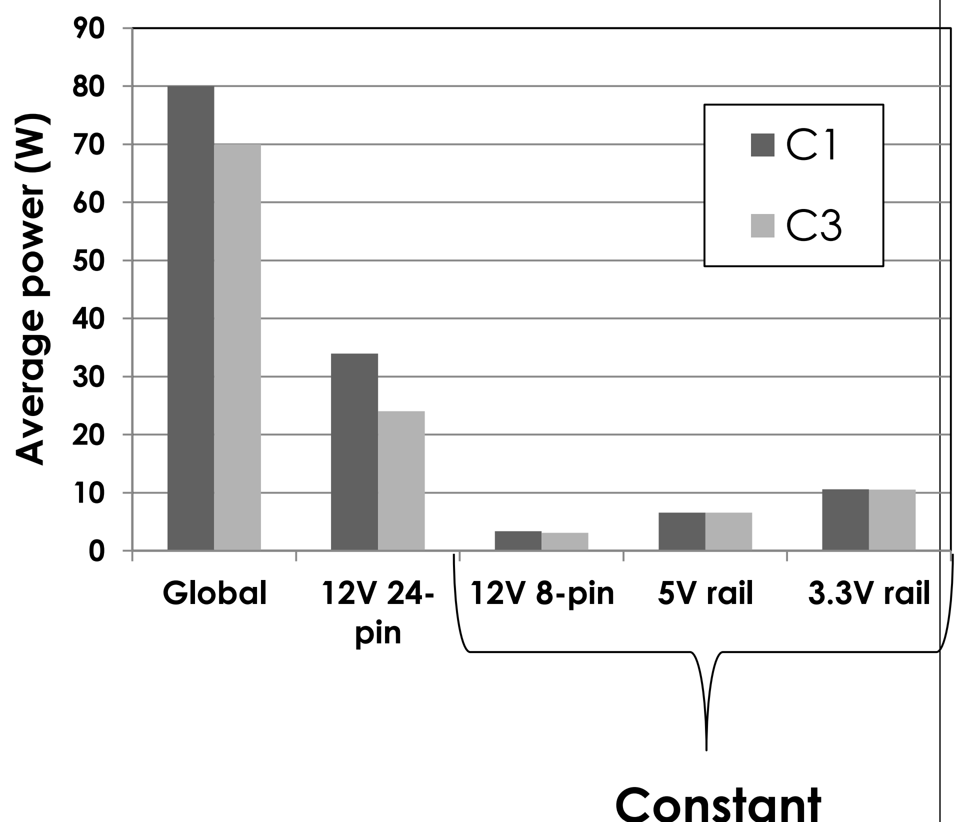

Energy profile of SR Data plane (2/2)

- The 12V 24-pin power remains almost constant
- The 12V 8-pin power changes with respect to the working frequency

Active power consumption of the 12V ATX rails for a subset of available frequencies and related maximum throughput

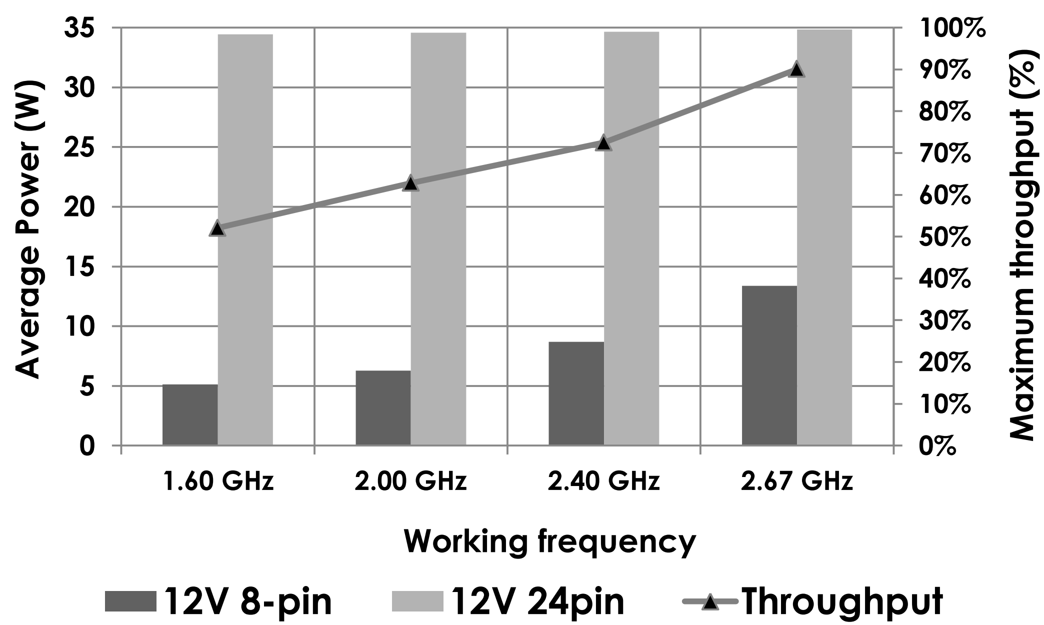

Dynamic Adaptation

- 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

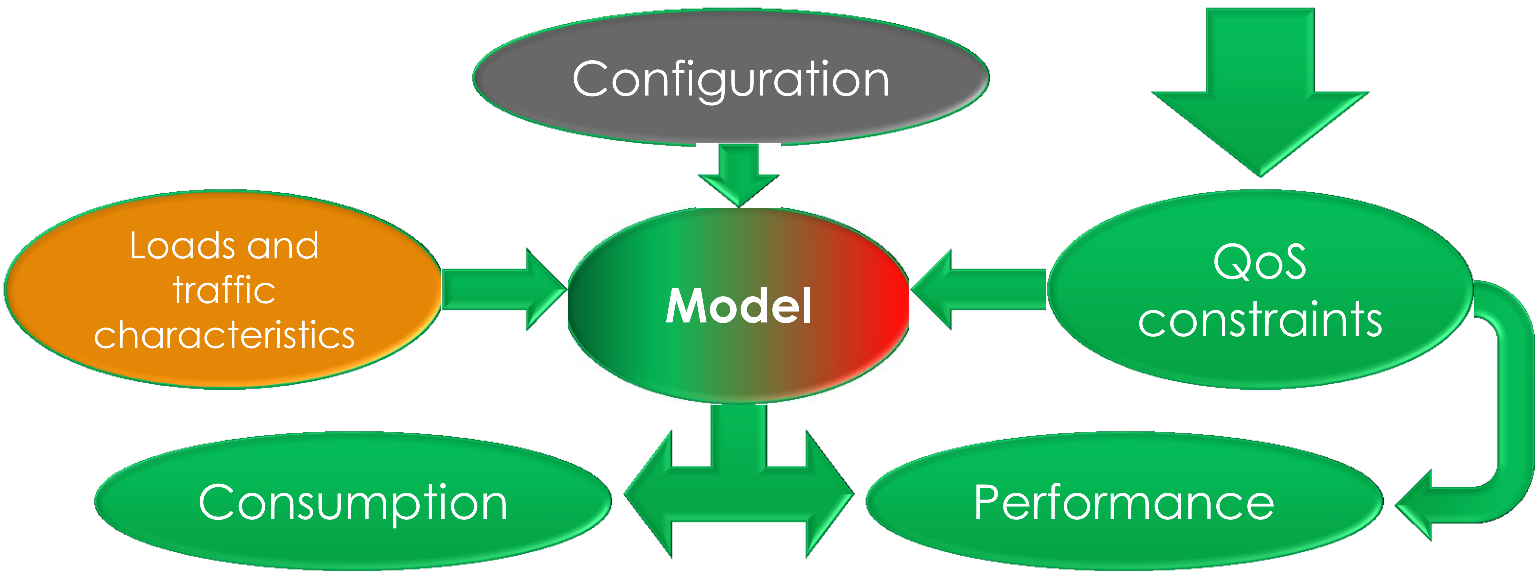

Idle Logic and Power Scaling

Idle vs Adaptive Rate – Examples of control strategies

S. Nedevschi, et. al., "Reducing Network Energy Consumption via Sleeping and Rate Adaptation", Proc. 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI'08), San Francisco, CA, April 2008.

- It is a seminal work exploring energy savings in networks.
- It explores (separately) and compares the effects of putting components to sleep when idle and also adapting the rate of "network operation" to workload.
- It tries to determine:
  - bounds and magnitudes for energy savings
  - where sleeping is best and where rate adaptation is best
- They start their analysis by stating that network devices should have similar power management primitives with respect to the ACPI of general purpose CPUs.

Dynamic Adaptation

Idle Logic and Power Scaling

Idle vs Adaptive Rate – Examples of control strategies

Understanding the Power-Performance Tradeoff

- A recently proposed simple model, based on classical queueing theory, allows representing the trade-off between energy and network performance in the presence of both AR and LPI capabilities.
- The model is aimed at describing the behaviour of packet processing engines.
- It is based on a $M^X/D/1/SET$ queueing system.

Source: R. Bolla, R. Bruschi, A. Carrega, F. Davoli, "Green Network Technologies and the Art of Trading-off," *Proc. IEEE INFOCOM 2011 Workshop on Green Communications and Networking*, Shanghai, China, April 2011, pp. 301-306.

R. Bolla, R. Bruschi, F. Davoli, P. Lago, "Trading off energy and forwarding performance in next-generation network devices" in J. Wu, S. Rangan, H. Zhang, Eds., *Green Communications: Theoretical Fundamentals, Algorithms and Applications*, CRC Press, Taylor & Francis, 2012, pp. 693-716.

R. Bolla, R. Bruschi, A. Carrega, F. Davoli, "Green networking with packet processing engines: Modeling and optimization", *IEEE/ACM Transactions on Networking*, 2013 (to appear); doi: 10.1109/TNET.2013.2242485.

Understanding the Power-Performance Tradeoff

Modeling and Control

- Service rate μ represents device capacity in terms of maximum number of packet headers that can be processed per second;
- Assumptions:
  - all packet headers require a constant service time;
  - Finite buffer of N packets is associated to the server for backlogging incoming traffic;
- We try to take into account the Long Range Dependency (LRD) and Multi-fractal traffic characteristics by using a Batch Markov Arrival Process (BMAP) with LRD batch sizes

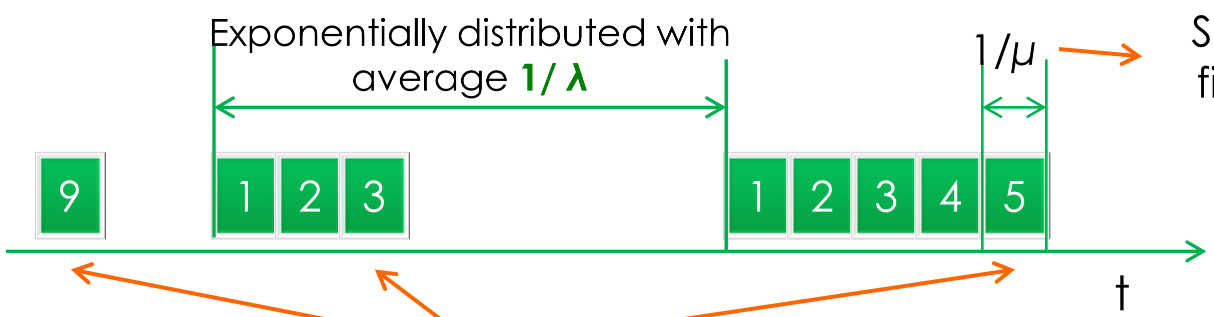

Served at a fixed rate ***μ***

# of pkts/batch follows **Zipf's law**

$$\Pr\{\text{incoming burst contains } j \text{ pkts}\} = \beta_j = \begin{cases} \frac{c}{j^s}, & j \le j_{\max} \\ 0, & j > j_{\max} \end{cases}$$

Normalization constant

Packet processing engine corresponds to a **M^x/D/1/N** queuing system.

Dynamic Adaptation

Understanding the Power-Performance Tradeoff

Modeling and Control

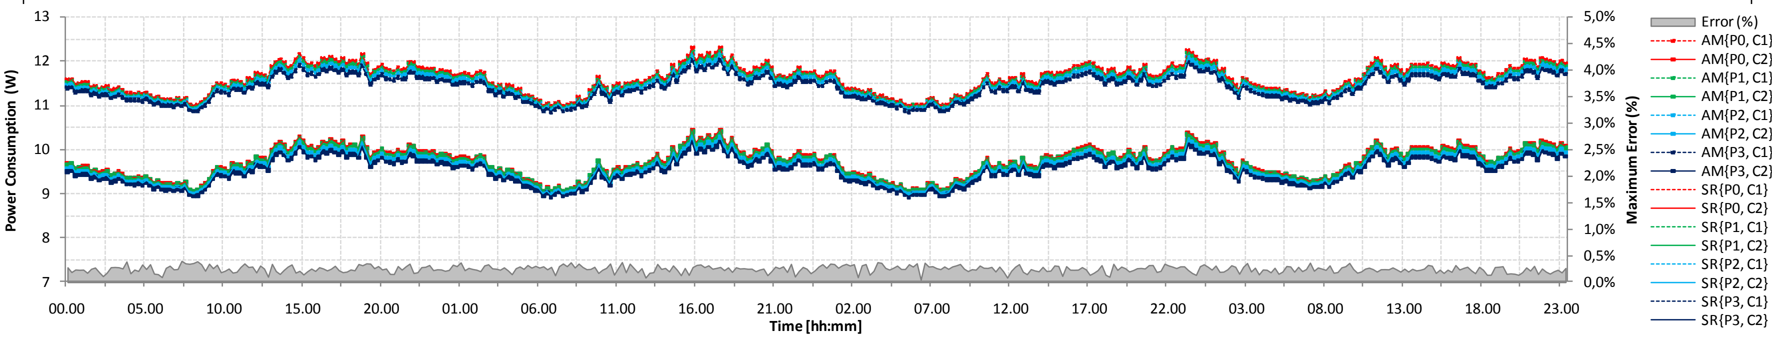

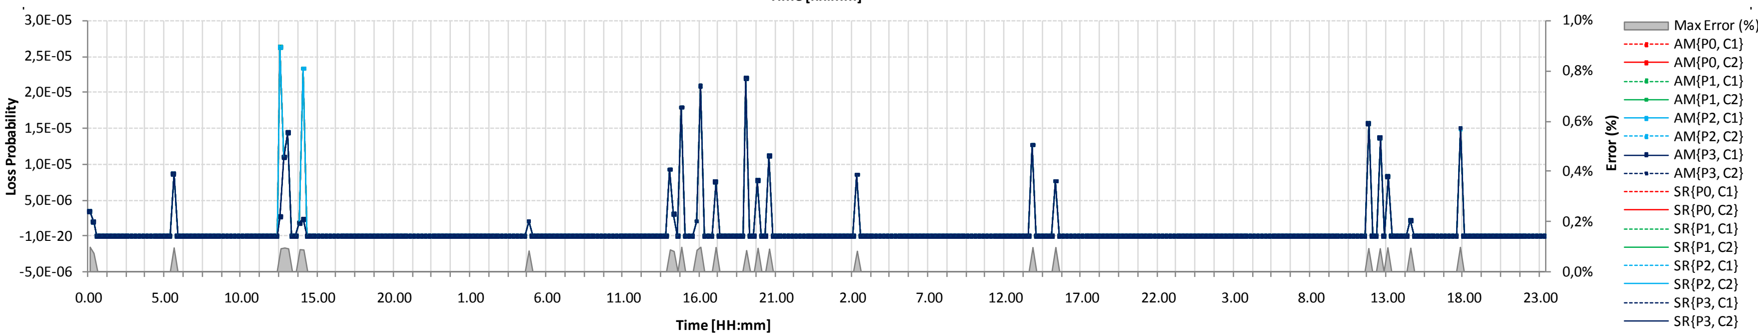

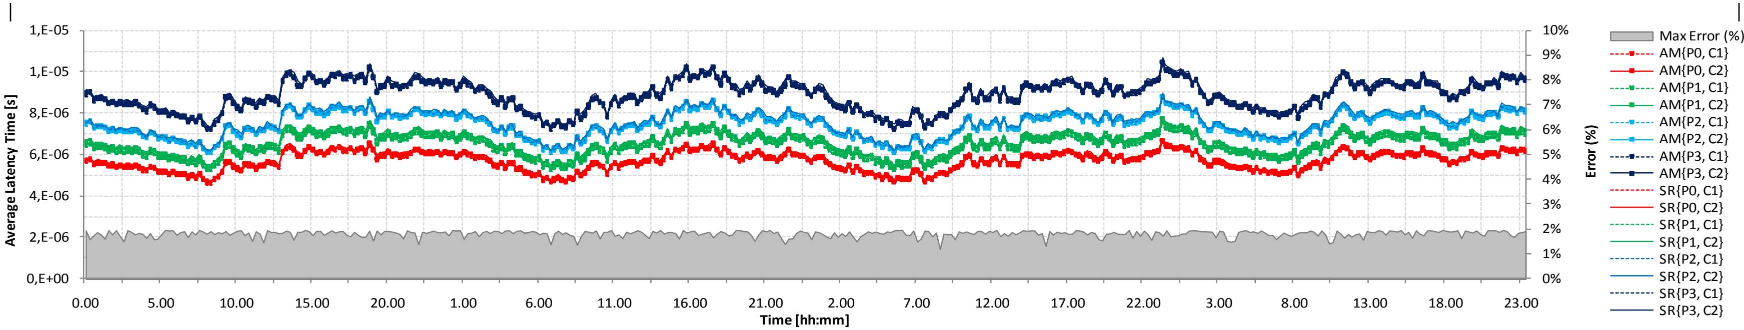

Energy-aware Load Balancing

- Focus on packet processing engines for network devices
  - highly parallel architectures
    - "divide and conquer" the traffic load incoming from a number of high-speed interfaces
    - Traffic flows enter and exit the engine by means of Serializer/Deserializer busses (SerDes)

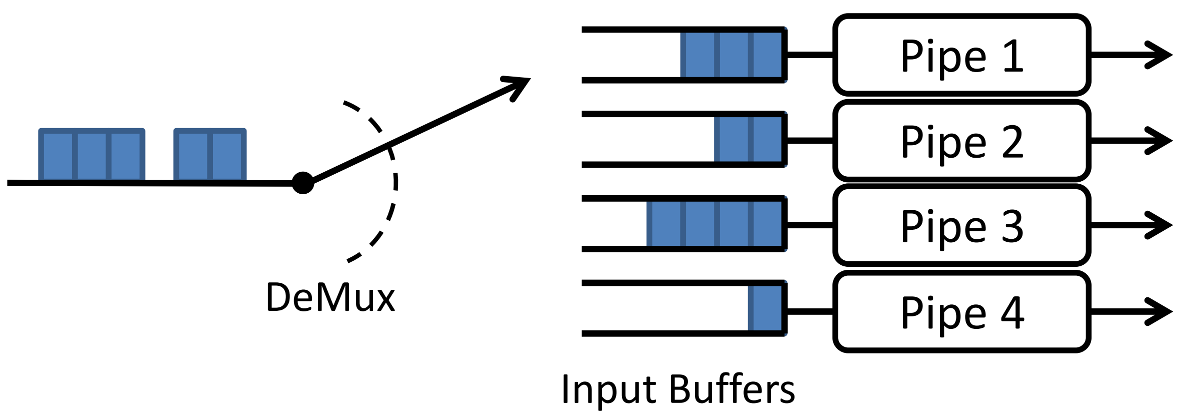

Energy-aware Load Balancing

Objectives

- The main goal is to dynamically manage the configuration of the packet processing engine, in order to optimally balance its energy consumption with respect to its network performance.
- To this purpose we want to dynamically act on:
  - how many pipelines have to actively work
  - their AR and LPI configurations
  - which share of the incoming traffic volume the load balancer module must assign to them
- Formulation of a general optimization problem to reflect different policies:
  - minimization of energy consumption for a certain constraint on packet latency time
  - maximization of network performance for a given energy cap
  - optimization of a given trade-off between the two previous objectives.

Energy-aware Load Balancing

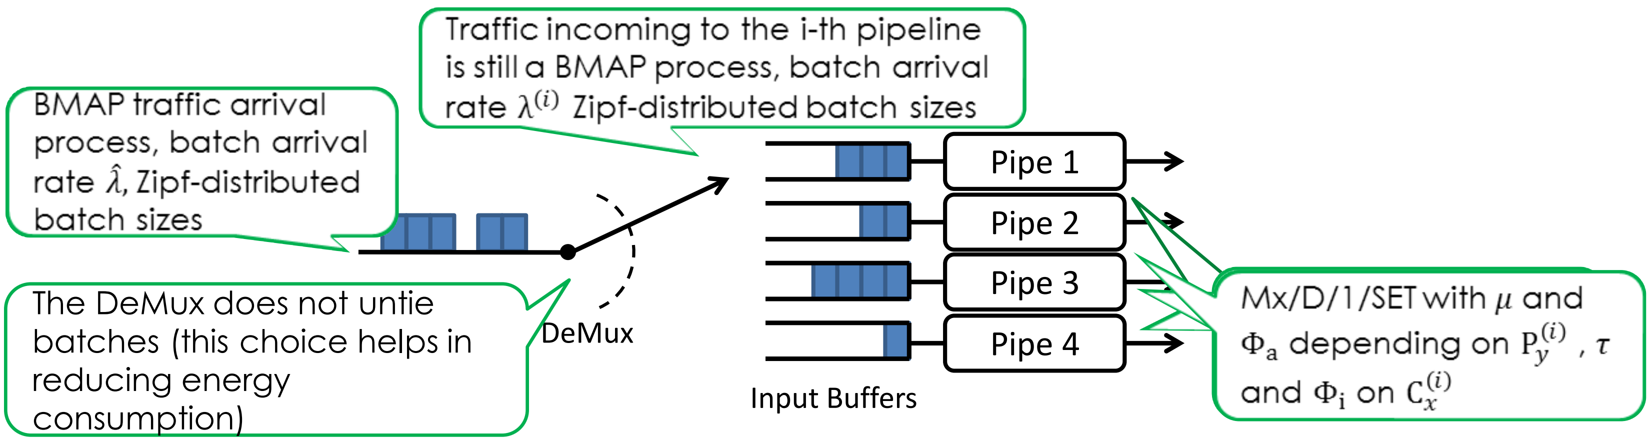

- Average power consumption: $\hat{\Phi} = \sum_{i=0}^{\Lambda-1} \tilde{\Phi}^{(i)}\left(\lambda^{(i)}, C_x^{(i)}, P_y^{(i)}\right)$
- Average packet latency: $\hat{W} = \sum_{i=0}^{\Lambda-1} \frac{\lambda^{(i)}}{\hat{\lambda}} \tilde{W}^{(i)}\left(\lambda^{(i)}, C_x^{(i)}, P_y^{(i)}\right)$
- We define our optimization problem as

$$\begin{cases} \min\limits_{\lambda^{(i)}, C_x^{(i)}, P_y^{(i)}} \gamma \frac{\hat{\Phi}}{\Phi^*} + (1-\gamma)\frac{\hat{W}}{W^*} \\ \hat{W} < W^* \\ \hat{\Phi} < \Phi^* \\ \sum_{i=0}^{\Lambda-1} \lambda^{(i)} = \hat{\lambda} \end{cases}$$

γ ($\in [0,1]$): "trade-off parameter" for modulating the weight of power consumption with respect to the network performance:

- $\gamma = 0$: maximization of network performance for given power consumption cap
- $\gamma = 1$: minimization of system power consumption constrained to maximum value of average latency

Analyzing the trade-off (1/2)

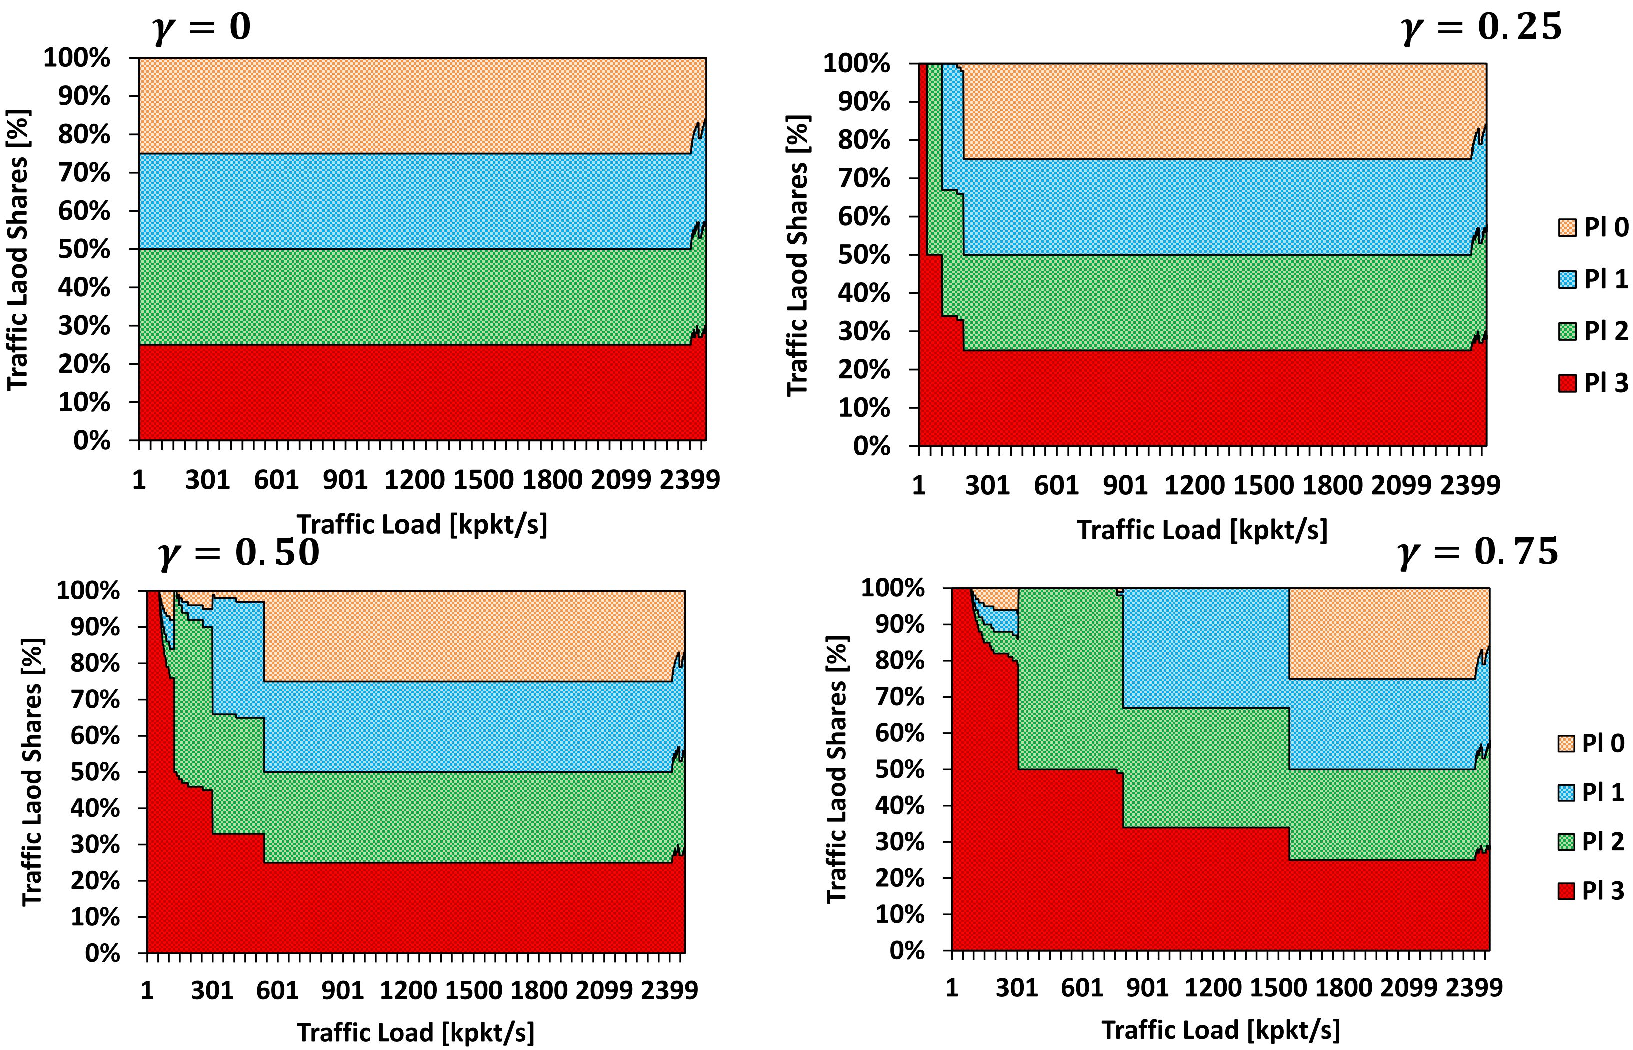

Analyzing the trade-off (2/2)

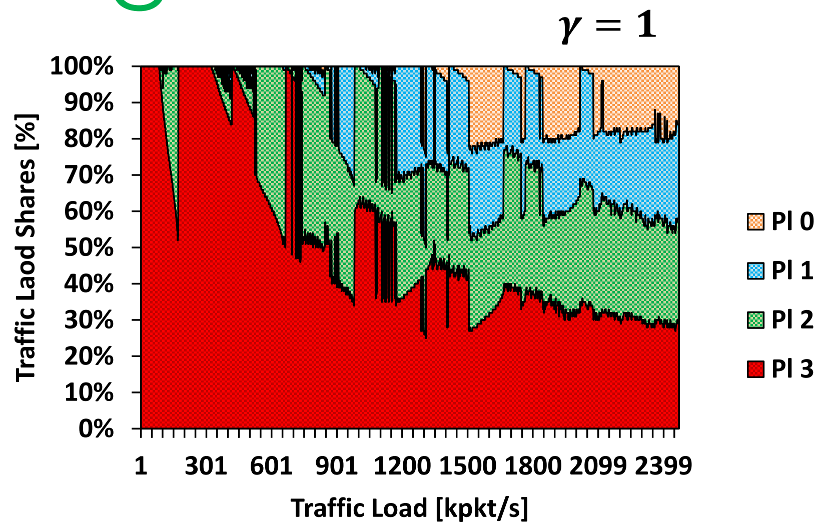

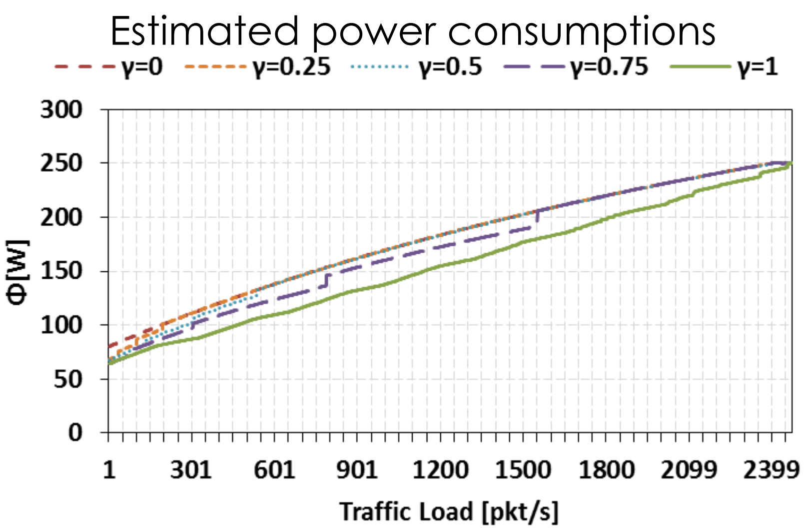

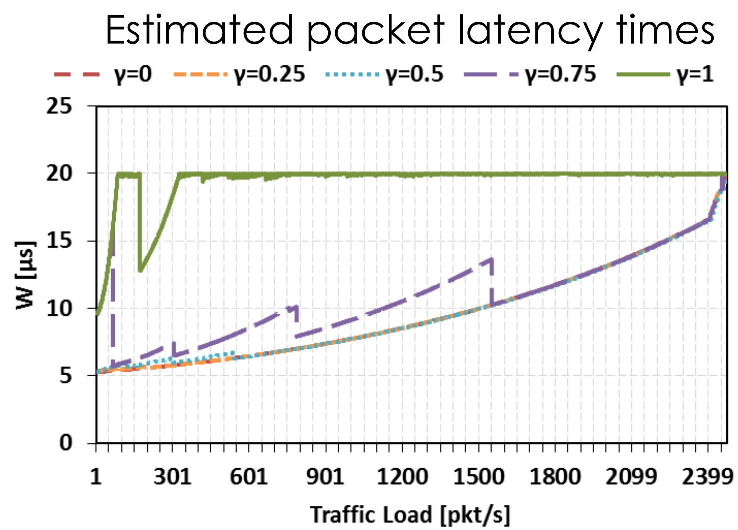

Modeling the IEEE 802.3az

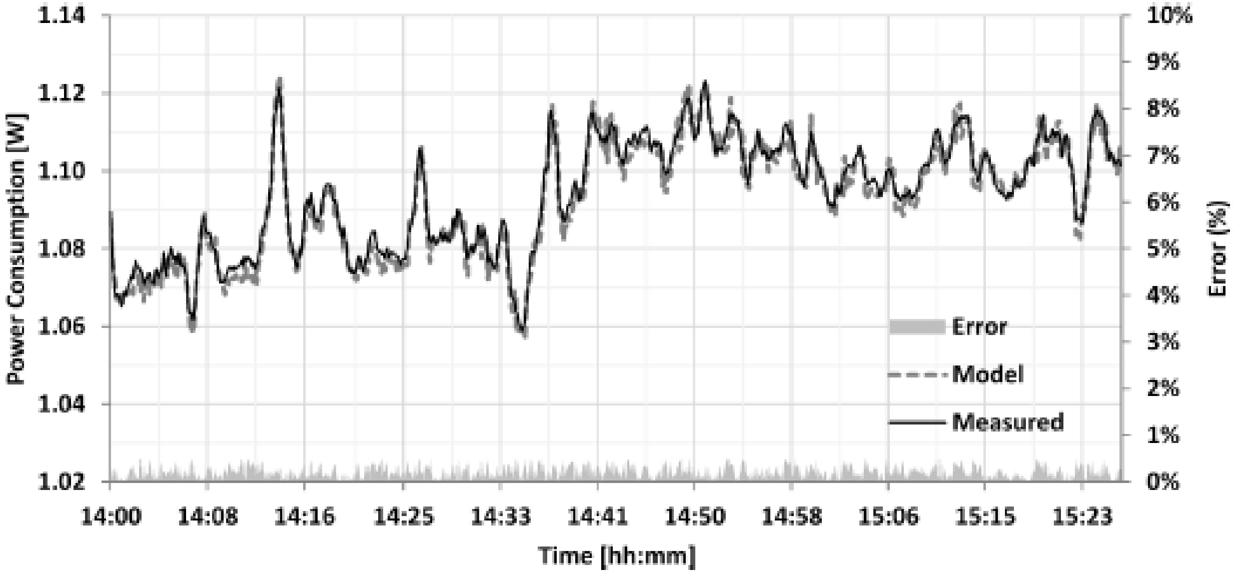

A similar model can be applied to assess the performance of IEEE 802.3az at 100 Mbps, 1 Gbps and 10 Gbps.

Source: 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, Jan. 2014.

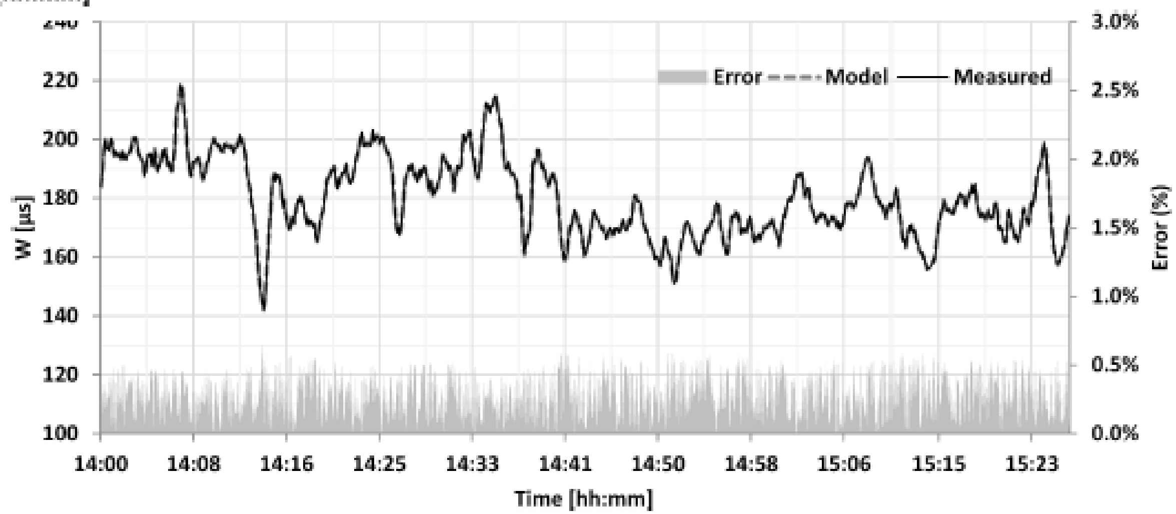

Sleeping/standby

- Sleeping/standby approaches are used to smartly and selectively drive unused network/device portions to low standby modes, and to wake them up only if necessary.
- However, since today's networks and related services and applications are designed to be continuously and always available,
  - standby modes have to be explicitly supported with special techniques able to maintain the "network presence" of sleeping nodes/components:

Solution: Proxying the network presence

- Moreover, additional techniques should be added to
  - enlarge as much as possible the number of "sleeping" parts or elements, but avoiding side effects or unacceptable performance reductions

Solution: Network virtualization

Solution: Energy aware traffic engineering and routing

Sleeping/standby

Proxying the Network Presence

- Proxying at the application layer

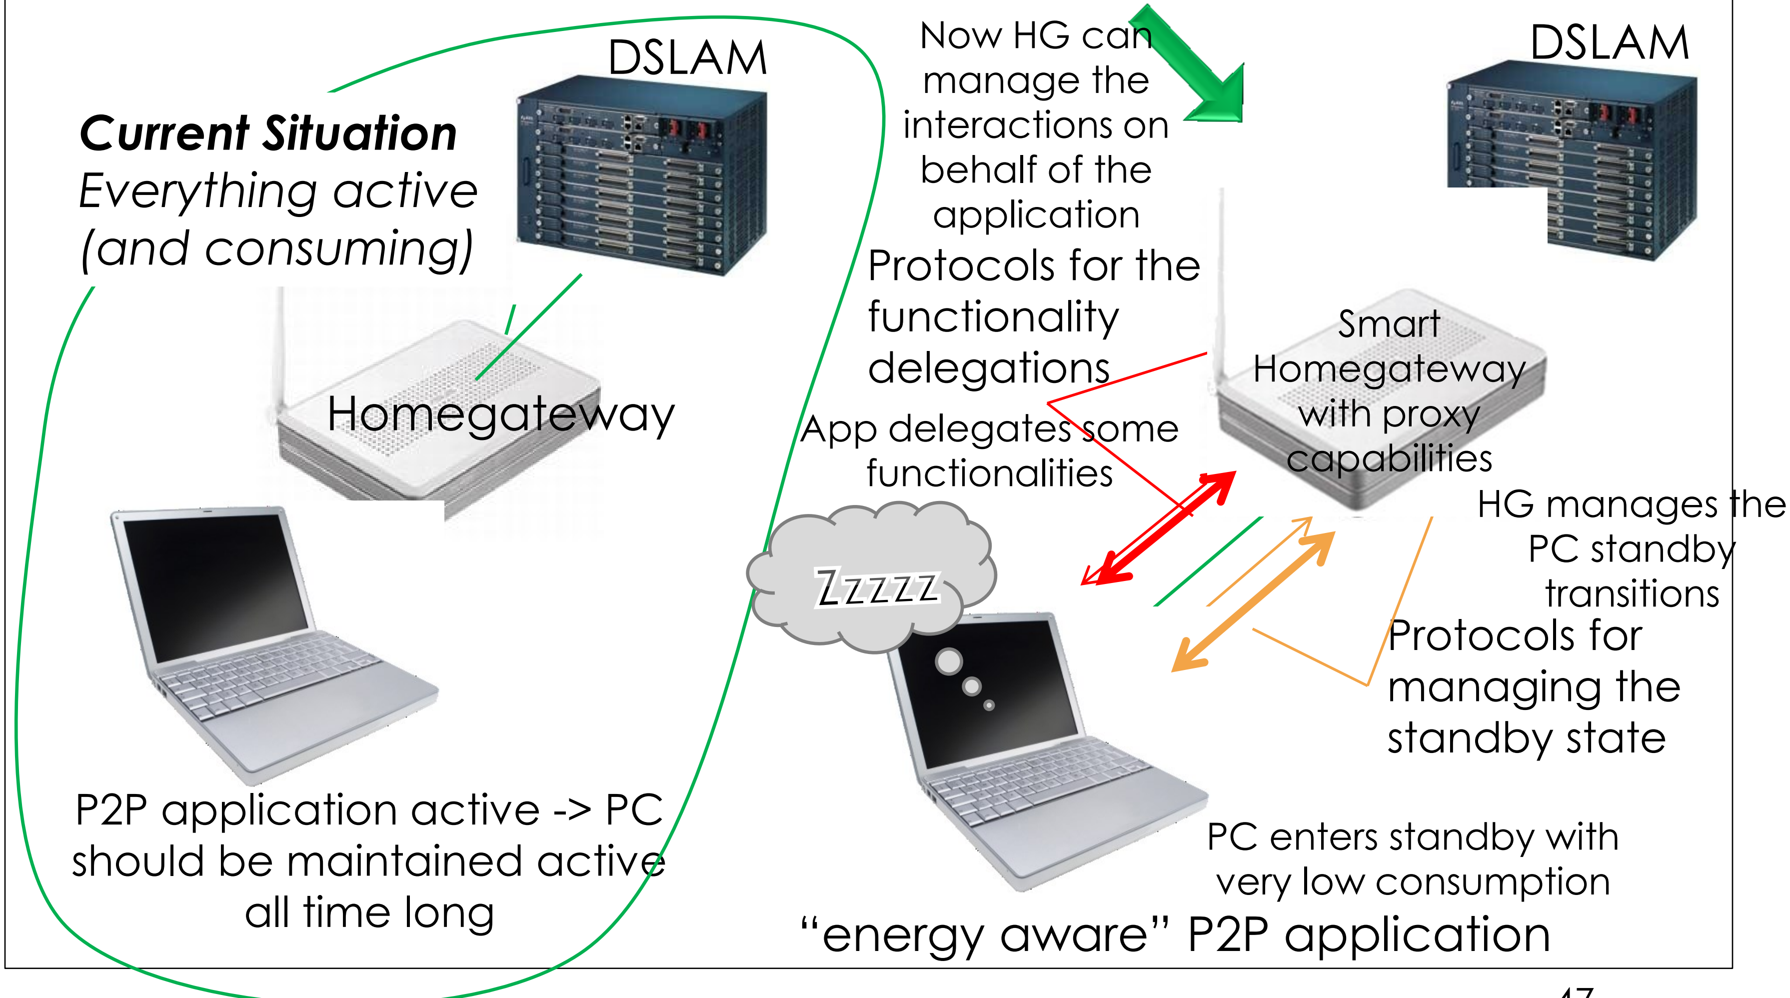

Supporting End-System Standby

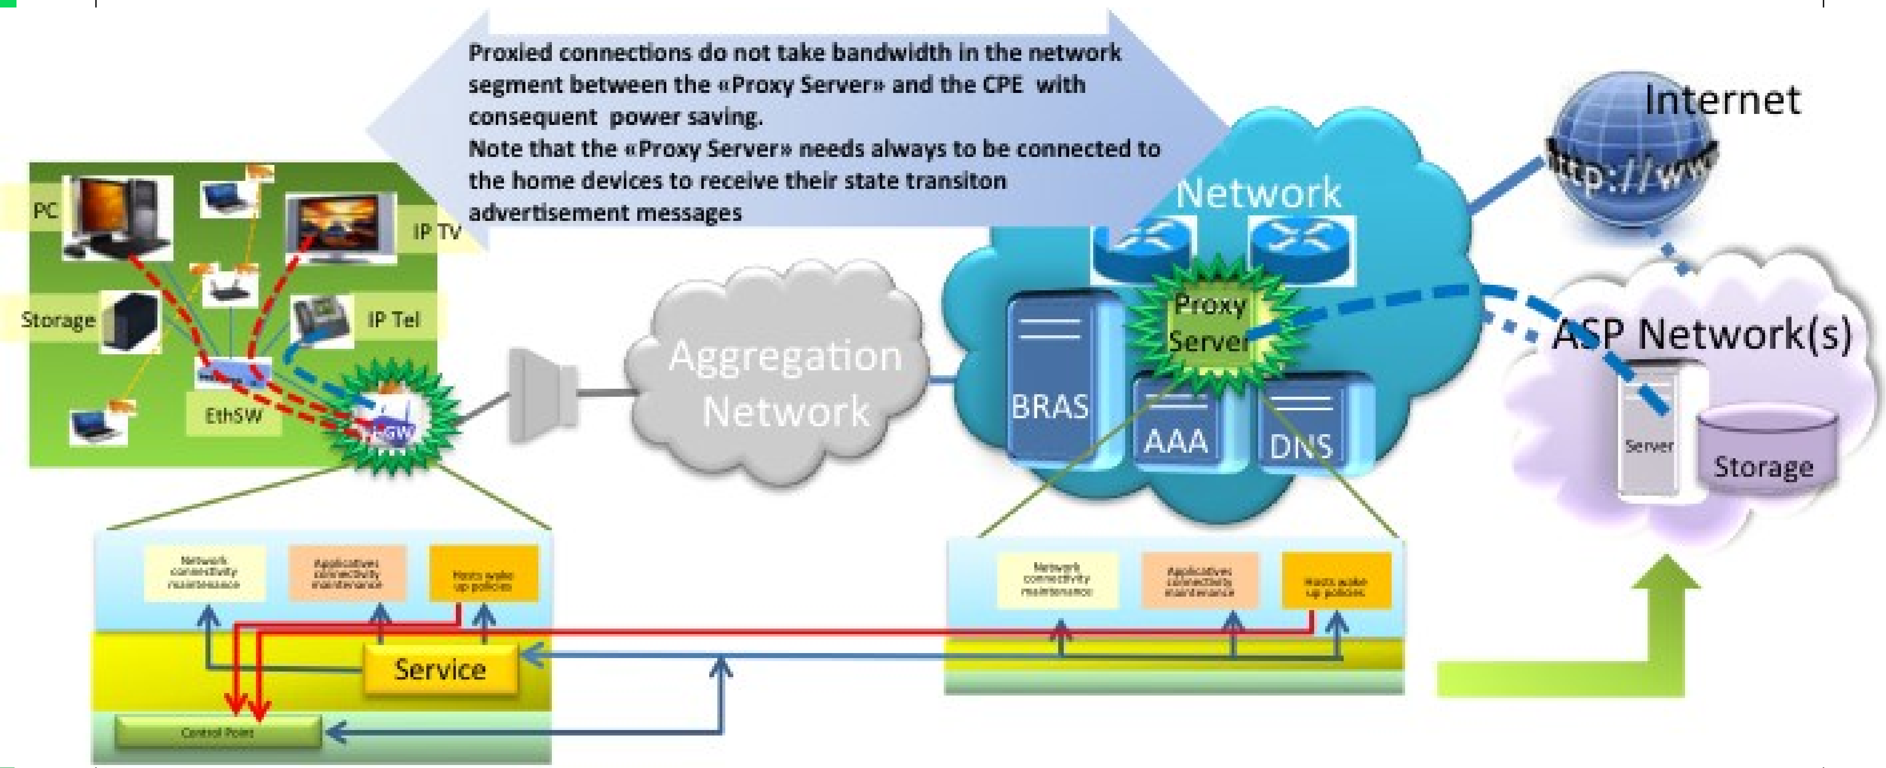

- Advertises itself as NCP to the devices in the home network
- Supports proxying of connections within the home network
- Exports towards the mate «external» NCP
  - requests for actions received from hosts related to connections over the «external» network
  - Hosts' power state
- Receives from mate «external» NCP requests for hosts' wake up

- Supports proxying of connections over the external network
- Collects from «internal» NCP
  - requests for actions received from hosts related to connection over the «external» network
  - Hosts' power state
- Forwards to «internal» NCP requests for hosts' wake up

Extending the reach by Network-wide Control

Only local control policies | **Local + network-wide control policies**

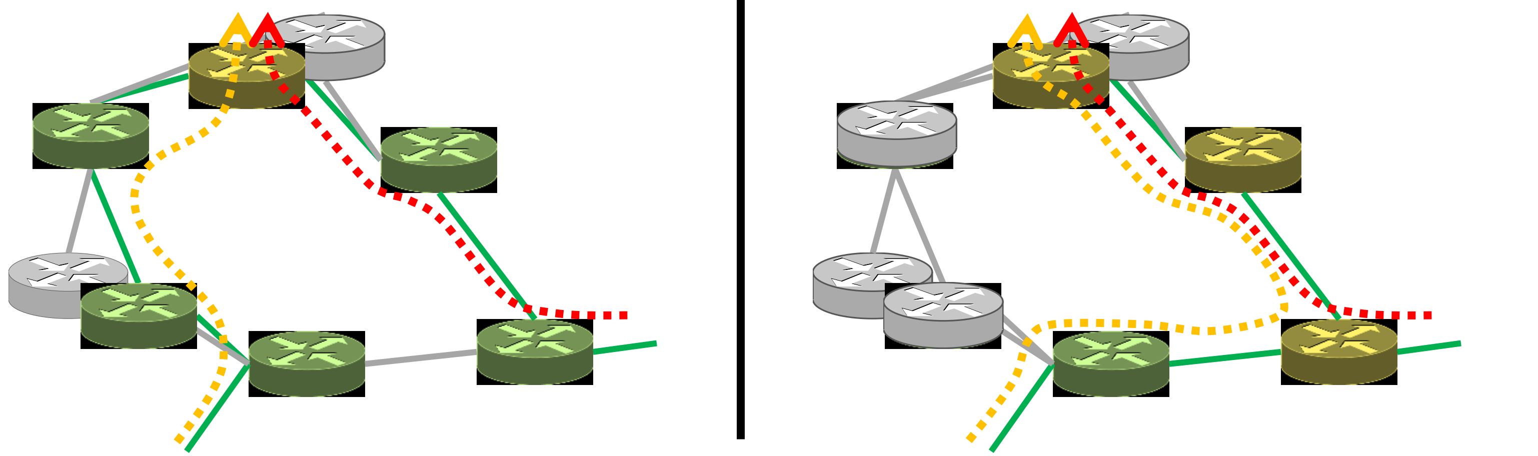

- Network-wide control strategies (i.e., routing and traffic engineering) give the possibility of moving traffic load among network nodes.
- **When a network is under-utilized, we can move network load on few "active" nodes, and put all the other ones in standby .**
- Different network nodes can have heterogeneous energy capabilities and profiles.
- Recent studies, obtained with real data from Telcos (topologies and traffic volumes) suggested that network-wide control strategies could cut the overall energy consumption by more than 23%.

Green network-wide control: traffic engineering and routing

- Current contributions in this area mainly focus on:
  - Putting links in standby modes -> calculating the minimal sub-topology for meeting QoS contraints.
    - L. Chiaraviglio, D. Ciullo, M. Mellia, M. Meo, “Modeling Sleep Modes Gains with Random Graphs”, *Proc. IEEE INFOCOM 2011 Workshop on Green Communications and Networking*, Shanghai, China, April 2011.
    - A. P. Bianzino, L. Chiaraviglio, M. Mellia, J.-L. Rougier, "GRiDA: Green Distributed Algorithm for Energy-Efficient IP Backbone Networks" *Computer Networks*, vol. 56, no.14, pp. 3219–3232, Sept. 2012.
    - A. Cianfrani, V. Eramo, M. Listanti, M. Polverini, "Introducing Routing Standby in Network Nodes to Improve Energy Savings techniques", *Proc. ACM e-Energy Conf.*, Madrid, Spain, May 2012.
  - Considering the energy profile of devices or their sub-components -> acting on routing/TE metrics in order to move flows towards “greener” alternative paths.
    - J. C. Cardona Restrepo, C. G. Gruber, C. Mas Machuca, “Energy Profile Aware Routing,” *Proc. Green Communications Workshop* in conjunction with *IEEE ICC'09 (GreenComm09)*, Dresden, Germany, June 2009.
    - P. Arabas, K. Malinowski, A Sikora, “On formulation of a network energy saving optimization problem”, Special Session on Energy Efficient Networking, ICCE 2012, Hue, Vietnam, Aug. 2012.
    - E. Niewiadomska-Szynkiewicz, A. Sikora, P. Arabas, J. Kołodziej, “Control system for reducing energy consumption in backbone computer networks”, *Concurrency and Computation: Practice and Experience,* vol. 25, no. 12, pp. 1738–1754, Aug. 2013; doi: 10.1002/cpe.2964.

Green network-wide control: traffic engineering and routing

J. Restrepo, C. Gruber, and C. Machoca, "Energy Profile Aware Routing," *Proc. IEEE GreenComm '09*, Dresden, Germany, June 2009.

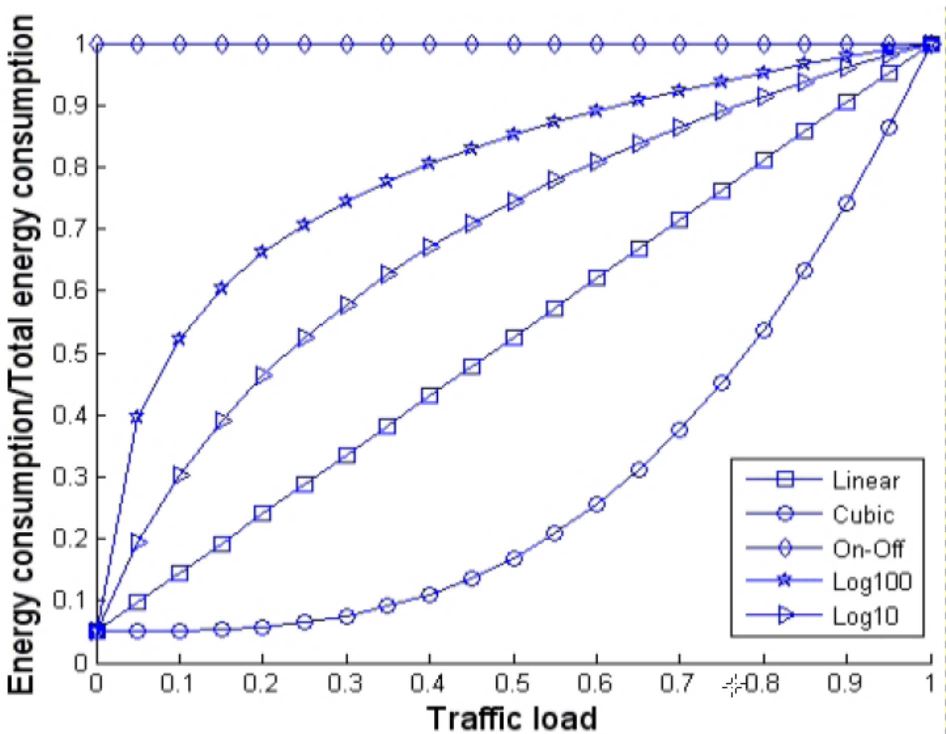

They showed the influence of different router energy profiles on the energy-aware routing problem solution.

Green network-wide control: traffic engineering and routing

W. Fisher, *et al.*, "Greening Backbone Networks: Reducing Energy Consumption by Shutting Off Cables in Bundled Links", *ACM SIGCOMM Workshop on Green Networking 2010*, New Delhi, India, Aug. 2010.

- They exploit the fact that many links in core networks are actually "bundles" of multiple physical cables and line cards that can be shut down independently.
- Since identifying the optimal set of cables to shut down is an NP-complete problem, the authors propose several heuristics based on linear optimization techniques.

Variable	Description
$G(V, E)$	backbone router and link representation
$\|V\|, \|E\|$	cardinality of set V and E, respectively
$c(u, v)$	capacity of edge (u, v)
B	number of cables in a bundle
D	set of demands
s_d	source of demand $d \in D$
t_d	destination of demand $d \in D$
h_d	flow requirement for demand $d \in D$
$f_d(u, v)$	flow on edge (u, v) of demand d
$f(u, v)$	total flow on edge (u, v)
n_{uv}	number of powered cables in link (u, v)

$$
\begin{aligned}
\min \quad & \textstyle\sum_{(u,v)\in E} f(u,v) \\
\text{s.t.} \quad & \textstyle\sum_{D} f_d(u,v) \le c(u,v) \quad \forall (u,v) \in E, \\
& \textstyle\sum_{v\in V} f_d(u,v) = \sum_{v\in V} f_d(v,u) \quad \forall d, u \neq s_d, t_d, \\
& \textstyle\sum_{v\in V} f_d(s_d,v) = \sum_{v\in V} f_d(v,t_d) = h_d \quad \forall d.
\end{aligned}
$$

Green network-wide control: traffic engineering and routing

G. Lin, *et al.,* "Efficient heuristics for energy-aware routing in networks with bundled links," *Computer Networks*, vol. 57, no. 8, June 2013, pp. 1774-1788.

- They considers the problem of shutting down a subset of bundled links during off-peak periods in order to minimize energy expenditure.
- Unfortunately, identifying the cables that minimize this objective is an **NP-complete** problem.
- Henceforththey propose several practical **heuristics** based on *Dijkstra*'s algorithm and *Yen*'s *k*-shortest paths algorithm.
- The authors propose an efficient approach – **Shortest Single Path First (SSPF)** to *power off redundant cables* as long as the remaining cables provide sufficient capacity to satisfy traffic demands.

Implementing it all: The Energy-aware Data-Plane

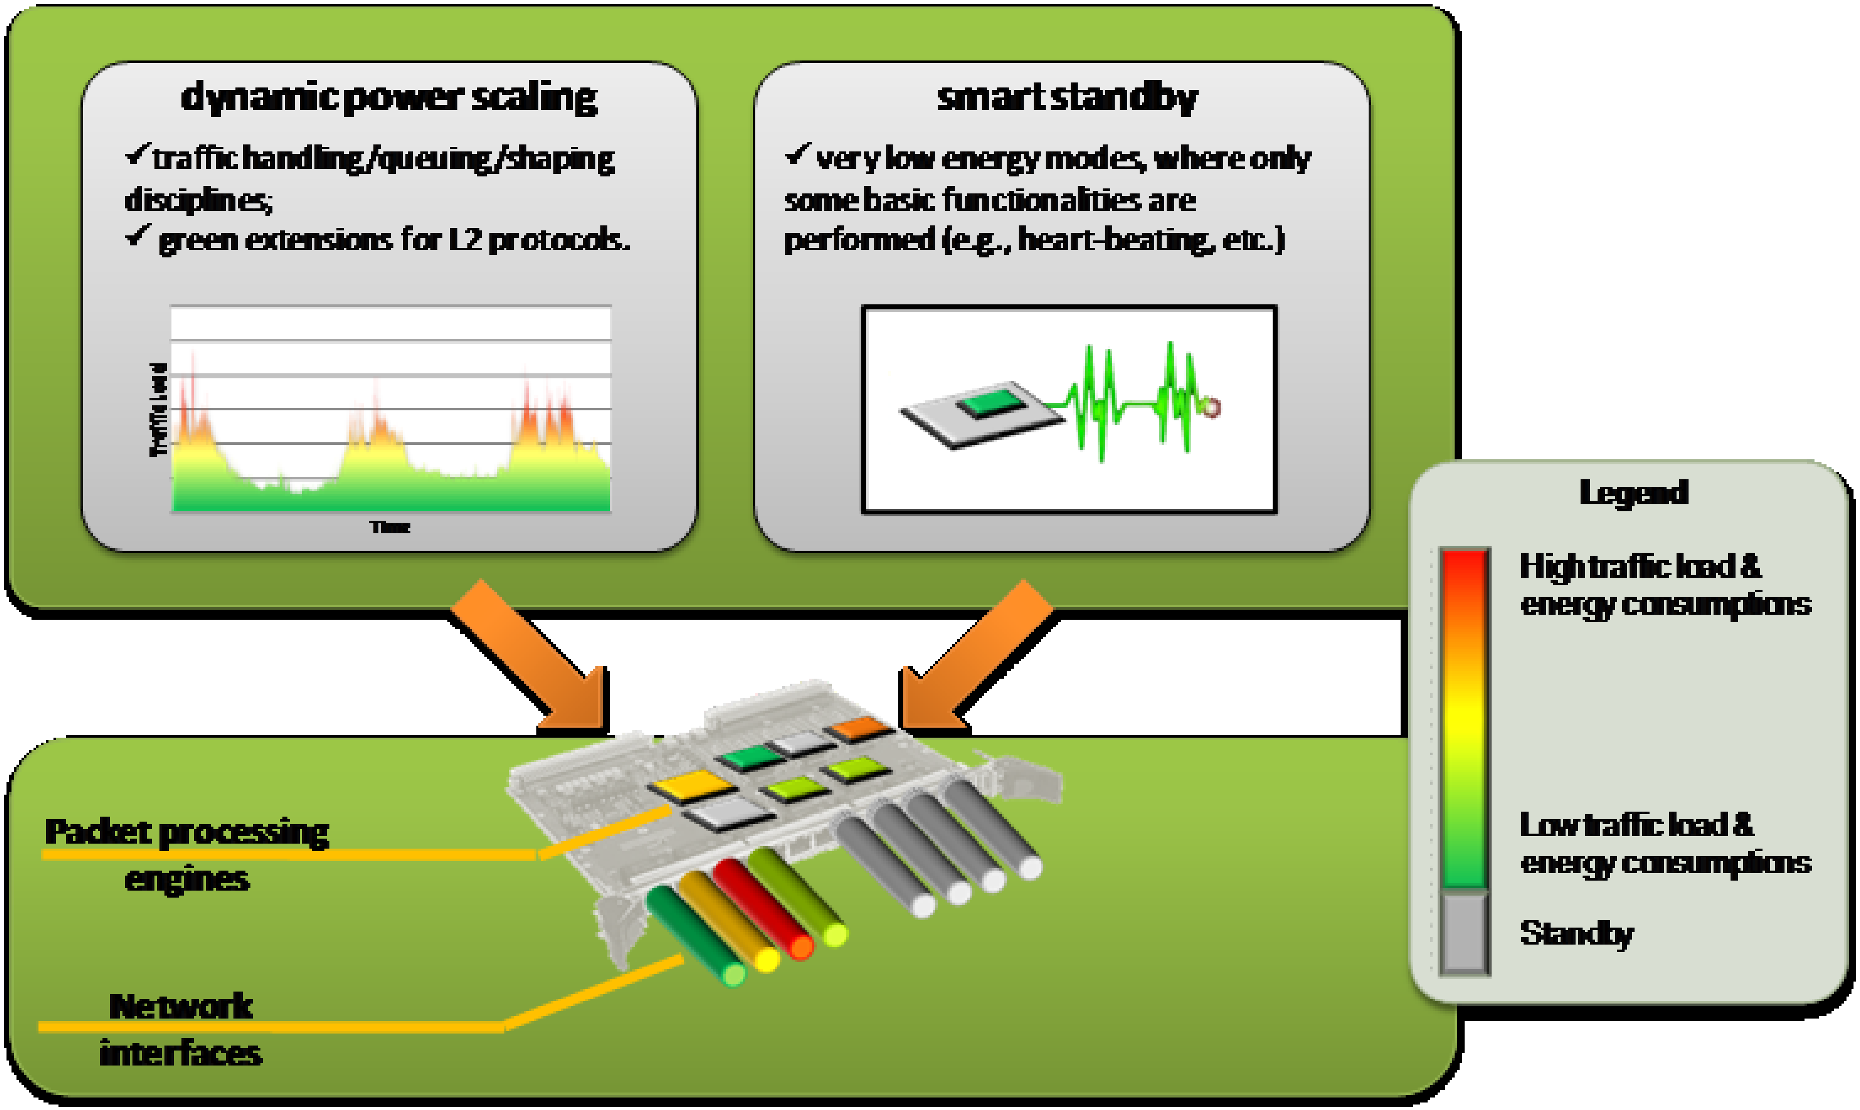

Implementing it all: The Control Plane

Autonomic and short-term on-line optimizations

Operator-driven long-term off-line optimizations

Local Control Policies (LCPs)

Given:
- the actual traffic workload from input links
- Local service requirements

Dynamically find the best energy-aware configuration

Routing & Traffic Engineering

Given:
- The traffic matrix
- Service requirements
- The energy-aware capabilities of network nodes and links

Dynamically move the traffic flows among network nodes in order to minimize the overall network consumption

Network-wide Monitoring

Given the history of measurements regarding:
- network performance
- energy consumption

The operator can explicitly plan and/or reconfigure the settings of:
- single device
- traffic engineering and routing.

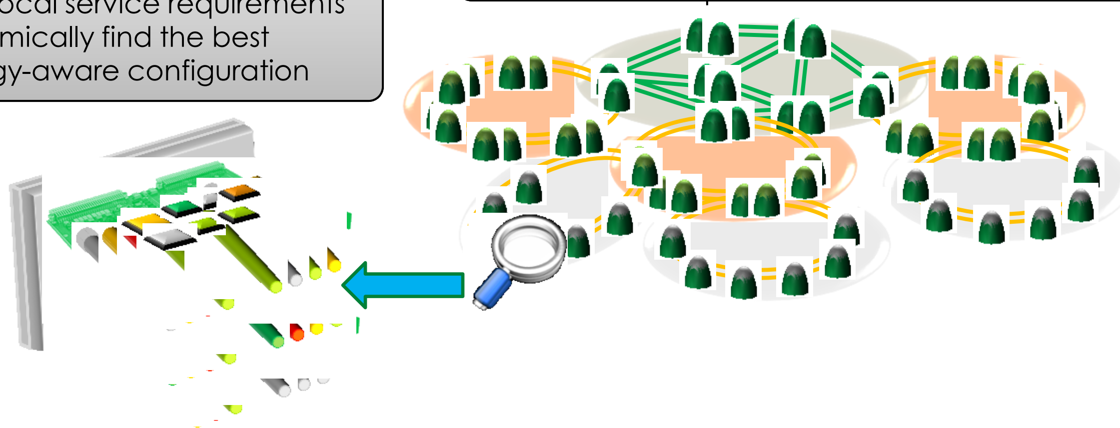

The Network Operation Center (NOC)

Network Control Policies (NCPs)

QoS vs Power Management

- **Energy consumption arises from the HW components inside the network device**
- Power management primitives can be natively applied to such HW components
- A network device can be composed by a huge number of HW components (**a device can be thought as a «network in a network»**)
- The overall configuration nee to be consistent (i.e., minimu power consumption & QoS constraint satisfation)

A Key Enabler: The Green Abstraction Layer

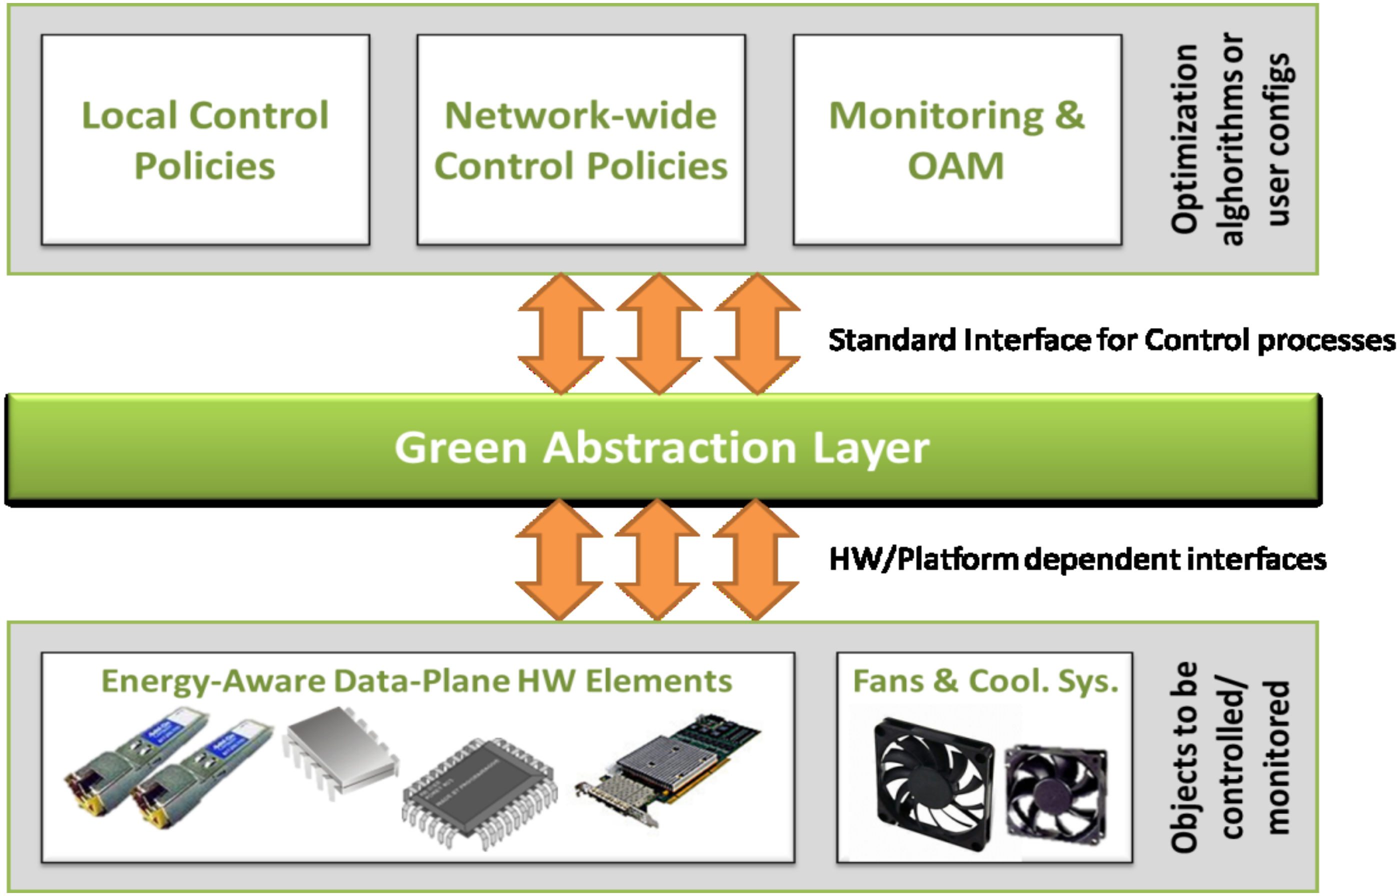

Smart standardization is required to enable efficient and network-wide dynamic power management

Device internal organization

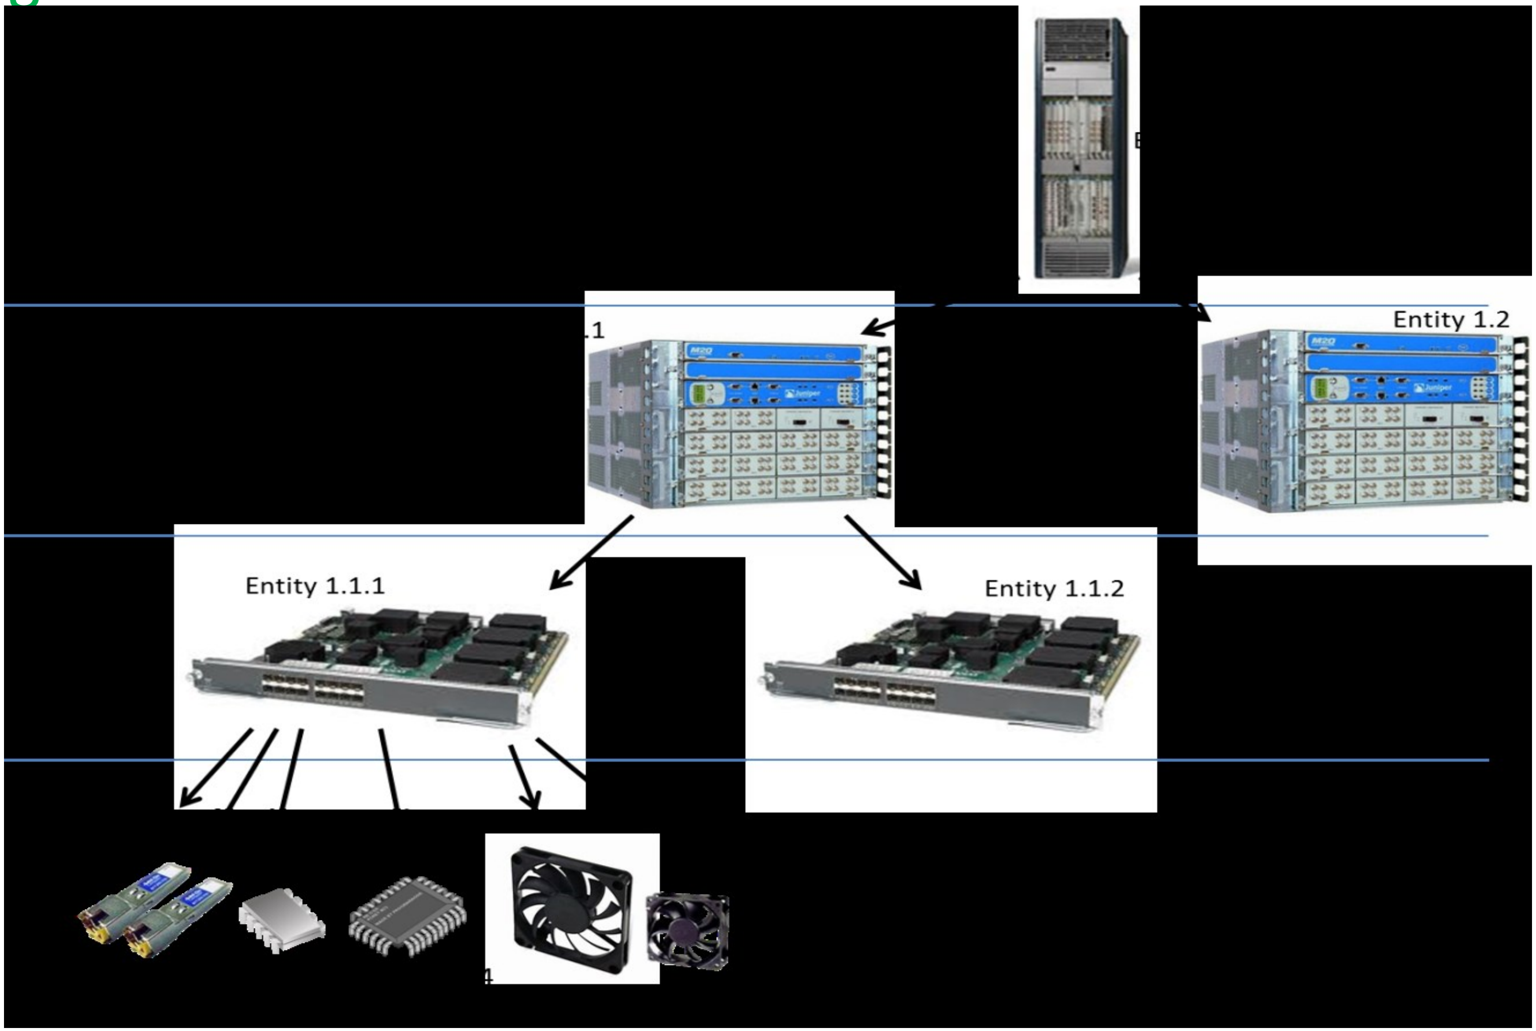

The GAL Hierarchical Architecture

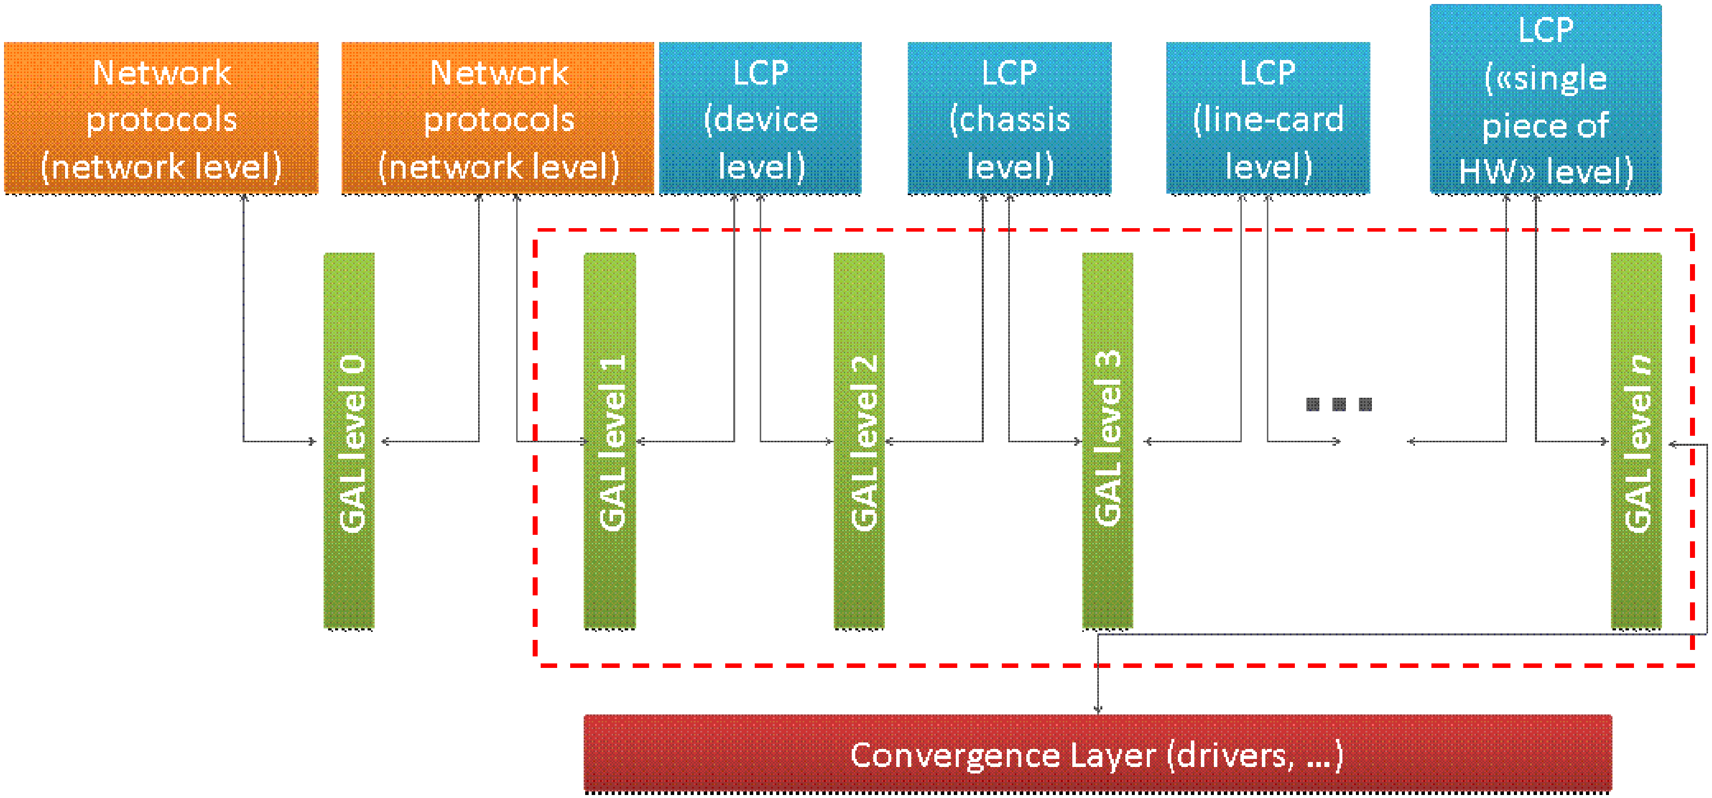

The GAL Hierarchical Architecture

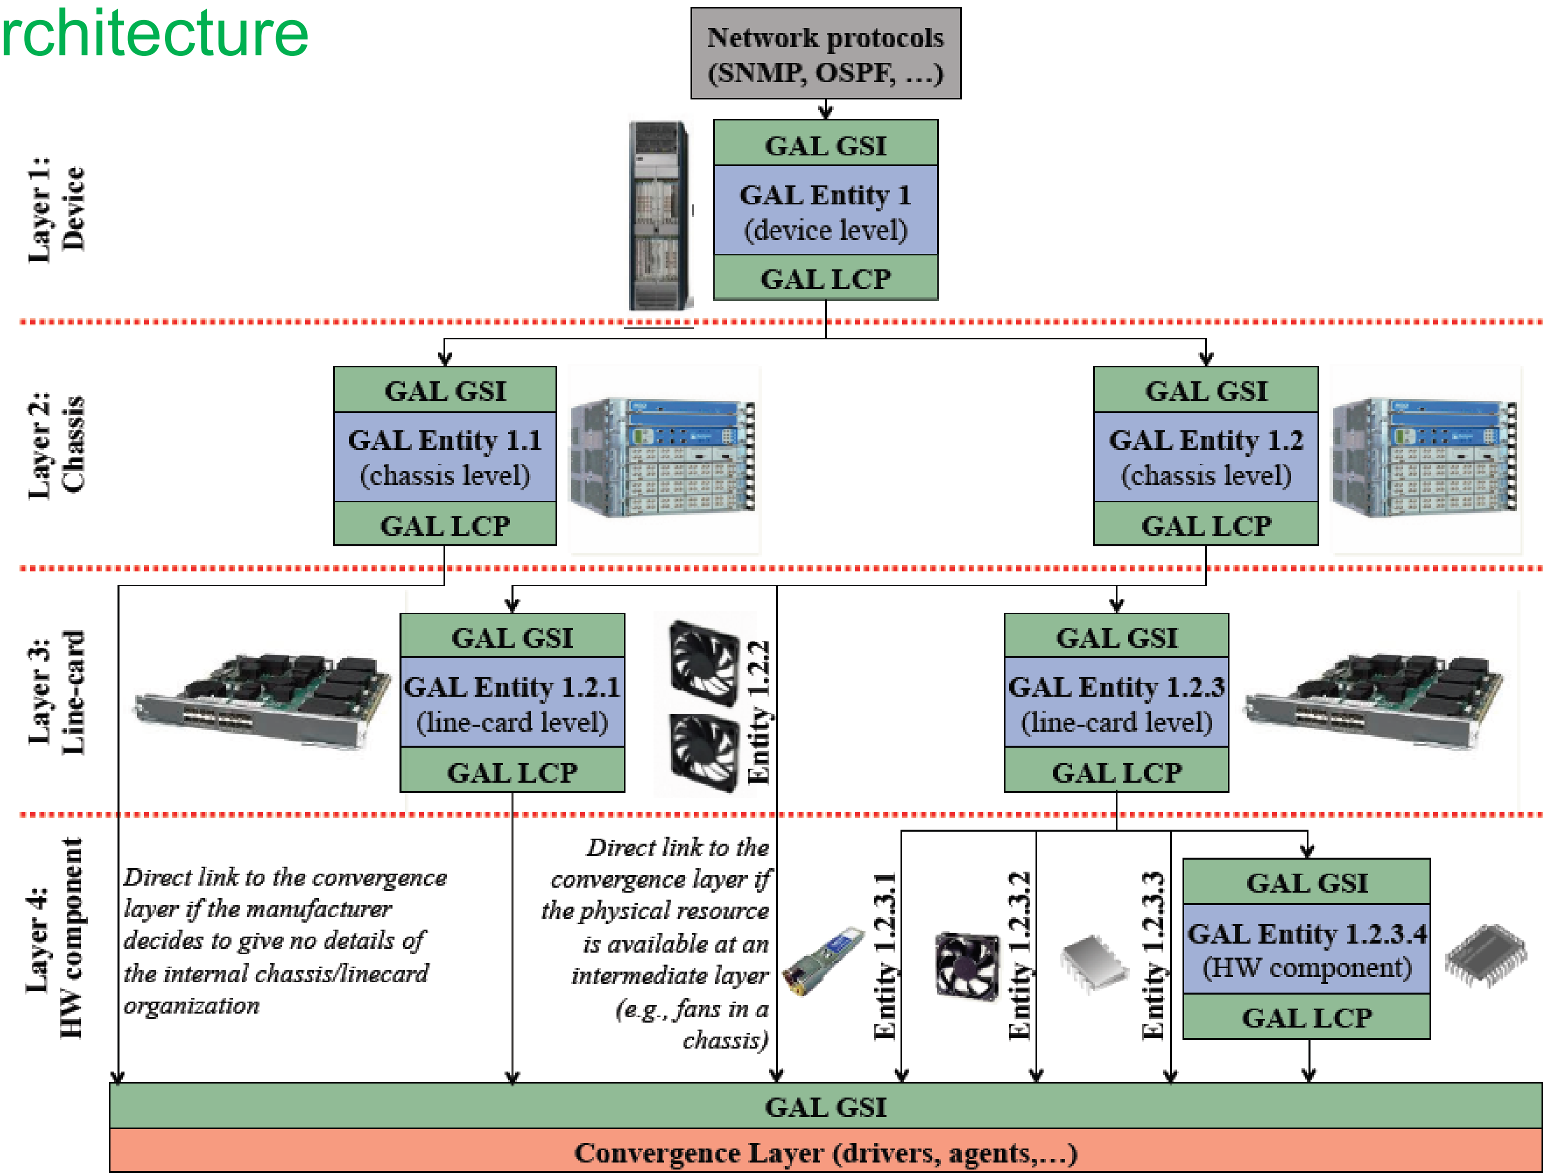

GAL - An example

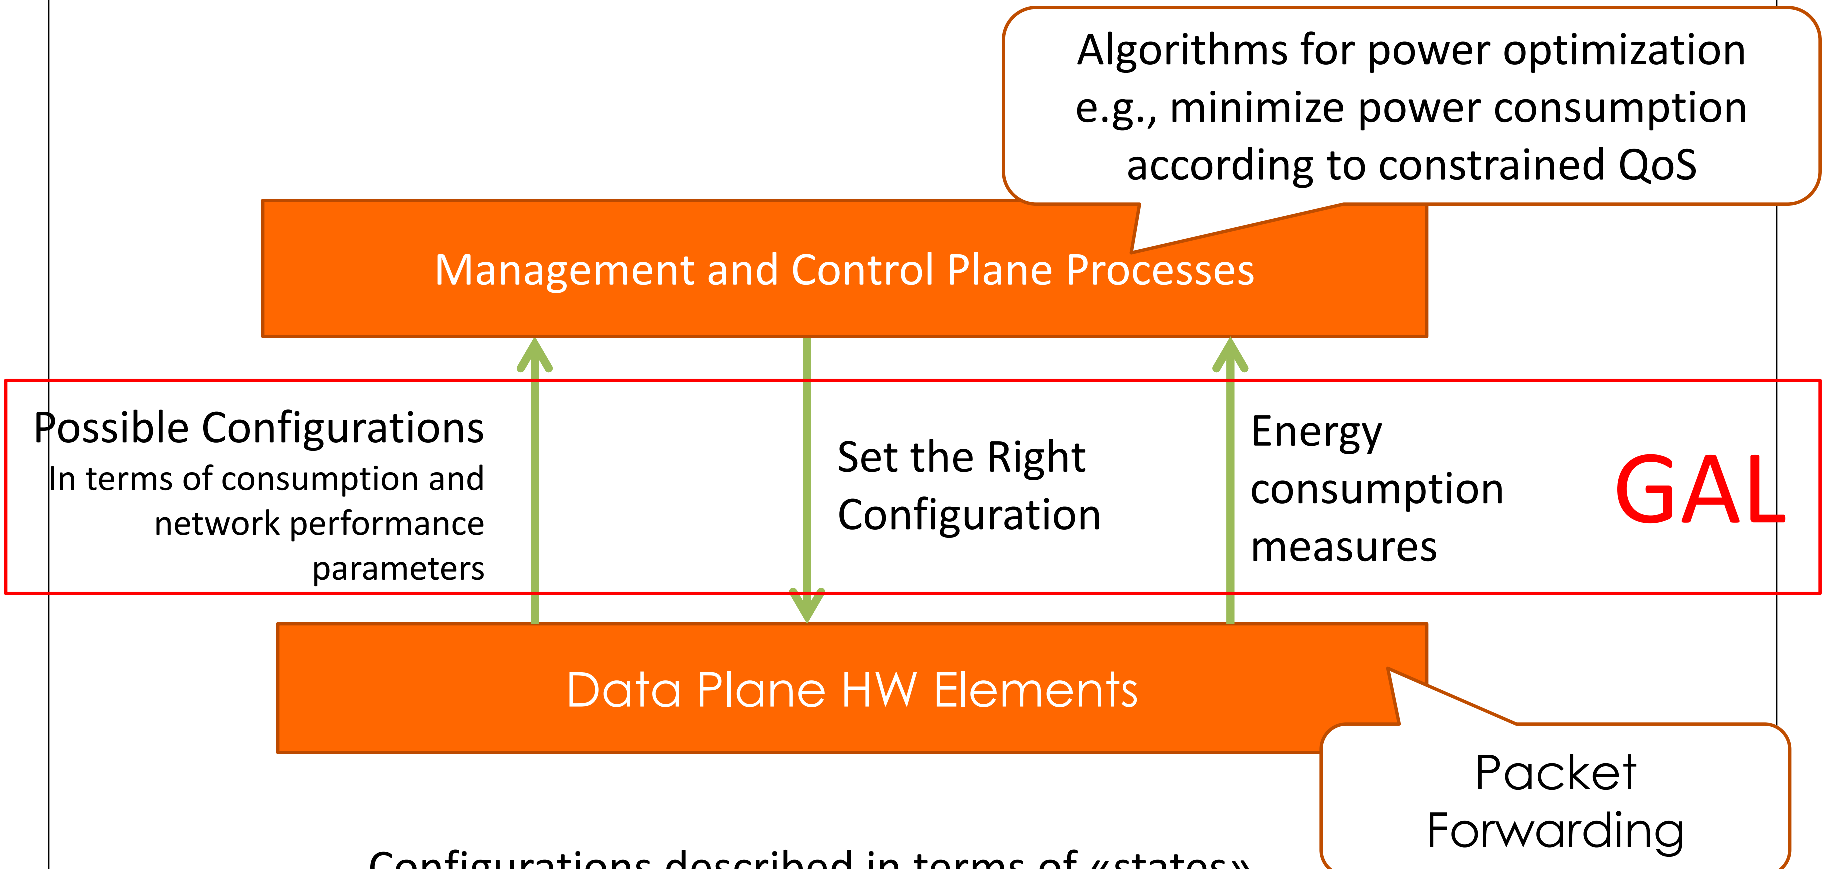

Configurations described in terms of «states»
A state is a stable configuration of the data plane HW

Energy-aware States

- An EAS can be modeled as a couple of energy-aware Primitive sub-States (PsS) related to the configuration of Standby and Power Scaling mechanisms:

$$EAS_n=\{P_j, S_k\} \text{ with } 0 \le j \le J \text{ and } 0 \le k \le K$$

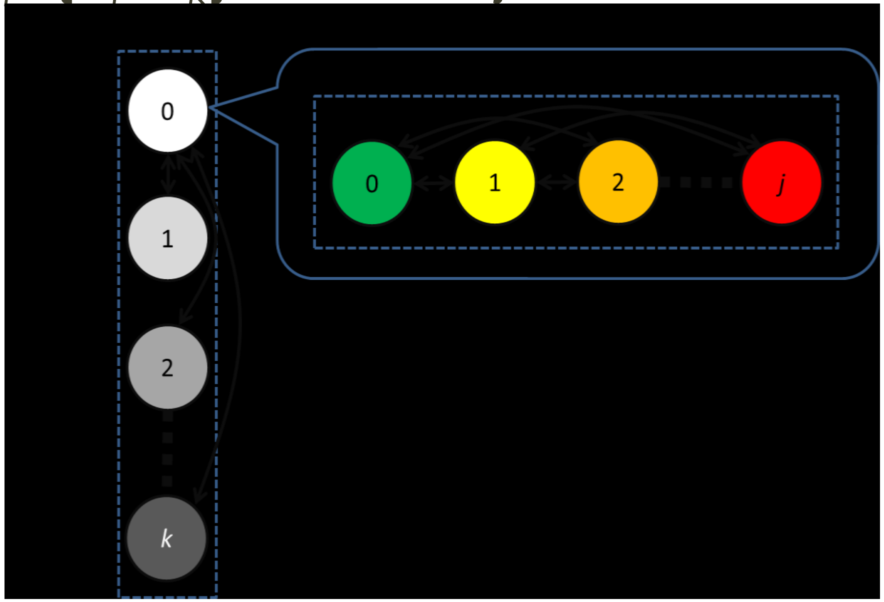

Definition of Energy-aware States

- An **Energy-Aware State (EAS)** was defined as an operational power profile mode implemented by the entity that can be configured by control plane processes.
- Energy-aware states are represented by a complex data type, which contains indications on the power consumption, the performance, the available functionalities, and the responsiveness of the entity when working in such configuration.
- Specific data types have been defined for the power scaling and standby capabilities, by taking into account different operational behaviours that can be provided by the implementations of such capabilities (e.g., autonomic or non-autonomic behaviours)

The GSI hierarchy

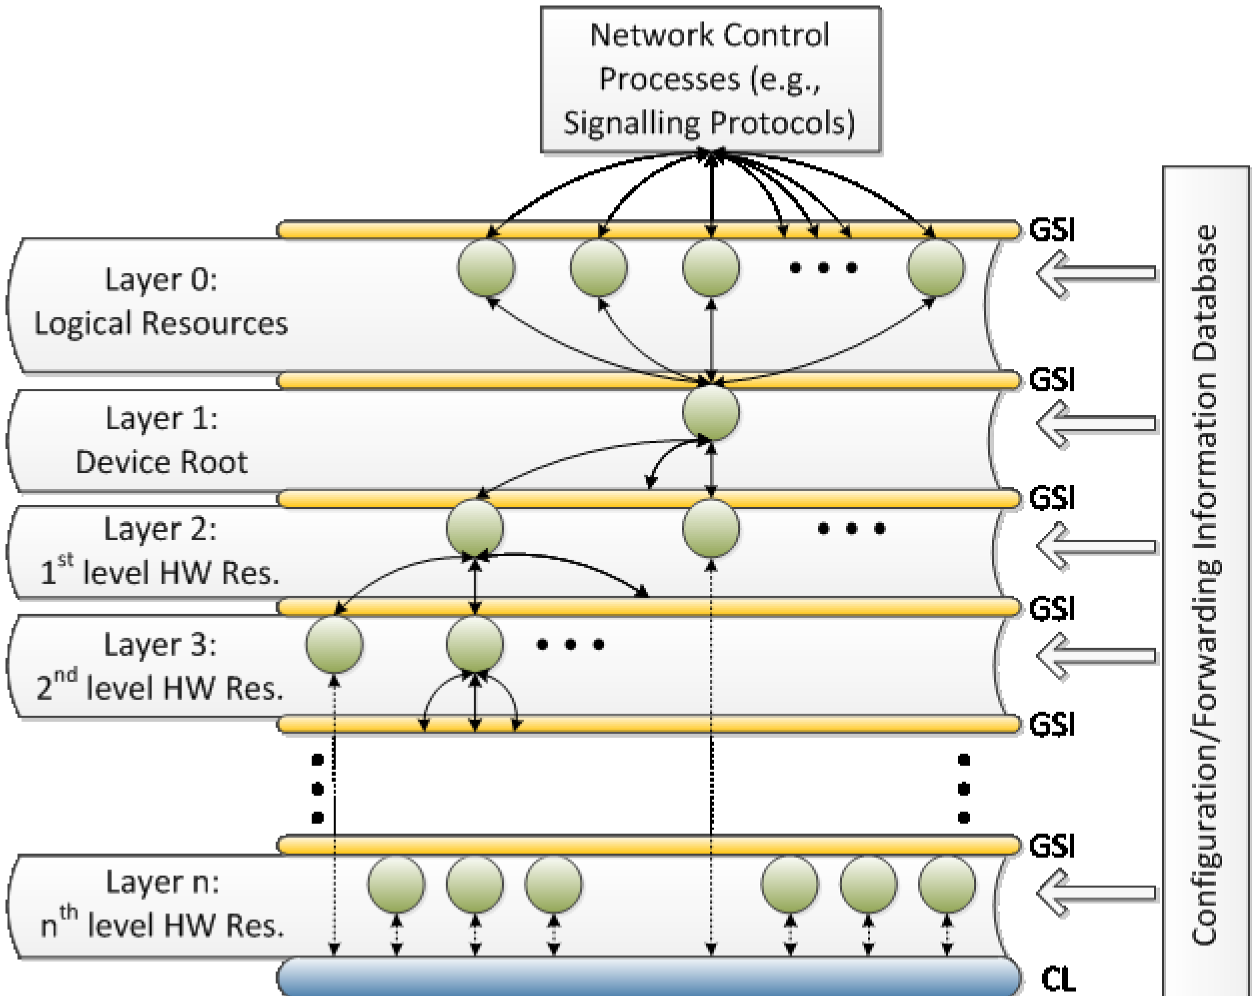

Toward standardization

- A specific Work Item on the GAL has been created by the ETSI Technical Committee on "Environmental Engineering" (TC-EE), named DES/EE-0030 "Green Abstraction Layer (GAL) power management capabilities of the future energy telecommunication fixed network nodes".
- The final definition of this standard is expected by the end of 2013.
- Coordination with the IETF EMAN group ongoing.

Source: *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.*

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, Jan. 2014 (to appear).

Back to SDN/NFV and the GAL

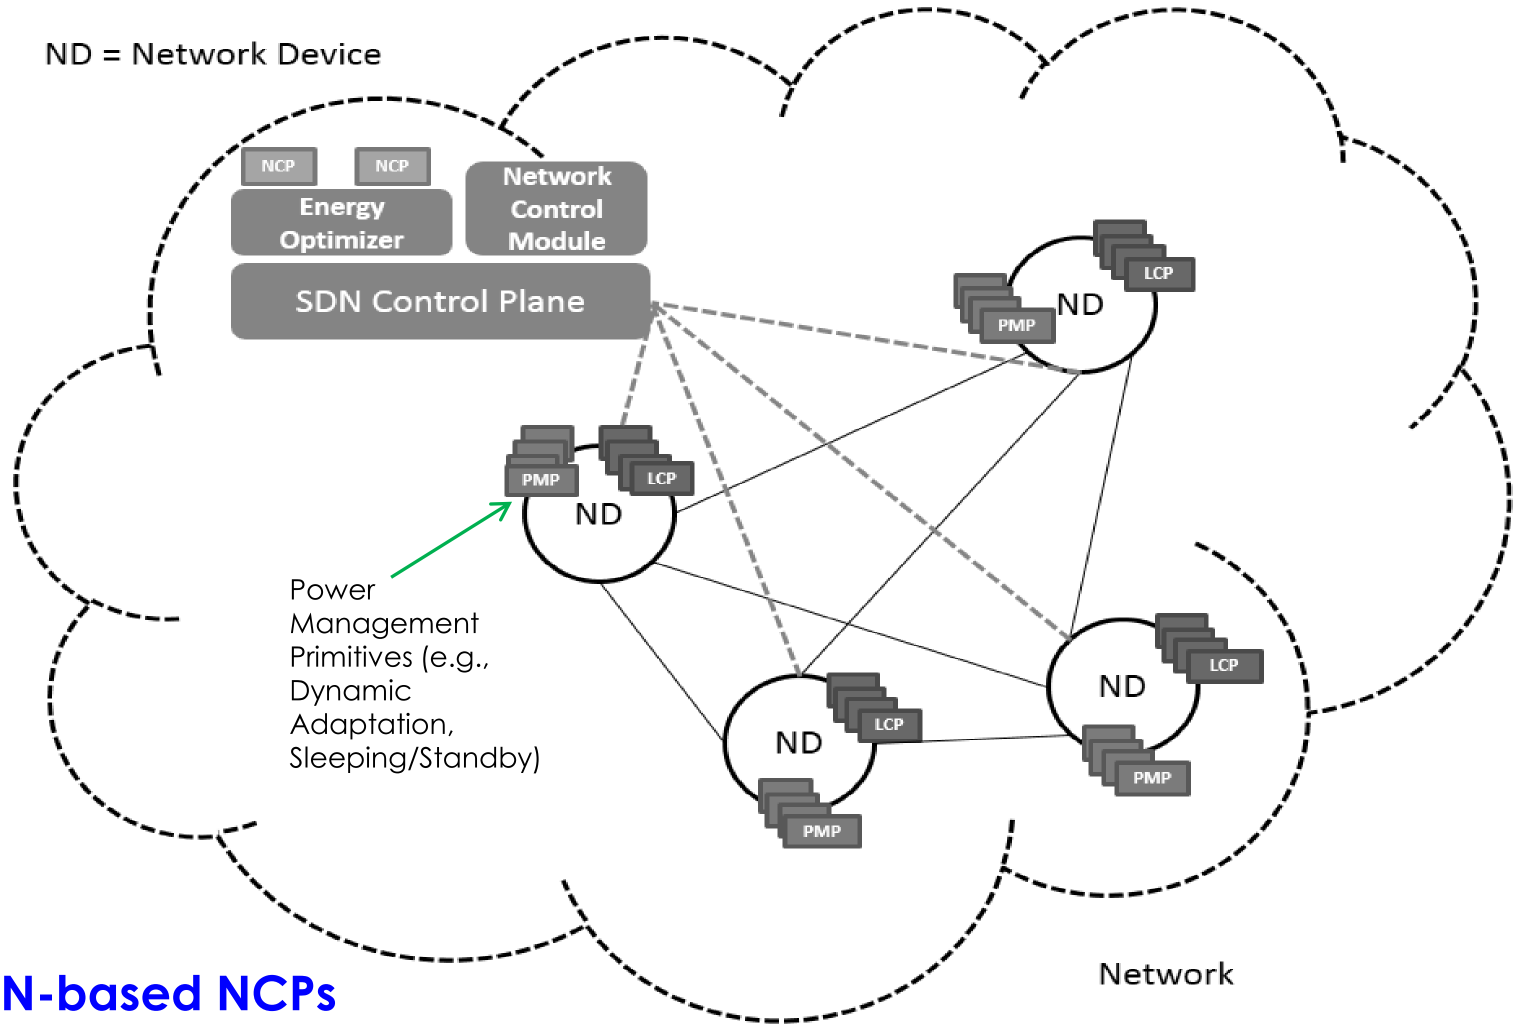

SDN-based NCPs

Back to SDN/NFV and the GAL

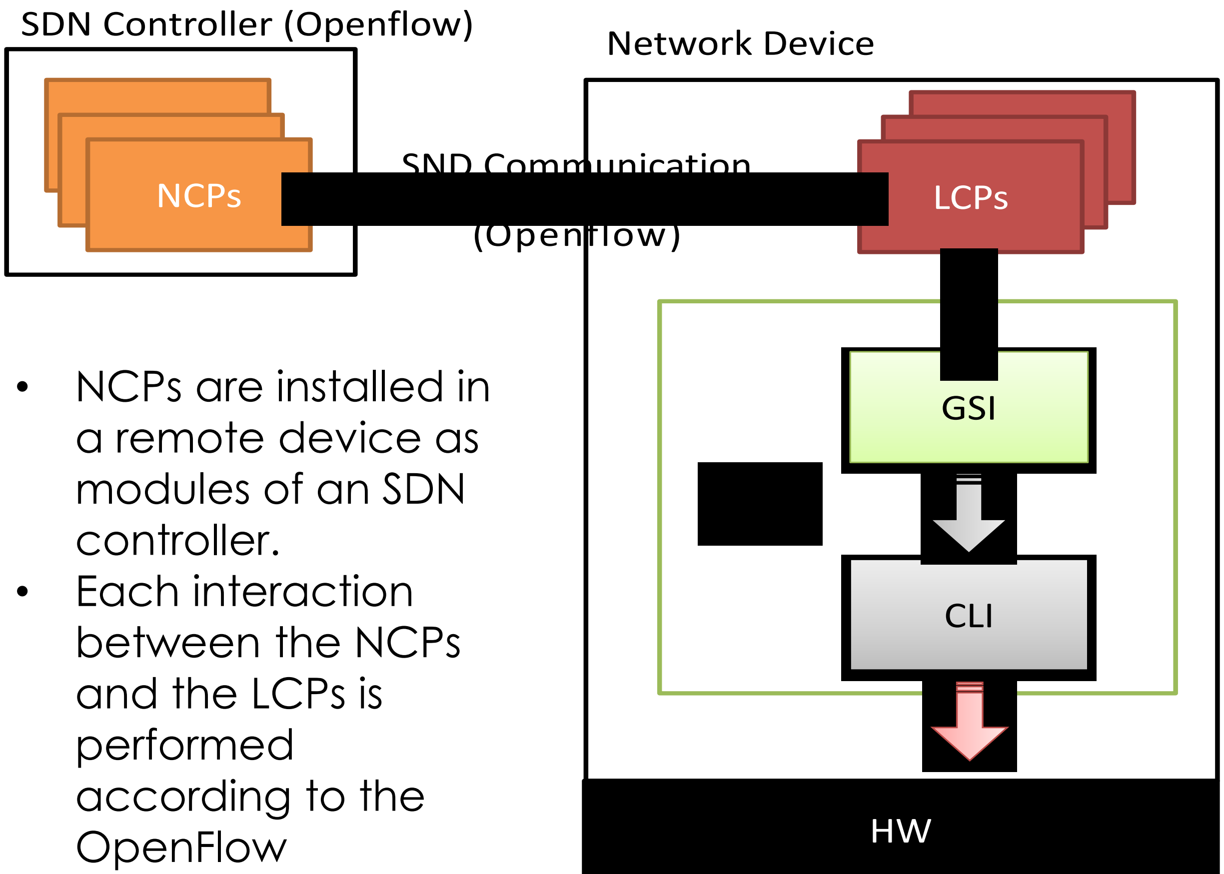

- NCPs are installed in a remote device as modules of an SDN controller.
- Each interaction between the NCPs and the LCPs is performed according to the OpenFlow Specification.

- LCPs set and get the energy-aware configuration by means of the 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

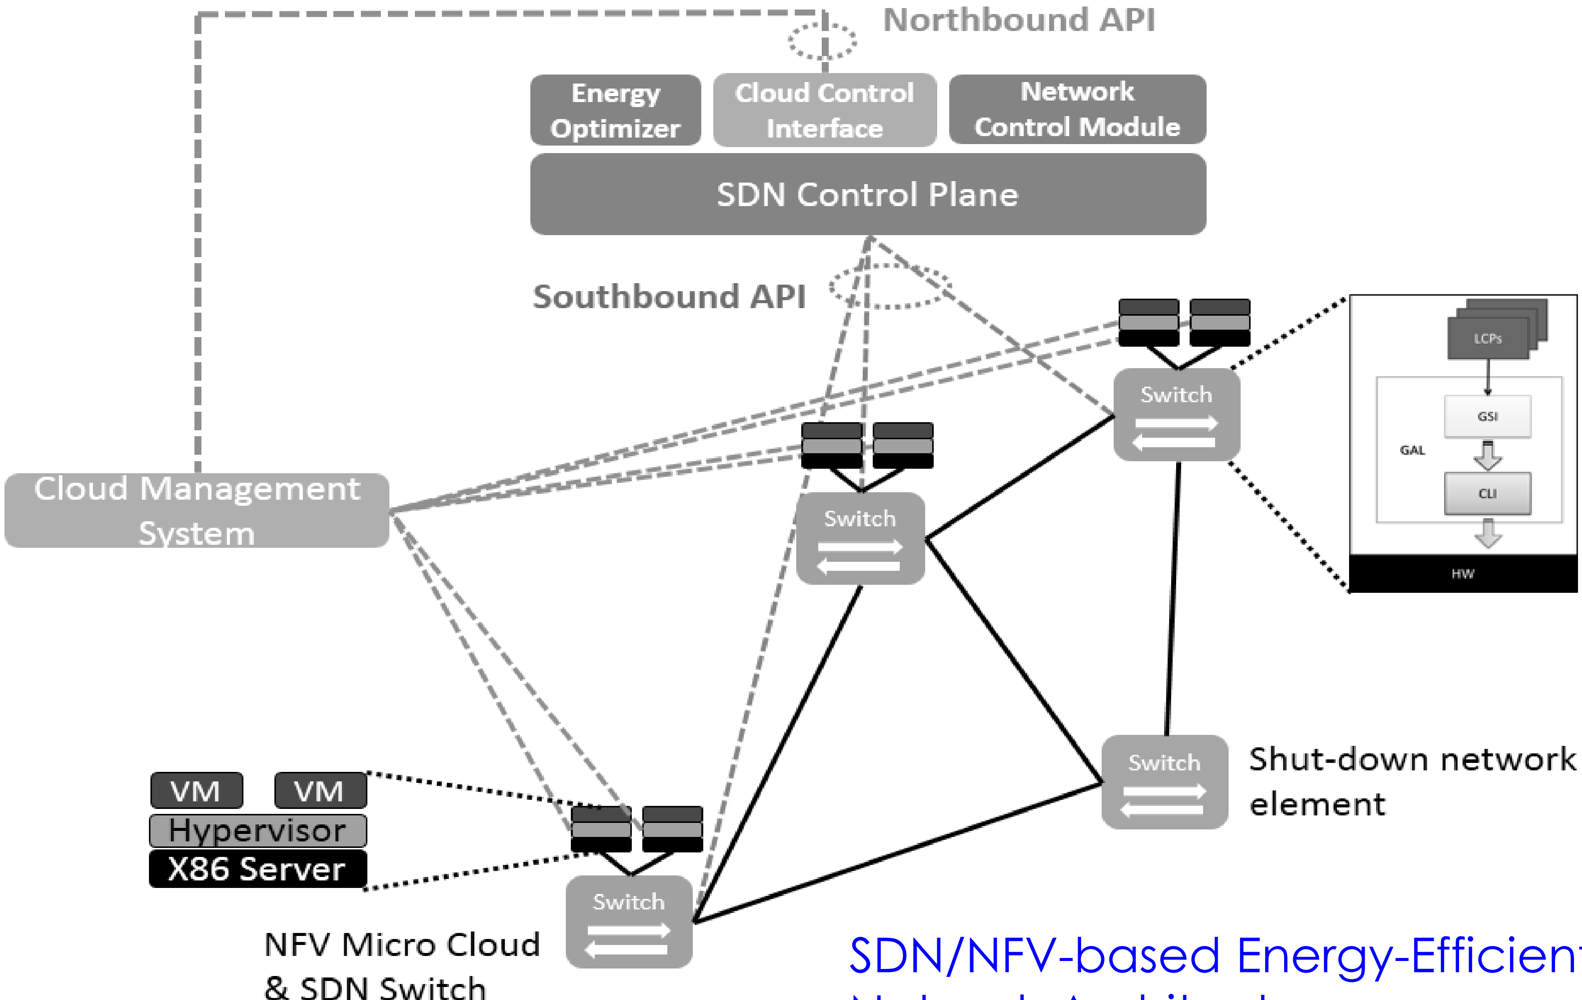

SDN/NFV-based Energy-Efficient Network Architecture

Back to SDN/NFV and the GAL

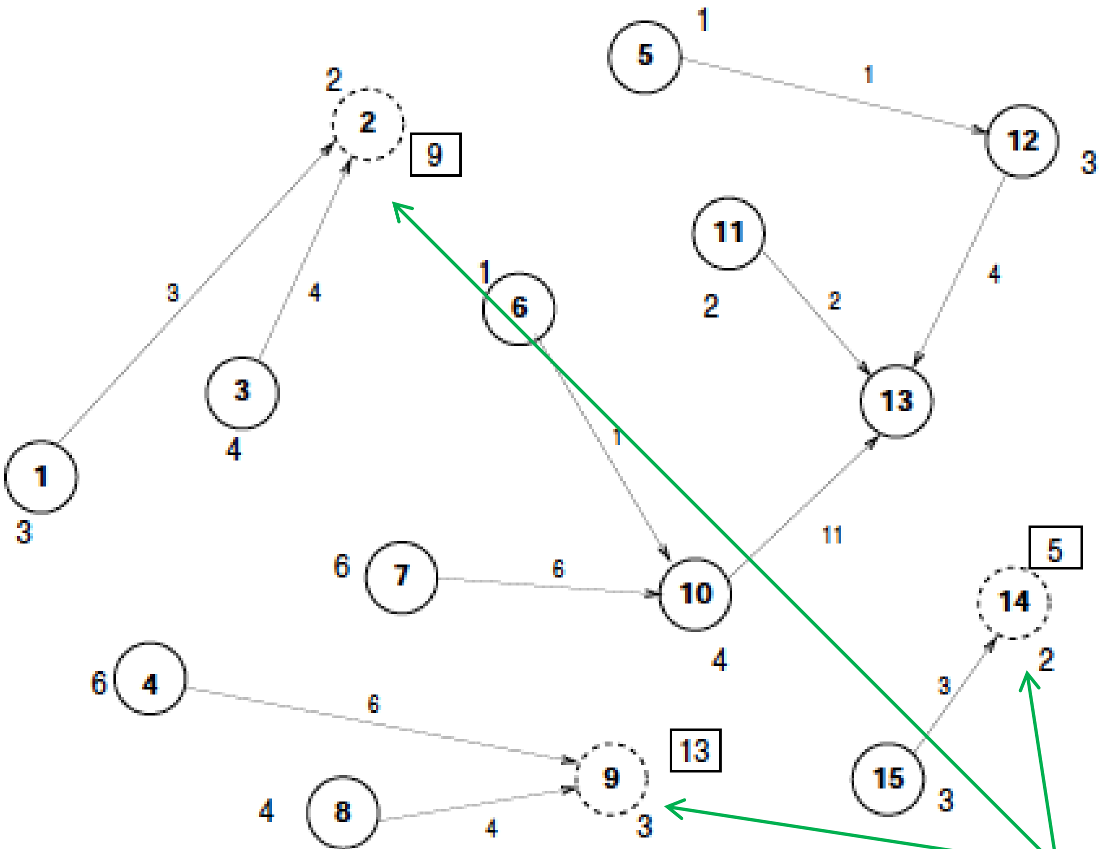

Independently Routable Traffic at the SDN-FEs

SDN-enabled nodes

Even partial SDN deployment may be beneficial

S. Agarwal, M. Kodialam, T. V. Lakshman, "Traffic Engineering in Software Defined Networks,", Proc. IEEE INFOCOM 2013, Torino, Italy, pp. 2211 – 2219.

Sleeping/standby

L2 Virtualization & Standby

- Consider a network scenario similar to the state-of-the-art backbone networks deployed by Telcos, where IP nodes have highly modular architectures, and work with a three-layer protocol stack.

Usually, IP does not work on physical resources, but on the terminations of L2 links…

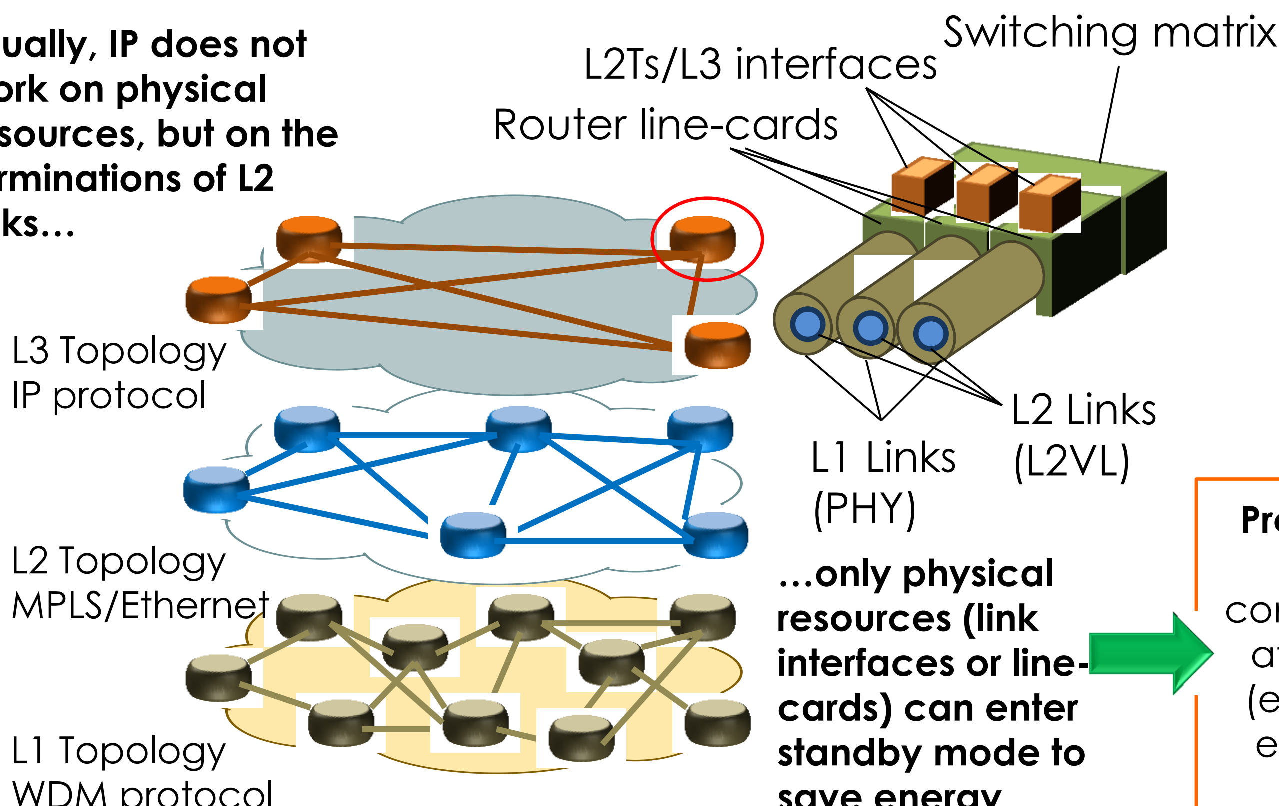

…only physical resources (link interfaces or line-cards) can enter standby mode to save energy

Problem: Network stability, convergence times at multiple levels (e.g., MPLS traffic engineering + IP routing)

L2 Virtualization & Standby

- ***Exploiting modularity***: making line-cards left active to “cover” sleeping parts, without the device losing any networking resource/functionality.
- ***Virtualizing (IP) network functionality***: before a line-card enters standby status, it has to transfer its resources and activated functionalities to other cards that will remain active.

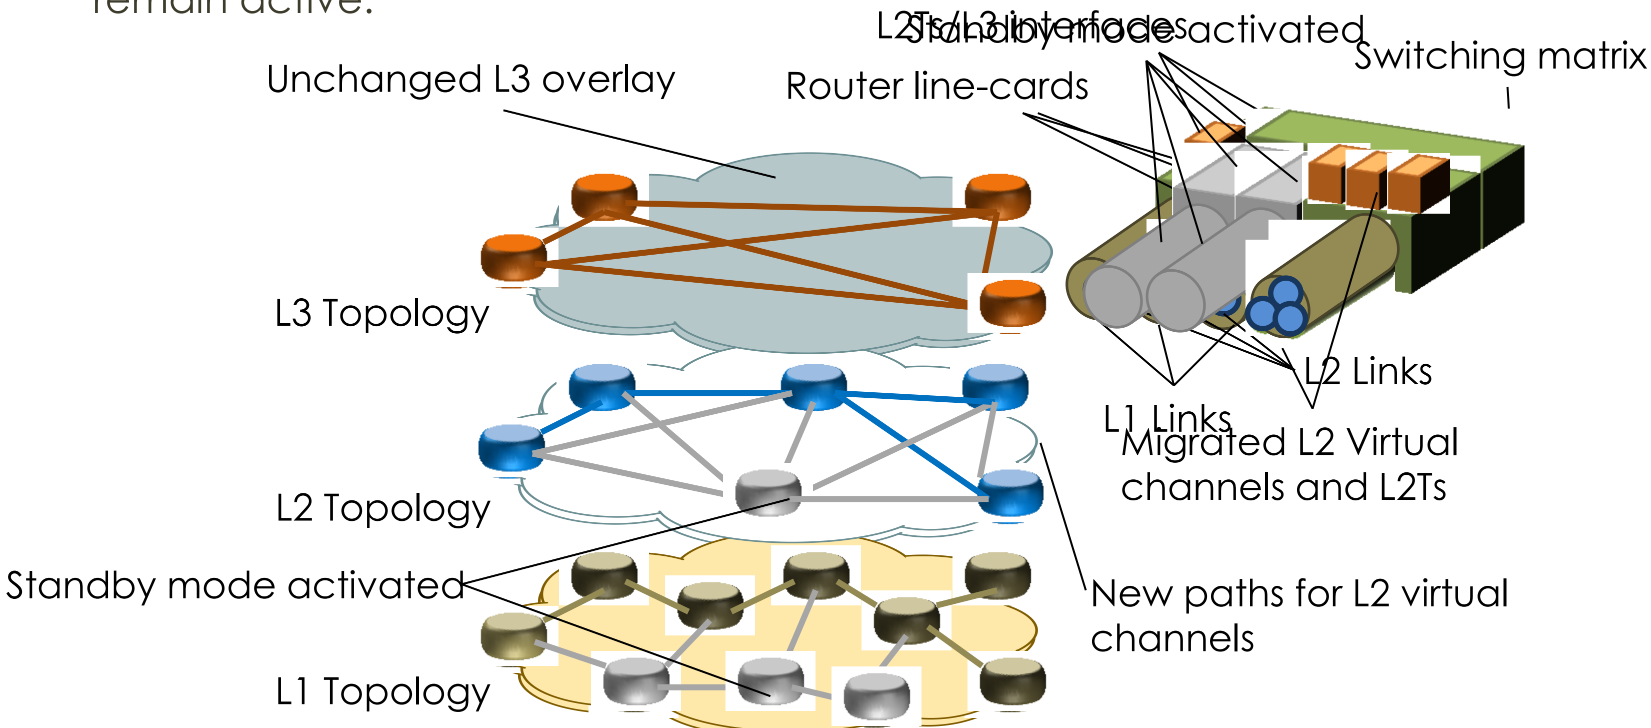

Source: R. Bolla, R. Bruschi, A. Cianfrani, M. Listanti, “Enabling Backbone Networks to Sleep”, *IEEE Network*, vol. 25, no. 2, pp. 26-31, March/April 2011.

The Distributed Router Open Platform (DROP) and NFV

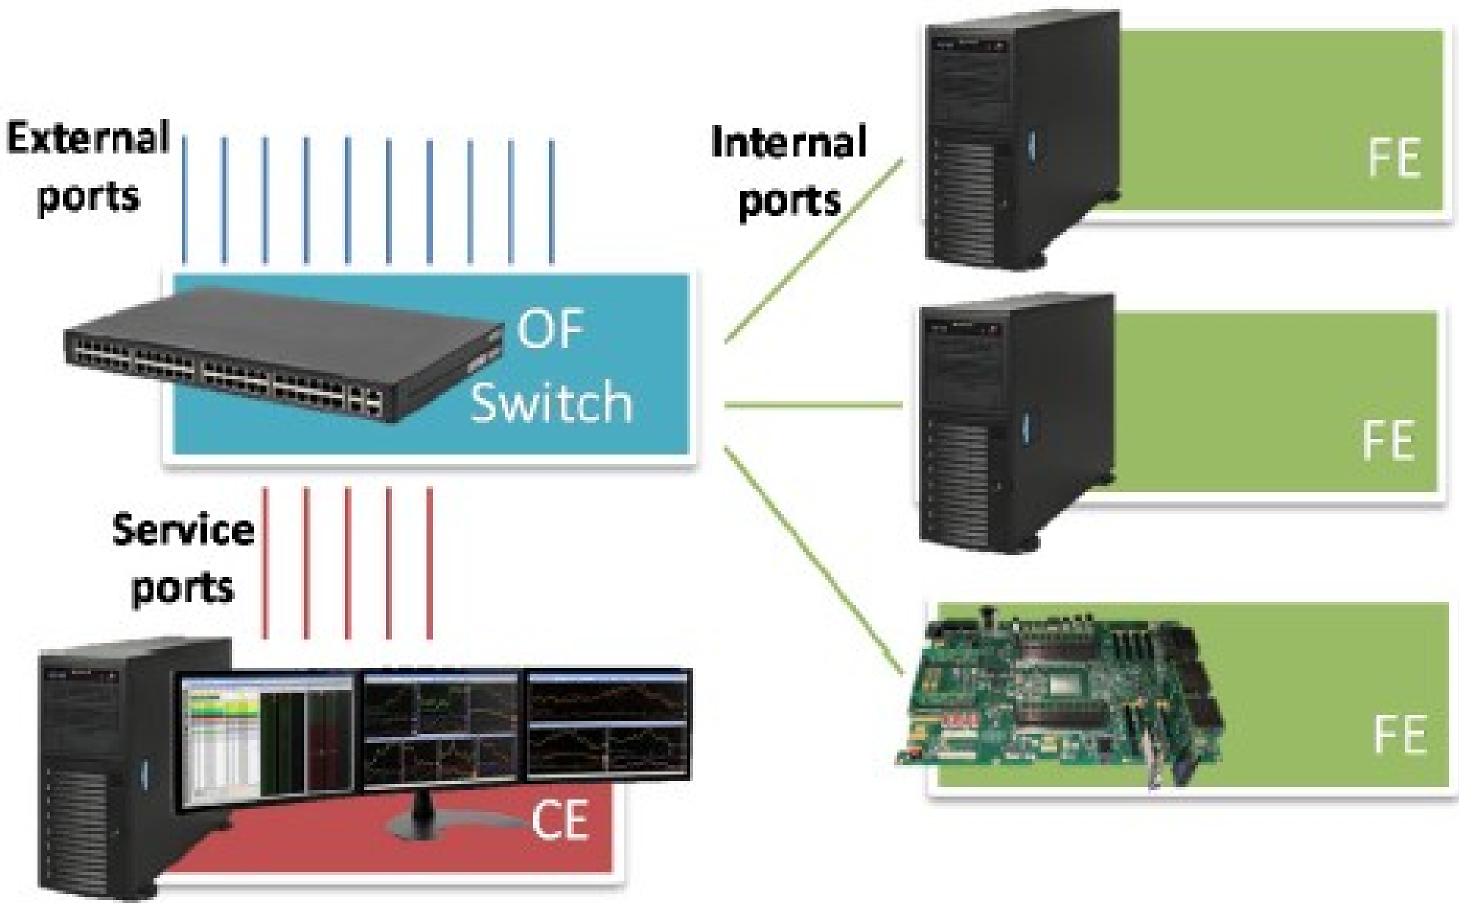

The Distributed Router Open Platform (DROP) and NFV

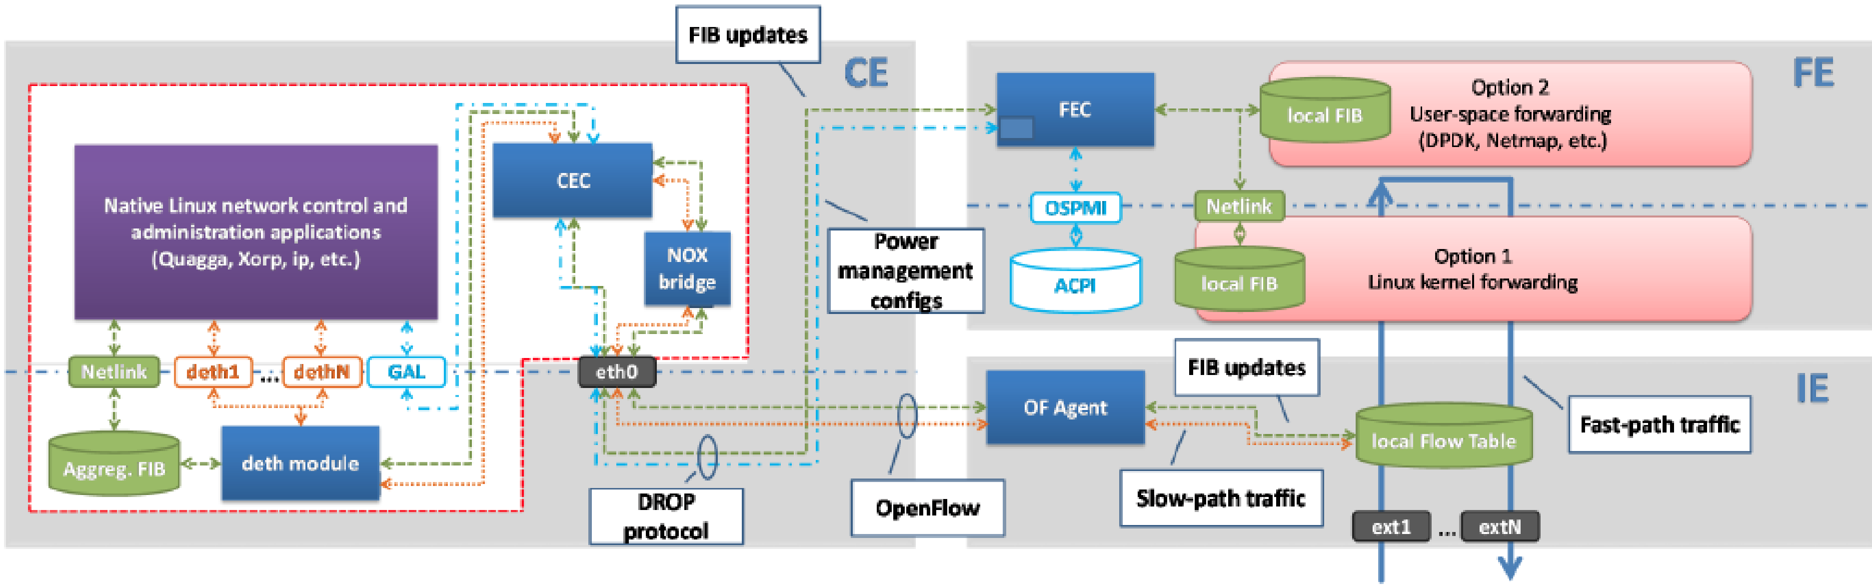

Source: *R. Bolla, R. Bruschi, C. Lombardo, S. Mangialardi, "DROPv2: Energy-Efficiency through Network Function Virtualization," IEEE Network (under review).*

The Distributed Router Open Platform (DROP) and NFV

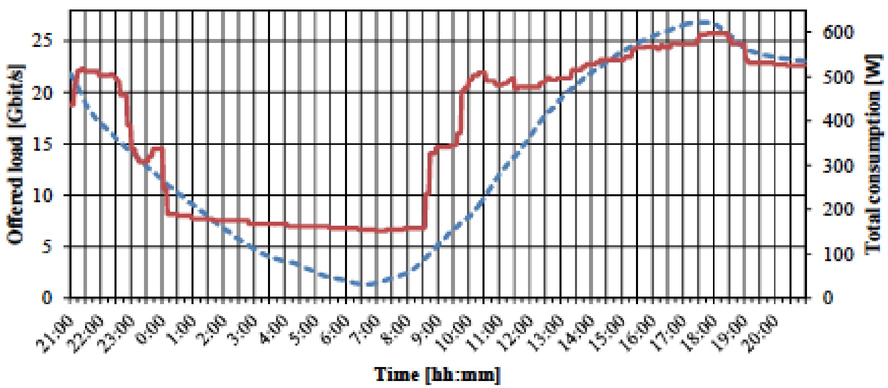

(a) Total

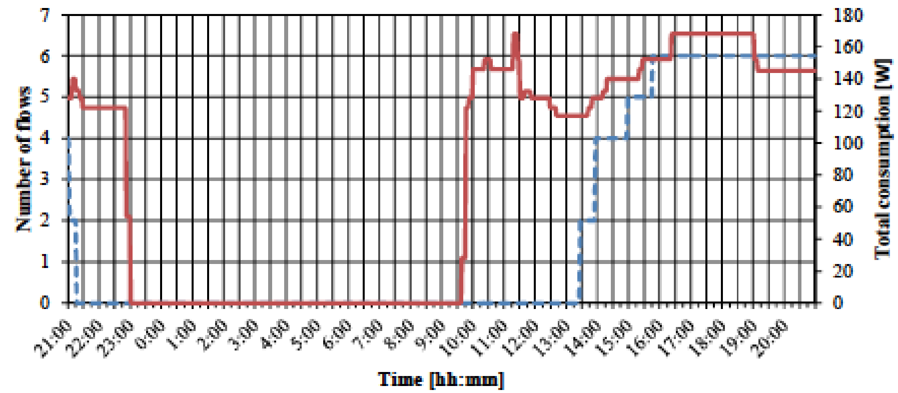

(b) First Xeon based FE

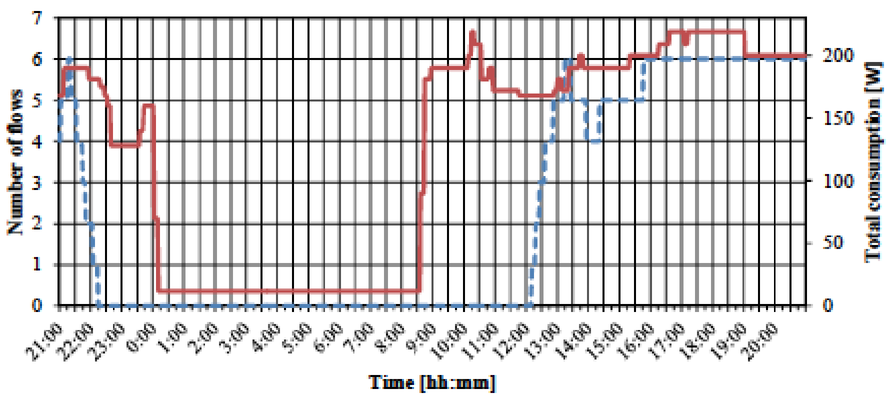

(c) Second Xeon based FE

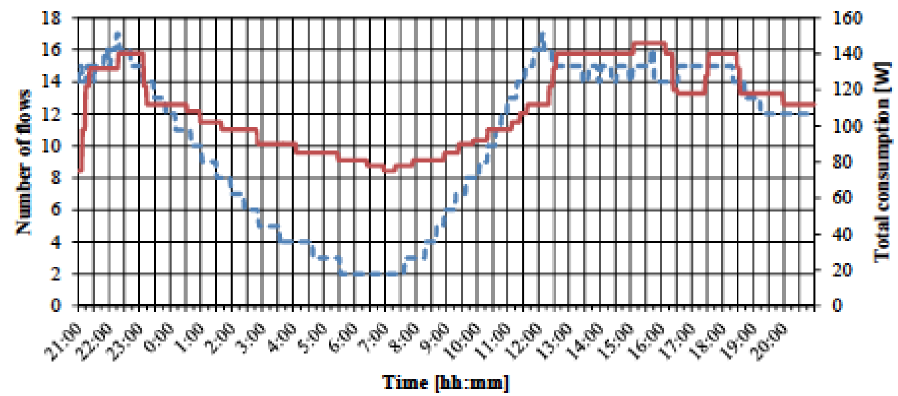

(d) XLP

Source: *R. Bolla, R. Bruschi, C. Lombardo, S. Mangialardi, "DROPv2: Energy-Efficiency through Network Function Virtualization," IEEE Network (under review).*

Conclusions - 1

- 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.

Conclusions - 2

- Several challenges to be faced
  - Scalability of the SDN environment
    - avoiding excessive flow table entries
    - avoiding Control – Data Plane communications overhead
    - managing short- & long-lived flows
  - Controller placement and (dynamic) allocation of switches
  - Cross-domain solutions
  - Defining Northbound APIs to enable real network programmability
  - More virtualization (multiple slices)?
  - Migrating virtual machines / consolidation across WAN domains