

# *The Green Abstraction Layer: A Standard Power Management Interface for Next- Generation Network Devices*

Raffaele Bolla<sup>1</sup>, Roberto Bruschi<sup>1</sup>, Franco Davoli<sup>1</sup>, Lorenzo Di Gregorio<sup>2</sup>, Pasquale Donadio<sup>3</sup>,  
Leonardo Fialho<sup>4</sup>, Martin Collier<sup>4</sup>, Alfio Lombardo<sup>5</sup>, Diego Reforgiato Recupero<sup>5</sup>, Tivadar Szemethy<sup>6</sup>.

1- CNIT, Genoa, Italy

2- Lantiq, Munich, Germany

3- Alcatel-Lucent, Italy

4- Dublin City University, Dublin, Ireland

5- Lightcomm, Catania, Italy

6- Netvisor, Budapest, Hungary

**Abstract** – The energy consumption of telecommunication networks can be reduced using dynamic power management techniques, which allow network performance to be traded off against energy consumption. Such techniques can be applied on devices by means of hardware-specific local control policies. Distributed power management across heterogeneous hardware requires a standardized representation of the capabilities of each hardware system and subsystem, to decouple distributed high-level algorithms from hardware specifics. We propose a Green Abstraction Layer (GAL), to provide such a standardized interface between the high-level algorithms and a lower level that represents the hardware and physical resources, where energy management actions are directly exerted. Since many network devices contain complex subsystems, this lower level can itself be subjected to a hierarchical decomposition. A key design constraint is for the GAL to be lightweight, so as not to incur a high energy penalty.

## *1 Introduction*

Improving network energy efficiency has become an important goal for Telco operators across all network segments, owing to the increasing power consumption exhibited over time [1] and to the increasing cost of energy. Estimates of public organizations, like the GeSI [2], provide alarming figures in the “Business As Usual” case; for the European Telecommunication companies alone, the estimated energy requirement in 2020 is 35.8 TWh, which would represent an increase of over 67% in a decade.

Such high power consumption arises from increasing user traffic and router capacities, which are not compensated for by a corresponding increase in silicon energy efficiency. Moreover, most network hardware is designed to operate constantly at maximum capacity, irrespective of the traffic load, even though their average utilization lies far below the maximum [3]. These observations have motivated adapting the network capabilities, and thus its energy requirements, to the actual traffic profiles [4,5]. This can be realized by providing *power management primitives* (PMPs) in the hardware platforms of networking devices, where energy absorption physically happens, similar to features of general-purpose computing systems (e.g., ACPI [6]), which interact with a *control policy* whose goal is to find a system configuration that consumes no more power than is necessary.

The main contribution of the present paper is the specification of a layer that provides a uniform way for the network control plane to access the green networking capabilities of devices. This novel standard interface is called the Green Abstraction Layer (GAL) [7]. The goal of the GAL is to provide:

- (i) a common and simple way for representing the power management capabilities available in heterogeneous data plane hardware;
- (ii) a framework for information exchange between power-managed data plane entities and control processes;
- (iii) a reference control chain allowing a consistent hierarchical organization of multiple local and network-wide energy management protocols.

Mediation by the GAL allows energy management protocols to determine in a consistent fashion which power management capabilities are available at the data-plane, their potential effect on both energy consumption and network performance, and how to interact with them. This allows the protocols to be insulated from the specific hardware features of the underlying devices, and yet enable them to choose the most suitable energy configuration.

## 2 Concepts in Network Energy Management

Dynamic adaptation of energy consumption is of particular benefit in network equipment, given their high *crest factor* (the ratio of peak to average load). The key concepts and facilities required to support this are considered below.

*PMPs* allow the energy consumption of networking devices or subsystems to be aggressively modulated, by putting them into standby states when not in use, or by decreasing their performance when incoming traffic is light. These are features locally available in network nodes, and their implementation and efficiency will be hardware-specific.

A means to provide the required Quality of Service (QoS) level to incoming traffic, whilst consuming the minimum power to achieve this, would invoke these primitives. One approach is to provide each network device with its own independent control algorithm, which effects a *local control policy (LCP)*; this algorithm may dynamically orchestrate the configuration of internal components (e.g., line-cards, link interfaces, network processors, etc.) to meet the desired QoS with the minimum consumption according to the current traffic. Superior results are likely to be obtained using a network-wide algorithm. This would implement a *network control policy (NCP)*.

Deploying such policies across heterogeneous equipment will require the representation of management and control actions and device/network status information in some standard abstract and implementation-agnostic form.

As in the ACPI [6] and the IETF EMAN [8] approaches, among others, such information can be abstractly represented as *energy-aware states (EASs)*. Both EMAN and ACPI provide pre-defined, non-extensible state sets. Furthermore, the power state description and characterization that existing standards provide is clearly insufficient for any remote control decision (e.g., lack of performance tradeoff indications in power state descriptions).

## 3 Power Management Primitives

PMPs allow dynamically trading off energy consumption and processing performance of hardware components. These primitives are usually realized at the hardware level, by introducing specific elements (e.g., clock frequency dividers, supply voltage regulators, sleeping transistors, etc.), either inside the power-managed component or close to it.

Such primitives can be classified on the basis of their effect on the working behavior of the managed component. According to the taxonomy proposed in [5], a basic classification is the following:

- *Standby primitives* allow devices or parts of them to enter very low energy states, where most of their functionalities are frozen. Wake-up times from standby modes are generally much faster than the times for recovering from off modes.
- *Power scaling primitives* allow the working rate of the component to be changed dynamically.

### 3.1 *Energy-Aware States*

An EAS is defined as a power setting that can be configured through the GAL, and that may impact on the power consumption, the performance, the available functionalities, and the responsiveness of the network device. To assure the correct interaction between the controlled entity and control plane processes, the EAS must report the entity's power consumption, performance, available functionalities, and its responsiveness when working in that state. These indexes are obviously necessary for control policies in order to decide which EAS provides the desired trade-off between power consumption and network performance.

Typically, the energy consumption and performance when in a given state will depend on the characteristics of arriving traffic and thus will not be known a priori. Hence, the most intuitive solution is to make such indexes describing the high-level effects introduced by power management techniques (i.e., speed reduction, wake-up times, etc.), which are independent of the system dynamics.

### 3.2 *Control Policies*

Recent suggestions for extending current routing and traffic engineering policies to explicitly consider network-wide energy consumption provide ways to implement a NCP [9-**Error! Reference source not found.**]. By moving traffic flows among alternative network paths, some nodes or subsystems may become inactive and can be put into standby mode [9], or power scaling primitives can reconfigure affected nodes [10]. Typical examples of these kinds of processes are traffic engineering protocols (e.g., OSPF-TE) with “green” extensions.

These are potentially more effective than LCPs. However, NCPs suffer from some drawbacks. They can exhibit much higher feedback/convergence delays. Routing and traffic engineering frameworks generally engage with logical network entities and are unaware how these are mapped onto physical resources, and thus to energy consumption.

Finally, NCPs often represent a network device simply as a node in a graph, whose arcs are the (virtual/physical) network links. This representation does not capture hardware nuances that may be very important for reducing consumption.

All these points motivate a strategy of jointly adopting LCPs and NCPs to optimize energy consumption in a hierarchical way. The ensuing complexity would be excessive in a network with heterogeneous equipment, without a standard interface to mediate between the various items of hardware and the entities seeking to control them.

## 4 *The GAL architecture and the Green Standard Interface*

The GAL is specifically designed to hide implementation details of energy-saving approaches, as well as to provide standard interfaces for interactions between green hardware and its control framework.

The GAL is provided with a hierarchical view of the device internal organization, which represents the components at the various levels as a tree of “*entities*”. A multi-chassis device can be represented through four layers (Figure 1).

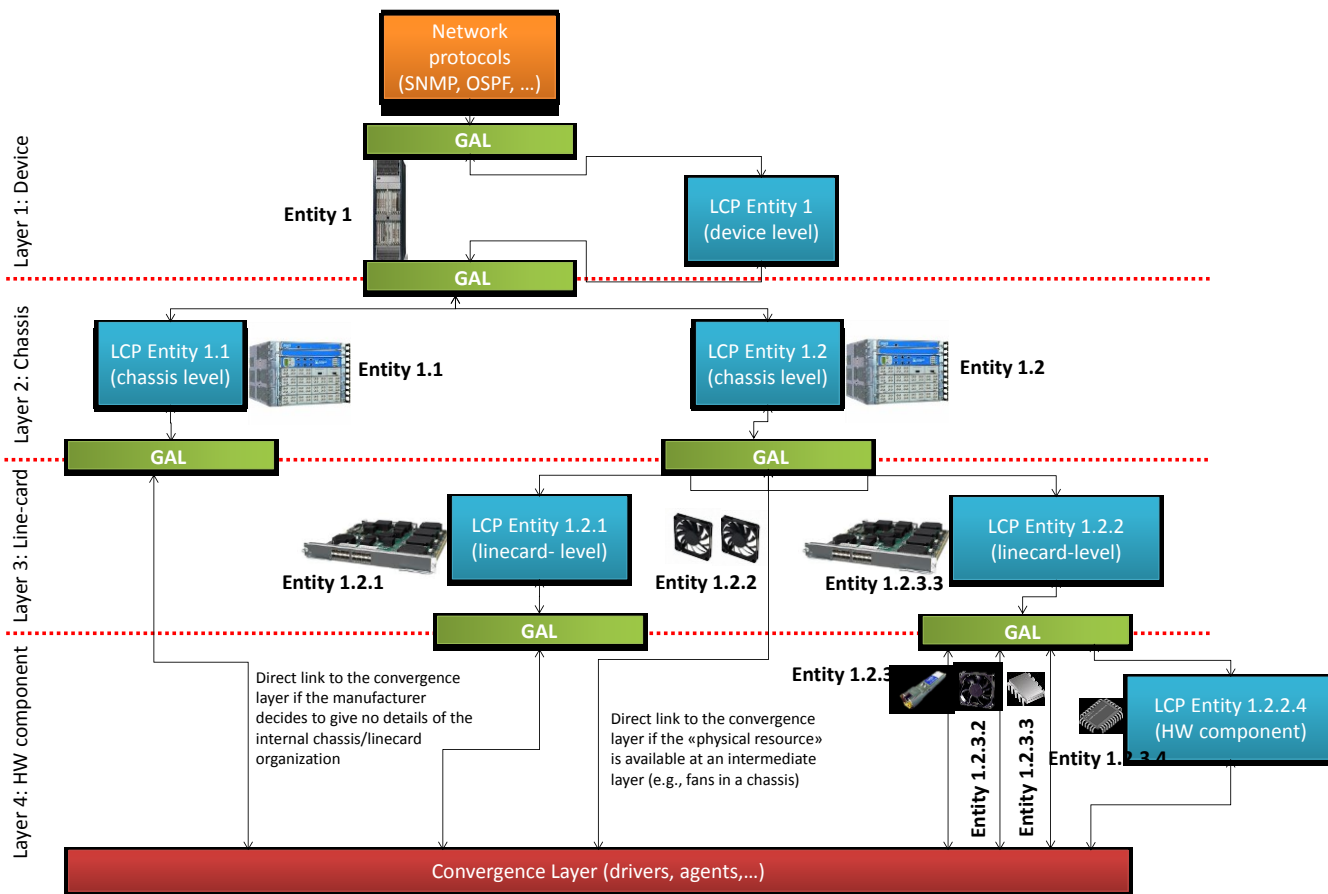


Figure 1: The hierarchical architecture of the GAL when applied to a multi-chassis network device.

The scope of the GAL embraces entities that provide energy-aware capabilities. An *energy-aware entity* is defined as an entity that can trade off its power consumption and its performance.

PMPs reside at the lowest levels of the hierarchy, and require specific LCPs to directly manage them. Thus, the top-level instance of the GAL should be used for interfacing such energy-aware entities at the lowest level with the relevant control-plane processes.

At intermediate levels, new LCPs (one for each entity at that level) are needed to orchestrate the settings and the operating behaviours of underlying energy-aware entities, and to expose a synthetic and aggregated set of operating characteristics and available configurations to higher levels. This process terminates at the device level, where the highest LCP orchestrates the high-level configuration of the device, and needs to expose a simplified view of it to NCPs.

This hierarchical approach results in a tree, where root nodes are LCPs and control applications, and leaf nodes correspond to hardware elements. The interface among the tree nodes is realized by means

of the GAL architecture and the *Green Standard Interface* (GSI) provided by the GAL. Leaf nodes may reside at different levels of the hierarchy for two main reasons:

- some energy-aware hardware entities may need to be accessible from higher hierarchical levels (e.g., fans in a chassis);
- some manufacturers shall prefer not to expose the internal organization and hardware architecture of the subcomponents.

The communication between the GAL and hardware components is realized through a “convergence layer” (external to the GAL). This layer includes all the drivers/agents or intercommunication means for setting or acquiring parameters from the supported energy-aware hardware elements

When a NCP (e.g., OSPF) decides to change the configuration of a logical link (e.g., the IP layer), by putting it into sleep mode, a command is sent to the LCP of the affected device. The (highest-level) LCP, in interpreting the command, first determines the physical resources bound to the logical link (e.g., which physical port).

For example, in Figure 1, the LCP of entity 1 selects the chassis containing the physical port bound to the logical link. Then, the chassis’ LCP (e.g., the entity 1.2) forwards the command to the correspondent line-card (entity 1.2.2), and (e.g.) reduces the speed of chassis fans.

If there are no further links active, the layer-3 LCP sends the entire line-card to sleep; otherwise, the LCP puts the physical interface into standby mode and reduces the performance of all the hardware components that process packets for that port. Complex hardware components (e.g., network processors) may require a further LCP layer in order to decide the best internal configuration.

The GSI is a lightweight interface to manage the energy-aware hardware entities. It is the northbound interface of a GAL instance instantiation. Although the exact implementation of the modules that interact with the hardware will depend on the devices (e.g. network elements, network cards, chips, fans, etc.), the GSI provides a set of functions and data types to enable the GAL interface to operate. The main functionalities offered by the GSI functions are:

- **Discovery:** It is used to retrieve information about:
  - available EASs and other information of the entity;
  - list of individually manageable components within the entity and their relations;
- **Provisioning:** Allows the setting of an EAS to an entity.
- **Monitoring:** of relevant parameters of the physical device.

The GSI can be interfaced with any network management protocol (e.g., SNMP, TMF). Thus, it can piggyback on the security offered by the selected protocol without incurring significant energy overhead.

## 5 Conclusions

We have proposed a novel flexible framework, the Green Abstraction Layer, to represent the energy-aware capabilities of networking devices to higher-level protocols. It features a multi-layered abstract model of the energy-saving capabilities of the devices for use in deploying local and network-wide control policies.

The abstract representation of the hardware as one or more EASs permits the tradeoff between network performance and energy saving to be evaluated across multiple hardware types. This will allow the GAL to support the minimisation of energy consumption and maintenance of QoS across a network with diverse and internally complex hardware.

## **References**

- [1] Bolla R., Bruschi R., Carrega A., Davoli F., Suino D., Vassilakis C., Zafeiropoulos A., “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, Jul. 2012.
- [2] Global e-Sustainability Initiative (GeSI), “SMART 2020: Enabling the Low Carbon Economy in the Information Age,” <http://www.theclimategroup.org/assets/resources/publications/Smart2020Report.pdf>.
- [3] Nedeveschi S., Popa L., Iannaccone G., Wetherall D., Ratnasamy S., “Reducing Network Energy Consumption via Sleeping and Rate-Adaptation,” *Proc. 5th USENIX NSDI Symp.*, San Francisco, CA, Apr. 2008, pp. 323-336.
- [4] Bolla R., Bruschi R., Christensen. K., Cucchietti F. Davoli, F., Singh S., “The Potential Impact of Green Technologies in Next Generation Wireline Networks - Is There Room for Energy Saving Optimization?,” *IEEE Communications Magazine*, vol. 49, no. 8, pp. 80-86, Aug. 2011.
- [5] Bolla R., Bruschi R., Davoli F., Cucchietti F., “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.
- [6] ACPI, *Advanced Configuration and Power Interface Specification*, Revision 5.0, 2011.
- [7] The ECONET project, <http://www.econet-project.eu>.
- [8] IETF Energy MANagement (EMAN) Working Group, <https://datatracker.ietf.org/wg/eman/charter/>
- [9] Cianfrani A., Eramo V., Listanti M., Polverini M., “An OSPF-Integrated Routing Strategy for QoS-aware Energy Saving in IP backbone Networks”, *IEEE Transactions on Network and Service Management*, vol. 9, no. 3, pp. 254-267, Sept. 2012.
- [10] Cardona Restrepo J.C., Gruber C.G., Mas Machuca C., “Energy Profile Aware Routing,” *Proc. IEEE ICC GreenComm Workshop*, Dresden, Germany, June 2009.

**Raffaele Bolla** is an associate professor of Telecommunication Networks at the University of Genoa. He is the Principal Investigator of the ECONET project.

**Roberto Bruschi** is with the CNIT as researcher. He is the coordination assistant of the ECONET project, and the principal investigator of the GreenNet project.

**Franco Davoli** is full professor of Telecommunication Networks at the University of Genoa, Italy. He is a member of the CNIT Scientific Board.

**Lorenzo Di Gregorio** is with Lantiq Deutschland GmbH. He led the development of an advanced packet processor, employed in most products of Lantiq.

**Pasquale Donadio** is a system engineer at Alcatel-Lucent Italia. He manages several research programs funded by the European Community.

**Leonardo Fialho** is a postdoctoral researcher at Dublin City University, Ireland.

**Martin Collier** is a Senior Lecturer at Dublin City University in Ireland, where he runs the Switching and Systems Laboratory.

**Alfio Lombardo** is full professor of Telematics at the University of Catania, Italy. He is responsible of the OpenLab at the University of Catania.

**Diego Reforgiato Recupero** is a postdoctoral researcher at the University of Catania. The support for this fellowship has been given by PROVIDEO, a Marie Curie International Grant.

**Tivadar Szemethy** is working as Chief Architect for Netvisor, Hungary, a company focused on network and systems operation and management.

