SustNMS: Towards Service Oriented Policy-Based Network Management for Energy-Efficiency

Carlos H. A. Costa, Marcelo C. Amaral,
Guilherme C. Januário, Tereza C. M. B. Carvalho
Escola Politécnica, University of São Paulo
São Paulo, Brazil
{chancosta,marceloamaral,gcjanuario,terezacarvalho}@usp.br

Catalin Meirosu
Ericsson Research, Packet Technologies,
Sockholm, Sweden
catalin.meirosu@ericsson.com

Abstract—**Considering the increasing importance of carbon free networking solutions, several energy-efficient networking systems have been proposed and evaluated in literature. One significant challenge is how to encompass actions targeting energy savings and possible corresponding degradations in network QoS (Quality of Service). This paper presents an energy-efficient policy-based network management architecture that enables the enforcement of energy-efficiency according to high-level business decisions on how to coordinate the trade-off between energy savings and the potential impact on the network operation. The system integrates a real-time energy-efficiency assessment with a dynamic evaluation of network availability and performance. Its operation is presented and advantages are discussed.**

I. INTRODUCTION

Building a network infrastructure that actively supports sustainability features is an important step towards enabling a low carbon economy in the information age. The forecast of the power consumption related to network infrastructure shows a considerable and increasing share of the total power consumption of ICT systems. It has been projected that the carbon footprint of the sector is likely to grow up to 1.4 $GtCO_2$ by 2020, of which 25% is related to telecom infrastructure and devices [1]. Many network device manufacturers incorporate power-saving technologies in their designs, allowing certain parts of the equipment to be switched to a power-saving mode, depending on their current utilization, such as [2], and [3].

The task of monitoring and administering network nodes has been traditionally the domain of Network Management Systems (NMS). An energy-efficient NMS would thus manage power saving-enabled network nodes at time intervals that are compatible with NMS interactions. In this context, an energy-efficient NMS should use network monitoring data for pinpointing underutilized nodes (e.g., routers or interfaces with low traffic demand) and calculating the network configuration that defines the optimal power state and traffic allocation for each node. The energy-efficient NMS should take into account the assurance of QoS and ensure fast response to either failures or sudden workload increases. For each decision on the state of the node (power-saving/normal mode), an online analysis of availability measures would provide further clarification on the potential consequences of the decision.

Business connectivity services are a multi-billion dollar market worldwide, offering packet-based connectivity to enterprises via different technologies. Service Level Agreements (SLAs) specify the performance objectives to be reached by the service. Availability of the service is one of the most important parameters included in such agreements.

The remaining part of this paper is organized as follows. Section II presents the related work, discusses requirements for an energy-efficient network management system and the reasons why they are not fulfilled by the existing approaches. Section III presents the novel proposal of an architecture for policy-based network management that supports energy-efficiency functions, and the system operation, describing how the modules are integrated and their interaction. In sections IV and V, we include a discussion and conclusions of this work.

II. RELATED WORK

Energy-efficient networking attempts to save energy by reducing powered-up capacity of routers in the network. Reducing the energy consumption allows the network operator to be environmentally friendly while also reducing OPEX. A standardization process is being conducted by the Internet Engineering Task Force (IETF) to define an Energy and Power Management Information Base (MIB) [4] that provides, in a standard format, power-related data from switches/routers. However, no standard is concluded yet. The architectures of existing solutions for energy-efficient networking can be typically divided into two categories: self-coordinated and policy-driven systems.

The self-coordinated approach relies on the interaction among nodes without a centralized entity in order to achieve energy efficiency. Some examples of this approach are the Dynamic Ethernet Link Shutdown (DELS) algorithm [5], [6], and the Adaptive Link Rate (ALR) [7]. Power-aware extensions to routing protocols also have been proposed as self-coordinated approaches. A fully OSPF complaint energy-aware routing protocol is proposed in [8]. Another approach is found in the “MiDORi” project, which achieves energy-efficiency by conducting traffic engineering to dynamically aggregate traffic on particular links [9].

In energy-efficient nodes, the amount of power consumed depends on the amount of traffic handled and processing applied to the packets in transit. Power profiles indicate a consumption baseline and a scaling factor that determines how the consumption increases with traffic. Linear scaling factors were determined experimentally for real routers [10], [11]. In [12], the authors consider multiple theoretical models that employ linear scaling factors of various slopes on top a consumption baseline value. Even simple power profiles such as those mentioned above play a crucial role in determining the most power-efficient operational state of a network composed of heterogeneous devices, produced by different manufacturers and potentially being at different stages of their lifetime.

Policy-based Network Management (PBNM) has been used as a flexible mechanism to allocate network resources such as traffic prioritization, bandwidth allocation, quality of service and security management [13]. An architecture and information model for policy-based management of Quality of Service (QoS) in IP networks has been defined by IETF through the Policy Framework Working Group (PFGW). Distributed Management Task Force (DMTF) also developed information models for network and policy management applications, and joined the IETF/PFWG to standardize the IETF policy information model. The standard IETF/DMTF policy-based architecture is based on four functional blocks: a policy management tool, a policy repository, a policy decision point, and a policy enforcement point.

A Policy-Based NMS has a constantly updated picture of the network status. Some examples are the Energy Management System (EMS) [14] and a NMS that continuously feeds the collected data into a Bayesian Belief Network (BBN) [15].

Even though these ideas have been discussed for almost a decade [16], [17], a policy-based framework that allows conserving energy, along with an understanding of the impacts on network performance and availability, has yet to be developed. This work presents a novel management system architecture that extends the IETF/PFGW in a way that yields dynamic energy-efficiency evaluation in a heterogeneous environment, concurrently with QoS monitoring, including a real-time network availability evaluation. Provisioning this data to a policy-driven decision point allows the enforcement of a management approach to deal with the trade-off between energy savings and quality of service assurance, reflected in each node.

There are at least two main problems with the existing solutions. The first one is related to the coordination between energy-efficiency features in a network with heterogeneous devices. Existing solutions do not take into account the role of power profiles, preventing the optimization of energy consumption in response to a varying traffic demand when redundant paths are available in a heterogeneous network. In addition, telecom operators might have legacy equipment that does not support energy-efficiency features but could still be temporarily shutdown in a coordinated way. The second problem is related to how service level is evaluated while actions targeting energy-savings are taken. Service-centric management enables operators to translate abstract SLA parameters into actions to be performed at the network. It also enables the operators to evaluate the impact of network events all the way up to the customer. The way an administrator or the NMS determines thresholds is crucial as the process ultimately depends on service constraints at the network level (desired fault-resilience and performance). An energy-efficient NMS has to enable the definition of a trade-off between energy savings and level of service, which can be dynamically changed reflecting the adaption of the management policy by the infrastructure operator.

III. SUSTAINABILITY ORIENTED POLICY-BASED NETWORK MANAGEMENT ARCHITECTURE

This work presents a novel architecture for an Energy-Efficient PBNM System (SustNMS) that support the following functionality: (1) energy-efficiency evaluation in a heterogeneous environment, (2) dynamic QoS evaluation, including network availability for dynamically adjustable topology, and (3) network configuration, optimized and managed through a management policy that determines the trade-offs between energy savings and QoS. The policy determines classes of services, and how those trade-offs should be enforced according to network domain.

The SustNMS extends the PBNM architecture defined by IETF by including three modules. The first module is a *Model Repository (MR)* that stores power and availability models. These models are composed by static parameters used in the evaluation of power consumption and network availability, which occurs according to the whole network and to its individual nodes states. The second is the *Quality of Service Monitor* (QoSM), which evaluates network QoS by a dynamic calculation of its availability and by gathering performance indicators, such as dropped packets, jitter, and delay. The third is the *Sustainability Monitor* (SM), which dynamically evaluates energy-efficiency indexes, such as the Watts/bps ratio. These modules rely on a *Device Updater* that collects network data and applies the actions determined by the *Policy Management Framework* (PMF), which could be, e.g., putting nodes in a reduced power mode, or even performing a set of actions that implement traffic engineering. The architecture might be also benefited by the inclusion of a *Traffic Load Forecast* (TLF) module that yields an expected tendency on traffic demand to be queried by the Policy Management Framework (PMF) while calculating an optimal network configuration.

The main advantage in the architecture consists in this integrated view of the network. It enables the enforcement of energy-efficiency oriented rules while assessing their impacts on network performance before they are enforced. The conflicts raised by quality of service constraints is reconciled through a policy that determines the priority for each class of service and network domain. The usage of power models guarantees that the system operates even in a legacy environment where there exist devices that are unable to provide power consumption information. The detailed SustNMS architecture and its implementation are defined as following.

A. System architecture

The SustNMS architecture is depicted in Figure 1. It presents the submodules for each main module, namely: (i) Policy Management Framework (PMF); (ii) Model Repository

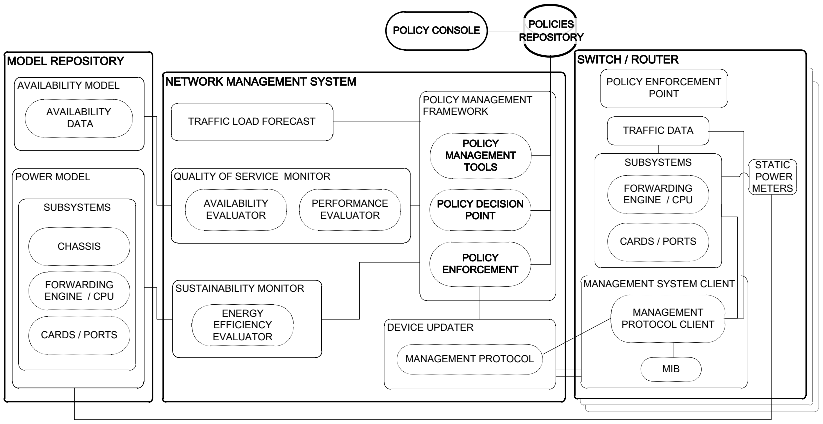

Fig. 1. Sustainability Oriented Policy-Based Network Management (SustNMS) Architecture

(MR); (iii) QoS Monitor (QoSM); (iv) Sustainability Monitor (SM); and (v) Device Updater (DU).

(i) The *Policy Management Framework* (PMF) is composed by the four modules of the IETF PBNM architecture [18]: the Policy Management Tools, the Policy Decision Point, the Policy Enforcement Points and the Policy Repository.

(ii) The *Model Repository* (MR) comprises two submodules, one for the storage of power models and another for availability models.

The *Power Model* is composed by sets of static parameters that define a power profile, i.e., a scaling factor for power consumption dependency on load. Such parameters are obtained through a one-time power measurement performed in the device with power meters. In this way, a power consumption profile would be determined by a function $f(load)$ defining how power consumption varies with the load handled by the device. For a power saving–enabled device, the model includes a set of variable related to: $P_{standby}$ (Reduced power mode), P_{idle} (Idle power consumption), and P_{full} (Full load power consumption). An example of a function $f(load)$ for a switch/router following a linear power consumption scaling is:

$$P(state, f(load)) = \begin{cases} P_{standby} & \text{if } state = standby \\ f(load) & \text{otherwise} \end{cases}$$

$$f(load) = P_{idle} + \frac{load}{capacity_{max}}(P_{full} - P_{idle}). \qquad (1)$$

In case a fine-grained control is available, such as the shutdown of individual line-cards in a blade architecture, a larger set of variables would allow a more accurate evaluation.

The *Availability Model* is composed by information of failure and repair rates for each network device type, such as mean time to failure (MTTF) and mean time to repair (MTTR).

(iii) The *Quality of Service Monitor* (QoSM) includes two modules for availability and performance evaluation respectively.

The *Availability Evaluator* dynamically evaluates network availability every time a new configuration is calculated by the PMF. One approach for dynamic evaluation is to first calculate availability for each node through a Markov model based on MTTF and MTTR stored in the MR [19]. In a second step, the standard standby modeling can be adapted to model devices in sleep mode. However, the default standby modeling does not take into account the time that a router takes to be switched back from a reduced power mode. An extended model that considers this delay should be applied, since a dynamic availability evaluation should capture the network unavailability related to power state switching [20].

The *Performance Evaluator* dynamically collects performance indicators for each node and for the whole network, such as the amount of packet losses, delay, or jitter. The indicators used depends on the management approach and the requirements defined in the SLAs.

(iv) The *Sustainability Monitor* (SM) is composed by a module responsible for energy efficiency evaluation. The *Energy Efficiency Evaluator* performs the evaluation of instant power consumption (given the time frame between probes of the system) for each node. The evaluation may be performed in two ways. In case the power consumption of the device is accessible through a management protocol, the evaluator just probes the device to collect this information and calculates an energy-efficiency ratio based on the current traffic observed on the node. In the lack of such feature, the module will retrieve from the MR the power model matching the device under evaluation. If a power model that matches the network device cannot be located in the MR, a pessimistic power model is loaded from the power model repository using a pre-defined lower bound for the power consumption variables. If a power model is found, consistency verification is performed checking the variables. The total power consumption, along with the load handled by the device, are used to evaluate the energy-efficiency ratio (W/bps).

(v) The *Device Updater* (DU) is the responsible for collecting data and applying changes to each device communicating with each instance of the Policy Enforcement Point (PEP). This module, using a standard network management protocol, probes network devices (or defines a trap) collecting traffic data, power state settings and network performance data, such

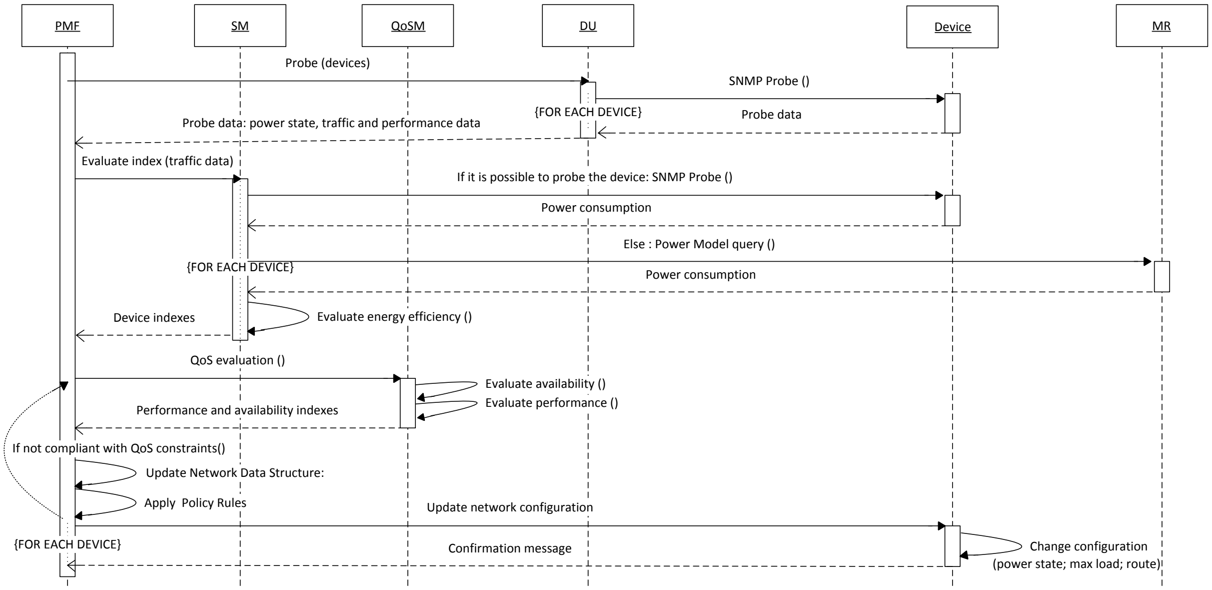

Fig. 2. SustNMS sequence diagram

as packet loss and the delay. The processed data provided as a network reconfiguration request is applied through a network management protocol (SNMP, NETCONF) or through a Command-Line Interface (CLI). The module implements power state changes and defines prioritized paths.

B. System Implementation

The system operates probing the network in a frequency defined by the operator. The managed devices are switches that have at least a power-saving feature defined as a sleep mode for which only parts of the node are powered on. These devices may be woken up within a given time frame. Each probe request is issued by the *Policy Management Framework* (PMF) and triggers actions in the sequence described in Figure 2. They are detailed in sets of actions as following.

(i) *Network state update request*: the Device Updater (DU) receives a request to collect network statistics. For each switch in the network, it issues a SNMP probe based on MIB-II (RFC1213) standards for data identifiers and collects traffic, state and performance data. This data is stored in the nodes of a graph that represents the network topology. The graph is sent to the requester.

(ii) *Efficiency indexes evaluation*: the Sustainability Monitor (SM) receives a request to evaluate efficiency indexes. For devices that support the Energy Monitoring MIB [4], instant power consumption (Watts) is obtained through the Power Monitor data identifiers. The remaining devices have power consumption evaluated through device-specific power models in MR, a database accessed through Simple Query Language (SQL) queries, defining the static parameters. The traffic and state data in each node are used to evaluate the power consumption function in Equation (1). The Network Energy Consumption Rate (NECR), the MilliWatts per Mbps ratio [21] is evaluated for each switch based on the traffic and power consumption in the nodes of the graph. The graph is updated to include the calculated indexes for each node.

(iii) *QoS evaluation*: the QoS Monitor (QoSM) receives a request to evaluate network availability and performance. The MTTF and MTTR of each switch are retrieved from MR and written in the respective node. This data along with the state of each node is used in a dynamic availability evaluation such as the one in [20]. The network availability is stored in a specific variable. Following, performance indexes are obtained through SNMP probes using MIB-II data identifiers for packet loss statistics for each switch. Delay and jitter are obtained through Internet Control Message Protocol (ICMP) request. The graph is updated with calculated indexes.

(iv) *Policy application*: the network status is handled to the PMF. The module relies on the general-purpose and distributed object management framework Ponder2 [22], used to create, store and handle management policies. The Policy Decision Point (PDP) performs a shortest path algorithm calculation that depends on which policies need to be applied to the current situation. Actions aiming the compliance with these policies are defined. The PMF module then checks all these actions and whether they can be applied by comparing the new configuration against its knowledge on hard constraints for the arrangement of the components and QoS requirements stated in the SLA. The network configuration may be recalculated to fulfill availability requirements, or the quality of service constraints are relaxed for energy efficiency improvement. This decision is defined in the policies. The output of the module is a new configuration for the network that is represented by specific configuration request stored in each node of the network graph.

(v) *Decision enforcement*: the decision on an updated network configuration is applied by the Device Updater (DU). The module translates the configuration request in each node into device-specific commands. They set the power state for each device or update switching-related data to enforce the prioritization of most efficient paths. Multi-protocol Label Switching (MPLS) is used to direct data among network nodes. Label Edge Routers (LERs) are set through CLI commands. The action is finalized with a confirmation on the updated configuration that is sent to the PMF.

The entire process is repeated for each probe issued by the SustNMS. The accuracy of the measures is defined by the rate at which these actions occur. The adequate probing rate can be defined through the energy-efficiency measurement since an increase of the probe rate would generate an increase of messages related to management operations, which in turn leads to an increase in energy consumption.

IV. DISCUSSION

The main advantages of the proposal are related to the support for agile operations via policies, the support for power profiles, and the support for real-time availability calculation.

The support for operations via policies decouple process-oriented behavior from system configuration. Therefore, the work of the human operators is considerably simplified. The policies may be added, removed and modified easily as new goals are specified or legislation is introduced. The policy refinement can take into account multi-layer aspects and optimize the overall system. Policies facilitate automation, thus creating an environment of increased efficiency.

The support for power profiles creates the premises for the operator to optimize the power consumption in accordance to the traffic to be supported by a certain set of network elements. Given a traffic profile, it allows the operator to forecast and accurately control the power consumption. This may be of particular significance in markets with spot prices for electricity or in smart grids scenarios.

The support for real-time availability calculation enables the operator to take service-centric decisions, rather than network-centric as proposed currently in the literature. In addition, it allows the operator to control SLA parameters and thus avoid unnecessary SLA-related penalties while saving energy. It also creates the potential to coordinate with OAM operations at the node level in order to reduce probing by OAM tools during energy saving intervals, thus increasing chances for physical layer energy efficiency functions to save power on idle links.

V. CONCLUSIONS AND FUTURE WORKS

Guaranteeing the maximum energy savings without compromising the networking performance remains a significant challenge when building an energy-efficient NMS solution. We discussed in this work why the existing approaches for energy-efficient networking still lack a management architecture that integrates an efficiency evaluation along with a dynamic QoS assessment. The proposal can be implemented using current technology and available equipment. The authors have an ongoing project evaluating the advantages of the architecture in scenarios of interest. A large-scale testbed is being developed to enable the obtainment of results showing the advantages of the proposal in relevant realistic network topology.

ACKNOWLEDGMENT

This work was supported by the Innovation Center, Ericsson Telecomunicações S.A., Brazil.

REFERENCES

[1] GeSI, "Smart 2020: Enabling the low carbon economy in the information age," The Climate Group on behalf of the Global eSustainability Initiative (GeSI), Tech. Rep., Jun. 2008.

[2] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, and J. Maestro, "IEEE 802.3az: The road to energy efficient ethernet," *IEEE Communications Magazine*, vol. 48, no. 11, pp. 50–56, 2010.

[3] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green ospf protocol," in *INFOCOM IEEE Conference on Computer Communications Workshops , 2010*, march 2010, pp. 1 –5.

[4] M. Chandramouli, B. Schoening, J. Quittek, T. Dietz, and B. Claise, "Power and energy monitoring mib," Network Working Group, IETF, Tech. Rep., May 2011.

[5] M. Gupta and S. Singh, "Dynamic Ethernet Link Shutdown for Energy Conservation on Ethernet Links," in *IEEE ICC'07*, 2007, pp. 6156–6161.

[6] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *Proc. of the NSDI'08*. USENIX, 2008, pp. 323–336.

[7] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption of ethernet with adaptive link rate (ALR)," *IEEE Transactions on Computers*, vol. 57, no. 4, pp. 448–461, April 2008.

[8] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *IEEE INFOCOM'10*, 2010, pp. 1–5.

[9] H. Takeshita, Y. Oikawa, H. Yonezu, D. Ishii, S. Okamoto, and N. Yamanaka, "Demonstration of the self organized dynamic link power management by MiDORi; energy optimal network topology design engine," in *OFC/NFOEC*, march 2011, pp. 1 –3.

[10] R. Bolla, F. Davoli, R. Bruschi, K. Christensen, F. Cucchietti, and S. Singh, "The potential impact of green technologies in next-generation wireline networks: Is there room for energy saving optimization?" *Comm. Magazine, IEEE*, vol. 49, no. 8, pp. 80 –86, august 2011.

[11] A. P. Bianzino, A. K. Raju, and D. Rossi, "Apples-to-apples: a framework analysis for energy-efficiency in networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 38, no. 3, pp. 81–85, Jan. 2011.

[12] S. Antonakopoulos, S. Fortune, and L. Zhang, "Power-aware routing with rate-adaptive network elements," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, dec. 2010, pp. 1428 –1432.

[13] J. Strassner, *Policy-Based Network Management: Solutions for the Next Generation*, 1st ed. Morgan Kaufmann, Sep. 2003.

[14] F. Pérez, D. Jorquera, and V. Iglesias, "Energy management system as an embedded service: Saving energy consumption of ICT," in *Proc. of the 22nd ARCS*. Springer-Verlag, 2009, pp. 195–206.

[15] A. Bashar, G. Parr, S. McClean, B. Scotney, M. Subramanian, S. Chaudhari, and T. Gonsalves, "Employing bayesian belief networks for energy efficient network management," in *2010 NCC*, 2010, pp. 1–5.

[16] M. Gupta and S. Singh, "Greening of the internet," in *Proc. of SIGCOMM*, 2003, pp. 19–26.

[17] K. Christensen, C. Gunaratne, B. Nordman, and A. George, "The next frontier for communications networks: power management," *Computer Communications*, vol. 27, no. 18, pp. 1758–1770, 2004.

[18] G. Waters, J. Wheeler, A. Westerinen, L. Rafalow, and R. Moore, "Policy framework architecture," Network Working Group, IETF, Tech. Rep., Feb. 1999.

[19] M. L. Shooman, *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*, 1st ed. Wiley-Interscience, Dec. 2001.

[20] M. C. Amaral, C. H. A. Costa, T. C. M. B. Carvalho, and C. Meirosu, "REASoN - REliability and/or availability evaluation for sustainable networking," in *4th Int. Work. on Reliable Networks Design and Modeling (RNDM'12)*, St. Petersburg, Russia, Oct. 2012.

[21] V. Manral, "Benchmarking power usage of networking devices," Network Working Group, IETF, Tech. Rep., May 2010.

[22] G. Russello, C. Dong, and N. Dulay, "Authorisation and conflict resolution for hierarchical domains," *Policies for Distributed Systems and Networks, IEEE International Workshop on*, vol. 0, pp. 201–210, 2007.