Computer Networks 65 (2014) 111–128

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Energy adaptation in multi-core software routers

Raffaele Bolla [a], Roberto Bruschi [b,*], Paolo Lago [a]

[a] Department of Electrical, Electronic and Telecommunication, and Naval Architecture (DITEN), University of Genoa, Via all'Opera Pia 13, 16145 Genova, Italy
[b] National Inter-University, Consortium for Telecommunications (CNIT), Research Unit of the University of Genoa, Via all'Opera Pia 13, 16145 Genova, Italy

ARTICLE INFO

Article history:
Received 13 June 2013
Received in revised form 18 December 2013
Accepted 24 February 2014
Available online 13 March 2014

Keywords:
Software routers
Green networking
Performance measurement

ABSTRACT

We consider a new generation of Commercial-Off-The-Shelf (COTS) Software Routers (SRs), able to effectively exploit multi-core/CPU hardware platforms. This choice was mainly driven from the fact that SRs are one of the few platforms including power management features, and that can be easily customized. Our main objective is to evaluate and to model the impact of power saving mechanisms, generally included in today's COTS processors, on the SR networking performance and behavior. To this purpose, we separately characterize the roles of both hardware and software layers through a large set of internal and external experimental measurements, obtained with a heterogeneous set of hardware platforms and SR setups. Starting from this detailed measure analysis, we propose a simple model, able to represent the SR performance with a high accuracy level in terms of packet throughput and related power consumption. The proposed model is thought to estimate how SRs behave under different network conditions and power management configurations. In fact, the model has been effectively applied inside a "green optimization" policy, which aims at minimizing power consumption, while maintaining a certain SR performance target. Finally, in order to provide an experimental evaluation of the proposed ideas, we apply such power management policies to a new generation software router platform, and we test their performance under real traffic traces.

1. Introduction

In the last few years, the research field of "green" and energy-efficient networking infrastructures has gained a great interest from both service/network providers and equipment manufacturers [1].

Besides a more widespread sensitivity to ecological issues, such interest springs from heavy and critical economical needs, since both energy cost and network electrical requirements have shown a continuous growth with an alarming trend over the past years. For example, as shown in [2] and in [3], energy consumption of Telecom Italia network in 2006 has reached more than 2 TW h (about the 1% of the total Italian energy demand), which was an increase of 7.95% over 2005. More recently, in 2009, this consumption increased to 2.14 TW h [4]. Similar trends can be generalized to a large part of other telecoms and service providers, since they essentially depend on data traffic volume increase and new services being offered.

To support new generation network infrastructures and related services for a rapidly increasing customer population, telecoms and service providers need a larger number of devices, with sophisticated architecture able to perform more and more complex operations in a scalable way. In such an environment, the challenge for network equipment manufacturers is to design new architectures able to scale their performance and functionalities preserving a high level of energy efficiency.

* Corresponding author. Tel.: +39 010 3532057; fax: +39 010 3532154.
E-mail addresses: raffaele.bolla@unige.it (R. Bolla), roberto.bruschi@cnit.it (R. Bruschi), paolo.lago@reti.dist.unige.it (P. Lago).

http://dx.doi.org/10.1016/j.comnet.2014.02.027
1389-1286/

It is well known that network links and devices are provisioned for busy or rush hour load, which typically exceeds their average utilization by a wide margin [5]. While this margin is seldom reached, nevertheless the power consumption is determined by it and remains more or less constant even in the presence of fluctuating traffic loads. Thus, the key of any advanced power saving criteria resides in dynamically adapting resources, provided at network, link or equipment levels, to current traffic requirements and loads [6–8].

In this sense, a clear example is provided by the new "green" Ethernet task force [9,10], which has explored in the last few years the possibility of saving energy by adopting two main schemes: the *Adaptive Rate* (AR) and *Low Power Idle* (LPI). The former allows dynamically modulating the capacity of a link, or of a processing engine, in order to meet traffic loads and service requirements; the latter, which has been selected in the final release of the IEEE 802.3az standard, forces links or processing engines to enter low power states when not sending/processing packets and to quickly switch to a high power state when sending one or more packets.

From a more general point of view, it is well-known that today's networks rely very strongly on electronics, despite the great progresses of optics in transmission and switching. Operational power requirements arise from all the hardware elements realizing network-specific functionalities, as the ones related to data- and control-planes, as well as from elements devoted to auxiliary functionalities (e.g., air cooling, power supply, etc.). In this respect, the data-plane certainly represents the most energy-starving and critical element in the largest part of network device architectures, since it is generally composed by special purpose hardware elements (packet processing engines, network interfaces, etc.) that have to perform per-packet forwarding operations at very high speeds.

In this sense, Tucker et al. [11,12] and Neilson [13] focused on high-end IP routers, and estimated that the data-plane accounts for 54% on the overall device power budget, vs. 11% for the control plane and 35% for power and heat management. The same authors further broke down energy consumption sources of the data-plane on a per-functionality basis. Internal packet processing engines require about 60% of the power consumption of the data-plane on a high-end router, network interfaces account for 13%, the switching fabric accounts for a further 18.5% and buffer management for 8.5%.

Starting from these considerations, our first objective is to evaluate the impact and the feasibility of adopting power scaling mechanisms on packet processing engines, which represent one of the most energy-hungry elements in many network platforms [14,15], in order to dynamically modulate energy absorption with respect to incoming traffic loads and service requirements.

In this scenario, we pursued a "bottom-up" approach: at first, we try to enable energy management capabilities in the data plane of an "open" reference platform for IP node, which allows us the possibility of enabling power management capabilities and to customizing their usage and control. Since the largest part of commercial network devices does not offer this possibility, and does not include any hardware primitives for power management, we decided to focus on new generation Software Routers (SR) [16–19,37], based on open source Software (i.e., Linux) and on multi-core Commercial Off-The-Shelves (COTS) hardware. Despite the possibility of customizing every aspect of the software implementing network functionalities, these platforms already provide advanced power management capabilities [20,21] by means of the Advanced Configuration and Power Interface (ACPI) technology [22]. Moreover, the latest generation of SR hardware platforms and relative building blocks are characterized by a so high parallelism and modularity that results and analyses obtained with these objects can be easily extended to extrapolate design feedback and guidelines for including power management capabilities in next-generation commercial network devices [23].

By enabling the power management on Software Routers at software and hardware levels, we had the opportunity to perform a complete experimental evaluation with a heterogeneous set of hardware platforms. The obtained measures provide insight into the relationship between SR internal dynamics and its external/power-related performance, and allow deriving a simple empirical model to estimate SR behavior with respect to different setups and traffic offered loads.

Starting from the analysis of experimental data, our next objective is to develop a model able to capture the impact of power management mechanisms by characterizing the trade-off between SR performance and energy requirements under different system configurations.

Finally, we propose an optimization policy (often referred to as "governor" in the Linux environment) that can be used for dynamically controlling the energy-aware capabilities with respect to traffic loads and traffic requirements. This paper is an extended version of previous works [24,25,35], where first approaches in evaluating energy-aware SRs and related control strategies were proposed.

The paper organization reflects the "bottom-up" nature of our approach. Section 2 introduces some background information and details on new generation SR architectures, specifically optimized for multi-core hardware systems, and on their standard support for power management. In Section 3, we present a large set of experimental results with the aim of investigating how power management mechanisms impact on SR performance. Then, starting from these results, we propose in Section 4 an empirical model, able to effectively estimate SR behavior in terms of throughput and power consumption under scenarios with increasing complexity (from the case of a single core and deterministic traffic to the one with multiple cores and stochastic traffic). In Section 5, the model is applied in an optimization policy for adapting the energy consumption with respect to incoming traffic load and service requirements. Section 6 reports the numerical results, Conclusions are drawn in Section 7. Appendices A–C are devoted to introduce demonstrations/derivation on statistical bounds we applied, and further details on the experimental testbed.