Modeling Power Consumption of the G-Lab Platform to Enable an Energy-Efficient Provision of Services

Gergö Lovász, Florian Niedermeier, Nasir Ali, Robert Basmadjian, and Hermann de Meer
University of Passau, Innstrasse 43, D-94032 Passau
{gergoe.lovasz|florian.niedermeier|nasir.ali|robert.basmadjian|demeer}@uni-passau.de

I. Introduction

The development of the future Internet is a challenging research topic: Not only restrictions of the current Internet need to be solved, also requirements of future services and protocols have to be anticipated. It needs to overcome the ossification and restrictions of the current Internet and has to be robust, reliable, and fault tolerant, for instance. One of the major issues, however, is to achieve an energy-efficient operation of the future Internet.

A lot of work regarding the future Internet is carried out on research networks like the G-Lab infrastructure[1]. While the aquisition of hardware is often financially supported, the cost caused by the hardware's energy consumption remains a constant financial burden, even after funding ends. To achive a more sustainable operation of the G-Lab network on the one hand, and an effective reduction of the energy consumption of future networks in general on the other hand, a comprehensive approach is needed. Hardware, services, and protocols need to collaborate to reduce energy consumption on a global scope, rather than achieving local optimizations.

To improve the energy-efficient provision of services within future networks, it is necessary to model the power consumption of services and the infrastructure that needs to provide these services. This abstract focuses on the development of load-based power consumption models for a standard G-Lab server as an important step of the holistic approach towards energy-efficient future networks and an energy-efficient G-Lab infrastructure.

II. Power Consumption Modeling

In general, servers, as we can find them at the different G-Lab sites, are composed of a variety of different components. Each component has its own characteristics in terms of technology and capacity. Therefore, capturing an entire server in a single model is a very complex and tedious task. However several suggestions for server power consumption models exist in literature. Some of them are very inaccurate due to the simplicity of the model, like the constant power model, presented in [2]. Fan et al. present in [1] a model for server power consumption that is based on the CPU utilization. Economou suggests in [3] to consider beside CPU and disk utilization also design and architecture properties of the server. All these models have in common that either they are too simple and inaccurate or very specific and complex. We decided to describe the different components of the G-Lab infrastructure in a hierarchical model. In our model every component is composed of a set of subcomponents. A subcomponent can have subcomponents itself. Components and subcomponents have static and dynamic properties that influence their overall power consumption. Static properties are constant factors which are defined at manufacturing time and do not change during runtime. Dynamic properties can change at runtime and are highly connected to the load on the component. On the one hand, such a hierarchical and modular description of a complex component with numerous subcomponents can reduce complexity, on the other hand it can also increase complexity for less complex components. Components with only few subcomponents that significantly contribute to the overall power consumption should instead be considered as 'black boxes' in order to reduce complexity. Therefore it is essential to decide for each (sub)component whether it should be regarded as a 'black box' or as a complex component with further subcomponents in order to reduce complexity.

The modeling of the service providing infrastructure enables a fine granular allocation of resources to services. Based on such models, resource allocation can be optimized with respect to energy consumption.

III. Modeling Example: Server

Figure 1 illustrates an excerpt of the server meta-model which represents an abstraction of a generic server. The root of the tree is the server itself for which the power consumption should be modeled. The nodes of the tree which are surrounded by a solid line represent the subcomponents whereas the leaves are the static (dashed nodes) and dynamic (dotted nodes) properties of the component and its subcomponents. For a server, several subcomponents can be identified (e.g. HDD, RAM, power supply unit (PSU), CPU, NIC). Some of them have further subcomponents themselves. E.g. the CPU could have several Cores with certain properties like the L1-cache size or the load.

Let SC be the set of subcomponents for a component. The index sp stands for the static parameters of a component and

[1] http://www.german-lab.de/

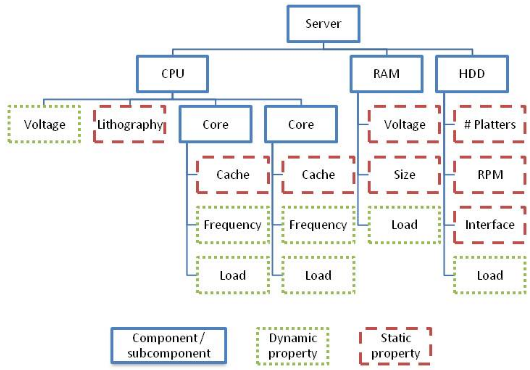

Fig. 1. Subcomponent tree of a server

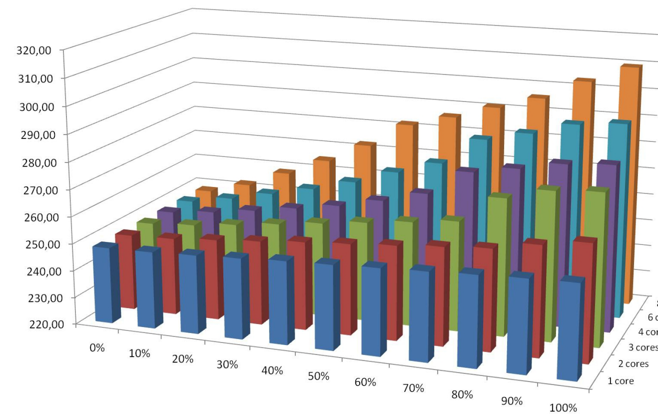

Fig. 2. Power consumption of a G-Lab server with different loads on the cores of its CPU

dp for the dynamic ones. Let $n = |SC|$. The function F_{SC} describes the correlation between the subcomponents in SC. The functions $F_1 ... F_n$ compute the power consumption of the subcomponents $c_1 ... c_n$. Then the power consumption of a server s is given by:

$$F(\vec{s_{sp}}, \vec{s_{dp}}, s) = \begin{cases} F_{SC=\emptyset}(\vec{s_{sp}}, \vec{s_{dp}}) & \text{if } SC=\emptyset \\ F_{SC}(\vec{s_{sp}}, \vec{s_{dp}}, F_1(\vec{c^1_{sp}}, \vec{c^1_{dp}}, c^1) & \\ , ..., F_n(\vec{c^n_{sp}}, \vec{c^n_{dp}}, c^n)), & \text{else} \end{cases}$$

In order to derive the functions $F_1, ..., F_n$ we apply a sensitivity analysis of both the static and dynamic parameters of the different components. Due to the limited space, in this abstract we do not describe our approach in detail. In order to get the power consumption of the server itself, the power consumption of the subcomponents cannot simply be added. Correlations between subcomponents have to be described by the function F_{SC}.

A. Measurements on a standard G-Lab Server

In order to determine the impact of CPU-load on the power consumption of a standard G-Lab server, we developed a benchmark that is able to load each single core of the CPU with a user-defined load. This is achieved by the determination of the number of operations (different types of mathematical and non-mathematical operations) that are needed to put 10%, 20%, ..., 100% load on a core of the CPU. Figure 2 shows the power consumption of a standard G-Lab server with different loads on the cores of its CPU. For this measurement a standard G-Lab server[2] with pack4Linux[3] as operating system was used. Pack4Linux supports real-time programming and allows assigning the CPU exclusively to a process. The experiment consisted of multiple measurements. In each measurement a certain number of cores were loaded with 10%, 20%, ..., 100%. The duration of every single measurement was 100 seconds. The figure shows the average power consumption for the measurement duration. This measurement shows the impact of load and number of loaded cores on the overall power consumption.

[2] http://www.sun.com/servers/x64/x4150/datasheet.pdf

[3] http://code.google.com/p/pack4linux/

B. Deriving an equation relying on available monitoring data

Due to the homogenious hardware used in the G-Lab infrastructure, the static properties listed in Figure 1 can be omitted. These static components will be used to extend the formula to be universally applicable for CPUs of different types and brands. By fitting a polynomial (which has as many independant variables as there are dynamic properties for the component) to the available measurement data, an easily computable equation can be derived. This approach also relies on the tree structuring of the components, as modelling an entire server or even datacenter in a single equation would increase the number of variables beyond feasability. The derived equation can now be called using available monitoring data like CPU load to predict the energy consumption for arbitrary load conditions.

IV. Future Work

Future work will include extending the equations to include the static properties of components, making them usable for a wider range of available hardware. Also, the impact of dynamic voltage and frequency scaling has to be integrated in the equation. Further measurements with longer durations will be carried out in order to get more accurate results

V. Acknowledgement

The authors would like to thank the German Federal Ministry for Education and Research (BMBF) that founds the Project G-Lab_Ener-G. This abstract is based on the work that was carried out in context of G-Lab_Ener-G.

References

[1] X. Fan, W. Weber, and L. Barroso, *Power provisioning for a warehouse-sized computer*, in Proceedings of the 34th annual international symposium on Computer architecture. ACM New York, NY, USA, 2007, pp. 13-23.

[2] S. Rivoire, P. Ranganathan, and C Kozyrakis, *A comparison of high-level full-system power models.* . s.l. : USENIX Association, 2008. HotPower.

[3] Economou, Dimitris, et al. *Full-system power analysis and modeling for server environments.* s.l. : Workshop on Modeling, Benchmarking, and Simulation (MoBS), 2006.