

Energy- and Performance-Aware Resource Management in G-Lab and Future Internet Infrastructures

Gergö Lovász, Florian Niedermeier, Andreas Berl, and Hermann de Meer
Chair for Computer Networks and Computer Communications
University of Passau
lovasz|niedermeier|demeer@fim.uni-passau.de

I. INTRODUCTION

The G-Lab infrastructure currently consists of about 170 nodes. Similar to other ICT infrastructures G-Lab is over-provisioned to deal with spontaneously occurring peak loads and future demands [1]. Therefore, the average utilization of a G-Lab node is about 10% – 20% [2] with regard to CPU and network load, which means that G-Lab is underutilized most of the time. Unfortunately, underutilized or idle (only the operating system is running) servers consume up to 70% of their maximum possible power consumption [3]. In the case of the Sun Fire X4150 [5], which is the standard node in the G-Lab infrastructure, the power consumption in idle state is approximately 250W whereas the power consumption at full load is 363W, which calculates to 69% of the maximum power consumption. This indicates a high energy-saving potential, within the G-Lab infrastructure as well as in Future Internet infrastructures in general. This abstract presents an energy- and performance-aware resource management that aims at the dynamic allocation of services to physical resources that goes beyond currently applied non-energy-aware, utilization-based consolidation approaches. The suggested resource management computes a resource allocation that is based on service requirement models on one hand and on the power consumption models of the physical resources on the other hand. This way, the overall power consumption of G-Lab (and other Future Internet infrastructures) will be minimized while service requirements are fully met. It is important to see that ICT infrastructures tend to consist of heterogeneous devices in terms of performance and energy consumption. This heterogeneity leads to a challenging situation with regard to the resource management which needs to solve a variant of the variable-sized multidimensional bin-packing problem. In contrast to this heterogeneity of ICT infrastructures, the nodes in the G-Lab infrastructure are largely homogeneous. Therefore G-Lab represents a simplified environment which facilitates the resource allocation. The virtualization of services is a key enabler of an energy-efficient resource

management, as resource virtualization allows a flexible and transparent allocation of physical resources to virtualized services. System virtualization is used within G-Lab to create virtual machines that are able to encapsulate services as, e.g., PlanetLab software. This virtualization allows the seamless migration of virtualized services and, therefore, enables the consolidation of several services on a single node. When services are consolidated on a small number of nodes, other nodes can be turned off to save energy. The main question that arises in this procedure is, on which host a certain virtualized service needs to be processed to achieve the minimum possible energy consumption within the supervised infrastructure (e.g., within a G-Lab data center). The infrastructure's energy consumption includes the energy consumption of servers and network equipment as well as the energy consumed by air-conditioning that is required to cool the utilized hardware. At the same time it has to be ensured that resource requirements of services are fully met. Such requirements are e.g. CPU time, RAM, I/O rate and network bandwidth.

II. ENERGY-AWARE RESOURCE MANAGEMENT

Apart from the virtualization approach that is required to enable the seamless migration and the consolidation of services within the managed ICT infrastructure, the suggested energy- and performance-aware resource management consists of three modules:

1. A monitoring/controlling module is needed that enables the monitoring of the energy-relevant parameters of virtualized services and physical hardware within the supervised ICT infrastructure. It provides mechanisms to initiate migrations of virtualized services, to shut down/hibernate hardware, and to wake up the hardware again, if needed.
2. The analyzer module interprets the current state of the ICT infrastructure and its virtualized services. If the changes exceed a specified threshold, the analyzer module reports the change to the optimizer module to initialize a change of the current resource allocation. Additionally, the analyzer module stores relevant

monitoring information and state changes within a dedicated data base. This data is used to build profiles of resource usage and dynamic device characteristics (e.g., heat or fan speed) that are based on historical and current states of the virtualized services and the physical infrastructure.

3. The optimizer module calculates energy-optimal allocations of physical resources to virtualized services that do not violate the resource requirements of the services. The resource allocation is modeled as a variant of the variable-sized multi-dimensional bin-packing problem. Services are represented as hypercubes and servers as hyperbins. Each edge of a hypercube represents a resource requirement of the service. The length of the edge indicates how much of the resource is required. Similarly, edges of the bins represent the server's resources, e.g, CPU cycles or RAM. The length of a bin's edge corresponds to the total amount of the physical resource it represents. To each bin a cost function is assigned that computes the current power consumption of the represented server. The goal of the packing problem is to pack the hypercubes in the hyperbins so that the sum of all cost functions is minimized. On one hand, service requirement models are needed to define the hypercubes. Based on the resource usage profile of a service and current monitoring data, the service requirement model estimates the future load generation of the service and determines the size of the representing hypercube. On the other hand, power consumption functions for the different server components are used as cost functions for the bins. They estimate the power that is consumed by the represented server based on its hardware characteristics and the hypercubes that are inside the bin. The size of the bin is determined by the represented server.

virtual counterpart of the physical infrastructure. The users of the infrastructure interact only with virtualized nodes and are not aware of the virtualization or the presence of an underlying physical infrastructure layer. The energy-aware management interacts with both, the physical and the virtual layer.

Particularly the optimizer component is in the focus of this work. Based on the power consumption models and the resource requirement models it has to be able to 1) estimate the power consumption of all possible resource mappings within the ICT infrastructure and 2) to choose an energy-optimal (or nearly energy-optimal) mapping of physical resources to virtualized services. To estimate the power consumption of different mappings, power consumption models and resource requirement models are used that model the physical and virtual layers of the ICT infrastructure. As an example, the optimizer has to decide whether it is more energy efficient to process 20 virtualized services on three highly energy efficient servers with moderate performance or on two less efficient servers with higher performance.

To find an energy-optimal mapping, heuristics are needed that solve a variant of the NP-hard variable-sized multi-dimensional bin-packing problem. It is important to see that the migration of virtualized services is costly: On the one hand it consumes performance in terms of CPU cycles and network load and on the other hand it costs additional energy. Therefore, it is necessary to achieve an energy-optimal mapping while producing the least possible amount of migrations within the system.

III. EXPECTED RESULTS

The expected result of the implementation of the Ener-G management framework is a significant energy consumption reduction of the supervised ICT infrastructure. The energy consumption is reduced in two steps: First, a straight forward consolidation is applied that consolidates services by using a first fit bin-packing approach. This kind of consolidation will already significantly reduce the energy consumption within G-Lab as a part of the servers can be turned off. In a second step, a full featured energy- and performance-aware management will be applied that considers service utilization as well as the energy consumption of the hardware. In this step, a further reduction of energy consumption is expected. Especially in heterogeneous infrastructures (as Future Internet infrastructures may be), a reduction of energy consumption is expected that significantly exceeds the energy consumption reduction that can be achieved by the first fit bin-packing approach.




Figure 1: Ener-G management framework

All of the described components are illustrated in Figure 1 in the context of the G-Lab infrastructure. The physical layer represents the physical hardware of the G-Lab infrastructure without virtualization and indicates the different energy-states of the server hardware. The virtual layer plays the role of a

- [1] Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e.V., Energieeffizienz im Rechenzentrum Band 2, 2008, p. 10
- [2] Paul Müller, Dennis Schwerdel and Robert Henjes, G-Lab Experimental Facility. G-Lab Status Meeting, February 2011, <https://files.server.german-lab.de/files/templates/Passau/ASP7.pdf>
- [3] David Meisner, Brian T. Gold, and Thomas F. Wenisch. PowerNap: eliminating server idle power. In Soffa and Irwin [4], pages 205-216.
- [4] Mary Lou Soffa and Mary Jane Irwin, editors. Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2009, Washington, DC, USA, March 7-11, 2009. ACM, 2009.
- [5] <http://www.sun.com/servers/x64/x4150/datasheet.pdf>, 31.05.2011
- [6] Bundesverband Informationswirtschaft, Telekommunikation und neue Medien e.V., Energieeffizienz im Rechenzentrum Band 2, 2008