

Demonstration of the Ener-G Resource Allocation System

Gergő Lovász*, Florian Niedermeier*, Hermann de Meer*, Daniel Günther[§], Paul Müller[§] and Stefan Dambeck[✧]

* University of Passau, Email: [lovasz|niederme|demeer}@fim.uni-passau.de](mailto:{lovasz|niederme|demeer}@fim.uni-passau.de); [§] University of Kaiserslautern, Email: [guenther|pmueller}@informatik.uni-kl.de](mailto:{guenther|pmueller}@informatik.uni-kl.de); [✧] :a:k:t: Informationssysteme AG, Email: stefan.dambeck@akt-infosys.de

I. INTRODUCTION

In order to cope with peak utilizations and unknown future demands, operators of communication and computational infrastructures oversize their IT equipment. However, such an over-provisioning is one of the main reasons for the unnecessarily high energy consumption of infrastructures like G-Lab and similar future Internet infrastructures. According to [1], the utilization of hardware resources in data centers is usually between 10% and 50%. In [2], Meisner et al. compare the power consumption of lowly utilized servers to their peak power consumption. They show that lowly utilized servers consume approximately 70% of their peak power consumption. Therefore, the reduction of over-provisioning has a huge potential with regard to energy saving.

A main goal of Ener-G is to reduce over-provisioning in G-Lab and other future Internet infrastructures by resource sharing and dynamic resource allocation. Sharing of resources allows for consolidating services on a subset of physical resources. Furthermore, if resource allocation is done dynamically, resources can be assigned to services based on their actual and not on their peak resource usage. Idle resources can be turned off to save energy and can be turned on again if the resource requirements increase.

The *Ener-G Resource Allocation System* uses resource sharing and dynamic resource allocation to optimize overall infrastructure energy consumption while considering performance constraints of services at the same time.

II. ENER-G RESOURCE ALLOCATION SYSTEM

This section gives an overview on the core concepts applied by the *Ener-G Resource Allocation System*.

A. Virtualization and consolidation

Virtualization and consolidation are key technologies to realize resource sharing and dynamic resource allocation. Services are operated on virtualized servers instead of physical ones. This enables a flexible and transparent resource management. When services are lowly utilized, they are consolidated on a subset of the physical servers so that unused hardware can be powered off. When utilization increases, servers can be powered on again and the virtualized servers are reallocated.

B. Prediction of infrastructure power consumption

The base for finding energy-minimal resource allocations is the ability to predict the infrastructure's power consumption

for different resource allocations. The *Ener-G Resource Allocation System* uses a linear, CPU load dependent server power consumption model suggested by Fan et al. [3]. The model allows for predicting server power consumption based on the actual CPU load. Although the model is simple, it provides sufficient accuracy for calculating energy-minimal resource allocations. Due to its simplicity, the computational overhead of the model is low. By summing up the power consumption of each server, the overall infrastructure power consumption can be determined. This enables comparing different resource allocations and selecting the one with the lowest power consumption.

C. Prediction of performance degradation

When multiple services are consolidated on the same physical server the probability of contention for hardware resources increases. This leads to a degraded performance, even if the hardware resources are not fully utilized. The *Ener-G Resource Allocation System* uses a performance degradation model for CPU-intensive services that allows for predicting the performance degradation caused by consolidation. As performance metric, the mean response time is chosen. Knowing the performance degradation of certain allocations beforehand, enables the consideration of performance constraints of services. The model calculates the expected performance degradation of a consolidated CPU-intensive service based on the number of consolidated services, the CPU load caused by the service itself and the overall CPU load of the hosting physical server. A detailed description of the model is given in [4].

D. Energy- and performance aware optimization model for resource allocation

The optimization model applied by the *Ener-G Resource Allocation System* is a modification of the variable-sized multidimensional vector packing problem. There are several reasons why a vector packing approach has been chosen. Foremost, vector packing is an approved model for task allocation. Furthermore, the vector packing approach has the ability to solve multidimensional problems. Therefore, the allocation of multiple resources to services with multiple resource requirements can be mapped to the dimensions of the vector packing. Finally, several heuristics exist that give good approximations of the NP-hard vector packing problem. In our model, the standard vector packing model is extended by a cost function for bins, representing server power consumption and by a cost function for vectors, representing service

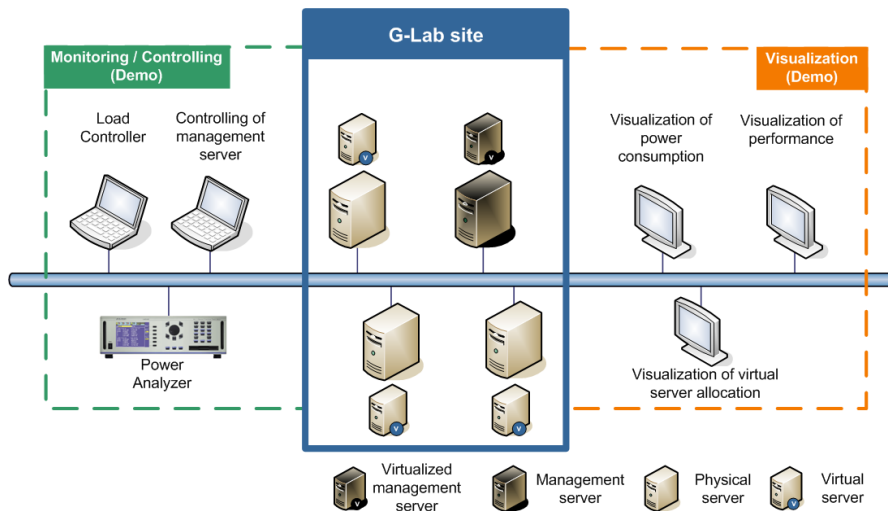


Figure 1: Ener-G demonstration setup

performance. The optimization goal of the modified vector packing is to minimize the sum of the cost functions of the bins (overall power consumption) while considering constraints on the cost functions for vectors (performance constraints of services). A detailed model description is given in [4].

III. DEMONSTRATION

Our demonstration illustrates the key features of the *Ener-G Resource Allocation System*: first, the minimization of infrastructure power consumption by virtualization and energy-aware consolidation and second, the consideration of performance constraints when virtualized servers are consolidated.

A. Demonstration setup

The demonstration setup of the *Ener-G Resource Allocation System* is shown in Figure 1:

The physical and virtual servers in the box *G-Lab site* form the core of the demonstration. One of the four physical servers takes over the role of the management server which corresponds to the head node of a real G-Lab site. Using the Proxmox Virtualization Platform¹, the physical servers are virtualized. The virtualized management server hosts the *Ener-G Resource Allocation System* that determines the allocation of the physical servers to the virtualized ones and is responsible for shutting down and powering up physical servers. The management server monitors the virtualized and physical servers using the Nagios² monitoring system. For the demonstration, on each virtualized server a computationally intensive service is running, utilizing the CPUs of the physical servers at different load levels.

The boxes *Monitoring / Controlling* and *Visualization* contain auxiliary components of the demonstration which are not part of the *Ener-G Resource Allocation System* but play an important role in the demonstration.

In the *Monitoring / Controlling* box a “Load Controller” sends computational requests to the services running on the virtualized servers of the simulated G-Lab site. By changing the frequency of the requests, the services can be operated at different load levels. Furthermore, a controlling entity is interconnected with the management server. The controlling entity allows for changing the allocation algorithms used by the *Ener-G Resource Allocation System* and for tuning performance constraints of the services at runtime.

Additionally, the physical servers of the simulated G-Lab site are connected to a power analyzer that monitors the power consumption of each physical server individually.

The *Visualization* box consists of three screens visualizing information on the virtualized and physical servers of the simulated G-Lab site. On one of the screens the actual mapping of the virtualized servers is shown. On a second screen the predicted and measured power consumptions of all four physical servers are depicted as graphs. The third screen visualizes the predicted and measured performance of the services running on the virtualized servers as graphs. All the visualizations are updated in real-time so that changes of the monitored values can be seen instantaneously.

B. Demonstration goals

The goal of the demonstration is to show the *Ener-G Resource Allocation System* in action. In the demonstration the load of the running services is modified, forcing the *Ener-G Resource Allocation System* to take actions. The energy consumption of the managed and non-managed infrastructure is compared by enabling and disabling the *Ener-G Resource Allocation System*. The visualization part of the demonstration provides an insight into the functioning of the different components of the *Ener-G Resource Allocation System*. Furthermore, the demonstration allows for visualizing the impact of different resource management strategies developed in Ener-G on service performance and infrastructure energy consumption.

ACKNOWLEDGEMENTS

The research leading to these results was supported by the German Federal Ministry of Education and Research (BMBF) in the context of the G-Lab_Ener-G project (ref. no. University of Passau: 01BK0910, ref. no. University of Kaiserslautern: 01BK0911 and ref. no. :a:k:t: Informationssysteme AG: 01BK0912).

REFERENCES

- [1] Barroso, L.A., Hözl, U.: The case for energy proportional computing. *IEEE Computer* 40(12), 33–37 (2007).
- [2] Meisner, D., Gold, B.T., Wensich, T.F.: Powernap: eliminating server idle power. In: *ASPLOS*, pp. 205–216. ACM (2009)
- [3] Fan, X., Weber, W.D., Barroso, L.A.: Power provisioning for a warehouse-sized computer. In: *Proceedings of the 34th Annual International Symposium on Computer Architecture, ISCA '07*, pp. 13–23. ACM, New York, NY, USA (2007)
- [4] Lovász, G., Niedermeier, F., De Meer, H.: Performance Tradeoffs of Energy-Aware Virtual Machine Consolidation. *Cluster Computing*, [NA]:[NA] (2012) DOI: [10.1007/s10586-012-0214-y](https://doi.org/10.1007/s10586-012-0214-y)

¹ <http://www.proxmox.com/>

² <http://www.nagios.org/>