WattDB – An Energy-proportional DBMS

Volker Höfner, Daniel Schall
Databases and Information Systems Group
University of Kaiserslautern, Germany
{hoefner, schall}@cs.uni-kl.de

Motivation

Recent experiments have shown that it is impossible to achieve energy efficiency in single-server environments using current, non-energy-proportional hardware [1]. State-of-the-art computer components use about 50% of their peak energy just for being idle.
Barroso et. al. have shown that servers typically operate under 10 – 50 % load [2], thus offering optimization potential in idle times. This topic recently gained a lot of attention and many papers have been published that try to improve the energy footprint of (database) servers. Nevertheless, no more than about 15% improvement in energy efficiency was gained, often by sacrificing performance.
Based on these findings, we conclude that a single, large database server as seen in today's enterprises will never be energy proportional. Due to its high idle power consumption and the lack of scalable hardware, single-server-based databases will waste most of their energy during times of low load.

The Vision

To overcome the drawbacks mentioned, we envision a locally distributed database system that runs on a cluster of lightweight nodes. While light-weighted servers show the same non energy-proportional behavior than larger ones, they have a much smaller energy footprint per node. Furthermore, smaller servers can be powered up and down independently.

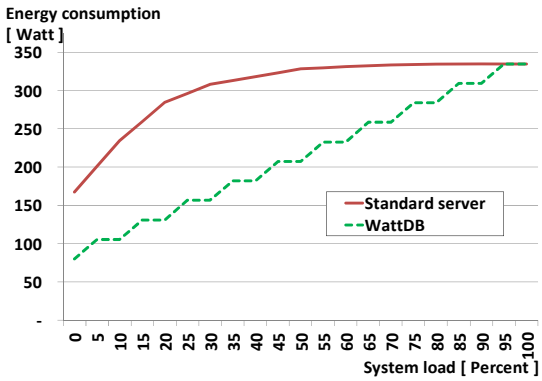

Figure 1 Energy proportionality

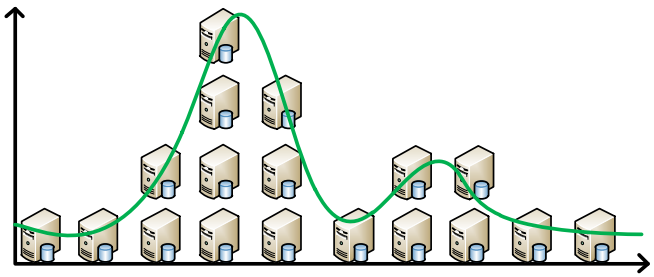

Figure 2 WattDB node provisioning

As depicted in Figure 1, by powering small-scale nodes according to the cluster's load, energy-proportional behavior can be approximated. The chart shows a server cluster, which consumes the same energy as a comparable, large server at peak. Especially in times of low utilization, the cluster's energy consumption can be much lower, due to powering down some of the nodes.

Clustering databases is a technique that is supported by all commercially available DBMS, e.g. to increase performance and fault tolerance. Nevertheless, removing nodes from the cluster to switch back to a more energy-efficient configuration is not supported.
Therefore, our goal is to develop a DBMS that balances power consumption proportionally to the system's load by dynamically powering its nodes individually up and down as sketched in Figure 2. We expect to combine fair peak performance with energy proportional behavior in order to save energy under changing workload situations.

Work in Progress

To bring our vision to life, we are going to implement a DBMS step-by-step from scratch, based on commodity, low-power hardware, e.g. using an Intel Atom CPU and mobile hard drives. Participating in the SIGMOD Programming Contest 2010, where we finished second [3], gave us first insights in programming distributed database systems. Based on this solely performance-centric implementation, we are going to implement *WattDB*.
Currently, we are developing a framework and additional tools for measuring the energy consumption of the cluster [4]. In June, we will present our first findings at the SIGMOD Demo Track [5].

About the Authors

Volker Höfner and Daniel Schall studied computer science at the University of Kaiserslautern. They already focused on energy efficiency in their diploma theses where they developed tools to measure the energy consumption and performance of solid-state drives.
In October 2009, they started their scientific career as PhD students under Prof. Dr. Dr. h. c. Theo Härder. Their first project was the participation in the SIGMOD Programming Contest 2010. Afterwards, they initiated the WattDB project to build an energy-proprtional DBMS.

References

[1] Tsirogiannis, Dimitris, Harizopoulos, Stavros, and Shah, Mehul A. Analyzing the Energy Efficiency of a Database Server. *SIGMOD Conference* (2010), 231-242.

[2] Barroso, Luiz André and Hölzle, Urs. The Case for Energy-Proportional Computing. *Computer*, 40, 12 (Dec. 2007), 33-37.

[3] Genzmer, Clément, Hudlet, Volker, Park, Hyunjung, Schall, Daniel, and Senellart, Pierre. The SIGMOD 2010 Programming Contest a Distributed Query Engine. *SIGMOD Rec.*, 39, 2 (June 2010), 61-64.

[4] Hudlet, Volker and Schall, Daniel. Measuring Energy Consumption of a Database Cluster. *Proc. 14-th GI-Conference Database Systems for Business, Technology and Web, LNI - P 180* (Feb. 2011), 734-737.

[5] Schall, Daniel and Hudlet, Volker. WattDB: An Energy-Proportional Cluster of Wimpy Nodes. *SIGMOD Conference – Demo Track*, to appear (June 2011).