Energy-Aware Computing in Heterogeneous Datacenters

Motivation

In recent years, power requirements have become a primary design objective for large scale compute systems. Energy has become a critical operating resource and plays a central role in the planning and operation of modern data centers. On the one hand, energy consumption forms significant part of the operating costs of a data center, which usually exceed its initial costs significantly. On the other hand, power dissipation has become a limiting factor on the performance and scalability of an IT system. The available power supply determines the size and number of machines that can be operated in an installation. The required cooling capacity is also closely linked to energy consumption, as the power dissipated by the compute system is converted into heat.

NVidia Jetson TX2 Board

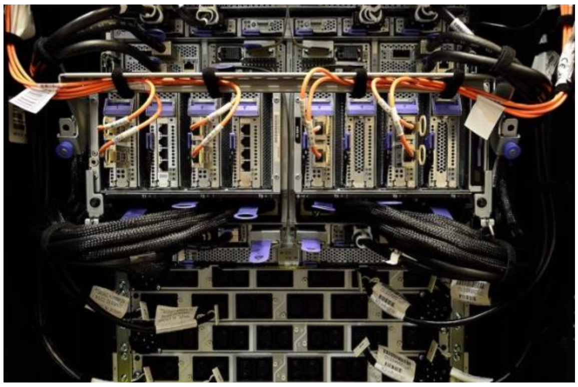

IBM E880 Server System

The situation is further complicated by an increasing level of heterogeneity in current system architectures. Besides traditional general-purpose processors, a heterogeneous architecture incorporates accelerators such as GPUs, FPGAs, or TPUs. This architectural trend results in a significantly increased power dynamism at runtime, because each type of processing element might run under different loads and have a distinct energy consumption curve. To cope with this complexity, empirical methods are indicated.

There is already an extensive body of cross-platform tools and to analyze performance characteristics of a system in terms of throughput, latency. Yet, current toolchains lack comparable instrumentation for analyzing a system's power consumption. Some hardware platforms do offer specific mechanisms to monitor and analyze power consumption at runtime, but different solutions vary widely in terms of measurement granularity and resolution. Consequently, there is a need for uniform power measurement methods.

From a sufficiently detailed trace of the system power consumption over time, an analysis can reveal valuable insights regarding the efficiency of the overall system operation as well as of specific workload executions. In addition, the gathered data can be used to optimize operation procedures or adjust workload execution modalities.

Project Goals

In this master project you will use and extend a platform-independent software suite for uniform power and energy measurements, as part of the ongoing research at our lab.

The objective is to gain insights on the energy characteristics of different workloads on different hardware platforms. You will be given access to a large variety of computing devices. With the collected data you can develop strategies to optimize job placement in data centers, towards a better energy efficiency.

We will assist you publishing your findings in a technical report or a scientific article.

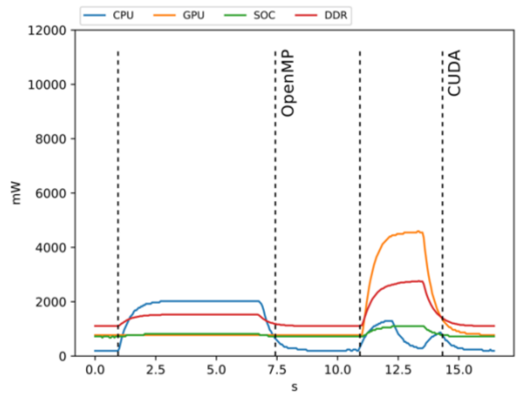

Sampled power draws over time on a Jetson TX2, showing the different behavior of the same physics simulation implemented with OpenMP, and CUDA.

```
$ pinpoint -r 4 -i 250 -- ./heatmap_sim random.csv
Energy counter stats:
[interval: 250ms, delay: 0 ms, runs: 4]
     5609.31 mJ CPU ( +- 5.37% )
     2311.00 mJ GPU ( +- 0.03% )
     2404.06 mJ SOC ( +- 1.06% )
     4496.69 mJ DDR ( +- 1.20% )
     18372.19 mJ IN ( +- 3.05% )
     2.87524 seconds time elapsed ( +- 1.09% )
```

Example run of pinpoint, our cross-platform energy measurement suite.

Contact

If you have any further questions, please do not hesitate to contact us.

Sven Köhler (C-1.11, sven.koehler@hpi.de)
Lukas Wenzel (C-1.11, lukas.wenzel@hpi.de)
Max Plauth (C-1.6, max.plauth@hpi.de)
Prof. Dr. Andreas Polze (C-1.8, andreas.polze@hpi.de)

Related Work

[1] K. W. Cameron, Rong Ge and Xizhou Feng, “High-performance, power-aware distributed computing for scientific applications,” in *Computer*, vol. 38, no. 11, pp. 40-47, Nov. 2005, doi:10.1109/MC.2005.380.
[2] Benedict Herzog, Max Plauth, Timo Hönig, Sven Köhler, Wolfgang Schröder-Preikschat, Andreas Polze. 2019. “Bridging the Gap: Energy-efficient Execution of Software Workloads on Heterogeneous Hardware Components.” In *Proceedings of the Tenth ACM International Conference on Future Energy Systems (e-Energy ’19)*, June 25–28, 2019, Phoenix, AZ, USA, doi:10.1145/3307772.3330176
[3] Pierre Olivier, A K M Fazla Mehrab, Stefan Lankes, and Mohamed Lamine Karaoui, Rob Lyerly, Binoy Ravindran. 2019. “HEXO: Offloading HPC Compute-Intensive Workloads on Low-Cost, Low-Power Embedded Systems.” In *The 28th International Symposium on High-Performance Parallel and Distributed Computing (HPDC ’19)*, June 22–29, 2019, Phoenix, AZ, USA, doi:10.1145/3307681.3325408