GreenClouds: Towards Green Computing on Clouds

Qingwen Chen*, Paola Grosso*, Rutger Hofman**, Cees de Laat* and Henri Bal**
* University of Amsterdam ** Vrije Universiteit Amsterdam

GreenClouds project

The GreenClouds project investigates energy-savvy solutions for modern High-Performance Computing (HPC) systems (like clouds) while taking into account both performance and energy consumption. The project will:

- Characterise the energy consumption of a broad range of applications on various types of hardware
- Provide global system-level optimizations through dynamic adaptation of computing and networking resources

The research results will be applied on the HPC systems of Dutch national HPC center (SARA) and DAS-4 clusters.

Research goal

GreenClouds focuses on three aspects that together provide the basic components for a system-level approach:

- **Hardware diversity**: we exploit the diversity of computing architectures (e.g. GPUs, multi-cores) to run computations on those architectures that perform them in the most energy-efficient way;
- **Elastic scalability**: we dynamically adapt the number of resources to the application needs accounting for computational and energy efficiency;
- **Hybrid networks**: we use optical and photonic networks to transport data and computations in a more energy-efficient way.

DAS-4 Clusters

DAS-4 (The Distributed ASCI Supercomputer 4)[2] is a six-cluster wide-area distributed system designed to provide a common computational infrastructure for scientific researchers within ASCI.

The two clusters hosted by Vrije Universiteit Amsterdam (VU) and University of Amsterdam (UvA) have the following features:

Cluster	Nodes	Type	Speed	Memory	Node HDDs
VU	74	dual-quad-core	2.4GHz	24GB	2*1TB
UvA	16	dual-quad-core	2.4GHz	24GB	1TB

Power measuring environment

Our power measuring setup consists of four componenets:

- *System under test (SUT)*: a single standard DAS-4 work node;
- *Workloads running on SUT*: LINPACK, which is used as a performance indicator for ranking supercomputers in TOP 500 list;
- *Data collection application*: Ganglia with our power metric extension;
- *Power meter*: a 32A PDU gateway from Schleifenbauer with the precision of 0.1 Watt.

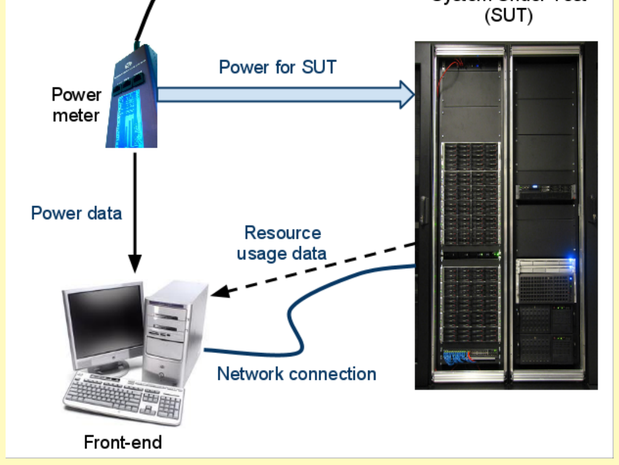

Power Model

The total power consumption of a SUT follows the power model described below:

$$P_{total} = P_0 + c_1 P_{CPU} + c_2 P_{cache} + c_3 P_{DRAM} + c_4 P_{HDD}$$

where P_0 is the idle power consumption of SUT, and $P_{CPU}, P_{cache}, P_{DRAM}$ as well as P_{HDD} are specific performance parameters for CPU, cache, DRAM and disk, and c_0, c_1, c_2, c_3, c_4 are weights.

In presence of accelerators, e.g. GPU or FPGA, the model is adjusted by adding the corresponding sub-components. Within our GreenClouds project, we will focus on the impact of variation in sub-components' performance with respect to the total power consumption of the SUT.

Results

Two series of power benchmarks are performed to characterise the power consumption of both the host and guest VMs running on it:

- Variatons of the host's CPU frequency;
- Changes of the number of a VM's virtual CPUs (vCPUs).

The LINPACK benchmark is used to stretch CPU and memory usage to their limits and report the result in terms of MFlops or GFlops.

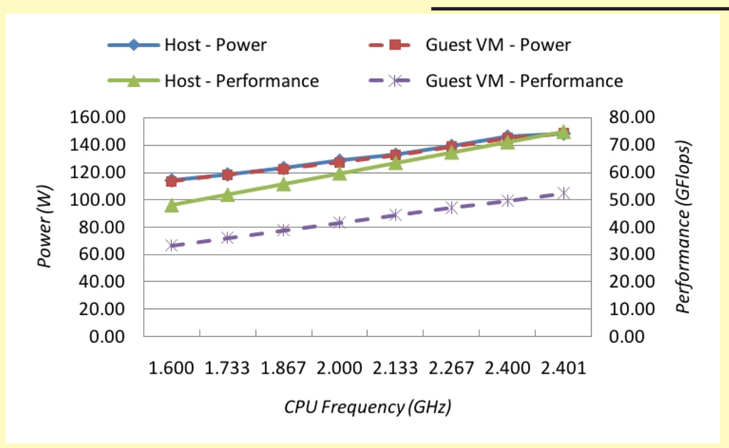

Fig. 1. Performance, power and energy consumption of LINPACK benchmark for various CPU frequencies.

The guest VM has nearly the same power consumption as the host does, but with worse performance. Moreover, the performance and power consumption increases almost linearly as the frequency scales up, so does the saved energy. Therefore, *it's 'greener' to overclock the CPU for CPU- and memory-intensive workload*.

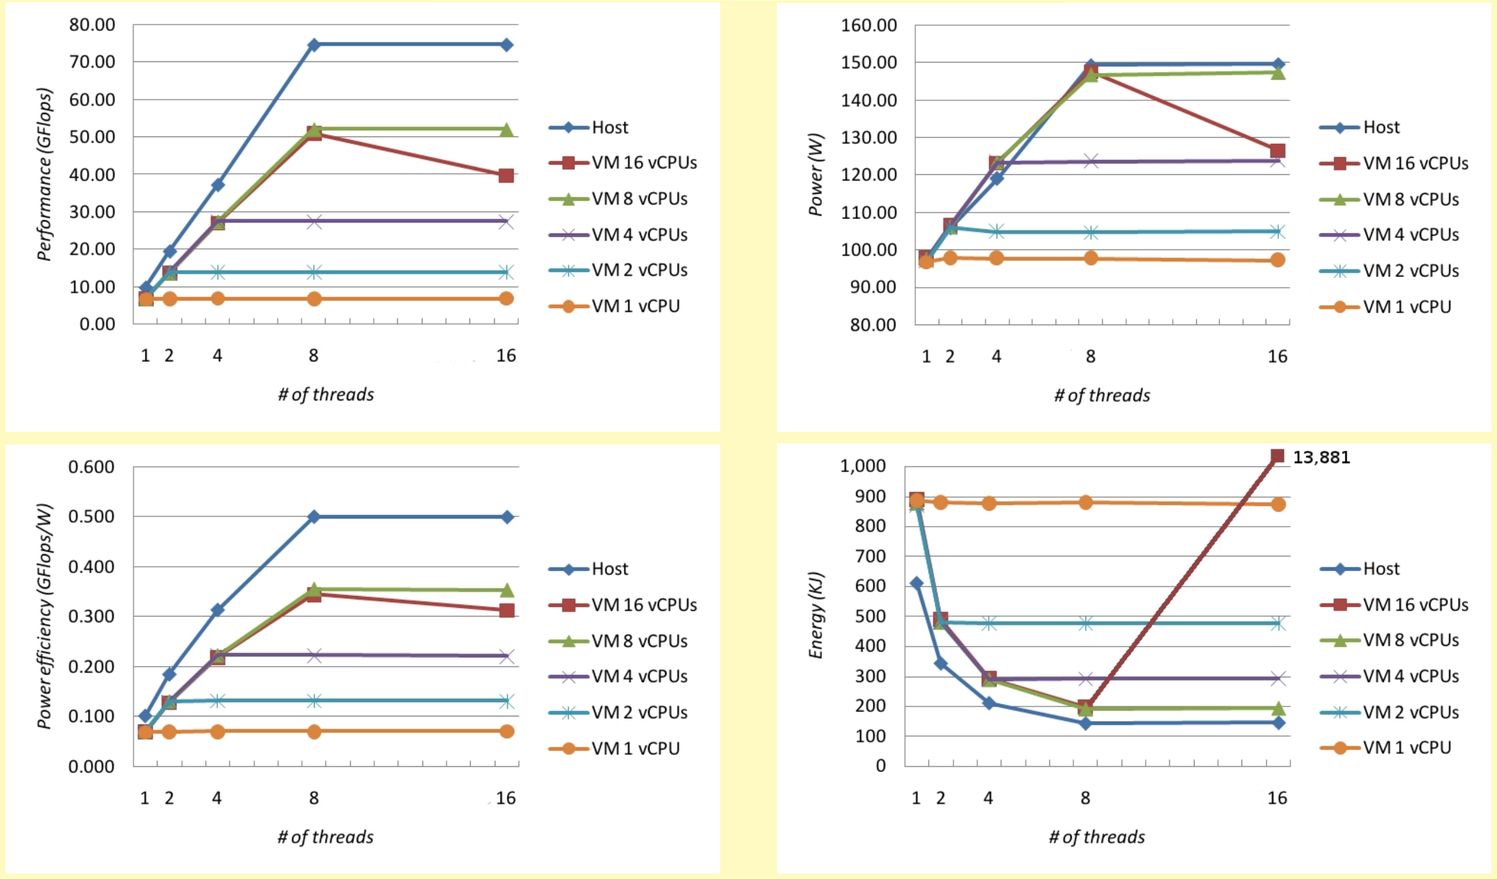

Fig. 2. Several VMs with the same configurations except the number of vCPUs.

All measured parameters increase until they reach a plateau when the number of threads is the same as the number of vCPUs for all non-overcommitted cases. Besides it,

- The host always has the best performance among all tests, and virtualisations normally result in ~30% degradation in overall performance.
- The overcommitted VM (i.e. the one with 16 vCPUs) consumes much more energy due to the extreme long execution time.

Funding

This work is sponsored by NWO - the Dutch Organization for Research.

References

[1] Green Clouds project: http://wiki.cs.vu.nl/greenclouds/index.php/Main_Page
[2] DAS-4 clusters: http://www.cs.vu.nl/das4