



enerGyPU for Monitoring Performance and Power Consumption on Multi-GPUs



John A. García. H.^{1,2}, Philippe O. A. Navaux², C.J. Barrios Hernández¹

^{1,2} john.garcia1@correo.uis.edu.co, ² navaux@inf.ufrgs.br, ¹ cbarrios@uis.edu.co

Super Computación y Cálculo Científico UIS

¹ High Performance and Scientific Computing Center, SC3UIS
Universidad Industrial de Santander - Bucaramanga, Colombia

² Parallel and Distributed Processing Group, GPPD
Universidade Federal do Rio Grande do Sul - Porto Alegre, Brasil

Abstract

enerGyPU is a tool that can be used to analyze multiple tests under different combinations of parameters to observe the key factors that determines the energy efficiency in terms of 'Energy per Computation' on Clusters with Multi-GPUs. enerGyPU is a monitor to centralize and automate the data captured in runtime while is executed in parallel with the scientific application and displays information via sequence plots, statistical tables, bar graphs and shows results in terms of energy efficiency.

Background and Motivation

The evaluation of performance and power consumption is a key step in the design of applications for large computation systems, such as supercomputers, clusters with nodes that have manycores and multi-GPUs. Researchers must design several experiments for workload characterization by observing the architectural implications of different combinations of parameters; such as problem size, number of cores per GPU, number MPI of process and observe the resulting clock frequency, memory usage, bandwidth usage, and power consumption, which are factors that determine the performance and energy efficiency of their workload implementation. The major problem in performance evaluation consists on experiment design given so many parameters. If the resulting data is not properly collected and the parameters of evaluation performance are not well organized, the design exploration may achieve wrong conclusions [1].

enerGyPU is developed to get accurate measurements of many experiments that are need for the project called "Energy Efficient Acceleration of Large-Scale Scientific Applications on Heterogeneous Architectures", where the goal is determine the distribution of tasks of large-scale scientific applications on a CPU-GPU heterogeneous architecture, for an efficient computational acceleration and low power consumption. This approach allows to obtain the optimal number of computational resources such as GPU, cores per GPU and number of Process MPI at a static time, in which a developer or framework of load balance can configure the number of computational resources to maximize the energy efficiency. This project is developed in collaboration by the High Performance and Scientific Computing Center, SC3UIS and Parallel Distributed Processing Group, GPPD.

enerGyPU Monitor Structure

enerGyPU is a type of batch monitor that is formed by two levels. One level is called *enerGyPU to data capture in runtime*, which is executed in parallel with the scientific application, and is composed by three events. A separate level called *enerGyPU to data visualization online*, that can be used later for analyses of results of each experiment in terms of energy efficiency, and is composed by four events.

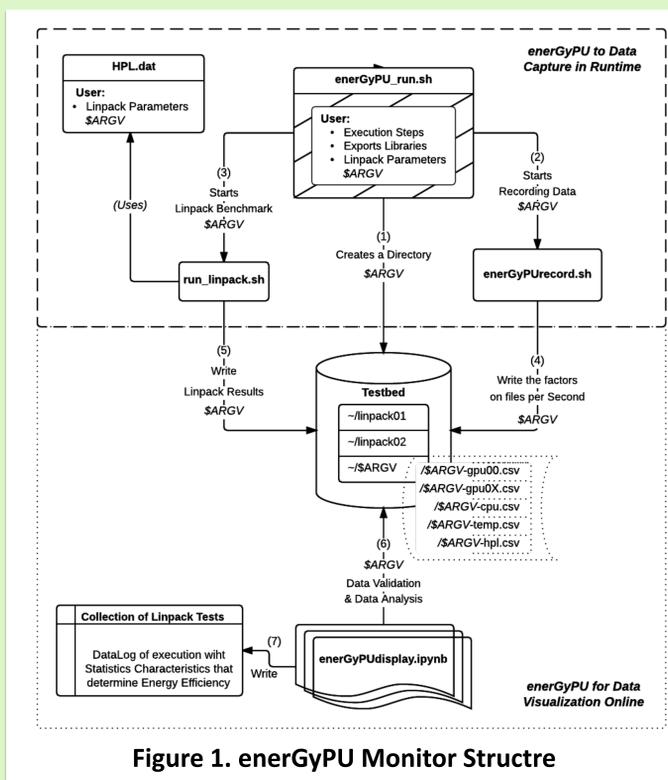


Figure 1. enerGyPU Monitor Structure

The first level is used to run a script called *enerGyPUn.sh* where execution steps have been declared, and the libraries needed to run scientific application, we used the Linpack Benchmark as study case. The user writes the same Linpack parameters defined in the file *HPL.dat* to execute the experiment. Those parameters are used as a unique key called ARGV to share the argument in all the following events. *enerGyPUrecord.sh* is executed in background for data extraction and writing: one file per GPU parameter (Streaming Multiprocessor Clock Frequency-SM, Memory Clock Frequency, Memory Usage, Power Consumption). New data registered every second.

The second level is used at post-processing for data visualization. It displays information via sequence plots, statistical tables, histograms and shows results in terms of energy efficiency. The event (6) displays information of tests through the IPython Notebook Viewer in a program called *enerGyPUDisplay.ipynb* that contains predefined code routines written in Python, where the researchers have to type the argument ARGV to identify the unique test of the Linpack benchmark.

Experimental Procedures and Results

This section presents the use of enerGyPU to observe the key factors that determine the energy efficiency in one test with Linpack benchmark. We used HPL.2.0 version configured for Tesla GPUs [3]. The Linpack parameters used are shown in Table 1. The computational resources used for this experiment are part of one node of the 'A' settings of cluster GUANE that consists in ProLiant SL390s-G7 computing nodes such as shown in Table 2.

Matrix size	49152
Block size	1024
GPU Used	4
Cores per GPU	3
Process MPI	4

Table 1. The Linpack parameters Used in experimental procedure.

Setting GUANE	A	B	C
Node type	SL390s	SL390s	SL390s
Number of nodes	8	3	5
Processor Intel	Xeon E5645	Xeon E5645	Xeon E5640
Processor Model	E5645	E5645	E5640
Processor by node (#)	2	2	2
Clock frequency (GHz)	2.40	2.40	2.670
Core/Processor (#)	6	6	4
Thread/Core (#)	2	2	2
GPUS Nvidia	Tesla M2075	Tesla M2050	Tesla M2050
GPUS Model	104	8	8
GPUs by node (#)	8	104	104
Memory DDR (GB)	200	200	200
SAS disk (GB)	200	200	200

Table 2. Settings GUANE Cluster Nodes.

This experiment procedure shows the 1 GPU of 4 GPU that are inactive, and other 2 GPUs of 4 GPUS actually working have constant frequencies of 1147MHz for SM clock and 1566MHz for memory clock, with a constant memory usage of 2128MiB, thus generating an average power of 120 watts between working GPUs. The results obtained for performance was 691.20 GFLOPS, for one node of Cluster GUANE achieved the energy efficiency of 1097.68 MFLOPS/W.

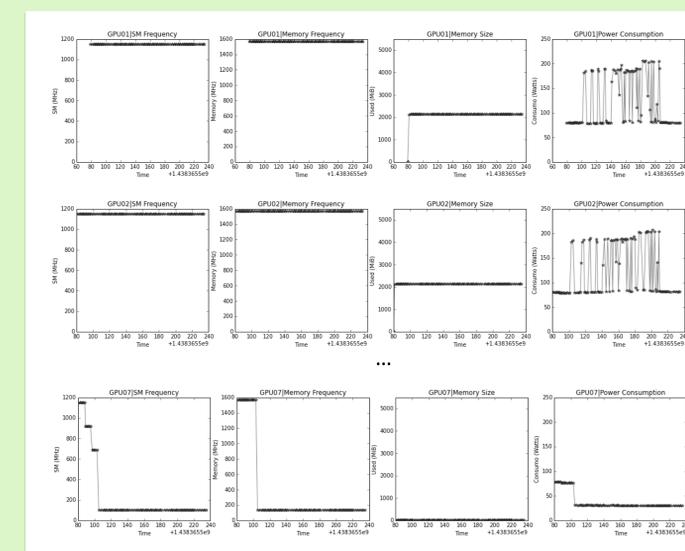


Figure 2. enerGyPU-Sequenceplot of Power Consumption by each GPU

References

- [1] Raj Jain. The art of computer systems performance analysis - techniques for experimental design, measurement, simulation, and modeling. Wiley professional computing. Wiley, 1991.
- [2] Jack J. Dongarra. Performance of Various Computers Using Standard Linear Equations Software. Electrical Engineering and Computer Science Department, University of Tennessee, University of Manchester, 2014.
- [3] Massimiliano Fatica. Accelerating linpack with CUDA on heterogenous clusters. In David R. Kaeli and Miriam Leeser, editors, GPGPU, volume 383 of ACM International Conference Proceeding Series. ACM, 2009.

More Information:
www.sc3.uis.edu.co/
www.inf.ufrgs.br/npd/

Project Repository:
<https://github.com/jagarcia/enerGyPU>

