Contribution ID: **11** Type: **Talk**

Energy-Time Analysis of Heterogeneous Clusters for EEG Classification

Tuesday, 5 November 2019 15:20 (30 minutes)

Power-aware computing introduces an additional dimension in the development of efficient parallel codes for heterogeneous computing architectures. Along with experimental frameworks that facilitate the accomplishment of experimental measures, there is a need for data analysis strategies and programming guidelines and strategies that jointly consider speed and consumption performance, among others. Our communication in the present edition of the PACO workshop describes the work developed on time-energy analysis of an electroencephalogram (EEG) classification problem implemented on heterogeneous clusters with nodes including CPU and GPU architectures, and the techniques proposed by our research group of the Department of Architecture and Computer Technology of the University of Granada in this topic [2, 3, 4, 5].

Our approach analyses the energy consumption and runtime behaviors of a parallel master-worker evolutionary algorithm according to the workload distribution among the GPU and CPU cores and their operating frequencies. In the kind of evolutionary procedure here considered, most of the workload corresponds to the fitness evaluation of the solutions that constitute the population evolved across the generations (iterations) of the algorithm. Thus, the way the solutions are allocated to GPU and CPU cores will determine the efficiency in performance and energy behavior of the parallel algorithm once a suitable workload distribution among the available cores is attained. We have defined models [3] for runtime and energy consumption that have been fitted to the experimental results by multiple linear regression with values of the R^2-statistics that show high statistical significance in all the cases.

The implemented program is an *OpenCL* (version 1.2) code (compiled with *GCC* 4.8.5) for a multi-objective feature selection problem corresponding to a BCI task [6] applied to a dataset recorded in the BCI Laboratory of the University of Essex. Each pattern is an electroencephalogram (EEG) described by 3600 features corresponding to 12 features per each of the 20 temporal segments and 15 electrodes [6].

From these models, it is possible to define a workload distribution that takes into account both the runtime and the energy consumption through a bi-objective cost function that properly weights both objectives. As many useful bioinformatics and data mining applications are tackled by programs with a similar profile to that of the here considered parallel master-worker procedure, the conclusions about energy consumption and speedups can be taken into account many times, and the proposed energy-aware scheduling approach could be frequently applied.

Regarding the operating mode control, the standard ACPI (Advanced Configuration and Power Interface) [9] includes mechanisms to manage and save energy, provides a suitable control of the BIOS functioning and gives information about the configuration and the control of the processor states related to the energy consumption $(C_0, C_1, C_2, C_3, \ldots, C_n)$ and performance $(P_0, P_1, \ldots, P_n)$. Likewise, the *Linux* kernel implements the CPUFreq (CPU Frequency scaling) subsystem [7], which allows the operating system, either automatically through the events generated by the ACPI, or through user program calls, to change the operating frequency of the processor for energy saving. The so-called governors [8] are included in the CPUFreq subsystem, to implement specific policies for controlling the processor clock. The interface to use these services at the user level can be found in the header file linux/cpufreq.h [1].

Acknowledgments
This work was supported by project PGC2018-098813-B-C31 (Spanish *Ministerio de Ciencia, Innovación y Universidades*), and by European Regional Development Funds (ERDF).

References

1. Brodowski, D.: The linux/cpufreq.h header file. https://github.com/torvalds/linux/blob/master/tools/power/cpupower/lib/cpufreq.h accessed: 2019-10-07
2. Escobar, J.J., Ortega, J., Damas, M., Savran, R., Gan, J.Q.: Energy-time analysis of convolutional neural networks distributed on heterogeneous clusters for EEG classification. In: Advances in Computational Intelligence. IWANN 2019. pp. 895–907. Lecture Notes in Computer Science, Springer, Gran Canaria, Spain (June 2019). https://doi.org/10.1007/978-3-030-20518-8_74
3. Escobar, J.J., Ortega, J., Díaz, A.F., González, J., Damas, M.: Energy-aware load balancing of parallel evolutionary algorithms with heavy fitness functions in heterogeneous CPU-GPU architectures. Concurrency and Computation. Practice and Experience 31(6), e4688 (2019). https://doi.org/10.1002/cpe.4688

4. Escobar, J.J., Ortega, J., Díaz, A.F., González, J., Damas, M.: A power-performance perspective to multiobjective electroencephalogram feature selection on heterogeneous parallel platforms. Journal of Computational Biology 25(8), 882–893 (2019). https://doi.org/10.1089/cmb.2018.0080
5. Escobar, J.J., Ortega, J., Díaz, A.F., González, J., Damas, M.: Time-energy analysis of multilevel parallelism in heterogeneous clusters: the case of EEG classification in BCI tasks. Journal of Supercomputing 75(7), 3397–3425 (2019). https://doi.org/10.1007/s11227-019-02908-4
6. Ortega, J., Asensio-Cubero, J., Gan, J.Q., Ortiz, A.: Classification of motor imagery tasks for BCI with multiresolution analysis and multiobjective feature selection. BioMedical Engineering OnLine 15(1), 73 (2016). https://doi.org/10.1186/s12938-016-0178-x
7. The Linux Kernel: CPU performance scaling. https://www.kernel.org/doc/html/v4.15/admin-guide/pm/cpufreq.html, accessed: 2019-10-07
8. The Linux Kernel: CPUFreq governors. https://www.kernel.org/doc/Documentation/cpu-freq/governors.txt, accessed: 2019-10-07
9. UEFI Forum, Inc.: Advanced configuration and power interface (ACPI) specification. https://uefi.org/sites/default/files/resources/A
accessed: 2019-10-07

Primary authors: ESCOBAR, Juan José (University of Granada, Granada, Spain); ORTEGA, Julio (University of Granada, Granada, Spain); DÍAZ, Antonio F. (University of Granada, Granada, Spain); GONZÁLEZ, Jesús (University of Granada, Granada, Spain); DAMAS, Miguel (University of Granada, Granada, Spain)

Presenter: ORTEGA, Julio (University of Granada, Granada, Spain)

Session Classification: Day I