

August 16, 2017

# Implementation of the K-means Algorithm on Heterogeneous Devices: a Use Case Based on an Industrial Dataset

Ying hao XU <sup>a,b</sup>, Miquel VIDAL <sup>a,b</sup>, Beñat AREJITA <sup>c</sup>, Javier DIAZ <sup>c</sup>,  
Carlos ALVAREZ <sup>a,b</sup>, Daniel JIMÉNEZ-GONZÁLEZ <sup>a,b</sup>, Xavier MARTORELL <sup>a,b</sup>,  
Filippo MANTOVANI <sup>a,1</sup>,

<sup>a</sup> *Barcelona Supercomputing Center (BSC)*

<sup>b</sup> *Universitat Politècnica de Catalunya, Barcelona Tech (UPC)*

<sup>c</sup> *Plethora IIoT*

**Abstract.** This paper presents and analyzes a heterogeneous implementation of an industrial use case based on K-means that targets symmetric multiprocessing (SMP), GPUs and FPGAs. We present how the application can be optimized from an algorithmic point of view and how this optimization performs on two heterogeneous platforms. The presented implementation relies on the OmpSs programming model, which introduces a simplified pragma-based syntax for the communication between the main processor and the accelerators. Performance improvement can be achieved by the programmer explicitly specifying the data memory accesses or copies. As expected, the newer SMP+GPU system studied is more powerful than the older SMP+FPGA system. However the latter is enough to fulfill the requirements of our use case and we show that it uses less energy when considering only the active power of the execution.

**Keywords.** FPGA, Arm, Clustering, Heterogeneous programming, OmpSs, FPGA automatic toolchain

## 1. Introduction

The complexity of current commercial and enterprise applications such as e-commerce, health monitoring, industrial production and financial data analysis, rely more and more on Machine Learning techniques to achieve their objectives. As a consequence, the popularity of programmable accelerators has grown in both the industry and the research community, which can provide large gains in efficiency and performance by performing specialized tasks. The use of field-programmable gate arrays (FPGAs) and general purpose graphic processing unit (GP-GPU) as programmable accelerators can significantly increase the performance and efficiency gains while retaining some of the flexibility of general-purpose processors [9].

Accelerating machine learning algorithms with reconfigurable hardware and comparing the achieved speed up results to the ones obtained with SMP and GP-GPU is not

---

<sup>1</sup>Corresponding Author: Filippo Mantovani, Barcelona Supercomputing Center, C/ Jordi Girona, 29 - 08034 Barcelona, Spain; E-mail: [filippo.mantovani@bsc.es](mailto:filippo.mantovani@bsc.es)