**Thesis Overview**

# Evaluation of High Performance Computing Platforms for Drug Discovery

University of Murcia, Computer Engineering and Technology Department
June, 27th 2014
Ginés D. Guerrero
gines.guerrero@ditec.um.es

In the first decade of the century, the Moore's Law, which has led the microprocessor design in the last fifty years, was put into question by the scientific community. This was mainly due to the physical limitations of silicon-based architectures, which caused a change in the trend of designing processors, guided by parallelism. This transition has placed (massively) programming parallel as the only way to extract the maximum performance to new consumer platforms; being this essential to address today's scientific challenges. Unfortunately, these challenges propose several issues whose computing needs are our of the scope for a single machine. Simulations, such as those discussed in this PhD Thesis, need to scale to large data centers; whose costs are only affordable for large institutions and governments. However, the current socio-economic situation requires an efficient use of resources. Tools such as cloud computing or volunteer computing offer an alternative to exploit computing resources in a flexible, fast, economical and environmentally friendly way.

In my dissertation, we evaluate the current landscape of computation, previously described, using as a case study a high-impact problem for society as virtual screening. Virtual screening is a computational tool extensively used for drug discovery. The study has covered all processing levels, starting with an extensive analysis of the different commercially available alternatives at chip level, through their evaluation in a cluster environment, to scale to cloud computing and volunteer computing levels.

This study concludes that GPUs are at the leading-edge of the development of scientific applications with massively parallel computing patterns and high computational demands, such as virtual screening. Moreover, this conclusion can be extended to other application fields with the same characteristics in its computation. However, migration to GPUs may cause an application redesign and even rethought, but this is actually part of computational thinking, which is now essential to develop scientific applications for the current state of high performance computing.

Different alternatives to the use of a cluster need to be evaluated, such as the use of cloud computing and volunteer computing for a larger-scale executions. Cloud computing can be an interesting option if the computation you may perform is executed periodically, as the non-use of local resources implies that the economic investment is not justified. On the other hand, the option of using a volunteer computing platform looks interesting for developing some kind of HPC applications, as it is offering huge amount of hardware resource at no cost.