Job Traffic-Aware Scheduling in Large-Scale Environments using GMMs

José Afonso Sanches[1], Lúcia Drummond[1], Claude Tadonki[2]

[1]Universidade Federal Fluminense, Instituto de Computação, Niterói, Brazil
{jalsanches, lucia}@ic.uff.br

[2]MINES-ParisTech, Centre de Recherche en Informatique, Fontainebleau, France
claude.tadonki@mines-paristech.fr

A *computing grid* can be defined as a system that coordinates distributed resources by using standardized, open and general-purpose protocols, and which offers non-trivial qualities of service. Due to the fact that this kind of system is multi-institutional and dynamic by nature [Foster et al. 2001], its characteristic issue is the *coordinated* resource sharing. According to [Iamnitchi and Foster 2004], the allocation and discovery of resources in large-scale computing environments - just as grids and clouds - remains an essential service. Grids' *resource* layer, as described by Foster, establishes protocols for operations on resources, such as their discovery and management.

[Foster et al. 2008] mention that the problems related to resources remain mostly the same in both grids and clouds. There exist some common needs here, which are: managing large structures and big data; defining methods by which the consumer may discover, request and use resources provided by central mechanisms; and finally implementing parallel computations on those resources. Whereas there exist differences in the details, it can be said that both communities share several common issues.

Rather than a single problem, scheduling in grid systems is a family of problems. The problem of finding optimal job scheduling in heterogeneous systems is mostly NP-hard. The dynamics of job submissions in a grid system may be described as follows: 1) an application can generate several jobs, which in turn can be composed of subtasks; 2) the grid system is responsible for sending each subtask to a resource, in order to be processed.

Although the simpler scenario in grid scheduling would be a user selecting the most adequate machines in the system and then submitting jobs to them, this is far from the standard case. Nowadays, a typical grid system must have meta-schedulers at disposal, in order to find and distribute jobs automatically and efficiently. There are several characteristics that make scheduling in grids more chalenging than in conventional distributed systems [Xhafa and Abraham 2010], among which we can mention: their dynamic structure, the high heterogeneity of resources, the job-resource requirements, and so on.

Makespan is one of the most popular and studied optimization criteria for job scheduling. It indicates the overall completion time for a group of tasks, and it is an indicator of the general productivity of a grid system. Another very popular one is *flowtime*, which is the sum of finalization times for all the tasks in a group. Minimizing the flowtime implies reducing the average response time of the Grid system. The minimization of both *makespan* and *flowtime* can be defined, respectively, as:

$$min_{S_i \in Sched}\{max_{i \in Tasks} F_i\} \quad \text{and} \quad min_{S_i \in Sched}\{\sum_{i \in Tasks} F_i\} \qquad \text{(1a, b)}$$

The goal of this PhD thesis is to reduce both *makespan* and *flowtime* for batches of jobs, by choosing a more adequate schedule with the aid of GMM predictions in what concerns some job traffic characteristics at a given site, like its workload, its network or its turnaround time. Our work is related to the extraction of refined job traffic information from sites on a large-scale environment, as an input for a job scheduling algorithm. A Gaussian Mixture Model, or GMM, is defined as "a parametric probability density function represented as a weighted sum of Gaussian component densities" [Reynolds 2009].

One of the greatest attributes of the GMM model is the ability to form smooth approximations of arbitrarily shaped densities. In fact, the classical uni-modal Gaussian model represents feature distributions by a position (mean vector) and an elliptic shape (covariance matrix) and then either a vector quantizer (VQ) or the "nearest neighbor" model will represent a distribution by a discrete set of characteristic templates. Hence, a GMM would act as a hybrid between these two models by using a discrete set of Gaussian functions, each one with their own *mean* and *covariance* matrices, so it will allow a better modeling capability.

The usage of GMMs is receiving increasing attention - the results in works such as [Ma et al. 2010] and [Prado et al. 2012] show that adopting GMMs may be good and flexible in modelling criteria like the shape of the network traffic for a grid site. GMMs are computationally affordable, since it contains Gaussian processes with a lower computational overhead and a simple underlying mathematics form. The estimation of GMM parameters can be made with the Expectation Maximization (EM) algorithm.

In order to better reproduce the effects of the network traffic in grid environments, we are in search for real-life systems, as many as possible. Currently, we are performing experiments in Grid5000, the national academic grid environment in France, in order to have a formal basis on the reason why the job traffic problems on grids/clouds can hamper jobs' allocation to resources, and therefore the overall efficiency.

References

Foster, I., Kesselman, C., and Tuecke, S. (2001). The anatomy of the grid - enabling scalable virtual organizations. *Int. Journal of Supercomputer Applications 15*.

Foster, I., Zhao, Y., Raicu, I., and Lu, S. (2008). Cloud computing and grid computing 360-degree compared. *Grid Computing Environments Workshop, CGE '08*.

Iamnitchi, A. and Foster, I. (2004). A peer-to-peer approach to resource location in grid environments. *Advances in Computer Science*.

Ma, T., El-khatib, Y., Mackay, M., and Edwards, C. (2010). Characterising a grid site's traffic. *ACM HPDC '10 Proceedings*.

Prado, R., Braun, J., Krettek, J., Hoffmann, F., García-Galán, S., Expósito, J. M., and Bertram, T. (2012). Gaussian mixture models vs. fuzzy rule-based systems for adaptive meta-scheduling in grid/cloud computing. *Management Intelligent Systems, AISC 171*.

Reynolds, D. (2009). Gaussian mixture models. In *Encyclopedia of Biometrics*. Springer.

Xhafa, F. and Abraham, A. (2010). Computational models and heuristic methods for grid scheduling problems. *Future Generation Computing Systems 26*.