
Minimizing rental cost for multiple recipe applications in the Cloud

Veronika Sonigo*¹

¹Département d'Informatique des Systèmes Complexes (FEMTO-ST/DISC) – Université de Technologie de Belfort-Montbéliard, Ecole Nationale Supérieure de Mécanique et des Microtechniques, Université de Franche-Comté, CNRS : UMR6174 – 24, rue Alain Savary 25000 - BESANCON, France

Abstract

Clouds are more and more becoming a credible alternative to parallel dedicated resources. The pay-per-use pricing policy however highlights the real cost of computing applications. This new criterion, the cost, must then be assessed when scheduling an application in addition to more traditional ones as the completion time or the execution flow. In this paper, we tackle the problem of optimizing the cost of renting computing instances to execute an application on the cloud while maintaining a desired performance (throughput). The target application is a stream application based on a DAG pattern, i.e., composed of several tasks with dependencies, and instances of the same execution task graph are continuously executed on the instances. We provide some theoretical results on the problem of optimizing the renting cost for a given throughput then propose some heuristics to solve the more complex parts of the problem, and we compare them to optimal solutions found by linear programming.

*Speaker