université PARIS DIDEROT PARIS 7

PREUVES PROGRAMMES SYSTÈMES

UMR 7126

Report on the PhD thesis of Krzystof Rządca

This thesis focuses on the efficient usage of large scale grids. This is a hot topic, with clear practical applications, and has been the subject of extensive study for over a decade, with a variety of proposed solutions.

What makes the issue quite complex is that there is not a unique, clear definition of what a grid really is, and what *efficiency* really means: the more one departs from the traditional case of a centrally controlled computing facility owned by a single organization, the more options for defining *efficiency* arise.

In this thesis, Krzystof Rządca proposes a new model of the grid, and then picks three concrete scenarii where this new model eases the performance analysis of the system, using notions and techniques from game theory and multi-criteria equitable optimization.

The thesis is accordingly organized in two parts: the first part introduces the multi-organizational grid model, and the second part contains the performance results (and scheduling algorithms) obtained using this model for the specific cases of grids performing specialised resource sharing, divisible load balancing and rigid load balancing.

Chapter 1 provides an introduction and recalls basic notions from scheduling, game theory and multi-criteria optimization.

Chapter 2 introduces the multi-organizational grid model; instead of simply separating resources and users, an approach that naturally leads to explore efficiency notions based on economical models of offer and demand, Krzystof Rządca proposes to consider explicitly the notion of *organization*, as the key focus for observing efficiency of the system. An organization is an entity that gives resources to the grid, and groups users willing to access the grid's global resources.
A grid is composed by a given set of organizations, which are independent and do not engage in any form of compensation: their only reason to join the grid is to increase

Adresse postale : Laboratoire PPS - Université Paris Diderot - Case 7014 - F-75205 Paris Cedex 13
Adresse physique : 175, rue du Chevaleret - F-75013 Paris
Tél : 33-1 44.27.44.27 - Fax : 33-1 44.27.86.54
Mél : Roberto.Di-Cosmo@pps.jussieu.fr Web : http://www.pps.jussieu.fr/

the performance of each organization users' jobs.
In this model, one considers three variations, with a mandatory centralised scheduler (using equitable multi-criteria optimization), a purely advisory centralized scheduler (using game-theory), or an intermediate solution allowing organization to quit a grid if they want, but forcing them to obey the central scheduler if they stay in (using constrained multi-criteria optimization).

Chapter 3 studies the case of grids built out of specialized resources, where each job needs to be executed on a specific resource, for each of the three variations. The main results here are the following:

- finding an equitably optimal centralised scheduling is NP-complete, and the set of all such schedules is exponential even with just two organizations
- a distributed schedule is analogous to the Prisoner's Dilemma, and ends up in a very inefficient resource allocation

Several algorithms are presented, exact and approximate, that are validated via a set of simulations.

Chapter 4 studies the case of general purpose grids executing divisible load tasks, making a clear set of assumptions, one of which is the fact that the finish time of a submitted job is guaranteed, and announced at submission time. This allows to formally obtain several interesting results :

- averaging load balancing is the socially optimal strategy, but
- if organization are able to decide pointwise whether to cooperate or not, no balancing will ever occur

When cooperation is enforced on a longer period, if the loads are similar, and the decisions are made without taking into account the expected queue lengths, then cooperation becomes an individually optimal strategy too.
An $O(n^2)$ modified load balancing algorithm is hence proposed, with a bias towards less loaded organizations: it allows to relax the condition on the similarity of loads, and is validated experimentally.

Chapter 5 studies the case of general purpose grids executing non divisible load tasks. Here again, the individual optimal strategies are non cooperative, and an algorithm is proposed that allow to preserve selfish goals, while 4-approximating the global optimum. The algorithm is validated experimentally, and evidence is shown that the

Adresse postale : Laboratoire PPS - Université Paris Diderot - Case 7014 - F-75205 Paris Cedex 13
Adresse physique : 175, rue du Chevaleret - F-75013 Paris
Tél : 33-1 44.27.44.27 - Fax : 33-1 44.27.86.54
Mél : Roberto.Di-Cosmo@pps.jussieu.fr Web : http://www.pps.jussieu.fr/

approximation constant can be significantly lower than 4 in practice.

Conclusion

This work provides a substantial contribution to a better understanding of some of the fundamental phaenomena that arise in large scale grids, via an original approach that more closely models the organisational structure of existing grids.

While the proposed solutions are validated only through simulations, I find the approach original and the analysis very compelling.

This work definitely qualify as a PhD thesis, and I strongly support its defense.

Paris, February 7th, 2008

Roberto Di Cosmo
Computer Science Professor
University Paris Diderot

Adresse postale : Laboratoire PPS - Université Paris Diderot - Case 7014 - F-75205 Paris Cedex 13
Adresse physique : 175, rue du Chevaleret - F-75013 Paris
Tél : 33-1 44.27.44.27 - Fax : 33-1 44.27.86.54
Mél : Roberto.Di-Cosmo@pps.jussieu.fr Web : http://www.pps.jussieu.fr/