The impact of data replication on job scheduling performance in the Data Grid

Ming Tang *, Bu-Sung Lee, Xueyan Tang, Chai-Kiat Yeo

School of Computer Engineering, Nanyang Technological University, Blk N4, #2a-32, Nanyang Avenue, Singapore 639798, Singapore

Abstract

In the Data Grid environment, the primary goal of data replication is to shorten the data access time experienced by the job and consequently reduce the job turnaround time. After introducing a Data Grid architecture that supports efficient data access for the Grid job, the dynamic data replication algorithms are put forward. Combined with different Grid scheduling heuristics, the performances of the data replication algorithms are evaluated with various simulations. The simulation results demonstrate that the dynamic replication algorithms can reduce the job turnaround time remarkably. In particular, the combination of shortest turnaround time scheduling heuristic (STT) and centralized dynamic replication with response-time oriented replica placement (CDR_RTPlace) exhibits remarkable performance in diverse system environments and job workloads.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Data replication; Grid scheduling; Data Grid; Simulation

1. Introduction

Nowadays, scientific research and commercial application generate large amount of data that are required by users around the world. A good example for this is the High Energy Physics (HEP) area, where a new particle accelerator, the Large Hadron Collider (LHC), will start to work at European Organization for Nuclear Research (CERN) in the year 2007, and several HEP experiments will produce Petabytes of data per year for decades [11]. The data captured from the experiments will be used by thousands of physicists, and the data are to be distributed to centers around the world for processing. As a specialization and extension of the computational Grid [8], the Data Grid is a solution for the above problem [4]. Essentially, the Data Grid is an infrastructure that manages large scale data files and provides computational resources across widely distributed communities.

The Grid resources, including computing facility, data storage and network bandwidth, are consumed by the jobs. For each incoming job, the Grid scheduler decides where to run the job based on the job

---

* Corresponding author. Tel.: +65 67904623; fax: +65 67926559.
E-mail addresses: mtang@pmail.ntu.edu.sg (M. Tang),
ebslee@ntu.edu.sg (B.-S. Lee), asxytang@ntu.edu.sg (X. Tang),
asckyeo@ntu.edu.sg (C.-K. Yeo).

requirements and the system status. In data-intensive applications, the locations of data required by the job impact the Grid scheduling decision and performance greatly. Creating data replicas can reroute the data requests to certain replica servers and offer remarkably higher access speed than a single server. At the same time, the replicas provide broader decision space for the Grid scheduler to achieve better performance from the perspective of the job.

Data replication is a practical and effective method to achieve efficient network performance in network bandwidth constrained environment, and it has been applied widely in the areas of distributed database and Internet [26,19]. New challenges are faced in the Data Grid, for example, huge data file sizes, system resources belonging to multiple owners, dynamically changing resources and complicated cost model.

In this paper, we propose a Data Grid architecture supporting efficient data replication and job scheduling. The computing sites are organized into individual domains according to the network connection, and a replica server is placed in each domain. Two centralized dynamic replication algorithms with different replica placement methods and a distributed dynamic replication algorithm are put forward. The Grid scheduling heuristics of shortest turnaround time, Least Relative Load and Data Present are proposed. In order to evaluate the performance of the scheduling heuristics combined with different replication algorithms, a Data Grid simulator called XDRepSim is developed. Various simulations are carried out with different system configurations and job workloads.

This paper is organized as follows. Section 2 presents the related works. Section 3 introduces the system model and performance metric for the Data Grid. Section 4 introduces our dynamic data replication algorithms. The Grid scheduling heuristics are put forward in Section 5. The simulation methods and results are described in Sections 6 and 7, respectively. Section 8 concludes the paper.

2. Related work

Several recent studies have taken into account both job scheduling and data replication in the Data Grid. In previous work [21], the External Scheduler is modelled to assign the job to specific computing site, and

the Data Scheduler running at each site is responsible for dynamically creating replicas for popular data files. Various combinations of scheduling and replication strategies are evaluated with simulations. Their results show that data locality is an important factor when scheduling the jobs. The simple scheduling policy of always assigning jobs to the sites that contain the required data works very well if the popular datasets are replicated dynamically. Takefusa et al. [24] also reported similar conclusions using the Grid Datafarm architecture and the Bricks Grid simulator.

Another set of closely related work is [2], which uses the Data Grid simulator OptorSim for studying scheduling and replication strategies. The simulated Grid architecture is similar to [21] in that a global Resource Broker schedules the jobs to the computing site and a Replica Optimiser in each site performs local replica optimization. The replication operation is independently determined by each site. For every data access required by the locally running job, the Replica Optimiser will determine whether the data should be replicated to local storage and which replicas should be removed if there is not enough space. The dynamic replication strategies used in the study evolve from traditional cache replacement methods. The economic replication strategies are put forward and they attempt to improve the profits brought by the replicas and decrease the cost of data management at the same time. The simulation results show that the scheduling strategy considering both the file access cost of the jobs and the workload of computing resources produces the shortest mean job execution time, and the economic replication strategy can improve the Grid performance tremendously.

Works of [20] and [13] study data replication strategies in the Grid. Ranganathan and Foster [20] put forward several straightforward dynamic replication strategies, including the Fast Spread and Cascading methods, for a hierarchical Data Grid. In [13], the Data Grid structure is a hybrid of tree and ring topologies, and data access among same tier nodes is allowed. A cost model is proposed to decide the replication. The model evaluates the data access gains by creating a replica and the costs of creation and maintenance for the replica. It is used by the Replica Manager in each intermediate storage site in a decentralized manner.

Job scheduling heuristics for the distributed computing systems have been extensively studied. Mah-