BIGHybrid: a toolkit for MapReduce in hybrid environments

Vinicius P. Perego, Kassiano J. Matteussi, Julio C.S. Anjos, Claudio Geyer
Institute of Informatics, UFRGS
Bento Gonąlves Avenue, 9500, Porto Alegre, RS, Brazil
vpperego@inf.ufrgs.br

Abstract

The high availability of Cloud Computing attracts users to deploy and run Big Data applications. Unfortunately, its usage implies in a huge price for the resources allocated to the job. Contrarily, heterogeneous environments like Desktop Grids also facilitates the processing of high datasets due to its large number of free resources. In order to reduce the costs with Cloud Computing, this work proposes a solution based on a hybrid architecture, which distributes data in the Cloud and Volunteer Computing. Alongside the hybrid model, a toolkit called BIGHybrid was developed to evaluate solutions in a hybrid environment.

1. Introduction

The increase in data consumption by modern applications requires that distributed systems have resources available to process it. In this scenario, Cloud Computing has become the most sought solution due to its availability. However, cloud service providers charge the users by resources used, which difficults the use of this type of infrastructure for processing high amounts of data. Alternately, Desktop Grid (DG) can be a good solution to execute these applications with additional power donated by idle computers (volunteer computing). The problem with this source of processing power is the volatility that needs an improved fault tolerance mechanism. To increase the resource pool, one could use a hybrid solution based on Cloud and Desktop Grid.

Hadoop [9] is a framework that implements MapReduce [5] , a programming model which eases the creation of distributed programs by abstracting the parallelization and achieving some degree of fault tolerance with its built-in algorithms.

This work proposes an update to the original Hadoop aiming to use it as a hybrid infrastructure based on Cloud Computing and DG. The changes in the Hadoop algorithms were made in a toolkit written over SimGrid.

SimGrid [4] is a toolkit for the simulation of distributed systems, enabling the user to define the execution platform, scheduling algorithms and take measures like the execution time.

2. Related Works

GroudSim [7] is a simulator toolkit for Grid and Cloud focusing on scientific applications. It is based on scalable discrete-event and supports traces from both application, host and event traces. Unlike BIGhybrid, GroundSim does not support hybrid or volatile simulations.

CloudSim[3], an extension of GridSim, focuses on Cloud simulation. It allows cloud service modeling such as service brokers, provisioning and allocation policies at a single physical node. However, CloudSim can not generate traces and simulate data-intensive applications, which are main features of BIGHybrid.

The AweSim simulator [8] is based on a network simulation framework that involves a fine-grained simulation for workflow computation and data movement across multiple Clouds. The implementation uses workload traces from a production data analysis service and is thus similar to the BIGhybrid simulator that adopts its behavior from traces in a volatile real-world environment.

3. Objectives

This work proposes a toolkit to simulate MapReduce in hybrid environments, adapting the original Hadoop MapReduce with enhances on its algorithms and additional inclusions to work with volunteer computing. BIGHybrid enables the evaluation of the impact of the changes on the job configuration or resources (e.g., network bandwidth, latency and processors flop power).

4. Model

BigHybrid [1] is organized in a layered architecture pattern as shown in Figure 1 . The System Config module allows the user to specify the workload (e.g., chunks size and data distribution among cloud and DG) and the platform for Simgrid of both cloud and DG environments.

In the lower layer, the middlewares BitDew-MapReduce and Cloud-BlobSeer control the DG and Hadoop resources, respectively. BlobSeer is a DFS that manages a huge amount of data in a flat sequence of bytes called BLOBs (Binary Large Objects). BlobSeer extends the Hadoop file system (HDFS) by enabling concurrent access to metadata and high throughput with concurrent reading, writing, and updating of the data. This incremental update is necessary for data management in a hybrid infrastructure.

Bit-Dew exploits protocols like P2P, HTTP, BitTorrent and FTP to data management in Grids. The BitDew-MapReduce, a MapReduce implementation adapted to a volatile environment, contains important features like fault-tolerance mechanism, data placement and incremental update which are important for hybrid infrastructure.

5. Prototype

BIGhybrid was implemented over the Simgrid toolkit. It relies on MRA++ [2] for BitDew-MapReduce and MRSG [6] for Cloud-BlobSeer. The original version of MRA++ was extended for the purpose of simulating Volunteer Computing, where the volatility module was added. All the simulators were developed with Simgrid through the MSG application programming interface.

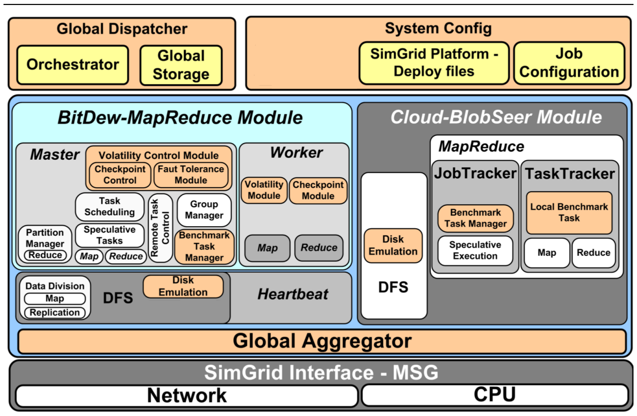

Figure 1. BIGHybrid architecture

6. Results

The hardware configuration is based on the computational capacity of processors equivalent to an Intel Xeon - 2 Cores, 4M Cache, 2.13 GHz 5E+09 Flops and the heterogeneous environment with its capacity distributed between 4E+09 to 6E+09 Flop. The networking bandwidth for Cloud is 1 Gbps and to DG is 100 Mbps.

An experiment was made in a low-scale hybrid environment with the results in Figure 2. The line red indicates time execution of 200 tasks (chunk size 64 MB) on 128 machines in a single Cloud deploying, equivalent to 503 seconds. Each execution distributes the machines and the workload in different scenarios. Executions A, B and C show a gain for the hybrid execution over the original cloud deployment.

7. Conclusion

In this study, the characteristics of a hybrid infrastructure were investigated, addressing problems to adapt the MapReduce to execute in these scenarios. It was possible to achieve an acceptable execution and even gain some performance with the low-scale scenario tested. Some features like I/O contention need to be added to increase the accuracy of the Cloud simulation. Although it is possible to reduce the overall resource allocation in the Cloud and consequently the cost, it is still necessary to measure the gains with the hybrid solution.

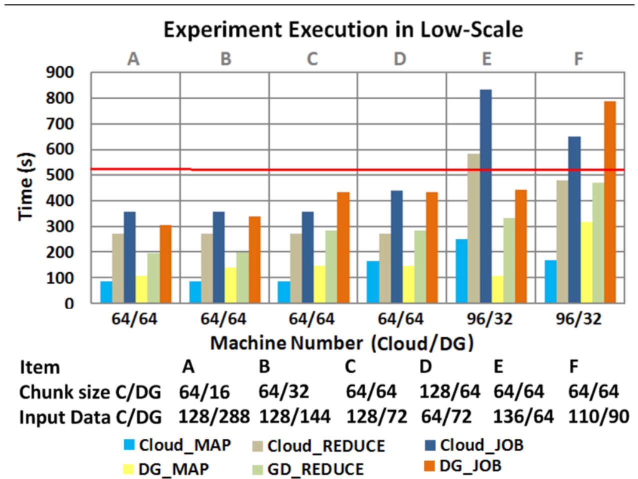

Figure 2. Low-Scale hybrid environment

References

[1] J. Anjos, G. Fedak, and C. F. Geyer. Bighybrid: a simulator for mapreduce applications in hybrid distributed infrastructures validated with the grid5000 experimental platform. *Concurrency and Computation: Practice and Experience*, 28(8):2416–2439, 2016.

[2] J. C. Anjos, I. Carrera, W. Kolberg, A. L. Tibola, L. B. Arantes, and C. R. Geyer. Mra++: Scheduling and data placement on mapreduce for heterogeneous environments. *Future Generation Computer Systems*, 42:22–35, 2015.

[3] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. De Rose, and R. Buyya. Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and experience*, 41(1):23–50, 2011.

[4] H. Casanova. Simgrid: A toolkit for the simulation of application scheduling. In *Cluster computing and the grid, 2001. proceedings. first ieee/acm international symposium on*, pages 430–437. IEEE, 2001.

[5] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[6] W. Kolberg, P. D. B. Marcos, J. C. Anjos, A. K. Miyazaki, C. R. Geyer, and L. B. Arantes. Mrsg–a mapreduce simulator over simgrid. *Parallel Computing*, 39(4):233–244, 2013.

[7] S. Ostermann, K. Plankensteiner, R. Prodan, and T. Fahringer. Groudsim: an event-based simulation framework for computational grids and clouds. In *European Conference on Parallel Processing*, pages 305–313. Springer, 2010.

[8] W. Tang, J. Jenkins, F. Meyer, R. Ross, R. Kettimuthu, L. Winkler, X. Yang, T. Lehman, and N. Desai. Data-aware resource scheduling for multicloud workflows: A fine-grained simulation approach. In *Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on*, pages 887–892. IEEE, 2014.

[9] T. White. *Hadoop: The definitive guide*. " O'Reilly Media, Inc.", 2012.