World Scientific News

An International Scientific Journal

Exploiting Dynamic Resource Allocation

G. S. Geethamani[1,*], M. Mayilvaganan[2]

[1]Department of Computer Science, Karpagam University, Coimbatore, Tamil Nadu, India

[1]Department of Information Technology, Hindusthan College of Arts and Sceience, Coimbatore, India

[2]Department of Computer Science, PSG College of Arts and Science,
Coimbatore - 641 028, Tamil Nadu, India

*E-mail address: gsgeethamani @gmail.com

ABSTRACT

In recent years ad-hoc parallel data processing has emerged to be one of the killer applications for Infrastructure-as-a-Service (IaaS) clouds. Major Cloud computing companies have started to integrate frameworks for parallel data processing in their product portfolio, making it easy for customers to access these services and to deploy their programs. However, the processing frameworks which are currently used have been designed for static, homogeneous cluster setups and disregard the particular nature of a cloud. Consequently, the allocated compute resources may be inadequate for big parts of the submitted job and unnecessarily increase processing time and cost. We discuss the opportunities and challenges for efficient parallel data processing in clouds and present our research project Nephele. Nephele is the first data processing framework to explicitly exploit the dynamic resource allocation offered by today's IaaS clouds for both, task scheduling and execution. Particular tasks of a processing job can be assigned to different types of virtual machines which are automatically instantiated and terminated during the job execution. Based on this new framework, we perform extended evaluations of Map Reduce-inspired processing jobs on an IaaS cloud system and compare the results to the popular data processing framework Hadoop.

Keywords: Single Processing, Multiprocessing, Cloud Computing, Neples Algorithm

1. INTRODUCTION

Today a growing number of companies have to process huge amounts of data in a cost-efficient manner. More companies have popularized an architectural paradigm based on a large number of commodity servers, problems like processing crawled documents or regenerating a web index are split into several independent subtasks, distributed among the available nodes, and computed in parallel.

In order to simplify the development of distributed applications on top of such architectures, many of these companies have also built customized data processing frameworks. Examples are Google's MapReduce, Microsoft's Dryad, or Yahoo!'s Map-Reduce-Merge. They can be classified by terms like high throughput computing (HTC) or many-task computing (MTC), depending on the amount of data and the number of tasks involved in the computation. Although these systems differ in design, their programming models share similar objectives, namely hiding the hassle of parallel programming, fault tolerance, and execution optimizations from the developer. Developers can typically continue to write sequential programs. The processing framework then takes care of distributing the program among the available nodes and executes each instance of the program on the appropriate fragment of data

2. OPTIMIZING MULTIWAY JOINS IN A MAP-REDUCE ENVIRONMENT

Implementations of map-reduce are being used to perform many operations on very large data. We examine strategies for joining several relations in the map-reduce environment. Our new approach begins by identifying the "map-key," the set of attributes that identify the Reduce process to which a Map process must send a particular tuple. Each attribute of the map-key gets a "share", which is the number of buckets into which its values are hashed, to form a component of the identifier of a Reduce process. Relations have their tuples replicated in limited fashion, the degree of replication depending on the shares for those map-key attributes that are missing from their schema. We study the problem of optimizing the shares, given a fixed number of Reduce processes. An algorithm for detecting and fixing problems where a variable is mistakenly included in the map-key is given. Then, we consider two important special cases: chain joins and star joins. In each case, we are able to determine the map-key and determine the shares that yield the least replication. While the method we propose is not always superior to the conventional way of using map-reduce to implement joins, there are some important cases involving large-scale data where our method wins, including: 1) analytic queries in which a very large fact table is joined with smaller dimension tables, and 2) queries involving paths through graphs with high out-degree, such as the Web or a social network.

3. MAP-REDUCE-MERGE: SIMPLIFIED RELATIONAL DATA PROCESSING ON LARGE CLUSTERS

We design and implement Mars, a Map Reduce framework, on graphics processors (GPUs). Map Reduce is a distributed programming framework originally proposed by Google for the ease of development of web search applications on a large number of commodity CPUs. Compared with CPUs, GPUs have an order of magnitude higher computation power and

memory bandwidth, but are harder to program since their architectures are designed as a special-purpose co-processor and their programming interfaces are typically for graphics applications. As the first attempt to harness GPU's power for Map Reduce, we developed Mars on an NVIDIA G80 GPU, which contains over one hundred processors, and evaluated it in comparison with Phoenix, the state-of-the-art Map Reduce framework on multi-core CPUs. Mars hides the programming complexity of the GPU behind the simple and familiar

4. IMPLEMENTATION

Cloud computing has emerged as a promising approach to rent a large IT infrastructure on a short-term pay-per-usage basis. Operators of so-called Infrastructure-as-a-Service (IaaS) clouds, like Amazon EC2, let their customers allocate, access, and control a set of virtual machines (VMs) which run inside their data centers and only charge them for the period of time the machines are allocated. In the proposed system Nephele, a new processing framework explicitly designed for cloud environments. Nephele is the first data processing framework to include the possibility of dynamically allocating or deallocating different compute resources from a cloud in its scheduling and during job execution. The actual execution of tasks which a Nephele job consists of is carried out by a set of instances. Each instance runs a so-called Task Manager (TM). A Task Manager receives one or more tasks from the Job Manager at a time, executes them, and after that informs the Job Manager about their completion or possible errors.

5. NEPHELE ALGORITHM

A data processing framework with support for dynamic allocation and de-allocation of different computational resources in the cloud. Compute resources available in a cloud environment are highly dynamic and possibly heterogeneous. In addition, the network topology is hidden so scheduling optimizations based on knowledge of the distance to a particular rack or server are impossible. The classic genotyping approach has been based on phylogenetic analysis, starting with a multiple sequence alignment. Genotypes are then established by expert examination of phylogenetic trees. However, such methods are suboptimal for a rapidly growing dataset, because they require significant human effort, and because they increase in computational complexity quickly with the number of sequences. This project uses a method for genotyping that does not depend on multiple sequence alignment. It uses the complete composition vector algorithm to represent each sequence in the dataset as a vector derived from its constituent k-mers, and affinity propagation clustering to group the sequences into genotypes based on a distance measure over the vectors. Our methods produce results that correlate well with expert-defined clades or genotypes, at a fraction of the computational cost of traditional phylogenetic methods. Increasingly, genotyping is coming into use for disease surveillance activities, as well as for microbial forensics. The classic genotyping approach has been based on phylogenetic analysis, starting with a multiple sequence alignment. Genotypes are then established by expert examination of phylogenetic trees. However, these traditional single-processor methods are suboptimal for rapidly growing sequence datasets being generated by next-generation DNA sequencing machines, because they increase in computational complexity quickly with the number of sequences.
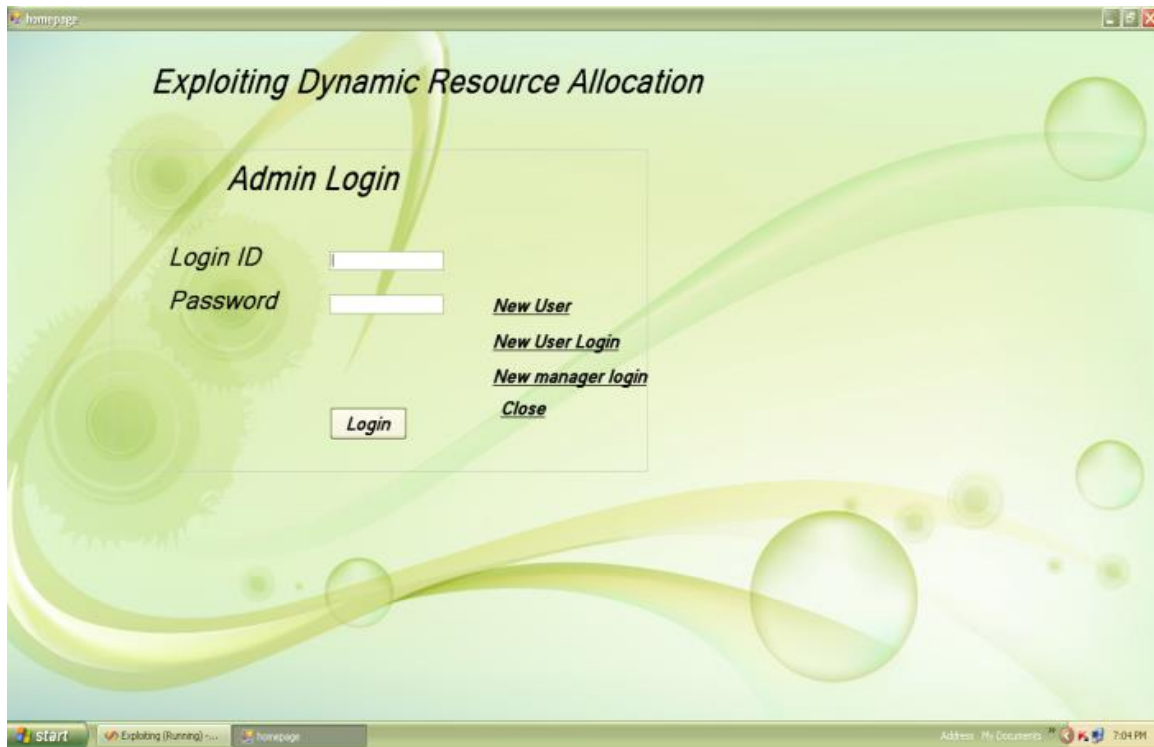
7. MODULES DEPLOYMENT

7. 1. Server Type Support
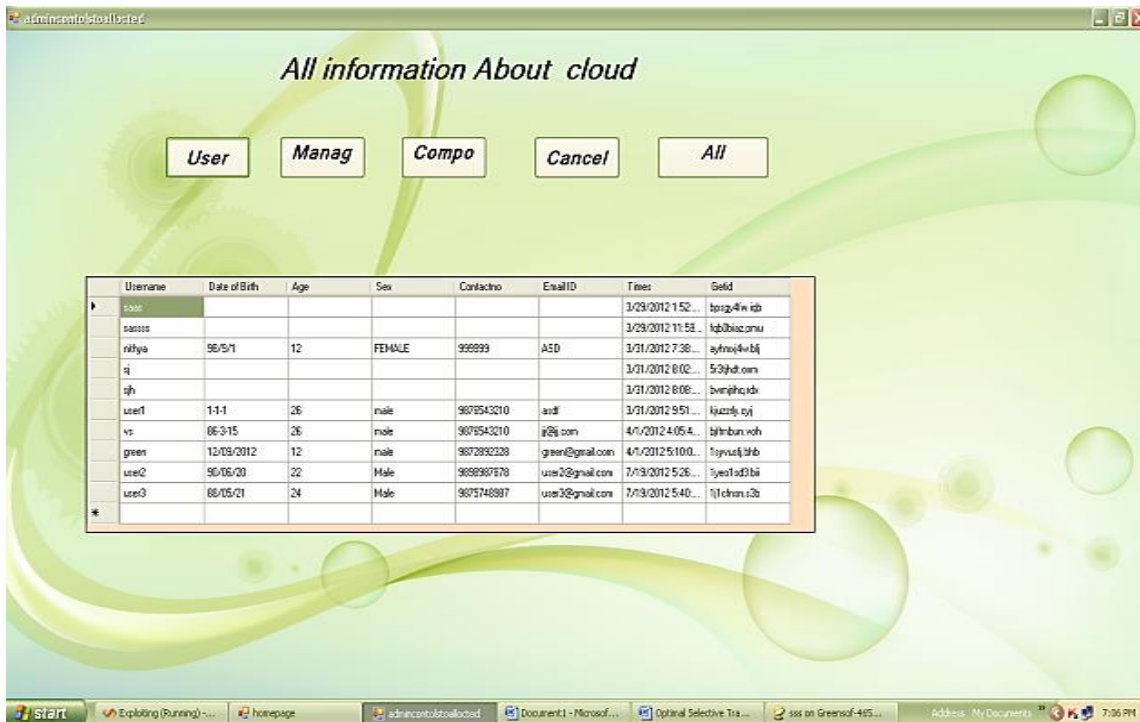
The scheduler is given a list of available server types and their cost per time unit. Each task can be executed on its own server type. An Execution Instance has an ID and a server type representing the hardware characteristics.

7. 2. Traffic Identifier

All edges of an Execution Graph are replaced by a channel before processing can begin. There are three channel types: A network channel is based on a TCP socket connection. Two subtasks connected via a network channel can be executed on different instances. Since they must be executed at the same time, they are required to run in the same Execution Stage.

7. 3. Storage Management System

After having received a valid Job Graph from the user, Nephele's Job Manager transforms it into a so-called Execution Graph. An Execution Graph is Nephele's primary PARALLEL AND data structure for scheduling and monitoring the execution of a Nephele job.

Each vertex of the Job Graph is transformed into one Execution Vertex. The default channel types are network channels. Each Execution Vertex is by default assigned to its own Execution Instance unless the user's annotations or other scheduling restrictions (e.g. the usage of in-memory channels) prohibit it. One fundamental idea to refine the scheduling strategy for recurring jobs is to use feedback data. With the collected data Nephele is able to detect both computational as well as I/O bottlenecks. We only use the profiling data to detect these bottlenecks.

8. CONCLUSIONS

The challenges and opportunities for efficient parallel data processing in cloud environments are discussed and presented Nephele, the first data processing framework to exploit the dynamic resource provisioning offered by today's IaaS clouds. Nephele's basic architecture and presented a performance comparison to the well-established data processing framework Hadoop.

The performance evaluation gives a first impression on how the ability to assign specific virtual machine types to specific tasks of a processing job, as well as the possibility to automatically Allocate/deallocate virtual machines in the course of a job execution, can help to improve the overall resource utilization and, consequently, reduce the processing cost. With a framework like Nephele at hand, there are a variety of open research issues, which we plan to address for future work. In particular, we are interested in improving Nephele's ability to adapt to resource overload or underutilization during the job execution automatically.

Our current profiling approach builds a valuable basis for this, however, at the moment the system still requires a reasonable amount of user annotations.

References

[1] Hung-chih Yang, Ali Dasdan, Ruey-Lung Hsiao, and D. Stott Parker. 2007. Map-reduce-merge: simplified relational data processing on large clusters. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data* (*SIGMOD '07*). Association for Computing Machinery, New York, NY, USA, 1029–1040. https://doi.org/10.1145/1247480.1247602

[2] L. D'Amore, G. Laccetti, D. Romano, G. Scotti & A. Murli (2015) Towards a parallel component in a GPU–CUDA environment: a case study with the L-BFGS Harwell routine. *International Journal of Computer Mathematics*, 92:1, 59-76, DOI: 10.1080/00207160.2014.899589

[3] L. D'Amore, D. Casaburi, A. Galletti, L. Marcellino, and A. Murli, *Integration of emerging computer technologies for an efficient image sequences analysis*, Integr. Comput.-Aided Eng. 18(4) (2011), pp. 365–378

[4] J.E. Dennis and J.J. More. Quasi-Newton methods, motivation and theory. *SIAM Rev.* 19 (1977), pp. 46–89. doi: 10.1137/1019005

[5] F. Gregoretti, G. Laccetti, A. Murli, G. Oliva, and U. Scafuri, *MGF: A grid-enabled MPI library*, Future Gener. Comput. Syst. 24(2) (2008), pp. 158–165. doi: 10.1016/j.future.2007.03.009

[6] D.C. Liu and J. Nocedal, *On the limited memory BFGS method for large scale optimization*, Math. Program. 45 (1989), pp. 503–528. doi: 10.1007/BF01589116

[7] L. Maddalena, A. Petrosino, and G. Laccetti, *A fusion-based approach to digital movie restoration*, Pattern Recognit. 42(7) (2009), pp. 1485–1495. doi: 10.1016/j.patcog.2008.10.026

[8] C. Zhu, R.H. Byrd, P. Lu, and J. Nocedal, Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization, *ACM Trans. Math. Softw*. 23 (1997), pp. 550–560. doi: 10.1145/279232.279236

[9] Chaiken, R., Bob Jenkins, P. Larson, Bill Ramsey, Darren Shakib, S. Weaver and Jingren Zhou. SCOPE: easy and efficient parallel processing of massive data sets. *Proc. VLDB Endow*. 1 (2008): 1265-1276.

[10] Dominic Battré, Stephan Ewen, Fabian Hueske, Odej Kao, Volker Markl, and Daniel Warneke. 2010. Nephele/PACTs: a programming model and execution framework for web-scale analytical processing. In *Proceedings of the 1st ACM symposium on Cloud Computing* (*SoCC '10*). Association for Computing Machinery, New York, NY, USA, 119–130. https://doi.org/10.1145/1807128.1807148