

Chapter 2

Towards Data Intensive Many-Task Computing

Ioan Raicu

*Illinois Institute of Technology, USA & Argonne
National Laboratory, USA*

Philip Little

University of Notre Dame, USA

Ian Foster

*University of Chicago, USA & Argonne National
Laboratory, USA*

Christopher M. Moretti

University of Notre Dame, USA

Yong Zhao

*University of Electronic Science and Technology
of China, China*

Amitabh Chaudhary

University of Notre Dame, USA

Alex Szalay

Johns Hopkins University, USA

Douglas Thain


University of Notre Dame, USA

ABSTRACT

Many-task computing aims to bridge the gap between two computing paradigms, high throughput computing and high performance computing. Traditional techniques to support many-task computing commonly found in scientific computing (i.e. the reliance on parallel file systems with static configurations) do not scale to today's largest systems for data intensive application, as the rate of increase in the number of processors per system is outgrowing the rate of performance increase of parallel file systems. In this chapter, the authors argue that in such circumstances, data locality is critical to the successful and efficient use of large distributed systems for data-intensive applications. They propose a "data diffusion" approach to enable data-intensive many-task computing. They define an abstract model for data diffusion, define and implement scheduling policies with heuristics that optimize real world performance, and develop a competitive online caching eviction policy. They also offer many empirical experiments to explore the benefits of data diffusion, both under static and dynamic resource provisioning, demonstrating approaches that improve both performance and scalability.

DOI: 10.4018/978-1-61520-971-2.ch002

Figure 1. Problem types with respect to data size and number of tasks



INTRODUCTION

We want to enable the use of large-scale distributed systems for task-parallel applications, which are linked into useful workflows through the looser task-coupling model of passing data via files between dependent tasks. This potentially larger class of task-parallel applications is precluded from leveraging the increasing power of modern parallel systems such as supercomputers (e.g. IBM Blue Gene/L (Gara et al, 2005) and Blue Gene/P (IBM BlueGene/P (BG/P), 2008)) because the lack of efficient support in those systems for the “scripting” programming model (Ousterhout, 1998). With advances in e-Science and the growing complexity of scientific analyses, more scientists and researchers rely on various forms of scripting to automate end-to-end application processes involving task coordination, provenance tracking, and bookkeeping. Their approaches are typically based on a model of loosely coupled computation, in which data is exchanged among tasks via files, databases or XML documents, or a combination of these. Vast increases in data volume combined

with the growing complexity of data analysis procedures and algorithms have rendered traditional manual processing and exploration unfavorable as compared with modern high performance computing processes automated by scientific workflow systems (Zhao, Raicu, & Foster, 2008).

The problem space can be partitioned into four main categories (see Figure 1). 1) At the low end of the spectrum (low number of tasks and small input size), we have tightly coupled Message Passing Interface (MPI) applications (white). 2) As the data size increases, we move into the analytics category, such as data mining and analysis (blue); MapReduce (Dean & Ghemawat) is an example for this category. 3) Keeping data size modest, but increasing the number of tasks moves us into the loosely coupled applications involving many tasks (yellow); Swift/Falkon (Zhao et al., 2007; Raicu, Zhao, Dumitrescu, Foster, & Wilde 2007) and Pegasus/DAGMan (Deelman et al., 2005) are examples of this category. 4) Finally, the combination of both many tasks and large datasets moves us into the data-intensive Many-Task Computing (Raicu, Foster, & Zhao, 2008) category (green);

44 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/towards-data-intensive-many-task/62821

Related Content

Event Models in Distributed Event Based Systems

Rolando Blanco and Paulo Alencar (2010). *Principles and Applications of Distributed Event-Based Systems* (pp. 19-42).

www.irma-international.org/chapter/event-models-distributed-event-based/44394

On Construction of a Multi-Grid Resource Selection Strategy on Grids

Chao-Tung Yang, Wen-Jen Huang and Kuan-Chou Lai (2014). *International Journal of Grid and High Performance Computing* (pp. 38-62).

www.irma-international.org/article/on-construction-of-a-multi-grid-resource-selection-strategy-on-grids/114712

Multi-Pattern GPU Accelerated Collision-Less Rabin-Karp for NIDS

Anas Abbas, Mahmoud Fayed and Heba Khaled (2024). *International Journal of Distributed Systems and Technologies* (pp. 1-16).

www.irma-international.org/article/multi-pattern-gpu-accelerated-collision-less-rabin-karp-for-nids/341269

Information Security in Data and Storage Grids through GS3

Vincenzo Daniele Cunsolo, Salvatore Distefano, Antonio Puliafito and Marco Scarpa (2012). *Computational and Data Grids: Principles, Applications and Design* (pp. 265-282).

www.irma-international.org/chapter/information-security-data-storage-grids/58749

Security

Valentin Cristea, Ciprian Dobre, Corina Stratan and Florin Pop (2010). *Large-Scale Distributed Computing and Applications: Models and Trends* (pp. 194-216).

www.irma-international.org/chapter/security/43108