Chapter 34
Parallel Data Reduction Techniques for Big Datasets

Ahmet Artu Yıldırım
Utah State University, USA

Cem Özdoğan
Çankaya University, Turkey

Dan Watson
Utah State University, USA

ABSTRACT

Data reduction is perhaps the most critical component in retrieving information from big data (i.e., petascale-sized data) in many data-mining processes. The central issue of these data reduction techniques is to save time and bandwidth in enabling the user to deal with larger datasets even in minimal resource environments, such as in desktop or small cluster systems. In this chapter, the authors examine the motivations behind why these reduction techniques are important in the analysis of big datasets. Then they present several basic reduction techniques in detail, stressing the advantages and disadvantages of each. The authors also consider signal processing techniques for mining big data by the use of discrete wavelet transformation and server-side data reduction techniques. Lastly, they include a general discussion on parallel algorithms for data reduction, with special emphasis given to parallel wavelet-based multi-resolution data reduction techniques on distributed memory systems using MPI and shared memory architectures on GPUs along with a demonstration of the improvement of performance and scalability for one case study.

INTRODUCTION

With the advent of information technologies, we live in the age of data – data that needs to be processed and analyzed efficiently to extract useful information for innovation and decision-making in corporate and scientific research labs. While the term of 'big data' is relative and subjective and varies over time, a good working definition is the following:

- **Big data:** Data that takes an excessive amount of time/space to store, transmit, and process using available resources.

One remedy in dealing with big data might be to adopt a distributed computing model to utilize its aggregate memory and scalable computational power. Unfortunately, distributed computing approaches such as grid computing and cloud computing are not without their disadvantages (e.g., network latency, communication overhead, and high-energy consumption). An "in-box" solution would alleviate many of these problems, and GPUs (Graphical Processing Units) offer perhaps the most attractive alternative. However, as a cooperative processor, GPUs are often limited in terms of the diversity of operations that can be performed simultaneously and often suffer as a result of their limited global memory as well as memory bus congestion between the motherboard and the graphics card. Parallel applications as an emerging computing paradigm in dealing with big datasets have the potential to substantially increase performance with these hybrid models, because hybrid models exploit both advantages of distributed memory models and shared memory models.

A major benefit of data reduction techniques is to save time and bandwidth by enabling the user to deal with larger datasets within minimal resources available at hand. The key point of this process is to reduce the data without making it statistically indistinguishable from the original data, or at least to preserve the characteristics of the original dataset in the reduced representation at a desired level of accuracy. Because of the huge amounts of data involved, data reduction processes become the critical element of the data mining process on the quest to retrieve meaningful information from those datasets. Reducing big data also remains a challenging task that the straightforward approach working well for small data, but might end up with impractical computation times for big data. Hence, the phase of software and architecture design together is crucial in the process of developing data reduction algorithm for processing big data.

In this chapter, we will examine the motivations behind why these reduction techniques are important in the analysis of big datasets by focusing on a variety of parallel computing models ranging from shared-memory parallelism to message-passing parallelism. We will show the benefit of distributed memory system in terms of memory space to process big data because of the system's aggregate memory. However, although many of today's computing systems have many processing elements, we still lack data reduction applications that benefit from multi-core technology. Special emphasis in this chapter will be given to parallel clustering algorithms on distributed memory systems using the MPI library as well as shared memory systems on graphics processing units (GPUs) using CUDA (Compute Unified Device Architecture developed by NVIDIA).

GENERAL REDUCTION TECHNIQUES

Significant CPU time is often wasted because of the unnecessary processing of redundant and non-representative data in big datasets. Substantial speedup can often be achieved through the elimination of these types of data. Furthermore, once non-representative data is removed from large datasets, the storage and transmission of these datasets becomes less problematic.

There are a variety of data reduction techniques in current literature; each technique is applicable in different problem domains. In this section, we provide a brief overview of sampling – by far the simplest method in implementation but not without intrinsic problems – and *feature selection* as a data

21 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/parallel-data-reduction-techniques-for-big-datasets/150191

Related Content

Acquiring Semantic Sibling Associations from Web Documents
Marko Brunzeland Myra Spiliopoulou (2007). *International Journal of Data Warehousing and Mining (pp. 83-98).*
www.irma-international.org/article/acquiring-semantic-sibling-associations-web/1795

A Survey of Managing the Evolution of Data Warehouses
Robert Wrembel (2009). *International Journal of Data Warehousing and Mining (pp. 24-56).*
www.irma-international.org/article/survey-managing-evolution-data-warehouses/1825

Dynamic Social Network Mining: Issues and Prospects
Luca Caglieroand Alessandro Fiori (2013). *Data Mining in Dynamic Social Networks and Fuzzy Systems (pp. 122-144).*
www.irma-international.org/chapter/dynamic-social-network-mining/77526

Social Network Synthesis: A Dynamic Approach for Building Distance Education Programs
E. Pnar Uça-Güneand Gülsün Eby (2017). *Social Media Data Extraction and Content Analysis (pp. 395-411).*
www.irma-international.org/chapter/social-network-synthesis/161971

Estimating the Number of Clusters in High-Dimensional Large Datasets
Xutong Zhuand Lingli Li (2023). *International Journal of Data Warehousing and Mining (pp. 1-14).*
www.irma-international.org/article/estimating-the-number-of-clusters-in-high-dimensional-large-datasets/316142