

Chapter 28

In-Memory Analytics

Jorge Manjarrez Sánchez
Instituto Tecnológico Superior de Jerez, Mexico

ABSTRACT

Analytics is the processing of data for information discovery. In-memory implementation of machine learning and statistical algorithms enable the fast processing of data for descriptive, diagnostic, predictive, and prescriptive analytics. In this chapter, the authors first present some concepts and challenges for fast analytics, then they discuss some of the most relevant proposals and data management structures for in-memory data analytics in centralized, parallel, and distributed settings. Finally, the authors offer further research directions and some concluding remarks.

INTRODUCTION

Processing speed and data set volume are a challenge for traditional analytics systems. Consider when do you want to make a recommendation for an online customer: right there while she is still browsing your catalogue or next time (if) she returns back? When do you want to detect a credit card is fraudulent: before a transaction is committed or after a client complain some days later? Score or decide when a credit card transaction is fraudulent or an incoming email is spam, is relatively easy from a computationally point of view provided the appropriate algorithms are implemented correctly. But in real-life systems where hundreds of thousands of such transactions arrive simultaneously and the correct decision must be taken instantaneously or in near real-time, speed is crucial. Even though only for the second example question an incorrect false negative, i.e., the no detection of a fraud, can have severe negative consequences, these examples highlight two main challenges in data analytics: the need to process ever larger data volumes and the need of obtaining fast or real time results. These are the concerns of In-memory analytics.

In what follows, we provide a definition and some concepts necessary to understand the problem and solution approach of in-memory analytics. Then we present a review of some important technological proposals to deal with two main challenges in data analytics: speed and scalability. The presentation is made from a data management perspective, independent of their nature (structured or unstructured), because an efficient data management strategy provides the infrastructure to handle and enable fast access to voluminous data sets for their analytical processing.

DOI: 10.4018/978-1-5225-7598-6.ch028

BACKGROUND

Analytics is the processing of data for facts and information discovery. Embedded within data there are facts about some events, perhaps several and a variety of them. These facts are knowledge helpful for taking decisions, recommendations systems, fraud detection, sentiment analysis in social networks and customer identification and many other applications. In a broad sense analytics can provide an answer for questions such as what happened? Why did it happen? What will happen? What should I do? And they correspond to the task of descriptive, diagnostic, predictive and prescriptive analytics respectively. The answer to these questions are obtained by the application of techniques from statistics, machine learning and computer science in the processing of data.

Analytics is the processing of data for making sense of all of it. Data sets are very large and to process them one has to cope with two orthogonal concerns: accuracy and speed of results. Accuracy needs to process as much data as possible, not just random samples. Not using complete datasets leads to loss of information (Wang, Callan, & Zheng, 2015) and obtained responses are approximations with certain level of error. Accuracy of the statistical model and henceforth that of predictions are proportional to samples size and quality. The bigger the better, but even that we can have today larger data sets, their processing takes considerable time. It consumes more processing time and disk input/output (IO) operations because data must be swapped back and forth from hard drives to RAM and CPU caches. But due to speed differences in data access and transfer between these devices, there is a processing bottleneck. One approach to avoid IO processing bottlenecks is to use as much as possible RAM memory, hence the name In-Memory Analytics, where data is loaded and kept into RAM for their fast consumption by the processing algorithms.

FAST ANALYTICS

In this section we address the speed challenge. When the speed of results is a factor to consider for efficient analytics, and taking into consideration that it must be used as much data as possible for the abovementioned reasons, then one must reduce or eliminate processing bottlenecks. They can be the IO latency and CPU usage. CPU is good at computing mathematical operations and multicore CPU's enhance their performance. IO depends on the type of disk and RAM. Solid state disks are faster than hard disk drives, but random access memory is faster. A typical hard drive has a latency of 5 ms while RAM has only 100 ns, it is 50000 times faster. But also the faster the more expensive. There is one more rapid type of memory: cache memory, which is a small and expensive memory within the CPU die. Main memory or In-memory refers to the processing of data performed in RAM.

The data management infrastructure provides the processing capabilities according to the workload type. The typical workloads in a database, can be Online Analytical Processing (OLAP) or Online Transaction Processing (OLTP). The difference between both are their usage patterns, size of result sets and processing speed. Traditionally OLTP deals with faster small transactions and OLAP seeks for answer precision on very large data sets, but nowadays speed is an important factor to provide for both of them, hence the use of in-memory analytics, the usage of RAM fast data access and to enhance analytics performance. But the issue is that real data sets are usually larger than the memory affordable and available in a system. Also note that not all available memory can be used for data storage, enough

7 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/in-memory-analytics/214629

Related Content

A Fast Image Encoding Algorithm Based on the Pyramid Structure of Codewords

Ahmed A. Radwan, Ahmed Swileman and Mamdouh M. Gomaa (2009). *International Journal of Mobile Computing and Multimedia Communications* (pp. 1-13).

www.irma-international.org/article/fast-image-encoding-algorithm-based/37452

Multi-Band Rectangular Zig-Zag-Shaped Microstrip Patch Antenna for Wireless Applications

Pranay Yadav, Bharat Bhushan Khare, Sudesh Gupta, Yash Kumar Kshirsagar and Swati Jain (2020). *Design and Optimization of Sensors and Antennas for Wearable Devices: Emerging Research and Opportunities* (pp. 66-86).

www.irma-international.org/chapter/multi-band-rectangular-zig-zag-shaped-microstrip-patch-antenna-for-wireless-applications/235783

Mobile Location Tracking: Indoor and Outdoor Location Tracking

Sima Nadler (2017). *Mobile Application Development, Usability, and Security* (pp. 194-209).

www.irma-international.org/chapter/mobile-location-tracking/169682

Multi-Criteria Recommender Systems: A Survey and a Method to Learn New User's Profile

Ferdaous Hdioud, Bouchra Frikh, Brahim Ouhbi and Ismail Khalil (2017). *International Journal of Mobile Computing and Multimedia Communications* (pp. 20-48).

www.irma-international.org/article/multi-criteria-recommender-systems/193258

Wireless-Enabled Fashion: Overall Supply Chain Impacts and Differentiating Technologies

L.F. Pau (2019). *International Journal of Mobile Devices, Wearable Technology, and Flexible Electronics* (pp. 37-56).

www.irma-international.org/article/wireless-enabled-fashion/268890