G Model
JOCS-670; No. of Pages 13
ARTICLE IN PRESS
Journal of Computational Science xxx (2017) xxx–xxx
Contents lists available at ScienceDirect
Journal of Computational Science
journal homepage: www.elsevier.com/locate/jocs

AQP++: a hybrid approximate query processing framework for generalized aggregation queries

Yuxiang Wang, Yixing Xia, Qiming Fang, Xiaoliang Xu*

School of Computer Science and Technology, Hangzhou Dianzi University, PR China

ARTICLE INFO

Article history:
Received 30 November 2016
Received in revised form 14 April 2017
Accepted 2 May 2017
Available online xxx

MSC:
00-01
99-00

Keywords:
Approximate query processing
Online aggregation
Bootstrap
Sampling

ABSTRACT

A sampling-based approximate query processing (AQP) method provides a fast way for users to obtain a trade-off between accuracy and time consumption by executing the query on a sample of data rather than the whole dataset. There are two major AQP methods: the (1) central limit theorem (CLT)-based online aggregation; and the (2) bootstrap method. The former is very efficient but is only suitable for simple aggregation queries, while the latter is quite general but has relatively high computational overhead. Both methods suffer from the possible estimation failure. However, there is no technology that can both support simple/complex queries within an acceptable time coupled with carefully considering the estimation failure. To make the current AQP method much more general and efficient, we propose a hybrid approximate query framework called AQP++ to combine the advantages of both methods and eliminate the limitations as far as possible. According to this hybrid framework, an estimation parameters adjustment method is presented for CLT-based online aggregation to improve its usability for much more complex aggregation queries. Then, an execution cost model is proposed to describe the computational overhead of the two AQP methods, which can be used to support our dynamic scheduling mechanism of AQP++ and make the whole system more efficient and flexible. Moreover, we have implemented our AQP++ prototype and conducted extensive experiments on the TPC-H benchmark for skewed data distribution. Our results demonstrate that our AQP++ can produce acceptable approximate results for both simple and complex queries within a much shorter time compared with the original CLT-based online aggregation and bootstrap method.

1. Introduction

Big data has been produced in various applications, including the user information from social networks, sensor data, scientific data, and a variety of log data, etc. And big data analytics is playing an important role in today's fast-paced data-driven businesses [1]. To solve the issues associated with the big data analytics, many distributed systems have been proposed such as Hadoop [2], Spark [3], Hyracks [4], etc., to provide more efficient and cost-effective solutions, which can be widely used for different areas such as task scheduling, data placement, replica management, and execution mechanisms [5–10].

Big data analytics applications typically need to process a tremendous amount of data on clusters of tens, hundreds, or thousands of machines to derive the latent useful information quickly. The large quantity of data and the relatively limited computation, storage, and network resources often make it infeasible to deliver exact answers at interactive speeds [11]. However, many applications can tolerate some degree of inaccuracy especially for exploratory queries, where the early approximate results that are accurate enough are often of much greater value to users than tardy exact results [12]. Sampling-based approximate query processing (AQP) provides a fast way for users to obtain a trade-off between accuracy and time consumption by executing each query on a sample of data rather than the whole dataset.

A large body of work has subsequently been proposed for AQP on centralized [13–17] and distributed (e.g. P2P and Cloud) [11,12,18–30] environments. The basic mathematical theory behind these researches can be divided into two major categories: the (1) central limit theorem (CLT)-based online aggregation (OLA); and the (2) bootstrap methods. Both methods allow users to observe the progress of a given query by showing iteratively refined approximate answers, and then stop the query execution

* Corresponding author.
E-mail addresses: lsswyx@hdu.edu.cn (Y. Wang), yixingx@hdu.edu.cn (Y. Xia), fangqiming@hdu.edu.cn (Q. Fang), xxl@hdu.edu.cn (X. Xu).

http://dx.doi.org/10.1016/j.jocs.2017.05.001
1877-7503/© 2017 Published by Elsevier B.V.

once the result achieves the desired accuracy (the confidence interval is narrowed according to the increasing sample size). For the CLT-based OLA method, a sequence of independent and identically distributed (i.i.d.) random variables are drawn from the population without replacement, and the statistical estimator is then calculated to obtain the approximate results. While he bootstrap method adopts a totally different way to achieve the same goal, that is a set of resampling procedures are conducted against to a given simple random sample (each resampling procedure generates a resample with replacement) to construct the "sample world", and then the statistical estimators that are computed based on such resamples can be viewed as an approximation to the "real world" (detailed information is given in Section 3).

Note that the CLT-based OLA is extremely efficient but lacks generality, which is restricted to very simple SQL queries (often with only a single layer of basic aggregates such as AVG, SUM, COUNT, VARIANCE, and STDEV with projections, filters, and a GROUP BY) [11]. This is because the estimators of complex aggregates cannot be given easily. On the other hand, the bootstrap method can be applied to arbitrarily complex SQL queries but with relatively higher computational overhead due to the resampling procedure. In addition, both AQP methods mentioned above suffer from the possible estimation failure, leading to the underestimation or overestimation of the actual error. Unfortunately, there is no technology can well support more categories of queries within an acceptable time coupled with carefully considering the estimation failure. To the best of the authors' knowledge, the best attempt to combine these two different AQP methods was made in [26], which can improve the overall performance by a dynamic scheme switching mechanism. However, the solution proposed in that paper does not solve the generalized problem of CLT-based OLA, and does not consider the estimation failure of the bootstrap method. This has motivated the new hybrid AQP framework called AQP++ that is proposed in this paper to find a balance between generality and efficiency.

In the AQP++ framework, an estimation parameter adjustment method is first presented for CLT-based OLA (which can be viewed as an extension of the original CLT-based OLA) to improve its usability for much more complex aggregation queries. Then, an execution cost model, which considers the effect of the estimation failure, is proposed to describe the computational overhead of the extended CLT-based OLA and original bootstrap method. Thirdly, a dynamic scheduling mechanism is given based on such cost model to make AQP++ much more flexible. Moreover, we have implemented our AQP++ prototype based on Hadoop and conducted extensive experiments on the modified TPC-H benchmark [31][1] to demonstrate the efficiency and effectiveness of AQP++.

The main contributions of this paper are as follows.

- We propose a new hybrid AQP framework called AQP++, which can obtain the generality and efficiency characteristics by combining the advantages of two AQP methods and eliminating the limitations as far as possible.
- We present an estimation parameters adjustment method for CLT-based OLA to improve its usability, which gives an alternative approach for complex aggregation queries.
- We derive an execution cost model by carefully considering the effect of the estimation failure, and then a dynamic scheduling mechanism is given according to this cost model to make our AQP++ more flexible.
- We implement our AQP++ prototype and conduct extensive experiments to demonstrate that it can well support more categories of queries and deliver reasonable precise online estimates within an acceptable time.

The remainder of this paper is organized as follows. In Section 2, we review related work firstly. In Section 3, we introduce the preliminaries of CLT-based OLA and bootstrap method, and then reveal the strengths and weaknesses of them by comparing from multiple aspects. In Section 4, we give an overview of AQP++. Section 5 proposes an estimation parameters adjustment method. Then we describes the cost model of AQP++ and the dynamic scheduling mechanism in Section 6. And in Section 7, we introduce the experimental setup and report results of the experimental evaluation. Finally, we conclude in Section 8.

2. Related work

Online aggregation (OLA) [13] is one commonly used AQP technique to provide a time-accuracy tradeoff for aggregation queries. In [14], Haas illustrates how the different theories can be used to derive formulas for both large-sample and deterministic confidence intervals. Meanwhile, the focus in [15,32] is on the OLA for join operations. However, all works in [13–15,32] are limited to single query processing. Therefore, Wu et al. proposed COSMOS to support multiple-query optimization for OLA [16].

However, these centralized OLA methods cannot be extended to a distributed manner easily, so well-designed distributed OLA methods have been proposed along with the development of peer-to-peer (P2P) and cloud computing [18–21,23–25]. Wu et al. extended OLA to a P2P context where sites are maintained in a distributed hash table (DHT) network [18,19]. In addition, in [20,21] the authors demonstrated a modified version of Hadoop that supports OLA, which can only return the query progress without any precision estimation. In [23], the authors formulated a statistical foundation that supported block-level sampling in OLA to improve the sampling efficiency. In [25], the authors proposed a fair-allocation strategy to guarantee the storage and computational load balancing for running OLA over a MapReduce-based cloud system. And in [25], the authors focused on the research of multiple-query sharing for OLA in a cloud system.

The basic mathematics theory of all works mentioned above is the CLT, which is quite efficient but only supports simple aggregation queries. To make OLA much more general, a new research area for OLA has emerged by adopting the bootstrap method to obtain the approximate query results [12,27–30,11]. In [12], the authors first introduced how the bootstrap method can be deployed on a basic cloud system such as Hadoop, which can make OLA available for arbitrary aggregation queries but pay less attention to the efficiency problem of bootstraps. In [27], the authors investigated the bootstrap performance issue for big data applications, and the bag of little bootstraps (BLB) approach was presented, which is well suited to modern parallel and distributed computing architectures and retains the generic applicability, statistical efficiency, and favorable theoretical properties of the original bootstrap. And then Kleiner et al. proposed a diagnostic method based on such a BLB approach to support estimation failure detection when the bootstrap is processing. However, this additional diagnostic procedure increased the overall overhead of query processing so that a set of optimizations (e.g. scan consolidation) were presented to decrease the side effects of the diagnostic procedure. Moreover, in [28] a novel probabilistic relational model was introduced for the bootstrap process, along with rigorous semantics and a unified error model, to further improve the bootstrap performance. In [29], the authors also presented an extended bootstrap method for interactive analysis on big data especially for the queries which have arbitrarily nested aggregates.

[1] We use the modified TPC-H benchmark which can generate the dataset with different data distribution.