

Estimating Aggregates in Time-Constrained Approximate Queries in Oracle

Ying Hu
ying.hu@oracle.com

Seema Sundara
seema.sundara@oracle.com

Jagannathan Srinivasan
jagannathan.srinivasan@oracle.com

ABSTRACT

The concept of time-constrained SQL queries was introduced to address the problem of long-running SQL queries. A key approach adopted for supporting time-constrained SQL queries is to use sampling to reduce the amount of data that needs to be processed, thereby allowing completion of the query in the specified time constraint. However, sampling does make the query results approximate and hence requires the system to estimate the values of the expressions (especially aggregates) occurring in the select list. Thus, coming up with estimates for aggregates is crucial for time-constrained approximate SQL queries to be useful, which is the focus of this paper. Specifically, we address the problem of estimating commonly occurring aggregates (namely, SUM, COUNT, AVG, MEDIAN, MIN, and MAX) in time-constrained approximate queries. We give both point and interval estimates for SUM, COUNT, AVG, and MEDIAN using Bernoulli sampling for various type of queries, including join processing with cross product sampling. For MIN (MAX), we give the confidence level that the proportion 100% of the population will exceed the MIN (or be less than the MAX) obtained from the sampled data.

1. INTRODUCTION

The growing nature of databases, compounded with the ability to formulate arbitrarily complex SQL queries, has led to the problem of long-running, complex SQL queries.

A solution being explored is to support time-constrained SQL queries [2], [3] that would complete in a specified time constraint either by computing the first few rows (top-K rows) or approximate results through sampling. Of the two approaches, the latter approach, namely approximate query processing, is very promising in that the query processing time could be reduced significantly by controlling the sample size. A practical application of time-constrained approximate query processing is queries involving aggregate functions.

Such queries are popular in applications such as OLAP and they tend to be long running as they compute aggregate values over large datasets. However, supporting time-constrained approximate SQL queries require work in two areas for them to become a practical and useful solution.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the ACM. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permissions from the publisher, ACM.

EDBT'09, March 24-26, 2009, Saint Petersburg, Russia.

Copyright 2009 ACM 978-1-60558-422-5/09/0003 ...\$5.00.

First, the user-specified time-constraint needs to be implicitly transformed to SAMPLE clauses on individual tables. This was addressed in [3], which presented estimation of sample sizes for queries involving various relational operations.

Second, the problem of estimating aggregates needs to be considered. Thus, in this paper, we focus on estimating aggregates in time-constrained approximate SQL queries. Since the aggregates in time-constrained approximate queries are computed only once, the result could vary significantly based on the chosen sample size, which warrants that additional measures are provided characterizing the goodness of the results. We consider commonly occurring aggregates, namely, SUM, COUNT, AVG, MEDIAN, MIN, and MAX. The measures (apart from the point estimate) that are useful are confidence intervals for aggregates SUM, COUNT, AVG, and MEDIAN, and confidence levels for aggregates returning extreme values (such as MIN and MAX) as tolerance limits. The aggregate estimation techniques are presented for join queries that employ cross-product sampling [1].

The rest of this paper is organized as follows: Section 2 describes the Bernoulli sampling scheme. Section 3 discusses the estimation for SUM, COUNT, and AVG. Sections 4 and 5 cover the estimation for MEDIAN, QUANTILE, MIN and MAX. The results in Sections 3, 4, and 5 are presented by assuming row sampling but could be extended to block sampling as well, as discussed in Section 6. Section 7 concludes the paper.

2. BERNOULLI SAMPLING

Oracle Database supports the Bernoulli (coin-flip) sampling scheme, where the sample percentage (f) indicates the probability of each row, or each cluster of rows in the case of block sampling, being independently selected as part of the sample. Because the database does not retrieve the exact sample size of the rows (blocks) of table, Bernoulli sampling is a variable size sampling scheme. The mean and variance of the random sample size n are given by $E(n) = fN$ and $V(n) = f(1-f)N$, where N is the population size, or the number of rows (blocks) in the case of row sampling (block sampling). In this paper we assume that the value of N is known from Oracle Database's object-level statistics, which includes the number of blocks and the number of rows in a table.

3. SUM, COUNT, AND AVG

We start with the formulas for the estimated COUNT, SUM, and AVG and their variance in join operations. We assume that there are k tables in the join operations and each table's sample percentage ($f_j, j = 1, \dots, k$) is calculated by an algorithm described in [3]. We also assume that the j -th sample S_j has n_j rows chosen from N_j rows of the j -th table R_j , where the value of N_j is known from table statistics. These assumptions also apply to the Sections 4 and 5. Note that under the Bernoulli (coin-flip) sampling, n_j is a random variable with mean $E(n_j) = f_j N_j$.

