A Unified Correlation-based Approach to Sampling Over Joins

Niranjan Kamat
The Ohio State University
kamatn@cse.ohio-state.edu

Arnab Nandi
The Ohio State University
arnab@cse.ohio-state.edu

ABSTRACT

Supporting sampling in the presence of joins is an important problem in data analysis, but is inherently challenging due to the need to avoid correlation between output tuples. Current solutions provide either correlated or non-correlated samples. Sampling might not always be feasible in the non-correlated sampling-based approaches – the sample size or intermediate data size might be exceedingly large. On the other hand, a correlated sample may not be representative of the join. This paper presents a *unified* strategy towards join sampling, while considering sample correlation every step of the way. We provide two key contributions. First, in the case where a *correlated* sample is *acceptable*, we provide techniques, for all join types, to sample base relations so that their join is *as random as possible*. Second, in the case where a correlated sample is *not acceptable*, we provide enhancements to the state-of-the-art algorithms to reduce their execution time and intermediate data size.

CCS CONCEPTS

•Information systems →Join algorithms;

KEYWORDS

Sampling, Join, Correlation, Random, Randomness

ACM Reference format:

Niranjan Kamat and Arnab Nandi. 2017. A Unified Correlation-based Approach to Sampling Over Joins. In *Proceedings of SSDBM '17, Chicago, IL, USA, June 27-29, 2017*, 12 pages. DOI: <http://dx.doi.org/10.1145/3085504.3085524>

1 INTRODUCTION

In cases where the data size restricts the hardware or software's ability to process it within reasonable time, sampling presents a pragmatic approach towards providing insights at scale. Joining multiple datasets helps incorporate related knowledge sources and develop deeper insights and has proven to be a valuable operation in data analysis in multiple scientific domains such as geo-spatial data [9], sensor networks [1], astronomy (SDSS [31]), etc. Performing statistical analyses using aggregations is a popular step following the joins. Sampling over joins is a compelling, yet challenging task. Performing a join can be expensive – sampling after materializing the join may not be a pragmatic approach. Initial efforts in join sampling

were directed towards obtaining *non-correlated samples*¹ [2, 8, 38]. As Chaudhuri et al. [8] demonstrated the inherent hardness of the problem and as online aggregation [21] started gaining momentum, efforts came to be directed towards obtaining correlated, probability samples for aggregation queries [13, 16, 24, 30, 45]. However, non-aggregation use cases (such as presenting a sample to user) have not been considered by this line of work. Correlated samples usually also have higher error estimates compared with non-correlated samples [10]. In this context, we aim to reduce correlation in correlated samples by maximizing *join randomness* (number of possible samples as a result of a join algorithm) – this has a side-effect of lower sampling error when compared with other correlated sampling techniques. Further, when non-correlated samples are mandated, we suggest enhancements to the state-of-the-art algorithms. To better understand the problem domain, we illustrate the primary challenge in non-correlated sampling of joins using an example.

1.1 Challenges in Avoiding Sample Correlation

Chaudhuri et al. [8] and Gibbons et al. [2] have provided excellent examples demonstrating one of the difficulties of join sampling – a relation has to be sampled in a biased fashion while considering the join key cardinality of the other. We look at another key challenge – that of the sample size far exceeding the relation size. Before doing so, we first demonstrate the cause of sample correlation, avoidance of which results in the aforementioned challenge.

1.1.1 Correlated Samples. To avoid correlation, a tuple being present in a sample should not affect the inclusion probability of another. For example, consider samples $S_1 = \{t_1, t_2\}$ and $S_2 = \{t_3, t_4\}$ of relations R_1 and R_2 respectively. Join between S_1 and S_2 produces the tuples $\{t_1 \circ t_3, t_1 \circ t_4, t_2 \circ t_3, t_2 \circ t_4\}$, where $t_i \circ t_j$ represents concatenation of t_i and t_j . Each tuple in S_1 and S_2 results in multiple join tuples, i.e., presence of a tuple, $t_i \circ t_j$, in the output increases the probability of other tuples having their R_1 component be t_i or R_2 component be t_j , resulting in a correlated sample.

1.1.2 Sample Inflation. We illustrate how avoidance of correlation can result in large samples. Consider joining two relations, each having a single key, and cardinalities of 50 and 100 respectively, causing join to have 5000 tuples. A 0.1 fraction of the join will consist of 500 tuples – far larger than the size of either relation. If a sampled tuple produces multiple joined tuples, the tuples will be correlated. To avoid correlation, a sampled tuple has to be restricted to *join only once*. This results in both samples having a size of 500, making sampling counter-productive. We define this need to avoid correlation between output tuples resulting in increased sample size as *sample inflation*. It is the primary reason behind sampling over joins being difficult, and an infeasible proposition at times.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SSDBM '17, Chicago, IL, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.
978-1-4503-5282-6/17/06...\$15.00
DOI: <http://dx.doi.org/10.1145/3085504.3085524>

¹We define a correlated sample as one where the inclusion probability of an item depends on that of another (details in Section 1.1.1) – samples having any degree of randomness in the sampling process can technically be called random samples [10].

1.1.3 Sample Inflation in Current Non-Correlated Sampling Approaches. Chaudhuri et al. [8] provide algorithms to obtain non-correlated samples for different index and statistics availabilities. We demonstrate how sample inflation can occur in each of these algorithms – looking at their experimental results, it appears that they avoided sample inflation through either careful implementation or query plan optimization. We provide enhancements to these non-correlated sampling algorithms to reduce the sample size and intermediate data size. However, the reduction might not be satisfactory due to sample inflation. Hence, the primary focus of this paper is to provide algorithms to reduce the correlation in correlated samples by maximizing the number of possible samples. This is based on our observation that non-correlated sampling techniques result in the maximum possible number of samples. Our efforts, thereby, in increasing the number of possible samples are directed towards making the samples as non-correlated as possible.

1.2 Overview & Contributions

This paper provides two key contributions – first, we look at techniques to maximize the number of possible samples in correlated sampling under the constraint of fixed sample size, when statistics over the join column are available. We provide strategies for allocating samples to different strata of multiple relations for different join types, including equi-join, outer join, self-join, non-equi-join, and theta join for the comparators $<$, $\leq$, $>$, $\geq$ (Sections 3 and 4). These techniques are derived mathematically in the Appendix. Although the derivations are complex, the resultant allocation strategies are simple, intuitive, and easy to compute. They have been experimentally validated to provide allocation close to optimal allocation that was found using a brute-force search. The sampling error of our techniques was found to be lower than that of other correlated join sampling techniques (Section 6.2.2).

Our second contribution is in non-correlated sampling, where we provide enhancements to the state-of-the-art algorithms [8]. When complete or partial statistics are available over a relation, our algorithms, GROUP-SAMPLE-ENHANCED (Section 5.1), and FREQUENCY-PARTITION-SAMPLE-ENHANCED (Section 5.2), access only half of the tuples in comparison, avoid the need to create large intermediate data which can be larger than the base relations, and remove a sampling step over the intermediate data. In the case where statistics and indexes are available over a single relation, we specify a filter-based criterion to decide whether sampling should be used, violating which can result in the sample size exceeding the relation size (Section 5.3). We also use this criterion to sample both relations, when statistics and indexes are available over both (Section 5.4).

2 RELATED WORK

As performing joins can be expensive, in addition to sampling, efforts have been directed towards accelerating them through various other means such as GPUs [20], MapReduce [6, 49], multi-threading [5, 40], networked execution [36], etc. Sampling over joins has been incorporated in numerous industrial and research systems such as *SQL Server*, *DB2*, *AQUA* [2], *Turbo-DBO* [12], *BlinkDB* [3], and *Quicksr* [30]. While these approaches target various layers of the database such as table scans, offline catalogs, and aggregation, there still exist several open opportunities.

In the context of non-correlated sampling, Olken et al. [38] show that it is possible to commute selection with sampling, but commuting projections and joins is harder, and sample a single relation. Chaudhuri et al. [8] provide better algorithms for joins and consider different availabilities of statistics and indexes. The *AQUA* system [2] obtains a simple random sample of a join in the primary key-foreign key scenario using join synopses, which *TuG* [43] extends to the many-to-many join scenario. While Gemulla et al. [15] aim to minimize the space overhead of join synopses, *CS2* [48] extends join synopses by proposing correlated sample synopsis, which instead of storing a sample of the join, stores a sample of the correlated tuples. *SciBORQ* [42] maintains correlation between join attributes using impressions. Bifocal sampling [14] recognizes multiplicative effect of strata² sizes and develops different sampling strategies based on strata sizes to estimate query size.

Research in obtaining non-correlated samples of joins stalled due to Chaudhuri et al. [8] showing the inherent hardness of the problem. However, considerable efforts have been expended towards obtaining correlated samples and their meaningful aggregation estimates through online aggregation, which allowed for error estimation during query execution [16–18, 21, 39]. Jermaine et al. [23, 25] removed dependence of ripple join algorithms on the data residing in memory to estimate error. The *DBO* [24] and *Turbo-DBO* [12] systems allowed processing of multiple relations in a scalable fashion. Nirkhiwale et al. [37] presented the sampling algebra inherent in these techniques. Our algorithm for sampling both sides of a join in the presence of indexes, STRATJOIN, is influenced by hash ripple joins [35], although they provide correlated samples. In *Wander Join*, Li et al. [33] improved ripple joins by sampling tuples from subsequent relations of the join which can join with the currently selected tuples. In this paper, we aim to improve correlated samples by reducing correlation in samples using strategies that maximize the number of possible samples. Our concept of maximizing the number of samples in joins is influenced by the notion of *sample randomness*, which has been introduced by Kateb et al. [4], who use it to improve stratified reservoir sampling.

Some correlated sampling-based approaches such as *End-Biased Sampling* [13], *Correlated Sampling* [45], and *Universe Sampler* [30] join *all* tuples whose hash of the join key lies in a randomly chosen range of the hashed join domain. Tuples whose key hash does not lie in the selected domain are discarded. This results in the samples and the join being cluster samples [10]. Cluster samples can be useful when join and measure columns are not correlated – our experiments, indeed, show that such approaches have a large error in the presence of correlation between the join and measure columns. In contrast, our approach provides representatives from all join keys and performs well in the presence of correlation.

3 MAXIMIZING RANDOMNESS FOR EQUI-JOIN

As mentioned earlier, we approach the correlated sampling problem from the perspective of maximizing join randomness. The problem statement we address can be given as follows: *For the case where correlated samples are acceptable, determine the sample allocation strategy to maximize join randomness, given the constraint of fixed sample size, in the presence of statistics over join columns.*

²Strata can be defined as the distinct keys in the join output.

Table 1: List of Notations

Symbol	Explanation
R_i	i^{th} Relation of the Join
N_i	Cardinality of R_i
S_i	Sample of R_i
n_i	Cardinality of S_i
f	Join Sampling Rate
A	Join Column
a	Join Value
$t.A$	Value of Column A in Tuple t
$m_i(a)$	Number of Tuples in R_i having Value a in A
m_i^j	Number of Tuples in R_i belonging to j^{th} stratum
$mm_i(a)$	Number of Tuples in S_i having Value a in A
mm_i^j	Number of Tuples in S_i belonging to j^{th} stratum
z	Number of Relations
n	Number of Strata
k	Total Sample Size
k^j	Number of Samples Allocated for Stratum j for all Relations

3.1 Join Randomness

We use *join randomness*, defined as number of samples possible in a join algorithm, to reduce correlation in a sample. To understand our motivation behind doing so, let us look at an example depicting the number of possible samples as a result of a few different sample allocations. Consider joining two relations with strata sizes, $mm_1^1 = 10$, $mm_2^1 = 10$, $mm_1^2 = 20$, and $mm_2^2 = 20$. Let the sample size be constrained at 30. The number of possible samples is $C_{mm_1^1}^{m_1^1} \times C_{mm_2^1}^{m_2^1} \times C_{mm_1^2}^{m_1^2} \times C_{mm_2^2}^{m_2^2}$. Table 2 shows that different sample allocations can result in the number of possible samples differing by multiple orders of magnitude.

Table 2: Effect of Sample Allocation on Number of Samples

Index	mm_1^1	mm_2^1	mm_1^2	mm_2^2	# Samples
1	5	5	10	10	2.2×10^{15}
4	3	7	8	12	2.3×10^{14}
2	2	3	12	13	5.3×10^{13}
3	1	1	14	14	1.5×10^{11}

If we are not restricted to sampling the input and then performing the join, but could sample the join of the relations (resulting in a non-correlated sample), the number of possible samples would have been exceedingly large, $C_{|S_1 \bowtie S_2|}^{|R_1 \bowtie R_2|} = C_{125}^{500}$ for the first configuration. Thus, our efforts can be perceived as being towards making the samples as non-correlated as possible, through the metric of number of samples. This view has been strengthened by our experiments, which show that our approach has lower error than other correlated sampling-based approaches – non-correlated samples result in lesser error than correlated samples theoretically.

3.2 Maximizing Randomness for Single Stratum

Consider allocating k^j tuples amongst relations $R_1, R_2 \dots R_z$, each having a single stratum, to maximize the number of possible samples, $\prod_{i=1}^z C_{mm_i^j}^{m_i^j}$. Appendix A shows that Equation 1 can be used for this purpose – Section 6.1.1 shows that it results in a low error, with the maximum difference in the sample stratum values as a result of our allocation and the optimal allocation found by searching through all possible allocations being 2.

$$mm_i^j = \text{round} \left(\frac{k^j \times m_i^j}{\sum_{i=1}^z m_i^j} \right) \quad (1)$$

3.3 Maximizing Randomness for Multiple Strata

We now provide the strategy to allocate a given sample size amongst different strata, in the general case of equi-join between multiple relations having multiple strata. Consider the problem of determining the sample allocation k^j for j^{th} stratum, for $j \in [1, n]$, with $k = \sum_{j=1}^n k^j$. Our goal is to maximize the number of possible samples, $\prod_{j=1}^n \prod_{i=1}^z C_{mm_i^j}^{m_i^j}$. Appendix B shows that we can use Equation 2 to do so – Section 6.1.2 shows that the allocation is close to optimal allocation, with the maximum difference being 1.

$$k^j = \text{round} \left(\frac{k \times \sum_{i=1}^z m_i^j}{\sum_{j=1}^n \sum_{i=1}^z m_i^j} \right) \quad (2)$$

3.4 Combined Algorithm – MAXRANDJOIN

Using equations 1 and 2, mm_i^j can be given as follows.

$$mm_i^j = \frac{m_i^j}{\sum_{i=1}^z m_i^j} \times \frac{k \times \sum_{i=1}^z m_i^j}{\sum_{j=1}^n \sum_{i=1}^z m_i^j} = \frac{k \times m_i^j}{\sum_{j=1}^n \sum_{i=1}^z m_i^j} \quad (3)$$

The resulting value will be rounded. We can now provide our overall algorithm to maximize randomness for equi-joins.

Algorithm 1 Maximize Randomness for Equi-Joins

1. Create z allocation tables, with each table having with n rows and 3 columns. The columns represent the stratum value, the target allocation mm_i^j calculated using Equation 3, and the current allocation (at a complexity of $O(n \times z)$).
 2. Sample the z relations in the following fashion. For the i^{th} relation, scan through its rows to create a reservoir sample without replacement [46] of size mm_i^j for every stratum $j \in [1, n]$ using the i^{th} allocation table (at a complexity of $\sum_{i=1}^z |R_i|$).
 3. Join the samples (variable complexity).
-

3.5 Applicability

MAXRANDJOIN needs statistics over the join columns. If they are unavailable, table scans will be needed in the case of row-stores. However, in the case of column-stores, it will only entail a comparatively inexpensive scan over the join columns. Complexity of the join step (Step 3) is dependent on the query optimizer, which will be able to choose the correct plan since precise join column statistics will be available at this step.

3.6 Derivation of Allocation Strategy

We use Lagrange multipliers, a popular tool in the statistical sampling community to find approximate optimal strata allocation in closed-form under space or cost constraints [11, 26, 32, 34, 44]. This includes different approaches for determining strata sizes in stratified random sampling such as Neyman allocation and cost-based allocation [34]. Lagrange multipliers provide us with critical points for maximum and minimum values, if they exist. Our functions will have a minimum value as the number of samples is non-negative. They will also possess a maximum value as the relation and sample sizes are bounded. The critical points have to be plugged into the function for the number of possible samples to determine if the resultant value is a maximum or a minimum. Our experiments (Section 6.1) show that suggested sample sizes are close to the optimal solution in practice – with a maximum difference of 2 in the case of single stratum allocation and 1 for multiple strata allocation. The derivations possess a couple of sources of potential error. They might result in rounding errors as they maximize the number of samples in the continuous domain, whereas the allocation occurs in the discrete domain. Another source of error can be our use of an approximation for the harmonic sum, which results in the assumption of mm_i and $m_i - mm_i$ being *large* (Equation 10). Finally, we note that our *simple*, intuitive, closed-form formulae maximize expressions consisting of factorials that are combinatorial in nature.

3.7 Underlying Intuition

While our allocation strategy for equi-joins is intuitive, its derivation is rather complex. However, the underlying intuition behind our strategy is interesting and straightforward. Consider allocating a fixed number of tuples randomly amongst different strata. Intuitively, proportional allocation is more likely to occur than other allocations. The number of ways to come across any particular allocation will equal the number of possible samples. As proportional allocation is the most likely occurrence, proportional allocation will also result in the maximum number of possible samples.

3.8 Comparison with Correlated Join Samplers

Our approach embraces correlation in the samples, and aims to reduce it by maximizing the join randomness. Section 6.2.2 shows that a side-effect of maximizing the join randomness is lower sampling error when compared with other correlated samplers.

End-Biased Sampling, *Correlated Sampling*, and *Universe Sampler* provide a cluster sample – they are ill-suited for handling correlation between join and measure columns. Such approaches circumvent the sample inflation issue by considering only those tuples that their hashing function accepts. This removes the need for histogram information as well. On the other hand, we provide a unified approach to both correlated and non-correlated sampling scenarios using strategies with sound mathematical origins that tackle sampling inflation head-on.

Ripple Join, *SMS Join*, and *Wander Join*, at every point in their execution, provide a correlated output as a result of a join between simple random samples of the relations. These approaches do not take strata-based skew into consideration and as a result have a higher error than our approach (Section 6.2.2).

MAXRANDJOIN is specifically designed to provide a sample that is *as random as possible* and takes all strata into consideration. In contrast, other correlated samplers have different goals such as responsiveness, streamability, scalability, removing the need for statistics and indexes, etc. Comparing MAXRANDJOIN to them is not straightforward – they have different objectives with the resultant benefits and drawbacks.

4 MAXIMIZING RANDOMNESS FOR OTHER JOIN TYPES

We have looked at maximizing randomness for the important case of equi-joins. We build upon the results and mathematical tools developed in the previous section to present techniques for maximizing randomness for all other join types.

4.1 Outer Join & Self-Join

The expressions for the number of samples that need to be maximized in the case of equi-join, $\prod_{i=1}^z C_{mm_i}^{m_i}$ (for single stratum) and

$\prod_{j=1}^n \prod_{i=1}^z C_{mm_i}^{m_i}$ (for multiple strata), are clearly applicable in the case of outer joins, such as full outer join, left outer join, and right outer join, as well (as $0! = 1$). Therefore, the techniques used for equi-joins can be used for outer joins as well. As self-join involves using the same relation on both sides of the join, the equi-join allocation strategy will be applicable for it as well.

4.2 Non-Equi-Join

In non-equi-joins ($\neq$), a tuple can be joined with all non-matching tuples of other relations. The number of possible samples can be given by $\prod_{j=1}^n \prod_{i=1}^z \left(C_{mm_i}^{m_i} \prod_{jj=1 \& jj \neq j}^n \prod_{ii=1 \& ii \neq i}^z C_{mm_{ii}}^{m_{ii}^{jj}} \right)$. Interestingly, derivation in Appendix C shows that the allocation strategy for equi-joins works in this case as well.

4.3 Theta Join

In theta joins, tuples are joined using the provided condition over the join columns. We provide allocation strategies, with derivation in Appendix D, for theta joins for the common comparators, $<$, $\leq$, $>$, and $\geq$. Algorithm 2 provides the strategy for the $\leq$ comparator.

Algorithm 2 Maximize Randomness for Theta Joins ($\leq$)

1. Number the strata from 1 to n in *descending* order.
 2. Determine approximate value of A using binary search, so that $k = \sum_{j=1}^n \left(A^{\frac{1}{j}} \sum_{i=1}^z m_i^j \right)$.
 3. Allocate the sample size using proportional allocation for a stratum j , $k^j = A^{\frac{1}{j}} \times \sum_{i=1}^z m_i^j$.
 4. Allocate k^j between the relations, $mm_i^j = k^j \times \frac{m_i^j}{\sum_{i=1}^z m_i^j}$.
-

Our derivation also shows that the only difference in the algorithm, when using the $<$ comparator instead of $\leq$ operator, is in finding A , such that $k = \sum_{j=1}^n \left(A^{\frac{1}{j-1}} \sum_{i=1}^z m_i^j \right)$. Other steps – 1, 3, and 4 – remain the same. These results can be extended for the comparators, $>$ and $\geq$, by reversing the strata order in Algorithm 2.

4.4 Cross Join

Cross join involves performing a cartesian product between the relations. The number of possible samples can be given by $\prod_{i=1}^z C_{|S_i|}^{|R_i|}$. This expression can be directly framed into the expression for maximizing randomness for a single stratum, $\prod_{i=1}^z C_{mm_i}^{m_i}$. The equation for space constraint can be reframed similarly as well. Hence, we use proportional allocation here, with the sample size given by $|S_i| = \frac{|R_i|}{\sum_{j=1}^n |R_j|}$.

5 NON-CORRELATED SAMPLING

We have looked at techniques to maximize randomness of *correlated* samples of joins. Here, we look at its complementary use case and present its problem statement as follows: *For the case where only non-correlated samples are acceptable, determine the sample allocation strategy to minimize sample size for the specified join sampling rate.* This is known to be a hard problem due to sample inflation (Section 1.1.2). Similar to Chaudhuri et al. [8], we consider different availabilities of statistics and indexes. We look at some of the issues in their state-of-the-art algorithms, and provide enhancements to them. We then provide an algorithm for the case where statistics and indexes are available over both relations (STRATJOIN).

5.1 Enhancement to Group-Sample

Group-Sample is the state-of-the-art algorithm for the case of statistics being available over one of the relations. We demonstrate its shortcoming (Section 5.1.1), provide an algorithm that rectifies it (Section 5.1.2), show the theoretical proof of its correctness (Section 5.1.3), provide its time (Section 5.1.4) and space (Section 5.1.5) complexities, and discuss a major enhancement if sorting were possible (Section 5.1.6).

5.1.1 Issues in Group-Sample. We briefly describe *Group-Sample* – please refer [8] for details. Using the statistics over R_2 , *Group-Sample* samples R_1 in a streaming fashion, by weighting each tuple in R_1 by the number of tuples in R_2 that can join with it, resulting in S_1 . Next, it joins S_1 with R_2 , generating a group of size $m_2(t_1.A)$ for each sampled tuple t_1 in S_1 . Finally, it chooses a single tuple from every group in a streaming fashion.

Joining S_1 with R_2 will result in the intermediate materialized data having a size of $f \times \sum_{a \in \text{Strata}} m_1(a) \times m_2(a) \times m_2(a)$ on average. This can result in the intermediate data size exceeding the size of join between non-sampled relations – $\sum_{a \in \text{Strata}} m_1(a) \times m_2(a)$ – which renders sampling counter-productive. Further, a scan is then needed over the intermediate materialized data to choose random tuples from each group, at a time complexity of $O(f \times \sum_{a \in \text{Strata}} m_1(a) \times m_2(a) \times m_2(a))$.

5.1.2 GROUP-SAMPLE-ENHANCED. Algorithm 3 eliminates the need to materialize large intermediate data and scan it. It can be extended to the case of multiple relations, by repeating steps 2 (b) through 2 (f) for the relations having statistics.

5.1.3 Proof of Correctness. We show that the probability of t_1 joining with any of the $m_2(t_1.A)$ tuples equals $\frac{1}{m_2(t_1.A)}$. Let the

Algorithm 3 GROUP-SAMPLE-ENHANCED

1. Obtain a with-replacement sample S_1 of R_1 , in a streaming fashion, by weighting a tuple t by $m_2(t.A)$.
 2. While tuples of S_1 are streaming by do:
 - (a) Extract the next tuple, t_1 , of S_1 .
 - (b) Start a scan of R_2 . Set i to 0.
 - (c) Extract the next tuple, t_2 , of R_2 .
 - (d) If $t_1.A \neq t_2.A$, go to Step 2 (c). Otherwise, increment i . Try to join t_1 with t_2 using $\text{Bernoulli}\left(\frac{1}{m_2(t_1.A) - i + 1}\right)$.
 - (e) If join is successful, output the joined tuple, and go to Step 2 (a).
 - (f) If join is unsuccessful, go to Step 2 (c).
-

i^{th} tuple amongst $m_2(t_1.A)$ tuples of R_2 be denoted by $m_2(t_1.A)[i]$. The probability of joining it with t_1 will be

$$\begin{aligned}
 & P(\text{reject } m_2(t_1.A)[1]) \times P(\text{reject } m_2(t_1.A)[2]) \cdots \\
 & \times P(\text{reject } m_2(t_1.A)[i-1]) \times P(\text{accept } m_2(t_1.A)[i]) \\
 & = \left(1 - \frac{1}{m_2(t_1.A)}\right) \times \left(1 - \frac{1}{m_2(t_1.A) - 1}\right) \cdots \\
 & \times \left(1 - \frac{1}{m_2(t_1.A) + 1 - (i-1)}\right) \times \frac{1}{m_2(t_1.A) - i + 1} \\
 & = \frac{m_2(t_1.A) - 1}{m_2(t_1.A)} \times \frac{m_2(t_1.A) - 2}{m_2(t_1.A) - 1} \cdots \times \frac{m_2(t_1.A) - i + 1}{m_2(t_1.A) - i + 2} \\
 & \times \frac{1}{m_2(t_1.A) - i + 1} = \frac{1}{m_2(t_1.A)}
 \end{aligned}$$

5.1.4 Time Complexity. As the initial sampling step is identical in both *Group-Sample* and *Group-Sample-Enhanced*, with a time complexity of $f \times \sum_{a \in \text{Strata}} m_1(a) \times m_2(a)$, we provide the average time complexity of our join step. First, we look at the probability of the join happening by the i^{th} tuple.

$$\begin{aligned}
 & P(\text{join does not occur by } m_2(t_1.A)[i]^{\text{th}} \text{ tuple}) \\
 & = P(\text{reject } m_2(t_1.A)[1]) \cdots \times P(\text{reject } m_2(t_1.A)[i]) \\
 & = \left(1 - \frac{1}{m_2(t_1.A)}\right) \times \left(1 - \frac{1}{m_2(t_1.A) - 1}\right) \cdots \\
 & \times \left(1 - \frac{1}{m_2(t_1.A) - i + 1}\right) \\
 & = \frac{m_2(t_1.A) - 1}{m_2(t_1.A)} \times \frac{m_2(t_1.A) - 2}{m_2(t_1.A) - 1} \cdots \times \frac{m_2(t_1.A) - i}{m_2(t_1.A) - i + 1} \\
 & = \frac{m_2(t_1.A) - i}{m_2(t_1.A)} \\
 & P(\text{join occurs by } m_2(t_1.A)[i]^{\text{th}} \text{ tuple}) \\
 & = 1 - P(\text{join does not occur by } m_2(t_1.A)[i]^{\text{th}} \text{ tuple}) \\
 & = \frac{i}{m_2(t_1.A)}
 \end{aligned}$$

The probability of a tuple joining by the halfway stage will be $\frac{1}{2}$, as $i = \frac{m_2(t_1.A)}{2}$. Thus, on average, half the tuples of a stratum will be accessed before the join occurs. As a result, assuming tuples are present in a random order in R_2 , on average, half the tuples in R_2 will be accessed for joining with every tuple from S_1 .

5.1.5 Space Complexity. Joining a tuple from S_1 with a random tuple from its corresponding stratum in R_2 does not materialize any intermediate data. Hence, in addition to the space required for S_1 and the output, $f \times \sum_{a \in \text{strata}} m_1(a) \times m_2(a)$, we will only need to keep a count of the number of tuples belonging to the stratum that have been accessed so far (i in Algorithm 3) in the join step.

5.1.6 Effect of Sorted R_2 . In the case that R_2 is not sorted, the time complexity of joining S_1 with R_2 in *Group-Sample* will be $O(f \times (\sum_{a \in \text{strata}} m_1(a) \times m_2(a)) \times |R_2|)$. While *GROUP-SAMPLE-ENHANCED* improves upon it, it still needs to perform a scan accessing half the tuples of R_2 to join with every tuple from S_1 . If it were possible to sort R_2 , the time complexities of both *Group-Sample* and *GROUP-SAMPLE-ENHANCED* will be greatly reduced – the number of tuples from R_2 that need to be accessed during the join step reducing from $O(|R_2|)$ to $O(\log(|R_2|) + m_2(t_1.a))$ for stratum a .

5.2 Enhancement to Frequency-Partition-Sample

Frequency-Partition-Sample is applicable for the case where statistics are available for the larger strata of a single relation. It uses *Group-Sample* for such strata. The strata lacking in statistics are first joined and then the output sampled. The space and time complexities are dominated by the larger strata. The shortcomings of *Group-Sample* will affect *Frequency-Partition-Sample* as well – the intermediate data can be expected to be large since *Group-Sample* is applied over the larger strata. Hence, our *FREQUENCY-PARTITION-SAMPLE* algorithm uses *GROUP-SAMPLE-ENHANCED* instead of *Group-Sample* for larger strata, and follows *Frequency-Partition-Sample*’s approach for smaller strata.

5.3 Enhancement to Stream-Sample

Stream-Sample is designed for the case where one of the relations, R_2 , has access to indexes and statistics. First, a with-replacement random sample, S_1 , is constructed over R_1 , by setting weight of a tuple t_1 to $m_2(t_1.A)$. Next, as the tuples of S_1 are streaming by, a tuple t_1 is joined with one of the random tuples from $m_2(t_1.A)$.

If $|S_1|$ is materialized, it will have a size $f \times |R_1| \propto |R_2|$ on average, which can be larger than $|R_1|$, rendering sampling counter-productive. This problem occurs at a stratum level – the sampling rate of a stratum a in R_1 will be $f_1(a) = \frac{f \times m_1(a) \times m_2(a)}{m_1(a)} = f \times m_2(a)$. Clearly, whenever $f_1(a) > 1$, sampling will be counter-productive – it will be prudent to not sample such strata, and only sample if it reduces the stratum size. Note that such an approach can be used to reduce the sample size in *GROUP-SAMPLE-ENHANCED* as well.

5.4 STRATJOIN – Sampling Both Relations

We now provide an algorithm for the case where indexes and statistics are available over both relations (Algorithm 4). A simple random sample of the desired size is generated for each stratum, resulting in the output being a stratified random sample. *STRATJOIN* can be easily extended to multiple relations.

6 EXPERIMENTAL EVALUATION

Our experiments were implemented in Java 8, and were run on an Ubuntu Linux 14.04.1 LTS system with a 24-core 2.4GHz Intel

Algorithm 4 STRATJOIN

1. *Sample:* For every stratum of both relations, create a with-replacement sample of size $f \times m_1(a) \times m_2(a)$ if $f \times m_{\text{other}}(a) \leq 1$ for a stratum a . Otherwise, use the entire stratum in the sample.
 2. *Join:* For a stratum, depending on whether none of the relations, one of the relations, or both the relations are sampled, join the samples as explained below.
 - (a) If both strata are sampled, randomly choose and join sampled tuples. Use a sampled tuple only once.
 - (b) If a single relation is sampled, while tuples from the sampled stratum are streaming by, join them with a randomly chosen tuple from the non-sampled stratum.
 - (c) If neither relation is sampled, join random tuples with-replacement from both strata till sampling rate is met.
-

Xeon CPU, 256GB DDR3 @ 1866 MHz memory, and a 500GB @ 7200 RPM disk. Section 6.1 studies effectiveness of our techniques at maximizing the number of samples. Section 6.2 compares randomness, and error in the presence of correlation between join and measure columns, between different join techniques. Our experiments use multiple datasets to illustrate the corresponding use cases well. Their size might be limited in some cases to determine the non-sampled ground truth to compare against.

6.1 Allocation Error

This section looks at the effectiveness of our sample allocation techniques at maximizing the number of samples in equi-joins. We validate our allocation strategy for both single stratum and multiple strata being involved in the join. To evaluate our accuracy, we find the best solution by searching through all possible allocations. This is computationally expensive and restricts size of datasets. We used the following metrics: mean squared error – $\frac{\sum_{i=1}^n (Y_i - \hat{Y}_i)^2}{n}$; mean squared relative error – $\frac{\sum_{i=1}^n \left(\frac{Y_i - \hat{Y}_i}{Y_i} \right)^2}{n}$; and maximum difference – $\text{Max}(\{|Y_i - \hat{Y}_i| : i \in [1, n]\})$; where Y_i is the actual value for optimal allocation found using brute-force search, $\hat{Y}_i$ is our predicted value, and n represents the number of observations.

6.1.1 Single Stratum Partitioning. Since our approach might not result in the optimal allocation (Section 3.6), we determine its accuracy by trying different relation and strata sizes, using 3 relations. The total *population size*, $|R_1| + |R_2| + |R_3|$, was varied from 150 to 1000. For each population size, all possible assignments to $|R_1|$, $|R_2|$, and $|R_3|$ were tried. We ensured that each relation had a minimum of 5 tuples. Sample size was varied from 15 to 100. We found the optimal allocation through brute-force search – searching for larger population and sample sizes was prohibitively expensive. Table 3 shows that all error metrics are low – validating our effectiveness in practice for the single stratum case.

6.1.2 Multiple Strata Partitioning. We again determined accuracy of our allocation strategy by varying the population and sample sizes. We used 3 relations having 3 strata each, with the minimum

Table 3: Error in Single Stratum Partitioning

Population	MSE	MSRE	Maximum Difference
150	0.1109	0.0069	1
200	0.1197	0.0094	2
300	0.1377	0.0135	2
400	0.1506	0.0162	2
500	0.1598	0.0180	2
600	0.1666	0.0196	2
700	0.1720	0.0208	2
800	0.1761	0.0218	2
900	0.1794	0.0226	2
1000	0.1821	0.0232	2

Table 4: Error in Multiple Strata Partitioning

Population	MSE	MSRE	Maximum Difference
40	0.1439	0.0183	1
45	0.1714	0.0221	1
50	0.1097	0.0081	1
55	0.1087	0.0128	1

stratum size being 3. In a similar fashion as above, the total population size was varied from 40 to 55 and the sample size was varied from 27 to 30 – again, these were the upper limits for which we could find the optimal solution through brute-force search (within a day in this case). Table 4 shows that all error metrics are low.

6.2 Comparison with Correlated Samplers

We also looked at different correlated sampling techniques in the presence of correlation between join and measure columns. Two relations with 8000 rows were generated – we have limited the data size to where a measurement could be obtained in around 10 minutes. Each relation had 2 columns – a join column and a measure column. We used two common distributions, Gaussian and Zipfian, to model the column values. We obtained similar results using the two – results using Gaussian distribution have been presented. The join columns were sampled from $\mathcal{N}(\mu = 100, \sigma = 10)$. The measure columns were sampled from $\mathcal{N}(\mu = 200, \sigma = 10)$ and $\mathcal{N}(\mu = 300, \sigma = 10)$, so that we obtain non-overlapping values. Correlation between join and measure columns was varied from 0.1 to 0.9³. We did not find a significant change for differing correlation values – results using a correlation of 0.5 have been presented. 50 runs were performed, with different relations and samples being created in each run – the median of the measurements has been presented. We use a modification of the L^1 -norm, the *average relative error* $\left(\sum_{i=1}^n \frac{|Y_i - \hat{Y}_i|}{Y_i}\right)$, to better represent the relative error.

In our context, *Ripple Join*, *SMS Join*, and *Wander Join* have identical semantics, as do different cluster samplers such as *End-Biased Sampling*, *Correlated Sampling*, and *Universe Sampler*. In the figures, *Stratified Random* represents stratified random sampling of the join, which results in a non-correlated sample. *Stratified Random* suffers from sample inflation – it provides us with the results for the best-case scenario from the perspective of correlation.

³A column having correlation ρ with X is generated from columns X and Y as $\rho \odot X \oplus (1 - \rho^2)^{1/2} \odot Y$. The $\odot$ operator represents multiplication of a number with every element of an array, and $\oplus$ operator represents pairwise addition of two arrays.

Figure 1: *Stratified Random* results in the most possible samples. *MAXRANDJOIN* always results in more samples than *Ripple Join*, while *Universe Sampler* results in the fewest.

Figure 2: *MAXRANDJOIN* has the least error amongst correlated samplers. Interestingly, *STRATJOIN* has lesser error than *MAXRANDJOIN* for the smallest sample size.

6.2.1 Number of Possible Samples. We look at the number of possible samples as a result of different techniques (Figure 1), i.e. given the samples, we find $\prod_{j=1}^n \prod_{i=1}^2 C_{mm_i}^{m_i^j}$. Note that the Y-axis is presented in *loglog* scale (base e) to better illustrate the growth pattern. *MAXRANDJOIN* consistently provides more samples than *Ripple Join* and *Universe*. *STRATJOIN* and *Stratified Random* result in the maximum possible value.

6.2.2 Sampling Error. We look at the average relative error for the measure sum (over rows) of product of the measure columns (Figure 2) – other column functions that we experimented with, such as sum of measure columns and product of measure columns, gave us similar results. Output for *Universe Sampler* has not provided as it resulted in larger errors – including it would obscure differences between the other algorithms from being seen. Ground truth for the measure value was determined after joining the relations (without sampling). *STRATIFIED RANDOM* had the least error, while *MAXRANDJOIN* consistently had lower error than *Ripple Join*. Interestingly, *STRATJOIN* had lower error than *MAXRANDJOIN* for the smallest sample size (125) indicating perhaps that avoiding correlation at the expense of smaller output size, which *STRATJOIN* does, might be the better option at lower sampling rates.

Figure 3: MAXRANDJOIN is resistant to white Gaussian noise with low standard deviation. At higher sample sizes, it can handle binomial perturbations to join count information.

6.2.3 Effect of Noise on MAXRANDJOIN. We investigate MAXRANDJOIN’s resistance to common types of noise – white Gaussian and Binomial (Figure 3). The standard deviation of the Gaussian noise was varied from 40 to 320. MAXRANDJOIN’s resistance starts waning with increasing standard deviation in the Gaussian noise. It can handle Binomial error for larger sample sizes.

6.3 Non-correlated Sampling

We have presented extensive theoretical results describing the improvements provided by GROUP-SAMPLE-ENHANCED over *Group-Sample* in Section 5.1. In this section, we demonstrate those benefits empirically, by looking at the number of intermediate data tuples created and the time taken to obtain the join result.

6.3.1 Experimental Setup. We use a similar setup to that used by Chaudhuri et al. [8]. Four tables were generated with 10000 tuples each. The join column in each table had counts modeled using a Zipfian distribution. The parameter z of the Zipfian distribution was varied from 0 to 3. Four other tables with 100000 tuples were generated similarly. Each row consists of three columns – RID (integer), JoinKey (integer), and Padding (integer). We have implemented the algorithms using our custom in-memory join system. By default, we discard the first run of each experiment and report the mean of the following three runs (the runs were nearly identical for all experiments). In the figures, LHS refers to the relation with 10000 tuples while RHS refers to the one with 100000 tuples. In the legends, the numbers following the algorithm name refer to the LHS and RHS skews (when available), respectively.

6.3.2 Intermediate Data Size. We look at the number of intermediate tuples that would be created by *Group-Sample* and GROUP-SAMPLE-ENHANCED for varying sampling rates (Figure 4) and RHS skews (Figure 5). They show that GROUP-SAMPLE-ENHANCED, whose intermediate data size is determined by the join size, can result in the intermediate data size being multiple orders of magnitude smaller than that of *Group-Sample*. Both algorithms exhibit a linear increase in the intermediate data size with increasing sampling rate. When an increase in the RHS skew does not increase the join size (LHS skew = 0), while GROUP-SAMPLE-ENHANCED intermediate data size does not increase, that of *Group-Sample* does.

Figure 4: While an increase in the sampling rate results in a linear increase in the intermediate data size, the size required by GROUP-SAMPLE-ENHANCED is multiple orders of smaller than that of *Group-Sample* due to the usage of reservoir sampling-based techniques.

Figure 5: GROUP-SAMPLE-ENHANCED intermediate data size depends on the join size, while *Group-Sample* sample size does not and can increase regardless (LHS skew = 0).

6.3.3 Execution Time. We also compared execution times for the two algorithms for varying sampling rates and RHS skews (Figures 6 and 7). The time limit for a run was capped at 4 hours, which resulted in some of the results being unavailable. The execution time of GROUP-SAMPLE-ENHANCED was usually an order of magnitude lower than that of *Group-Sample*.

7 CONCLUSION & FUTURE WORK

Join sampling is an interesting and important area of research. We have presented techniques to sample joins, in the context of correlated and non-correlated sampling – illustrating the benefits and drawbacks in doing so. We provided novel techniques for increasing randomness of joins when correlated samples are acceptable. We showed that our techniques to maximize join randomness were effective over varying population and sample sizes. While other correlated samplers are applicable only in the case of aggregate, approximable queries, our techniques are application-agnostic – they still result in the sample having a lower aggregate measure error when compared with other state-of-the-art correlated samplers. It affirms our intuition that increasing the number of possible samples reduces the correlation in the samples. In the context

Figure 6: GROUP-SAMPLE-ENHANCED provides results around an order of magnitude sooner. It is also able to finish execution within the time limit in one of the cases where *Group-Sample* is not ($f = 1\%$).

Figure 7: GROUP-SAMPLE-ENHANCED handles increasing skew better than *Group-Sample*.

of non-correlated sampling, we provided major improvements to the state-of-the-art algorithms. We also provided an algorithm for sampling both sides of a join in the presence of statistics and indexes over both. Finally, while our algorithms use statistics over join columns, this dependency cannot be done away with for high performance querying. While one approach for gathering such statistics has been to run regularly scheduled *cron jobs*, another recent one has been to collect statistics over first pass over the data and use them over the subsequent passes [7, 30].

Going forward, several interesting avenues of research are enabled by our findings. We wish to investigate applicability of search-based solutions to obtain theoretically perfect allocation [47]. We would like to develop an estimation framework for computing variance for our correlated sampling-based approaches [28]. We wish our work influences others to take correlation into consideration while designing algorithms. Finally, we hope other research areas that use correlated samples, in both computer science and statistics communities, use strategies to maximize sample randomness.

ACKNOWLEDGMENTS

We would like to thank Ymir Vigfusson, our research group, and the reviewers for their extremely insightful and helpful comments. This work has been supported by the National Science Foundation, under grants 1422977, 1453582, and 1527779.

A MAXIMIZING RANDOMNESS FOR SINGLE STRATUM EQUI-JOIN

The notation given in Table 1 is used in the Appendix – superscript and subscript denote the stratum index and the relation index, respectively. Superscript has been omitted in this section, since a single stratum is considered. The number of possible samples can be given by

$$f(mm_1, \dots, mm_z) = \prod_{i=1}^z C_{mm_i}^{m_i} = \prod_{i=1}^z \frac{m_i!}{mm_i!(m_i - mm_i)!} \quad (4)$$

The fixed sample size constraint can be given by

$$g(mm_1, mm_2, \dots, mm_z) = \sum_{i=1}^z mm_i = k^1 \quad (5)$$

g denotes the constraint function. Lagrange multipliers gives a system of z equations, with i^{th} ($i \in [1, z]$) equation obtained by taking partial derivative of equation 4 with respect to mm_i

$$f_{mm_i} = \left(\prod_{j=1 \& j \neq i}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \right) \times m_i! \times \frac{\partial}{\partial mm_i} \left(\frac{1}{mm_i! \times (m_i - mm_i)!} \right) \quad (6)$$

Using the basic Lagrange multipliers equation of $f_{mm_i} = \lambda \times g_{mm_i}$ and equation 5, we also get

$$f_{mm_i} = \lambda \times g_{mm_i} = \lambda \times \frac{\partial \left(\sum_{i=1}^z mm_i \right)}{\partial mm_i} = \lambda \quad (7)$$

Using the Gamma function to represent factorial, we get

$$(n!) = \Gamma'(n+1) = n! \times \left(-\gamma + \sum_{a=1}^n \frac{1}{a} \right) \quad (8)$$

where γ is Euler-Mascheroni constant. $\frac{\partial}{\partial mm_i} \left(\frac{1}{mm_i! \times (m_i - mm_i)!} \right)$ can be simplified as

$$\begin{aligned} & \left(mm_i! \times \frac{\partial((m_i - mm_i)!)}{\partial mm_i} \right) + \left((m_i - mm_i)! \times \frac{\partial(mm_i!)}{\partial mm_i} \right) \\ &= - \frac{mm_i! \times (m_i - mm_i)! \times \left(-\gamma + \sum_{a=1}^{m_i - mm_i} \frac{1}{a} \right) \times (-1)}{(mm_i! \times (m_i - mm_i)!)^2} \\ &= - \frac{(m_i - mm_i)! \times mm_i! \times \left(-\gamma + \sum_{a=1}^{mm_i} \frac{1}{a} \right)}{(mm_i! \times (m_i - mm_i)!)^2} \\ &= - \frac{mm_i! \times (m_i - mm_i)! \times \left(-\sum_{a=1}^{m_i - mm_i} \frac{1}{a} + \sum_{a=1}^{mm_i} \frac{1}{a} \right)}{(mm_i! \times (m_i - mm_i)!)^2} \\ &= - \frac{\left(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a} \right)}{mm_i! \times (m_i - mm_i)!} \end{aligned} \quad (9)$$

Plugging this into equation 6 and using equation 7, we get

$$\begin{aligned}\lambda &= \left(\prod_{j=1 \& j \neq i}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \right) \times m_i! \times \\ &\quad - \left(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a} \right) \\ &\quad \frac{mm_i! \times (m_i - mm_i)!}{mm_i! \times (m_i - mm_i)!} \\ &= - \left(\prod_{j=1}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \right) \times \left(\sum_{a=1}^{mm_i} \frac{1}{a} - \sum_{a=1}^{m_i - mm_i} \frac{1}{a} \right)\end{aligned}$$

The sum of harmonic series can be approximated by

$$\begin{aligned}\sum_{a=1}^{mm_i} \frac{1}{a} &= \ln(mm_i) + \gamma + \frac{1}{2 \times mm_i} - \frac{1}{12 \times mm_i^2} \\ &\approx \ln(mm_i) + \gamma\end{aligned}\quad (10)$$

for large mm_i , where γ is the Euler-Mascheroni constant. Using the above two equations, we get

$$\begin{aligned}-\lambda &= \prod_{j=1}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \times (\ln(mm_i) - \ln(m_i - mm_i)) \\ &= \prod_{j=1}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \times \ln\left(\frac{mm_i}{m_i - mm_i}\right)\end{aligned}\quad (11)$$

Note that the term $\left(\prod_{j=1}^z \frac{m_j!}{mm_j! \times (m_j - mm_j)!} \right)$ is present in all $i \in [1, z]$ equations – we denote it by C , which gives us

$$\frac{-\lambda}{C} = \ln\left(\frac{mm_i}{m_i - mm_i}\right)\quad (12)$$

Consider the equation for another relation p

$$\frac{-\lambda}{C} = \ln\left(\frac{mm_p}{m_p - mm_p}\right)\quad (13)$$

From the above two equations, we get

$$\frac{mm_i}{m_i - mm_i} = \frac{mm_p}{m_p - mm_p} = A\quad (14)$$

for a constant A . Therefore,

$$mm_i = \frac{A \times m_i}{1 + A}\quad (15)$$

Using equations 5 and 15, we get

$$k^1 = \sum_{i=1}^z mm_i = \sum_{i=1}^z \frac{A \times m_i}{1 + A} = \frac{A}{1 + A} \sum_{i=1}^z m_i$$

Using above two equations, we get

$$mm_i = \left(\frac{A}{1 + A} \right) \times m_i = \frac{k^1 \times m_i}{\sum_{a=1}^z m_a}\quad (16)$$

B MAXIMIZING RANDOMNESS FOR MULTIPLE STRATA EQUI-JOIN

The constraint of fixed sample size for equi-joins with multiple strata can be given as follows, given $k^j = \sum_{i=1}^z mm_i^j$.

$$k = \sum_{j=1}^n k^j = g(k^1, k^2 \dots k^n)\quad (17)$$

The number of possible samples can be given by

$$\begin{aligned}f(k^1, k^2 \dots k^n) &= \prod_{j=1}^n \prod_{i=1}^z C_{mm_i^j}^{m_i^j} \\ &= \prod_{j=1}^n \prod_{i=1}^z \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \\ &= \prod_{j=1}^n \prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \times \frac{m_z^j!}{mm_z^j! \times (m_z^j - mm_z^j)!} \\ &= \prod_{j=1}^n \prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \times \\ &\quad \frac{m_z^j!}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!}\end{aligned}\quad (18)$$

Equation 16 provides optimal allocation amongst relations given total allocation to a stratum. Using it in the above equation results in the number of samples increasing or staying the same. Hence, k^j ($j \in [1, n]$) are the variables in this case. Lagrange multipliers gives us the j^{th} equation as

$$\begin{aligned}f_{k_j} &= \prod_{j=1}^n \prod_{i=1}^{z-1} \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \times m_z^j! \times \\ &\quad \frac{\partial}{\partial k^j} \left(\frac{1}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!} \right)\end{aligned}\quad (19)$$

We simplify $\frac{\partial}{\partial k^j} \left(\frac{1}{(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!} \right)$ as,

$$\begin{aligned}&\frac{-1}{((k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!)^2} \times \\ &\left(k^j - \sum_{i=1}^{z-1} mm_i^j \right)! \times \left(-\gamma + \sum_{l=1}^{k^j - \sum_{i=1}^{z-1} mm_i^j} \frac{1}{a} \right) \times \\ &\left(m_z^j - \left(k^j - \sum_{i=1}^{z-1} mm_i^j \right) \right)! \times \left(k^j - \sum_{i=1}^{z-1} mm_i^j \right)! \times \\ &\left(m_z^j - \left(k^j - \sum_{i=1}^{z-1} mm_i^j \right) \right)! \times \left(-\gamma + \sum_{l=1}^{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)} \frac{1}{a} \right) \\ &= \frac{-1 \times (k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!}{((k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!)^2} \times \\ &\left(-\gamma + \sum_{l=1}^{k^j - \sum_{i=1}^{z-1} mm_i^j} \frac{1}{a} + \gamma - \sum_{l=1}^{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)} \frac{1}{a} \right) \\ &= \frac{\ln(k^j - \sum_{i=1}^{z-1} mm_i^j) + \gamma - \ln(m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)) - \gamma}{-(k^j - \sum_{i=1}^{z-1} mm_i^j)! \times (m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j))!}\end{aligned}$$

$$= \frac{\ln \left(\frac{m_z^j - (k^j - \sum_{i=1}^{z-1} mm_i^j)}{k^j - \sum_{i=1}^{z-1} mm_i^j} \right)}{mm_z^j! \times (m_z^j - mm_z^j)!} \quad (20)$$

Substituting above equation in equation 19, we get

$$f_{k_j} = \prod_{j=1}^n \prod_{i=1}^z \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \times \ln \left(\frac{m_z^j - k^j + \sum_{i=1}^{z-1} mm_i^j}{k^j - \sum_{i=1}^{z-1} mm_i^j} \right)$$

From the standard Lagrange multipliers equation of $f_{k_j} = \lambda \times g_{k_j}$ and equation 17, we get

$$f_{k_j} = \lambda \times g_{k_j} = \lambda \times \frac{\partial \left(\sum_{j=1}^n k^j \right)}{\partial k_j} = \lambda \quad (21)$$

From the above two equations, we get

$$\lambda = \left(\prod_{j=1}^n \prod_{i=1}^z \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \right) \times \ln \left(\frac{m_z^j}{k^j - \sum_{i=1}^{z-1} mm_i^j} - 1 \right)$$

Consider the equation for another stratum p

$$\lambda = \left(\prod_{j=1}^n \prod_{i=1}^z \frac{m_i^j!}{mm_i^j! \times (m_i^j - mm_i^j)!} \right) \times \ln \left(\frac{m_z^p}{k^p - \sum_{i=1}^{z-1} mm_i^p} - 1 \right)$$

From the above 2 equations, given a constant A, we get

$$\frac{k^j - \sum_{i=1}^{z-1} mm_i^j}{m_z^j} = \frac{k^p - \sum_{i=1}^{z-1} mm_i^p}{m_z^p} = A \quad (22)$$

Using equation 16 and the above equation, we get

$$A = \frac{k^j - \sum_{i=1}^{z-1} \frac{k^j \times m_i^j}{\sum_{a=1}^z m_a^j}}{m_z^j} = \frac{k^j \times \left(1 - \sum_{i=1}^{z-1} \frac{m_i^j}{\sum_{a=1}^z m_a^j} \right)}{m_z^j} \quad (23)$$

$$= \frac{k^j \times \left(\sum_{a=1}^z m_a^j - \sum_{i=1}^{z-1} m_i^j \right)}{m_z^j \times \sum_{a=1}^z m_a^j} = \frac{k^j}{\sum_{a=1}^z m_a^j}$$

$$k^j = A \times \sum_{a=1}^z m_a^j \quad (24)$$

Using above equation and equation 17, we get

$$k = \sum_{j=1}^n \left(A \times \sum_{a=1}^z m_a^j \right) \quad (25)$$

$$A = \frac{k}{\sum_{j=1}^n \sum_{a=1}^z m_a^j} \quad (26)$$

Using equations 24 and 26, we get

$$k^j = \frac{k \times \sum_{i=1}^z m_i^j}{\sum_{j=1}^n \sum_{i=1}^z m_i^j} \quad (27)$$

C MAXIMIZING RANDOMNESS FOR NON-EQUI-JOINS

Number of samples for non-equi-joins can be given by

$$\prod_{j=1}^n \prod_{i=1}^z \left(C_{mm_i^j}^{m_i^j} \prod_{jj=1}^n \prod_{ii=1 \text{ \& } ii \neq i}^z C_{mm_{ii}^{jj}}^{m_{ii}^{jj}} \right)$$

$$= \prod_{j=1}^n \prod_{i=1}^z \left(C_{mm_i^j}^{m_i^j} \frac{C_{mm_i^j}^{m_i^j} \prod_{j=1}^n \prod_{i=1}^z C_{mm_i^j}^{m_i^j}}{\prod_{a=1}^n C_{mm_a^j}^{m_a^j} \prod_{b=1}^z C_{mm_b^j}^{m_b^j}} \right)$$

$$= \prod_{j=1}^n \prod_{i=1}^z \left(C_{mm_i^j}^{m_i^j} \frac{C_{mm_i^j}^{m_i^j} K}{\prod_{a=1}^n C_{mm_a^j}^{m_a^j} \prod_{b=1}^z C_{mm_b^j}^{m_b^j}} \right)$$

$$= K^{nz} \prod_{j=1}^n \prod_{i=1}^z \left(C_{mm_i^j}^{m_i^j} \frac{C_{mm_i^j}^{m_i^j}}{\prod_{a=1}^n C_{mm_a^j}^{m_a^j} \prod_{b=1}^z C_{mm_b^j}^{m_b^j}} \right)$$

$$\dots \text{ where } K = \prod_{j=1}^n \prod_{i=1}^z C_{mm_i^j}^{m_i^j}$$

$$= K^{nz} \prod_{j=1}^n \prod_{i=1}^z C_{mm_i^j}^{m_i^j} C_{mm_i^j}^{m_i^j} \prod_{j=1}^n \prod_{i=1}^z \frac{1}{\prod_{a=1}^n C_{mm_a^j}^{m_a^j} \prod_{b=1}^z C_{mm_b^j}^{m_b^j}}$$

$$= K^{nz} K^2 \frac{1}{\prod_{j=1}^n \prod_{i=1}^z \prod_{a=1}^n C_{mm_a^j}^{m_a^j}} \frac{1}{\prod_{j=1}^n \prod_{i=1}^z \prod_{b=1}^z C_{mm_b^j}^{m_b^j}}$$

$$= K^{nz} K^2 \frac{1}{\prod_{j=1}^n K} \frac{1}{\prod_{i=1}^z K} = K^{nz} K^2 \frac{1}{K^n} \frac{1}{K^z} = K^{nz+2-n-z}$$

K can be maximized using proportional allocation (Appendix B). Hence, we use the equi-join allocation strategy here as well.

D MAXIMIZING RANDOMNESS FOR THETA JOINS

For $\leq$ comparator, the number of samples can be given by

$\prod_{i=1}^z \prod_{j=1}^n C_{mm_i^j}^{m_i^j} \prod_{ii=1 \text{ \& } ii \neq i}^z \prod_{jj=1}^n C_{mm_{ii}^{jj}}^{m_{ii}^{jj}}$ and its log by

$$\sum_{i=1}^z \sum_{j=1}^n \left(\ln \left(C_{mm_i^j}^{m_i^j} \right) + \sum_{ii=1 \text{ \& } ii \neq i}^z \sum_{jj=1}^n \ln \left(C_{mm_{ii}^{jj}}^{m_{ii}^{jj}} \right) \right)$$

$$= \sum_{i=1}^z \sum_{j=1}^n \ln \left(C_{mm_i^j}^{m_i^j} \right) + \sum_{i=1}^z \sum_{j=1}^n \sum_{ii=1 \text{ \& } ii \neq i}^z \sum_{jj=1}^n \ln \left(C_{mm_{ii}^{jj}}^{m_{ii}^{jj}} \right)$$

$$= \sum_{i=1}^z \sum_{j=1}^n \ln \left(C_{mm_i^j}^{m_i^j} \right) + \sum_{i=1}^z \sum_{j=1}^n \sum_{ii=1}^z \sum_{jj=1}^n \ln \left(C_{mm_{ii}^{jj}}^{m_{ii}^{jj}} \right) - \sum_{i=1}^z \sum_{j=1}^n \sum_{jj=1}^n \ln \left(C_{mm_i^{jj}}^{m_i^{jj}} \right)$$

$$= \sum_{i=1}^z \sum_{j=1}^n \ln \left(C_{mm_i^j}^{m_i^j} \right) + (z-1) \times \sum_{i=1}^z \sum_{j=1}^n \sum_{jj=1}^n \ln \left(C_{mm_i^{jj}}^{m_i^{jj}} \right)$$

$$= \sum_{i=1}^z \sum_{j=1}^n \ln \left(C_{mm_i^j}^{m_i^j} \right) + (z-1) \times \sum_{i=1}^z \sum_{j=1}^n (n-j+1) \ln \left(C_{mm_i^j}^{m_i^j} \right)$$

$$\begin{aligned}
&= (z-1) \times \sum_{i=1}^z \sum_{j=1}^n (n-j+1 + \frac{1}{z-1}) \times \ln \left(C_{mm_i^j}^{m_i^j} \right) \\
&\approx (z-1) \times \sum_{i=1}^z \sum_{j=1}^n (n-j+1) \times \ln \left(C_{mm_i^j}^{m_i^j} \right) \\
&= (z-1) \times \sum_{i=1}^z \sum_{j=1}^n j \times \ln \left(C_{mm_i^j}^{m_i^j} \right)
\end{aligned}$$

Strata order was reversed in the last step. We now have to minimize $\sum_{i=1}^z \sum_{j=1}^n j \times (\ln(mm_i^j!) + \ln((m_i^j - mm_i^j)!))$, with $k = \sum_{i=1}^z \sum_{j=1}^n mm_i^j = g(mm_1^1 \dots mm_1^n)$. Lagrange multipliers gives us the $\{i, j\}^{th}$ ($i \in [1, z], j \in [1, n]$) equation as

$$\begin{aligned}
f_{mm_i^j} &= \lambda \frac{\partial g}{\partial mm_i^j} = \lambda = \frac{\partial}{\partial mm_i^j} (j(\ln(mm_i^j!) + \ln((m_i^j - mm_i^j)!))) \\
&= j \left(\frac{1}{mm_i^j!} mm_i^j! \left(-\gamma + \sum_{a=1}^{m_i^j} \frac{1}{a} \right) \right) + \\
& j \left(\frac{1}{(m_i^j - mm_i^j)!} (m_i^j - mm_i^j)! \left(-\gamma + \sum_{a=1}^{m_i^j - mm_i^j} \frac{1}{a} \right) (-1) \right) \\
&= j \left(\sum_{a=1}^{mm_i^j} \frac{1}{a} - \sum_{a=1}^{m_i^j - mm_i^j} \frac{1}{a} \right) = j(\ln(mm_i^j) + \gamma - \ln(m_i^j - mm_i^j) - \gamma) \\
&= j \left(\ln \left(\frac{mm_i^j}{m_i^j - mm_i^j} \right) \right) = -j \left(\ln \left(\frac{m_i^j}{mm_i^j} - 1 \right) \right) \approx -j \left(\ln \left(\frac{m_i^j}{mm_i^j} \right) \right)
\end{aligned}$$

For every equation, $(m_i^j/mm_i^j)^j$ will be constant. For any stratum as j is constant, m_i^j/mm_i^j will be constant – showing that proportional allocation should be used to allocate samples for a stratum between relations. Therefore, $mm_i^j = k^j m_i^j / \sum_{i=1}^z m_i^j$. For all strata, using above two equations, $(k^j / \sum_{i=1}^z m_i^j)^j$ will be constant (denoted by A). Therefore, $k^j = A^{\frac{1}{j}} \sum_{i=1}^z m_i^j$ and $k = \sum_{j=1}^n k^j = \sum_{j=1}^n A^{\frac{1}{j}} \sum_{i=1}^z m_i^j$. A is the only unknown and can be found using binary search within error. Finding A lets us find k^j , which leads us to finding mm_i^j . For the $<$ comparator, following a similar line of reasoning, the only difference in the algorithm is in determining A , given by $k^j = A^{\frac{1}{j-1}} \sum_{i=1}^z m_i^j$.

REFERENCES

- [1] Intel Lab Data. In URL: <http://db.csail.mit.edu/labdata/labdata.html>.
- [2] Swarup Acharya, Phillip B Gibbons, Viswanath Poosala, and Sridhar Ramaswamy. 1999. Join Synopses for Approximate Query Answering. In *SIGMOD*.
- [3] Sameer Agarwal, Barzan Mozafari, and others. 2013. BlinkDB: Queries With Bounded Errors and Bounded Response Times on Very Large Data. In *EuroSys*.
- [4] Mohammed Al-Kateb, Byung Suk Lee, and X Sean Wang. 2007. Adaptive-size Reservoir Sampling Over Data Streams. In *SSDBM*.
- [5] Spyros Blanas, Yanan Li, and Jignesh M Patel. 2011. Design and Evaluation of Main Memory Hash Join Algorithms for Multi-core CPUs. In *SIGMOD*.
- [6] Spyros Blanas, Jignesh M Patel, Vuk Ercegovic, and Jun Rao. 2010. A Comparison of Join Algorithms for Log Processing in Mapreduce. In *SIGMOD*.
- [7] Sunil Chakkapen, Thierry Cruanes, Benoit Dageville, and others. 2008. Efficient and Scalable Statistics Gathering for Large Databases in Oracle 11g. In *SIGMOD*.
- [8] Surajit Chaudhuri, Rajeev Motwani, and Vivek Narasayya. 1999. On Random Sampling over Joins. In *SIGMOD*.
- [9] Eunjoon Cho, Seth A Myers, and Jure Leskovec. 2011. Friendship and Mobility: User Movement in Location-based Social Networks. In *SIGKDD*.
- [10] William G Cochran. 2007. Sampling Techniques. John Wiley & Sons.
- [11] José A Diaz-Garcia and Ma Magdalena Garay-Tapia. 2005. Optimum Allocation in Stratified Surveys. In *Comunicaciones del Cimat*.
- [12] Alin Dobra, Chris Jermaine, Florin Rusu, and Fei Xu. 2009. Turbo-Charging Estimate Convergence in DBO. In *VLDB*.
- [13] Cristian Estan and Jeffrey F Naughton. 2006. End-Biased Samples for Join Cardinality Estimation. In *ICDE*.
- [14] Sumit Ganguly, Phillip B Gibbons, Yossi Matias, and Avi Silberschatz. 1996. Bifocal Sampling for Skew-resistant Join Size Estimation. In *SIGMOD*.
- [15] Rainer Gemulla, Philipp Röscher, and Wolfgang Lehner. 2008. Linked Bernoulli Synopses: Sampling along Foreign Keys. In *SSDBM*.
- [16] Peter J Haas and others. 1999. Ripple Joins for Online Aggregation. In *SIGMOD*.
- [17] Peter J Haas and Peter J Haas. 1996. Hoeffding Inequalities for Join-Selectivity Estimation and Online Aggregation. In *Research Report*. IBM.
- [18] Peter J Haas, Jeffrey F Naughton, S Seshadri, and others. 1996. Selectivity and Cost Estimation for Joins Based on Random Sampling. In *JCSS*.
- [19] Dirk Habich, Wolfgang Lehner, and Alexander Hinneburg. 2005. Optimizing Multiple Top-K Queries over Joins.. In *SSDBM*.
- [20] Bingsheng He, Ke Yang, Rui Fang, Mian Lu, Naga Govindaraju, Qiong Luo, and Pedro Sander. 2008. Relational Joins on Graphics Processors. In *SIGMOD*.
- [21] Joseph M Hellerstein and others. 1997. Online Aggregation. In *SIGMOD*.
- [22] Prasanth Jayachandran and others. 2014. Combining User Interaction, Speculative Query Execution and Sampling in the DICE System. In *VLDB*.
- [23] Christopher Jermaine and others. 2006. The Sort-Merge-Shrink Join. In *TODS*.
- [24] Christopher Jermaine, Subramanian Arumugam, Abhijit Pol, and Alin Dobra. 2008. Scalable Approximate Query Processing with the DBO Engine. In *TODS*.
- [25] Christopher Jermaine, Alin Dobra, Subramanian Arumugam, Shantanu Joshi, and Abhijit Pol. 2005. A Disk-based Join with Probabilistic Guarantees. In *SIGMOD*.
- [26] Joseph B Kadane. 2005. Optimal Dynamic Sample Allocation among Strata. In *Journal of official Statistics*.
- [27] Niranjana Kamat, Prasanth Jayachandran, Karthik Tunga, and Arnab Nandi. 2014. Distributed and Interactive Cube Exploration. In *ICDE*.
- [28] Niranjana Kamat and Arnab Nandi. 2016. A Closer Look at Variance Implementations In Modern Database Systems. In *SIGMOD RECORD*.
- [29] Niranjana Kamat and Arnab Nandi. 2017. A Session-Based Approach to Fast-But-Approximate Interactive Data Cube Exploration. In *TKDD*.
- [30] Srikanth Kandula, Anil Shanbhag, Aleksandar Vitorovic, and others. 2016. Quickr: Lazily Approximating Complex AdHoc Queries in BigData Clusters. In *SIGMOD*.
- [31] Stephen M Kent. 1994. Sloan Digital Sky Survey. In *Science with Astronomical Near-infrared Sky Surveys*.
- [32] Kyriaki Kitikidou. 2012. Optimizing Forest Sampling by using Lagrange Multipliers. In *AJOR*.
- [33] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2016. Wander Join: Online Aggregation via Random Walks. In *SIGMOD*.
- [34] Sharon L Lohr. 2010. Sampling: Design and Analysis.
- [35] Gang Luo and others. 2002. A Scalable Hash Ripple Join Algorithm. In *SIGMOD*.
- [36] Mohamed F Mokbel, Ming Lu, and Walid G Aref. 2004. Hash-merge Join: A Non-blocking Join Algorithm for Producing Fast and Early Join Results. In *ICDE*.
- [37] Supriya Nirakhiwale, Alin Dobra, and Christopher Jermaine. 2013. A Sampling Algebra for Aggregate Estimation. In *VLDB*.
- [38] Frank Olken. 1993. Random Sampling from Databases.
- [39] Chengjie Qin and Florin Rusu. 2014. PF-OLA: A High-performance Framework for Parallel Online Aggregation. In *Distributed and Parallel Databases*.
- [40] Suprio Ray, Bogdan Simion, Angela Demke Brown, and Ryan Johnson. 2014. Skew-resistant Parallel In-memory Spatial Join. In *SSDBM*.
- [41] Hyoseop Shin, Bongki Moon, and Sukho Lee. 2000. Adaptive Multi-stage Distance Join Processing. In *SIGMOD*.
- [42] Lefteris Sidiropoulos, Martin L Kersten, and Peter A Boncz. 2011. SciBORQ: Scientific Data Management with Bounds On Runtime and Quality. In *CIDR*.
- [43] Joshua Spiegel and Neoklis Polyzotis. 2009. TuG Synopses for Approximate Query Answering. In *TODS*.
- [44] Pandurang Sukhatme. 1957. *Sampling Theory of Surveys with Applications*. ISAS.
- [45] David Vengerov, Andre Cavalheiro Menck, Mohamed Zait, and Sunil P Chakkapen. 2015. Join Size Estimation Subject to Filter Conditions. In *VLDB*.
- [46] Jeffrey S Vitter. 1985. Random Sampling with a Reservoir. In *TOMS*.
- [47] Benjamin W Wah and Zhe Wu. 1999. The Theory of Discrete Lagrange Multipliers for Nonlinear Discrete Optimization. In *CP*. Springer.
- [48] Feng Yu, Wen-Chi Hou, Cheng Luo, Dunren Che, and Mengxia Zhu. 2013. CS2: A New Database Synopsis for Query Estimation. In *SIGMOD*.
- [49] Xiaofei Zhang, Lei Chen, and Min Wang. 2012. Towards Efficient Join Processing over Large RDF Graph using Mapreduce. In *SSDBM*.
- [50] Weijie Zhao, Florin Rusu, Bin Dong, and Kesheng Wu. 2016. Similarity Join over Array Data. In *SIGMOD*.