

The Sort-Merge-Shrink Join¹

Christopher Jermaine, Alin Dobra, Subramanian Arumugam, Shantanu Joshi, Abhijit Pol
CISE Department, University of Florida, Gainesville
{cjermain, adobra, sa2, ssjoshi, apol}@cise.ufl.edu

Abstract

One of the most common operations in analytic query processing is the application of an aggregate function to the result of a relational join. We describe an algorithm called the *Sort-Merge-Shrink (SMS) Join* for computing the answer to such a query over large, disk-based input tables. The key innovation of the SMS join is that if the input data are clustered in a statistically random fashion on disk, then at all times, the join provides an online, statistical estimator for the eventual answer to the query, as well as probabilistic confidence bounds. Thus, a user can monitor the progress of the join throughout its execution and stop the join when satisfied with the estimate's accuracy, or run the algorithm to completion with a total time requirement that is not much longer than that of other common join algorithms. This contrasts with other online join algorithms, which either do not offer such statistical guarantees, or can only offer guarantees so long as the input data can fit into main memory.

1 Introduction

One promising approach to data management for analytic processing is *online aggregation* [9][10] (OLA). In OLA, the database system tries to quickly discover enough information to immediately give an approximate answer to a query over an aggregate function such as COUNT, SUM, AVERAGE, and STD_DEV. As more information is discovered, the estimate is incrementally improved. A primary goal of any system performing OLA is providing some sort of statistical guarantees on result quality, usually in the form of statistical *confidence bounds* [2]. For example, at a given moment, the system might say, "0.32% of the banking transactions were fraudulent, with $\pm 0.027\%$ accuracy at 95% confidence." A few seconds later, the accuracy guarantee might be $\pm 0.013\%$. If the user is satisfied with the accuracy or determines from the initial results that the query is uninteresting, then computation can be terminated by a mouse click, saving valuable human time and computer resources.

In practice, one of the most significant technical difficulties in OLA is relational join processing. In fact, join processing is even more difficult in OLA than in traditional transaction processing systems. In addition to being computationally efficient, OLA join algorithms must have meaningful and statistically-justified confidence bounds. The current state-of-the-art for join processing in OLA is the *ripple join* family of algorithms, developed by Haas and Hellerstein [6]. However, there is one critical drawback with the ripple join and others that have been proposed for OLA:

¹A preliminary version of this paper appeared as "A Disk-Based Join With Probabilistic Guarantees" in the SIGMOD 2005 Conference [22].

All existing join algorithms for OLA (including the ripple join) are unable to provide the user with statistically meaningful confidence bounds as well as efficiency from start-up through completion, *if the total data size is too large to fit in main memory*.

This is unfortunate, because a ripple join may run out of memory in a few seconds, but a sort-merge join or hybrid hash join [14] may require hours to complete over a very large database. Using current technology, if the user is unsatisfied with the ripple join's accuracy, the only option is to wait until an exact answer can be computed.

It turns out to be exceedingly difficult to design disk-based join algorithms that are amenable to the development of statistical guarantees. The fundamental problem is that OLA join algorithms must rely on *randomness* to achieve statistical guarantees on accuracy. However, randomness is largely incompatible with efficient disk-based join processing. Efficient disk-based join algorithms (such as the sort-merge join and hybrid hash join) rely on careful organization of tuples into blocks, so that when a block is loaded into memory, *all* of its tuples are used for join processing before the block is unloaded from memory. This careful organization of tuples is the antithesis of randomness, and makes statistical analysis difficult. This is the reason, for example, why the scalable version of the ripple join algorithm by Luo et al. [13] can no longer maintain any statistical guarantees as soon as it runs out of memory.

Previously, Haas and Hellerstein have asserted that the lack of an efficient, disk-based OLA join algorithm is not a debilitating problem, because often the user is sufficiently happy with the accuracy of the estimate that the join is stopped long before the buffer memory is exhausted [10]. While this argument is often reasonable, there are many cases where convergence to an answer will be very slow, and buffer memory could be exhausted long before a sufficiently accurate answer has been computed. Convergence can be slow under a variety of conditions, including:

- if the join has high selectivity;
- if a select-project-join query has a relational selection predicate with high selectivity;
- if the attributes queried over have high skew;
- if the query contains a GROUP BY clause with a large number of groups or if the group cardinalities are skewed; or
- if the database tuples or key values are large (for example, long character strings) and so they quickly consume the available main memory.

Our Contributions

In this paper, we develop a new join algorithm for OLA that is specifically designed to be amenable to statistical analysis at the same time that it provides out-of-core efficiency. Our algorithm is called the *SMS join* (which stands for Sort-Merge-Shrink join). At a high level, the SMS join is best viewed as a variation of the sort-merge join [14] or more specifically, the Progressive Merge Join [4], with

the addition of significant statistical computations that allow the final result of the query to be guessed at throughout execution. The SMS join represents a contribution to the state-of-the-art in several ways:

- The SMS join is totally disk-based, yet it maintains statistical confidence bounds from start-up through to completion.
- Because the first phase of the SMS join consists of a series of hashed ripple joins, the SMS join is essentially a generalization of the ripple join. If a satisfactory answer can be computed in the few seconds before the join runs out of core memory, the SMS join is equivalent to a ripple join.
- Despite the fact that it maintains statistical guarantees, the SMS join is competitive with traditional join algorithms such as the sort-merge join, even with the extra computation required to maintain statistical estimates throughout the process. Our prototype implementation of the SMS join is not significantly slower than our own sort-merge join implementation. Our experiments show that if the SMS join is carefully implemented, the join is only around 10% slower than the traditional sort-merge join. Given this, we argue that if the SMS join is used to answer aggregate queries, online statistical confidence bounds can be maintained *virtually for free*, with the only additional cost being a small performance hit required to complete the entire join compared to a more traditional algorithm.

Paper Organization

The remainder of the paper is organized as follows. Section 2 gives some background on the problem of developing out-of-core joins for OLA. [Section 3 gives an overview of the SMS join, and the subsequent three sections describe the details of the algorithm.](#) Section 7 considers how the SMS join can be implemented so that it runs with high efficiency. [Our benchmarking is described in Section 8, related work is described in Section 9,](#) and the paper is concluded in Section 10. [The paper’s appendices discusses statistical considerations, including how the SMS join can be applied if data are not clustered randomly on disk.](#)

2 Disk-Based Joins for OLA: Why Are They So Difficult?

2.1 A Review of the Ripple Join Algorithm

We begin with a short review of the *ripple join* family of algorithms, which constitute the state-of-the-art in join algorithms for OLA [6]. We will describe the ripple join in the context of answering the query: “Find the total sales per office for the years 1980-2000” over the database tables EMP (OFFICE, START, END, NAME) and SALES (YEAR, EMP_NAME, TOT_SALES). The first table describes the time period that an employee is associated with an office, and the second table gives the total sales attributed to an employee per year. If an employee can only be associated with one office at a time, SQL for this query is given below:

```
SELECT SUM (s.TOT_SALES) , e.OFFICE , s.YEAR
FROM EMP e , SALES s
```

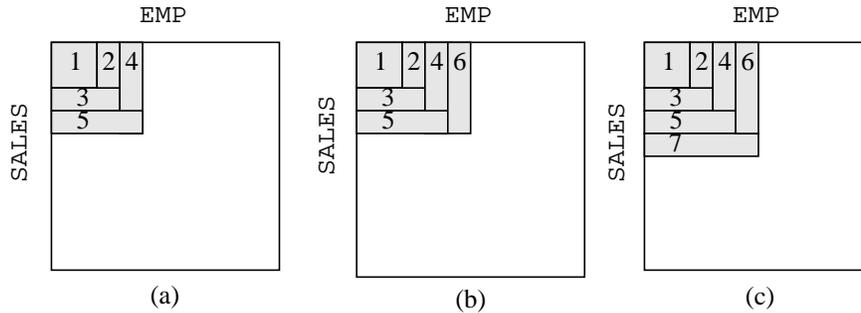


Figure 1: Performing sampling steps during a ripple join.

```

WHERE e.NAME = s.EMP_NAME AND s.YEAR BETWEEN e.START
AND e.END AND s.YEAR BETWEEN 1980 AND 2000
GROUP BY e.OFFICE, s.YEAR

```

To answer this query, a ripple join will scan the two input relations EMP and SALES in parallel, using a series of operations known as *sampling steps*. At the beginning of each sampling step, a new set of tuples of size n_{EMP} is loaded into memory from the input relation EMP. These new tuples from EMP are joined with all of the tuples that have been encountered thus far from the second input relation, SALES. This is depicted as the addition of region 6 to Figure 1 (b) from Figure 1 (a). Next, the roles of EMP and SALES relations are swapped: n_{SALES} new input tuples are retrieved from SALES, and joined with all existing tuples from EMP (shown as the addition of region 7 in Figure 1 (c)). Finally, both the estimate for the answer to the query and the statistical confidence interval for the estimate's accuracy are updated to reflect the new tuples from both relations, and then output to the user via an update to the graphical user interface. The estimate and bounds are computed by associating a normally distributed random variable N with the ripple join; on expectation, the ripple join will give a correct estimate for the query, and the error of the estimate is characterized by the variance of N , denoted by σ^2 . A key requirement of the ripple join is that all input relations must be clustered randomly on disk, so that there is *absolutely* no correlation between the ordering of the tuples on disk and the contents of the tuples. Even the slightest correlation can invalidate the statistical properties of the join, leading to inaccurate estimates and confidence bounds.

The ripple join is actually a *family* of algorithms and not a single algorithm, since the technique does not specify exactly how the new tuples read during a sampling step are to be joined with existing tuples. This can be done using an index on one or both of the relations, or in the style of a nested-loops join, or using a main-memory hash table to store the processed tuples.

In their work, Haas and Hellerstein showed that the hashed version of the ripple join generally supports the fastest convergence to an accurate answer [6]. To perform a hashed ripple join, a single main memory hash table is used to hold tuples from both relations. When a new tuple is read from input relation EMP, it is added

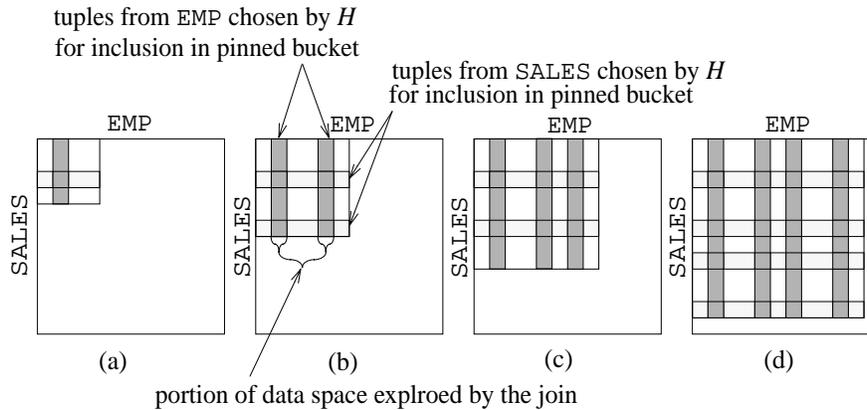


Figure 2: Progression of the first pass of symmetric version of the hybrid hash join.

to the table by hashing the key(s) that the input relations are joined on. Any tuple s from the other input relation SALES which matches e will be in the same bucket, and the tuples can immediately be joined.

Haas and Hellerstein’s work demonstrated that other types of ripple joins are usually less efficient than the hashed ripple join. A *nested-loops ripple join* gives very slow convergence, and is expectedly even slower to run to completion than the traditional (and often very slow) nested-loops join. The reason for this is that since the ripple join is symmetric, each input relation must serve as the inner relation in turn, leading to redundant computation in order to maintain statistics and confidence intervals, compared with even a traditional nested-loops join. An *indexed ripple join* can be faster than a nested-loops ripple join, but since random clustering of the input relations is an absolute requirement for estimation accuracy, the index used *must* be a secondary index, and an indexed ripple join is essentially nothing more than a slower version of the traditional indexed, nested-loops join. Thus, it still requires at least one or two disk seeks per input tuple to perform an index lookup to join each tuple. Even using modern hardware, one would expect a processing rate no faster than a few thousand tuples per second per disk under such a regime.

2.2 OLA and Out-Of-Core Joins

While it is very fast in main memory, the hashed ripple join becomes unusable as soon as the central hash table becomes too large to fit into main memory. The problem is that when a large fraction of the hash table must be off-loaded to disk, then each additional tuple encountered from the input file will require one or two random disk head movements in order to add it to the hash table, join it with any matching, previously-encountered tuples, and then update the output statistics. At a few milliseconds per seek using modern hardware, this equates at most to only a few thousand tuples processed per second per disk.

Furthermore, designing a join algorithm that can operate efficiently from disk as well as provide tight, meaningful confidence bounds is a difficult problem. Consider a hybrid hash join [14] over a query of the form:

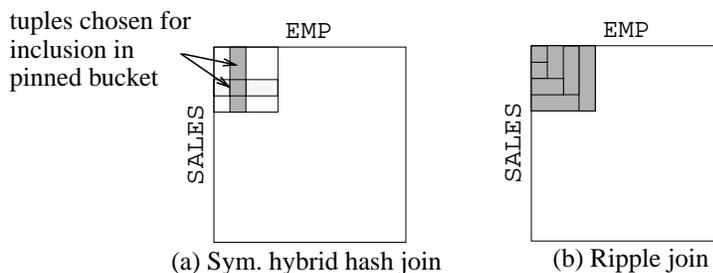


Figure 3: Comparison of a symmetric hybrid hash join and a ripple join, with enough memory to buffer four tuples from each input relation.

```

SELECT expression (e, s)
FROM EMP e, SALES s
WHERE predicate (e, s)

```

where *predicate* is a conjunction containing the clause $e.k = s.k$. A hybrid hash join would operate by first scanning the first input relation EMP, hashing all of EMP's tuples on the attribute EMP.*k* to a set of buckets that are written out to disk as they are constructed, except for a single, "special" bucket that is pinned in main memory. After hashing EMP, the second relation SALES is then scanned and hashed on the attribute SALES.*k*. Any tuples from SALES that hash to the "special" pinned bucket are joined immediately, and the results are output online. All other tuples are written to disk with the other tuples from their respective buckets. In a second pass over the buckets from both EMP and SALES, all matching buckets are joined, and resulting tuples are output as they are produced.

One could very easily imagine performing a ripple-join-like symmetric variation on the hybrid hash join, where the first passes through both SALES and EMP are performed concurrently, and all additions to the "special" pinned bucket are joined immediately and used to produce an estimate to the query answer. If a hash function is used which truly picks a random subset of the tuples from each relation to fall in the special pinned buckets, then the resulting join is amenable to statistical analysis. As the relation EMP is scanned, a random set of tuples from EMP will be chosen by the hash function. Because any matching tuples from SALES will be pinned in memory, any tuples from EMP chosen by the hash function can immediately be joined with all tuples encountered thus far from SALES. If the join attribute is a candidate key of the EMP relation, the set of tuples chosen thus far from EMP is a true random sample of EMP, and this set is independent from the set of all tuples that have been read from SALES. Thus, a statistical analysis very similar to the one used by Haas [7], Haas et al. [8] and Haas and Hellerstein [6] can be used to describe the accuracy of the resulting estimator. The progression of the join with enough memory to buffer four tuples from each input relation is shown in Figure 2.

However, there are several serious problems with this particular online join, including:

- *Early on, the resulting estimator will have less accuracy than a ripple join.* This is particularly worrisome, because if OLA is used for exploratory data analysis, it is expected that many queries will be aborted early on. Thus, a fast and accurate estimate is essential. In the case of the symmetric hybrid hash join, the issue is that the pinned bucket containing tuples from EMP will not fill until EMP has been entirely scanned during the first pass of the join. Contrast this with a ripple join: if we have enough memory to buffer n_{EMP} tuples from EMP, then after n_{EMP} tuples have been read from EMP, a ripple join will use all n_{EMP} tuples immediately to estimate the final answer to the join. On the other hand, after scanning n_{EMP} tuples, the symmetric hybrid hash join would use only around $n_{\text{EMP}}/|\text{EMP}|$ tuples to estimate the final answer. A comparison between the two algorithms after four tuples from each input relation have been read is shown in Figure 3. While the symmetric, hybrid hash join has ignored three of the four tuples from each input relation, the ripple join has fully processed all four.
- *If the join attribute is not a candidate key of EMP, then statistical analysis of the join is difficult.* The problem in this case is that the tuples chosen by the hash function will not be a random sample of the tuples from EMP, since the function will choose all tuples with a given key value k . The statistical properties of the join under such circumstances are exceedingly difficult to analyze.
- *Extending the statistical analysis of the symmetric hybrid hash join to the second phase is difficult as well.* From a statistical perspective, the issue is that each pair of buckets from EMP and SALES joined during the second phase are not independent, since they contain tuples hashed to matching buckets. Taking into account this correlation is difficult. Furthermore, the tuples from each on-disk bucket from SALES have already been used to compute the estimator associated with the growing region of Figure 2, so each bucket join is not independent of the initial estimator computed during the hash phase.

3 The SMS Join: An Overview

The difficulties described in the previous section illustrate why developing out-of-core joins with statistical guarantees is a daunting problem. We now give a brief overview of the SMS join, which is designed specifically to maintain statistically meaningful confidence bounds in a disk-based environment. The SMS join is a *generalization* of the ripple join, in the sense that if it is used on a small data set or if it is terminated before memory has been exhausted, then it is identical to the ripple join. However, the SMS join continues to provide confidence bounds throughout execution, even when it operates from disk. Intuitively, the SMS join is best characterized as a disk-based sort-merge join, with several important characteristics that differentiate it from the classic sort-merge join.

3.1 Preliminaries

In our description of the SMS join, we assume two input relations, EMP and SALES (as described in Section 2.1) which are joined using the following generic query:

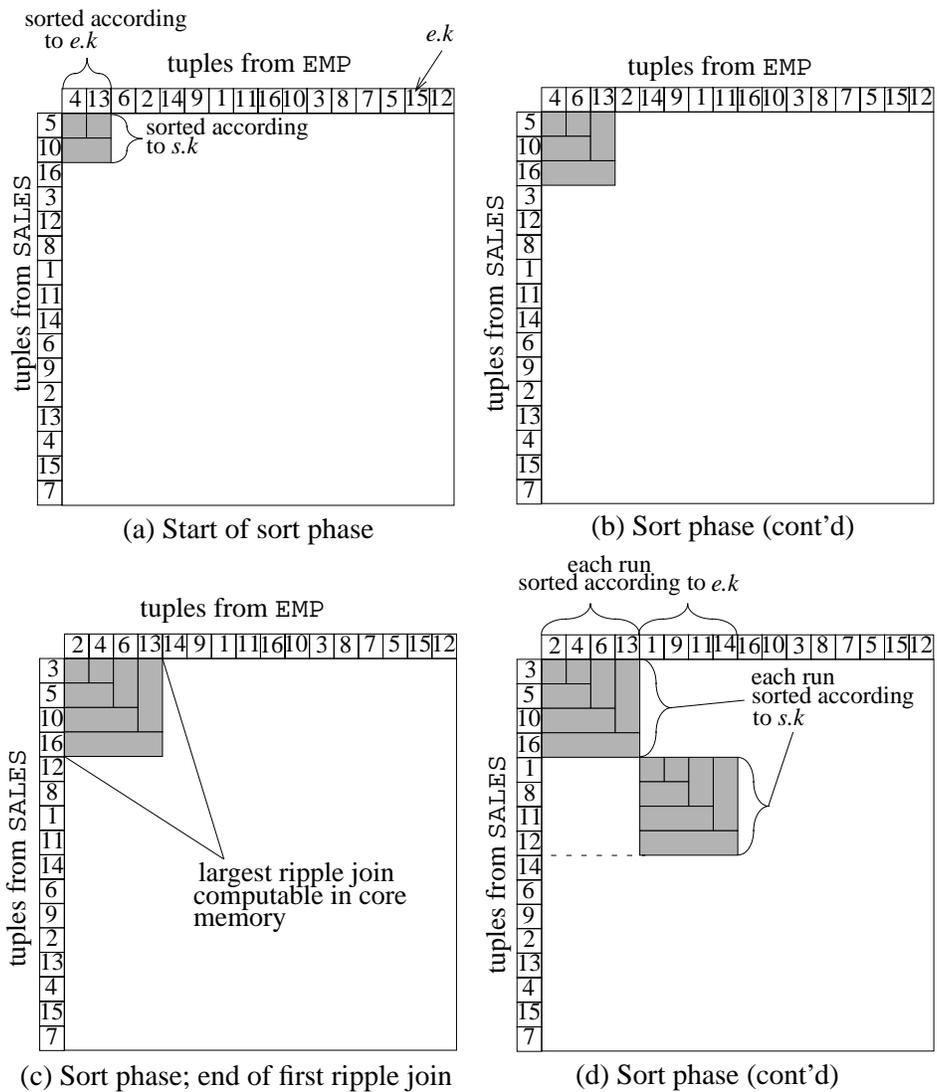


Figure 4: Outline of an SMS join, assuming enough main memory to buffer four tuples from each input relation. The remainder of the join is depicted in Figure 5.

```

SELECT SUM  $f(e, s)$ 
FROM EMP as  $e$ , SALES as  $s$ 

```

We restrict our discussion to the case of two input relations, though just like a sort-merge join, the SMS join can handle additional input relations as long as the same sort order can be used to join all of them. Just as in the case of the ripple join, in order to provide our statistical confidence bounds, we will require that input rela-

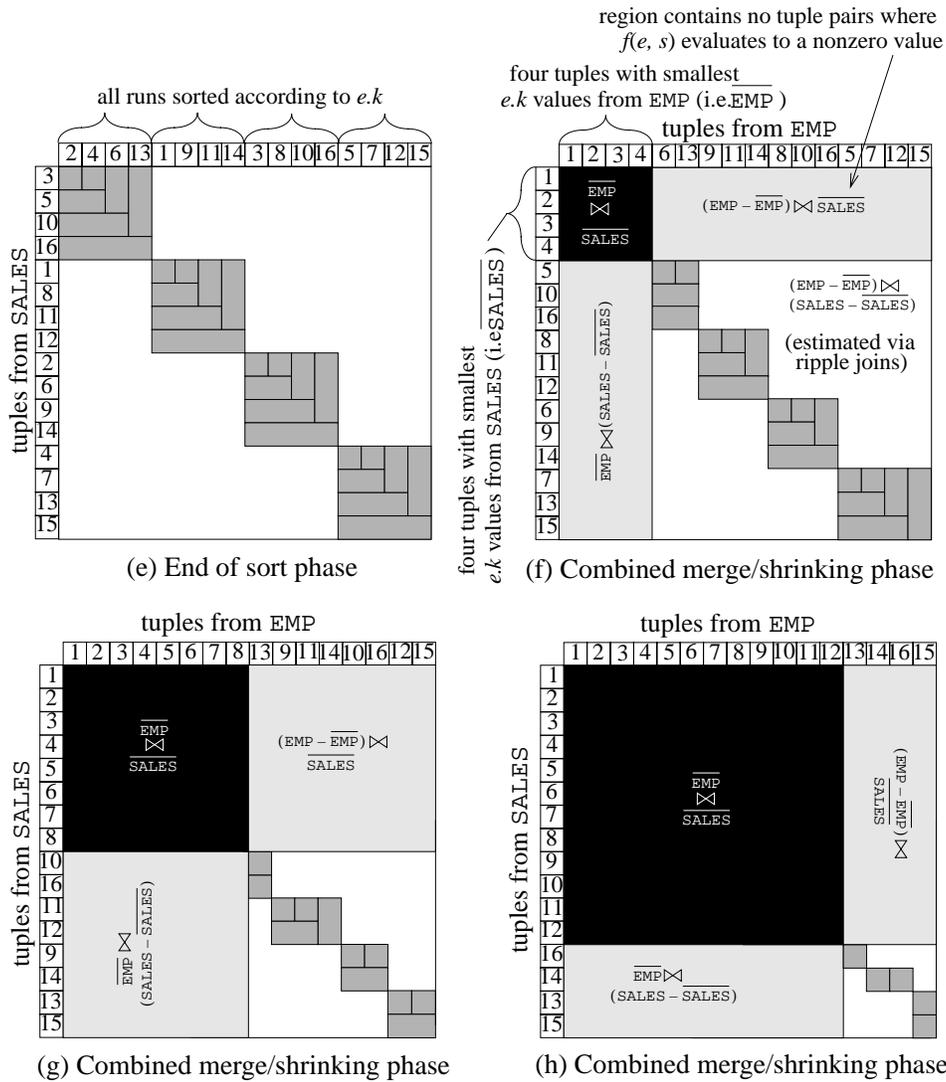


Figure 5: Outline of an SMS join, continued.

tions EMP and SALES be clustered in a statistically random fashion on disk¹. One way to achieve the desired randomization in a scalable fashion is to associate a pseudo-random number with each tuple. Then, a scalable disk-based sorting algorithm (such as a two-phase, multi-way merge sort) is used on the tuples to order them based upon the associated pseudo-random numbers. Once the tuples are ran-

1. If such an ordering is not available or is not practical, the SMS join can still be used; see Appendix C for a discussion.

domized, the numbers can then be discarded. The function f can be any function over a pair of tuples from the two input relations, and can encode arbitrary relational selection predicates or join conditions over the input tuples. We assume that f evaluates to zero unless a conjunction of boolean expressions over the two input tuples evaluates to true, and that one of these boolean expressions is an equality check of the form $(e.k = s.k)$, so that a classical hash join or sort-merge join could be used over the attribute k . We subsequently refer to the attribute k as the *key*. In our discussion, we also assume that the aggregate function evaluated is SUM; AVERAGE or STD_DEV queries can be handled by treating the query as an expression over multiple queries that are evaluated concurrently. Since the associated confidence bounds must simultaneously take into account the potential for error in each of the constituent queries, one of several methods must be chosen to combine the estimates. [One method suggested previously is the use of Bonferroni’s inequality \[8\].](#)

In general, we will assume that B blocks of main memory are available for buffering input tuples while computing the join. For simplicity, we assume that additional main memory (in addition to those B blocks) is [available for buffering output tuples and storing the required statistics](#); see Section 6.4 for a detailed discussion of the additional memory required by the SMS join compared to the sort-merge join. We will use the notations $|EMP|$ and $|SALES|$ to denote the size of the EMP and SALES input relations measured by the number of tuples, respectively. $|EMP|_b$ and $|SALES|_b$ denote the number of blocks used to store these relations.

3.2 Three Phases of the SMS Join

Given these preliminaries, the SMS join has three phases:

(1) The *sort phase* of the SMS join corresponds to a modified version of the sort phase of a sort-merge join. At a high level, the sort phase of the SMS join closely resembles the sort phase of the Progressive Merge Join [4] in that all of the data currently in memory from both relations are used to search for early output results. The process of reading in and sorting each pair of runs from EMP and SALES is treated as a series of hashed ripple joins, each of which is used to provide a separate estimate for the final result of the join. Using the techniques of Haas and Hellerstein, it is possible to characterize the estimate corresponding to the i th ripple join as a random variable N_i whose distribution becomes normal or Gaussian with increasing sample size. As described subsequently, all of the estimators can be combined to form a single running estimator N , where on expectation N provides the correct estimate for the query result. The sort phase of the SMS join is illustrated in Figure 4 (a) through (d) and Figure 5(e).

(2) The *merge phase* of the SMS join corresponds to the merge phase of a sort-merge join. Just as in a sort-merge join, the merge phase pulls tuples from the runs produced during the sort phase, and joins them using a lexicographic ordering. The key difference between the SMS join and the sort-merge join or the Progressive Merge Join is that in the SMS join, the sort phase runs concurrently with (and sends information to) the *shrinking phase* of the join.

Algorithm 1: Sort Phase of the SMS Join

Input: Amount of available main memory B , parameter to control the interleaving of reads and writes m

- (1) Let $m_{EMP} = (B \times |EMP|_b) / (|EMP|_b + |SALES|_b)$
- (2) Let $m_{SALES} = (B \times |SALES|_b) / (|EMP|_b + |SALES|_b)$
- (3) Let $n_{EMP} = m_{EMP} / m, n_{SALES} = m_{SALES} / m$
- (4) Perform a ripple join of the first m_{EMP} and m_{SALES} blocks from EMP and SALES, respectively; these blocks will be EMP_1 and $SALES_1$
- (5) Sort EMP_1 and $SALES_1$ on $EMP.k$ and $SALES.k$
- (6) For int $i = 1$ to $(|EMP|_b + |SALES|_b) / B - 1$
- (7) While $EMP_i \neq \emptyset$
- (8) Write n_{EMP} blocks from EMP_i to disk
- (9) Write n_{SALES} blocks from $SALES_i$ to disk
- (10) Read n_{EMP} blocks from EMP into EMP_{i+1}
- (11) Read n_{SALES} blocks from SALES into $SALES_{i+1}$
- (12) Join the new blocks from SALES with EMP_{i+1}
- (13) Join the new blocks from EMP with $SALES_{i+1}$
- (14) Sort EMP_{i+1} and $SALES_{i+1}$

(3) The *shrinking phase* of the SMS join is performed concurrently with the merge phase. Conceptually, the merge phase of the SMS join divides the data space into four regions, shown in Figure 5 (f) through (h). If \overline{EMP} and \overline{SALES} refer to the subsets of EMP and SALES that have been merged thus far, then as depicted in Figure 5, the four regions of the data space are $(EMP - \overline{EMP}) \bowtie (SALES - \overline{SALES})$, $(EMP - \overline{EMP}) \bowtie \overline{SALES}$, $\overline{EMP} \bowtie (SALES - \overline{SALES})$, and $\overline{EMP} \bowtie \overline{SALES}$. Assuming that the query to be evaluated is SUM or COUNT, then the final answer to our query is simply the sum over each of these four regions.

Just as in a classical sort-merge join, the merge phase of the SMS join allows us to compute the value of SUM $f(e, s)$ for the latter three joins exactly (note that in the case of $(EMP - \overline{EMP}) \bowtie \overline{SALES}$ and $\overline{EMP} \bowtie (SALES - \overline{SALES})$, this value will always be zero). However, in the classical sort-merge join, the region corresponding to $(EMP - \overline{EMP}) \bowtie (SALES - \overline{SALES})$ in Figure 5 would be ignored. Because the goal of the SMS join is to produce a running estimate for the eventual query answer, the value for SUM $f(e, s)$ over $(EMP - \overline{EMP}) \bowtie (SALES - \overline{SALES})$ will be continually *estimated* by the shrinking phase of the join algorithm.

The specific task of the shrinking phase is handling the removal of tuples from the sort phase ripple join estimators, and updating the estimator N corresponding to

the value of $\text{SUM } f(e, s)$ over the join $(\text{EMP} - \overline{\text{EMP}}) \bowtie (\text{SALES} - \overline{\text{SALES}})$. Tuples that are consumed by the merge phase must be removed from each N_i that contributes to N so that we ensure that each of the four regions depicted in Figure 5 (f) through (h) remain disjoint. As long as this is the case, the value of $\text{SUM } f(e, s)$ over each of the four regions can simply be summed up to produce a single estimate for the eventual query result.

We now describe each of the phases of the SMS join in more detail.

4 The Sort Phase

The *sort* phase of the SMS join is very similar to the first phase of a sort-merge join. The phase is broken into a series of steps, where at each step a subset of each input relation of total size small enough to fit in main memory is read from disk, sorted, and then written back to disk. Despite the similarity, some changes will be needed to provide an online estimate for the eventual answer to our query.

4.1 The Basic Algorithm

Pseudo-code for the sort phase of the SMS join is given above in Algorithm 1. During each step of the sort phase, a subset of the blocks from each input relation is first read in parallel and joined in memory using a hashed ripple join. After the blocks from each relation have consumed all of the available buffer memory, the tuples of each subset (called a “run”) are then sorted on the join attribute(s). The operation of the sort phase is visualized above in Figure 4 (a) through (d) and Figure 5 (e).

4.2 Statistical Analysis

We now address the problem of providing online estimates for the eventual answer to the query during the sort phase of the SMS join, along with statistical guarantees on the accuracy of those estimates.

4.2.1 Overview

It is possible to characterize the estimate provided by each of the individual sort phase estimators by any of a number of methods, including making use of the analysis provided by Haas and Hellerstein [6] as a starting point. However, the problem we face is significantly different than that faced by Haas and Hellerstein: we have *many* sort phase ripple join estimators, and not a single one. The question we must tackle when analyzing the SMS join is: How can these estimates be combined in a meaningful way, so as to produce a single estimator that is better than any of its constituent parts?

The basic tactic we use is as follows. First, let N_i be a random variable corresponding to the estimate provided by the i th sort phase ripple join. Note that each individual N_i is identical to Haas and Hellerstein’s ripple join estimator, and is computed in exactly the same way. However, to extend our analysis to the *many* ripple joins used by the SMS join, we associate a weight w_i with each individual estimate, subject to the constraint that $\sum_i^n w_i = 1$. Assuming a total of n ripple

joins in the sort phase, the random variable $N = \sum_i^n N_i w_i$ is then itself an unbiased estimator for the eventual query answer. By carefully choosing the weights associated with the various ripple joins, we can maximize the accuracy of the estimator. The remainder of this subsection focuses on the three-part problem of first characterizing the accuracy of N , then tuning the weights within N so as to maximize its accuracy, and finally estimating the accuracy of N on-the-fly.

4.2.2 Characterizing the Accuracy of N

We weight each ripple join because a careful choice of the various weights will tend to increase the accuracy of the estimate provided by N . A new ripple join that has just begun is likely to have a large degree of inaccuracy, and should thus be given less weight than the sort phase ripple joins that have completed. In the remainder of the paper, we use the standard statistical technique of measuring the accuracy of an unbiased estimator by measuring its *variance*. Given a set of weights, the general formula for the variance of N is:

$$Var(N) = \sum_i^n w_i^2 Var(N_i) + \sum_i^n \sum_{j \neq i}^n w_i w_j Cov(N_i, N_j) \quad (1)$$

In the above formula, the function *Cov* refers to the *covariance* of the two ripple join estimators. The covariance appears in the formula for the variance of N because the various ripple join estimators are not independent: if a tuple is used during the i th ripple join, then it cannot be used during the j th ripple join.

Unfortunately, deriving formulas for $Var(N_i)$ and $Cov(N_i, N_j)$ is a difficult task. We could use Haas and Hellerstein’s derivation for a single ripple join as the basis for our derivation of these terms, but this would present two primary difficulties:

- (1) First, Haas and Hellerstein’s derivation approximates sampling without replacement using sampling with replacement. As a result, in their analysis all samples can be assumed to be independent. In the application domain imagined for the original ripple join, this is reasonable since it was assumed that the join will be terminated after only a rough estimate of the eventual answer has been obtained. [This means that the join probably will end](#) before a significant fraction of either input relation has been processed. In such a situation, the difference between sampling with replacement and sampling without replacement is negligible. In our own application where each join may cover a significant fraction of the database (and all tuples will eventually be part of exactly one ripple join), such an approximation is very hard to justify.
- (2) Second, the Haas and Hellerstein’s derivation was non-trivial to say the least, and can be traced through a series of papers by Peter Haas and his co-authors [8][7][6]. Extending this analysis to the more complicated estimator described in this paper seems like a very daunting task, specifically because the covariance between the various ripple-join estimators must be considered as well.

As a result, we give an exact, permutational derivation for these values in Appendix A (the results are given as Theorem A.1, Corollary A.2, Theorem A.3, and

Corollary A.4). The permutational analysis given in the paper is arguably simpler than the without-replacement analysis that led to a characterization of the original ripple join; this is the main reason that it is even feasible to give an analysis of the more complicated estimator described in this paper. While the exact formulas and their derivation are not important to the exposition presented here, the salient properties of the formulas are:

- (1) They give an exact characterization of the accuracy of our estimators for any data sets and queries of the type considered in this paper.
- (2) Given a few sufficient statistics about the data sets used by the join (or estimates for the value of these statistics), it is an easy matter to use the formulas to compute the variance of N , as will be described in Section 4.2.4.

4.2.3 Tuning the Weights for Maximum Accuracy

As discussed above, the value of N and the variance of N both depend on the weights that are chosen for each of the constituent ripple joins. To maximize the accuracy of the estimate provided by N , it is necessary to compute the various weights so as to minimize the variance of N .

Performing this minimization is a straightforward task using basic calculus. Before we begin, we make note of two important observations that will guide the calculation. Assuming that we are currently computing the k th sort phase ripple join where $1 \leq k \leq n$, then the following observations simplify the problem:

- (1) First, we know that $Var(N_i) = Var(N_k)$ for all $1 \leq i < k$, so it must be the case that $w_i = w_k$ for all $1 \leq i < k$ in any optimal weighting scheme. This is because (by definition) all of the sort phase ripple join estimators that have completed make use of the same number of tuples, and so the properties of their estimators are identical (this assumes that we disregard the tiny amount of bias associated with the fact that the number of ripple joins may not divide evenly into the number of tuples in either input relation, and we assume a serial and not parallel execution of the join).
- (2) Second, we know that $Var(N_i) = \infty$ for all $i > k$, because no estimate has been produced by any ripple join past the k th one. Thus, it must be the case that $w_i = 0$ for all $i > k$.

Using these two facts, along with the fact that the formula for the covariance is independent of the amount of data used to compute each ripple join (Corollary A.4), Equation 1 can be re-written as:

$$Var(N) = (k-1)w_1^2 Var(N_1) + w_k^2 Var(N_k) + (1 - (k-1)w_1^2 - w_k^2)Q_U$$

In this expression, Q_U is the part of the expression for the covariance that is independent of the weighting that is chosen (see Appendix A and the proof of Theorem A.3 for an exact expression for Q_U). To minimize this value and maximize the accuracy of N , we make use of the fact that all weights must add to one, so $(k-1)w_1 + w_k = 1$. Let $j = k - 1$. Substituting, we have:

$$\text{Var}(N) = jw_1^2 \text{Var}(N_1) + (1 - jw_1)^2 \text{Var}(N_k) + (1 - jw_1^2 - (1 - jw_1)^2) Q_U$$

We wish to minimize this value. Taking the partial derivative with respect to w_1 gives:

$$\frac{\partial}{\partial w_1} \text{Var}(N) = 2jw_1 \text{Var}(N_1) + (-2j + 2j^2 w_1) \text{Var}(N_k) + (-2jw_1 + 2j - 2j^2 w_1) Q_U$$

Setting this value to zero and solving gives the optimal value for w_1 :

$$w_1^{opt} = \frac{\text{Var}(N_k) - Q_U}{\text{Var}(N_1) + j\text{Var}(N_k) - Q_U - jQ_U}$$

Given the above formula for the optimal weighting of w_1 , the weight w_k can then be computed using simple substitution.

One concern is that while this solution does find a “flat spot” on the variance function with respect to w_1 , it is not obvious that this must be a minimum. To address this, we note that the second partial derivative with respect to w_1 is always constant when the other variables are fixed:

$$\frac{\partial^2}{\partial^2 w_1} \text{Var}(N) = 2j \text{Var}(N_1) + 2j^2 \text{Var}(N_k) + (-2j - 2j^2) Q_U$$

This means that the curve has no local minima (or maxima) that is not also a global minima (or maxima). Since the feasible values of w_1 are all from 0 to 1 inclusive, we can then conclude that the minimum variance value must be either at $w_1 = 0$, 1, or w_1^{opt} . Checking all three values is then guaranteed to yield an optimal solution.

4.2.4 Estimating the Parameters Needed to Compute N , $\text{Var}(N)$

Of course, the above exposition assumes that we have access to $\text{Var}(N_k)$, $\text{Var}(N_1)$, and Q_U . It is an often-overlooked fact that given most sampling-based estimators, it is not possible to compute the variance of an estimator exactly, since computing the variance requires that we have access to all of the data. It is not realistic to assume that we have easy access to all of the data; if we did, then there would not be a need to sample! As a result, the variance of most sampling-based estimators must be estimated as well, and the SMS join is no different.

Looking over the formulas given in Appendix A for the various components of the variance of the SMS join, there are four separate quantities that must be estimated to compute the various components of the variance of N . These quantities are:

- Q , which is the eventual answer to the query;

- $\sum_{e \in \text{EMP}} (\sum_{s \in \text{SALES}} f(e, s))^2$, which we refer to as c_1 ;
- $\sum_{s \in \text{SALES}} (\sum_{e \in \text{EMP}} f(e, s))^2$, which we refer to as c_2 ;
- and $\sum_{(e, s) \in \text{EMP} \times \text{SALES}} f^2(e, s)$, which we refer to as c_3 .

Once each of these quantities is known or estimated, then, using the formulas given in Appendix A, it is an easy matter to compute $\text{Var}(N)$.

Accurately estimating each of these values requires some careful consideration. For example, consider c_1 . To compute a high-quality estimator for c_1 , we first produce a series of individual c_1 estimators, where each individual estimator is computed using the data that are loaded into memory as each of the ripple joins is processed (the process of computing each of these individual estimates is described in detail below). At any given time during the execution of the SMS join, we can then estimate c_1 as a simple, linear combination of the estimators computed along with each individual ripple join. Let $\hat{c}_{1,i}$ denote the value of c_1 estimated using the data input into the i th ripple join. After k ripple joins have been computed,

we can simply use $\hat{c}_1 = \sum_{i=1}^{k-1} \frac{\hat{c}_{1,i}}{k-1}$ as an estimate for c_1 . It is important to

note that \hat{c}_1 does not make use of $\hat{c}_{1,k}$, where $\hat{c}_{1,k}$ is the value of c_1 estimated by the ripple join that is currently being executed. The reason for this is that $\hat{c}_{1,k}$ is likely to have poorer accuracy than the other estimates (this is because if the j th ripple join is still being processed, then $\hat{c}_{1,k}$ will be estimated using fewer data points). Given a formal characterization of the effect of the reduced data on the accuracy of $\hat{c}_{1,k}$, it would allow each $\hat{c}_{1,i}$ to be weighted properly. In the absence of such a characterization, as long as $k > 1$ it is preferable to ignore $\hat{c}_{1,k}$ during the computation.

Estimating \hat{Q}_i and $\hat{c}_{3,i}$

Estimating each of Q , c_1 , c_2 , and c_3 as a mean of estimates over each individual ripple join requires that we have a principled way to compute \hat{Q}_i , $\hat{c}_{1,i}$, $\hat{c}_{2,i}$, and $\hat{c}_{3,i}$ over each ripple join. For \hat{Q}_i this is easy to do: since we do not know the eventual answer to the query, it is natural to use $\hat{Q}_i = N_i$. For $\hat{c}_{3,i}$ this is easy as well. Let EMP_i denote the set of tuples from EMP that has been assigned to the i th ripple join, and let SALES_i denote the analogous set for SALES. Let α_i denote the frac-

tion of tuples from EMP processed by the i th ripple join, and let β_i denote the analogous fraction for SALES. Then:

$$\hat{c}_{3,i} = \frac{1}{\alpha_i \beta_i} \sum_{(e,s) \in \text{EMP}_i \times \text{SALES}_i} f^2(e,s)$$

is an unbiased estimator for c_3 (this is true because $\hat{c}_{3,i}$ does nothing more than run a ripple join estimator using f^2 rather than f , so if the ripple join is itself unbiased, then so is $\hat{c}_{3,i}$). Since \hat{c}_3 is then a linear combination of unbiased estimators, it is unbiased as well.

Estimating $\hat{c}_{1,i}$ and $\hat{c}_{2,i}$

Things are not quite so simple for c_1 and c_2 , and care must be taken in developing estimators for these values. For example, consider the following, natural estimator for c_1 over the i th ripple join:

$$\hat{c}'_{1,i} = \frac{1}{\alpha_i} \sum_{e \in \text{EMP}_i} \left(\frac{1}{\beta_i} \sum_{s \in \text{SALES}_i} f(e,s) \right)^2$$

Unfortunately, this is *not* an unbiased estimator for $c_{1,i}$; in particular, the expected value or mean of this estimator is approximately:

$$E[\hat{c}'_{1,i}] \approx \frac{1 - \beta_i}{\beta_i} \sum_{e \in \text{EMP}_i} \sum_{s \in \text{SALES}_i} f^2(e,s) + c_{1,i}$$

(see Appendix B for a derivation of this expression). The quantity inside the parentheses is larger than $c_{1,i}$ by the amount:

$$E[\hat{c}'_{1,i}] - c_{1,i} = \frac{1 - \beta_i}{\beta_i} \sum_{e \in \text{EMP}_i} \sum_{s \in \text{SALES}_i} f^2(e,s)$$

This quantity is the *bias* of $\hat{c}'_{1,i}$, and we must compensate for this bias. Thus, the estimator for c_1 that we use in our implementation of the SMS join is:

$$\hat{c}_{1,i} = \frac{1}{\alpha_i} \sum_{e \in \text{EMP}_i} \left[\left(\frac{1}{\beta_i} \sum_{s \in \text{SALES}_i} f(e,s) \right)^2 + \left(\frac{1}{\beta_i} - \frac{1}{\beta_i^2} \right) \sum_{s \in \text{SALES}_i} f^2(e,s) \right]$$

By simply switching the relations EMP and SALES, a formula for $\hat{c}_{2,i}$ is easily derived:

$$\hat{c}_{2,i} = \frac{1}{\beta_i} \sum_{e \in \text{SALES}_i} \left[\left(\frac{1}{\alpha_{i,s}} \sum_{s \in \text{EMP}_i} f(e, s) \right)^2 + \left(1 - \frac{1}{\alpha_i} \right) \frac{1}{\alpha_{i,s}} \sum_{s \in \text{EMP}_i} f^2(e, s) \right]$$

Given these various estimators, it is then an easy matter to compute N and an estimator for the variance of N .

4.3 Of Aspect Ratios and Performance Hiccups

To ensure a smooth increase of the estimation accuracy that allows the statistical information given to the user to be constantly updated as the sort phase progresses, one requires that the processing of all of the input relations during the sort phase should progress at the same rate and that the processing of each relation should finish at the same time. This is the reason for the calculations of the first few lines of Algorithm 1, which together guarantee that every time we begin the next iteration of the loop of line (6), the same fraction of each input relation has been processed.

Why would it be a problem if we find ourselves in a situation where EMP has been totally processed and organized into runs, while SALES has only been half processed (or vice versa)? The issue is that if one relation were completely processed while the other was not, there would be a ‘hiccup’ where the estimate and accuracy could not be updated as the remainder of the other relation was sorted into runs as a prelude to the merge phase of the join.

This is one of the reasons that we do not consider the possibility of altering the relative rate at which the various input relations are processed, or ‘aspect ratio’ of the sort phase, as was suggested by Haas and Hellerstein in their work on the ripple join. The aspect ratio of the individual ripple joins *could* be altered by incorporating the adaptive techniques described by Haas and Hellerstein into lines (10) and (11); but these techniques should only be applied locally, to each individual ripple join. To provide smooth, consistent increases in estimation accuracy, the requirement is that after each ripple join that makes up the sort phase has been completed, the same fraction of each input relation must have been consumed.

5 The Merge Phase of the SMS Join

After the sort phase of the SMS join completes, the *merge phase* of the SMS join is invoked. The merge phase of the SMS join is similar to the merge phase of a classical sort-merge join. At all times, the merge phase of the join maintains the value:

$$total = \sum_{e \in \overline{\text{EMP}}} \sum_{s \in \overline{\text{SALES}}} f(e, s)$$

$total$ is simply the total value of the aggregate function f applied to the subsets $\overline{\text{EMP}}$ and $\overline{\text{SALES}}$ of the input relations, as depicted in Figure 4 and Figure 5.

To process the input relations, the merge phase repeatedly pulls tuples off the sort phase runs and joins them. Just as in a classical sort-merge join, two sets of

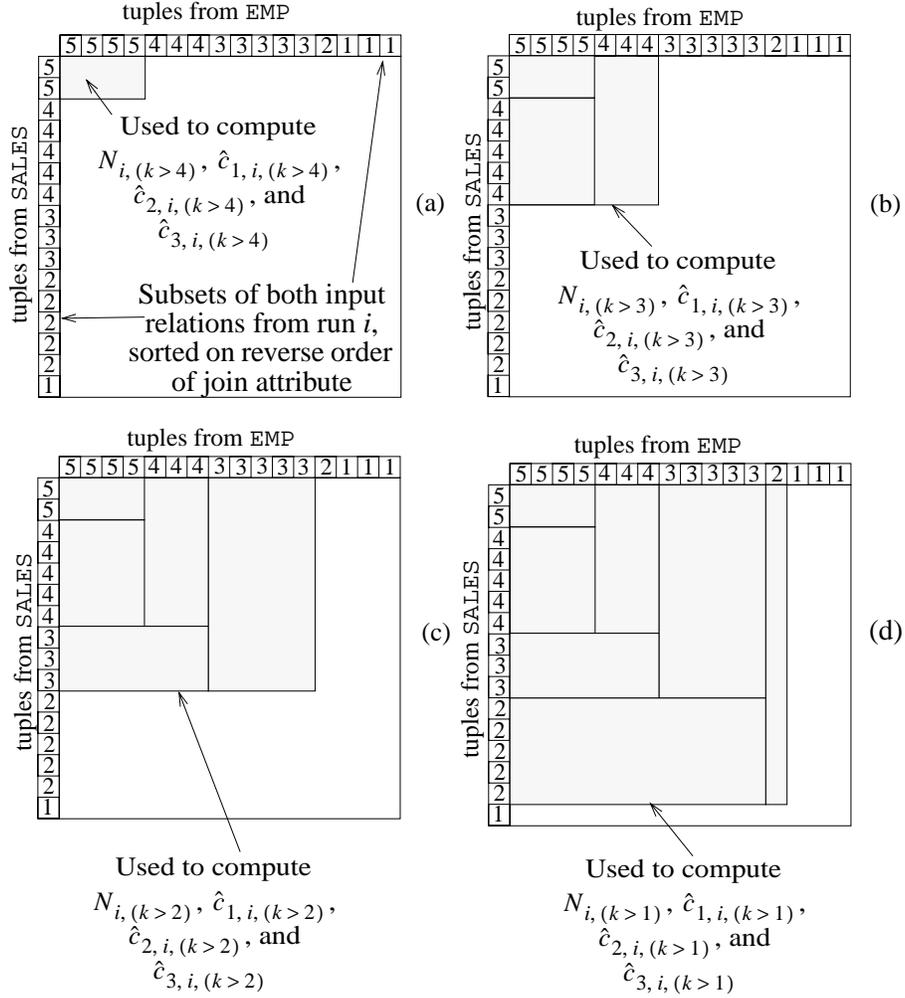


Figure 6: Reverse ripple join used to compute statistics used by the shrinking phase.

tuples $E \subset \text{EMP}$ and $S \subset \text{SALES}$ are pulled off the head of each run produced by the sort phase, such that for all $e \in E$ and $s \in S$, $e.k = s.k$. Assuming that the aggregate function in question is SUM or COUNT, the merge phase then computes the value:

$$v = \sum_{e \in E} \sum_{s \in S} f(e, s)$$

Given v and $total$, then the value of f applied to $(\overline{\text{EMP}} \cup E) \bowtie (\overline{\text{SALES}} \cup S)$ is then $total = v + total$.

However, unlike in a classical sort-merge join, the merge phase of the SMS join is assigned one additional task; it must ask the shrinking phase of the join to com-

Algorithm 2: Running a Reverse Ripple Join

```

(1) Let  $\alpha = \frac{|\text{EMP}_i|}{|\text{EMP}|}$ ,  $\beta = \frac{|\text{SALES}_i|}{|\text{SALES}|}$ 
(2) While (true)
(3)   Let  $\mathbf{K}$  be  $\max(\text{head}(\text{EMP}_i), \text{head}(\text{SALES}_i))$ 
(4)   While ( $\text{head}(\text{EMP}_i) == \mathbf{K}$ )
(5)      $e \leftarrow \text{Pop}(\text{EMP}_i)$ 
(6)     Add  $e$  to the ripple join  $RJ$ 
(7)   While ( $\text{head}(\text{SALES}_i) == \mathbf{K}$ )
(8)      $s \leftarrow \text{Pop}(\text{SALES}_i)$ 
(9)     Add  $s$  to the ripple join  $RJ$ 
(10)  If ( $\text{head}(\text{EMP}_i) == \text{head}(\text{EMP}_i) == \text{null}$ ) /* ith ripple join done */
(11)  Break
(12)  Again let  $\mathbf{K}$  be  $\max(\text{head}(\text{EMP}_i), \text{head}(\text{SALES}_i))$ 
(13)   $N_{i, (k > \mathbf{K})} = \frac{1}{\alpha\beta} \sum_{\{e, s\} \in RJ} f(e, s)$ 
(14)   $\hat{c}_{1, i, (k > \mathbf{K})} = \frac{1}{\alpha} \sum_{e \in RJ} \left[ \left( \frac{1}{\beta} \sum_{s \in RJ} f(e, s) \right)^2 + \left( 1 - \frac{1}{\beta} \right) \frac{1}{\beta} \sum_{s \in RJ} f^2(e, s) \right]$ 
(15)   $\hat{c}_{2, i, (k > \mathbf{K})} = \frac{1}{\beta} \sum_{s \in RJ} \left[ \left( \frac{1}{\alpha} \sum_{e \in RJ} f(e, s) \right)^2 + \left( 1 - \frac{1}{\alpha} \right) \frac{1}{\alpha} \sum_{e \in RJ} f^2(e, s) \right]$ 
(16)   $\hat{c}_{3, i, (k > \mathbf{K})} = \frac{1}{\alpha\beta} \sum_{\{e, s\} \in RJ} f^2(e, s)$ 

```

pute an estimate over the portion of the data space that has not yet been incorporated into *total*. At the same time that the merge phase is executed, the shrinking phase is charged with maintaining the estimator for the aggregate value associated with the region $(\overline{\text{EMP}} - \overline{\text{EMP}}) \bowtie (\overline{\text{SALES}} - \overline{\text{SALES}})$. This estimator is then added to *total* to produce an unbiased estimate for the eventual answer to the query.

To produce its estimate, the merge phase asks the shrinking phase to remove $\overline{\text{EMP}}$ and $\overline{\text{SALES}}$ from the computation of the estimator associated with $(\overline{\text{EMP}} - \overline{\text{EMP}}) \bowtie (\overline{\text{SALES}} - \overline{\text{SALES}})$ via the function call *shrink*(\mathbf{K}), where \mathbf{K} is the most recent key value to be removed from both relations. In response to this function call, *shrink* returns the value and variance of a random variable $N_{(k > \mathbf{K})}$ to the merge phase, where $N_{(k > \mathbf{K})}$ estimates the sum of f over the portion of the data space that has not yet been merged. Given this information, the current estimate for the answer to the query is simply $\text{total} + N_{(k > \mathbf{K})}$, and the variance of this estimate is $\text{Var}(N_{(k > \mathbf{K})})$. The next section considers the problem of how the shrinking phase can provide $N_{(k > \mathbf{K})}$ as well as an estimate for $\text{Var}(N_{(k > \mathbf{K})})$ to the merge phase of the SMS join.

6 The Shrinking Phase of the SMS Join

As the portion of the data space associated with the merge phase grows, the task of the shrinking phase is to update the estimator N associated with the region $(\overline{\text{EMP}} \bowtie (\text{SALES} - \overline{\text{SALES}}))$ by incrementally updating each of the estimators N_1, N_2, \dots, N_n , to reflect the fact that tuples are constantly being removed from the portion of the data space covered by the sort phase estimators. Updating N to reflect the continual loss of the tuples to the sets $\overline{\text{EMP}}$ and $\overline{\text{SALES}}$ is equivalent to removing these tuples from each of the ripple joins, recomputing the statistics associated with each individual ripple join, and then recombining the estimators N_1, N_2, \dots, N_n (as described in Section 4.2) to produce a new estimator associated with the reduced portion of the data space. The shrinking phase systematically undoes the work accomplished during the sort phase by removing tuples from each of the sort phase ripple joins.

6.1 Sufficient Statistics for the Shrinking Phase

When the merge phase begins, keys are removed one at a time from both EMP and SALES , in lexicographic order. For example, imagine that the smallest key in the data set has the value 1. The first step undertaken by the merge phase will be the removal of all tuples with key value 1 from both EMP and SALES . Thus, the shrinking phase must have access to statistics that are sufficient to estimate the result of the join $(\text{EMP} - \{\text{all tuples with key value 1}\}) \bowtie (\text{SALES} - \{\text{all tuples with key value 1}\})$. We will use the notation $N_{(k>1)}$ to denote this new estimator.

If the next key value is 2, then the next step undertaken by the merge phase is removal of all tuples with key value 2 to compute $N_{(k>2)}$. Thus, the shrinking phase should also have access to statistics that are sufficient to recompute the value of N over the relations $(\text{EMP} - \{\text{all tuples with key value 1 or 2}\}) \bowtie (\text{SALES} - \{\text{all tuples with key value 1 or 2}\})$ as well.

Just as N is computed as a combination of the individual estimators N_1, N_2, \dots, N_n , $N_{(k>\mathbf{K})}$ for the key value \mathbf{k} is computed as a combination of the individual estimators $N_{1, (k>\mathbf{K})}, N_{2, (k>\mathbf{K})}, \dots, N_{n, (k>\mathbf{K})}$, the variance of $N_{(k>\mathbf{K})}$ is estimated just as is described in Section 4, and the formulas to estimate the variance of $N_{(k>\mathbf{K})}$ is almost identical to the estimator for $\text{Var}(N)$. However, the one issue that does require careful consideration is how to *efficiently* compute these values for each $N_{i, (k>\mathbf{K})}$.

6.2 The Sort Phase Revisited

Computing $N_{(k>\mathbf{K})}$ using the algorithms of Section 4.2 requires that we have access to $N_{i, (k>\mathbf{K})}, \hat{c}_{1, i, (k>\mathbf{K})}, \hat{c}_{2, i, (k>\mathbf{K})}$ and $\hat{c}_{3, i, (k>\mathbf{K})}$ for every value of i . However, it is impractical to re-read each of the runs from disk every time that *shrink()* is called in order to compute these values. This would render the SMS join

far slower than classical disk-based join algorithms like the sort-merge join and hybrid hash join. To avoid the necessity of re-reading each of the runs, we will instead compute $N_{i, (k > \mathbf{K})}$ as well as the statistics necessary to compute the variance of $N_{(k > \mathbf{K})}$ during the sort phase of the join, at the same time that each run is in memory.

To do this, we make the following modification to Algorithm 1. After the runs EMP_i and $SALES_i$ from EMP and $SALES$ have been read into memory and joined via a ripple join, but before the runs have been written back to disk, both are sorted in *reverse order* according to $EMP.k$ and $SALES.k$. Then Algorithm 2 is executed. The execution of this algorithm is shown pictorially in Figure 6. Essentially, this algorithm operates exactly as a normal ripple join does, except that it operates over sets of tuples with the same key values, and it operates in reverse order. As this “reverse ripple join” is computed, the values for each $N_{i, (k > \mathbf{K})}$, $\hat{c}_{1, i, (k > \mathbf{K})}$, $\hat{c}_{2, i, (k > \mathbf{K})}$ and $\hat{c}_{3, i, (k > \mathbf{K})}$ are computed and stored in memory for use during the shrinking phase of the join.

Note that in general, for relatively large \mathbf{K} , only a small portion of the data space is used to compute $N_{i, (k > \mathbf{K})}$. The expected value of this estimate will become smaller with increasing \mathbf{K} . Assuming that the relative accuracy of the estimate is relatively constant over all \mathbf{K} (that is, assuming that $\frac{E[N_{i, (k > \mathbf{K})}]}{stddev(N_{i, (k > \mathbf{K})})}$ does not vary much with \mathbf{K}) $stddev(N_{i, (k > \mathbf{K})})$ for large \mathbf{K} will then be less than for small \mathbf{K} . This implies that the variance of the SMS join’s estimate decreases continuously as the shrinking phase progresses. Eventually, as the join completes, $N_{(k > \mathbf{K})}$ will always be zero and the answer to the query depends only on the value *total*. This reflects the fact that as the relations $(EMP - \overline{EMP})$ and $(SALES - \overline{SALES})$ lose more and more tuples during the merge phase of the join, most of the estimate for the eventual answer to the query is provided by the exact quantity *total* maintained by the merge phase of the join.

6.3 Reducing Main Memory Requirements

One potential problem with the tactic of pre-computing and storing these values is that if there is a large number of distinct key values in the database and/or a large number of runs needed to perform the join, then storing the statistics necessary to compute $N_{(k > \mathbf{K})}$ may require more core memory than is available. The solution to this problem is actually rather simple. We note that we do not actually need to store these statistics for *every* key value, since the statistics are used exclusively to update the running estimate for the eventual query result that is presented to the user. Extremely frequent updates of the estimate are probably overkill and will be ignored by the user anyway. In other words, if the key value 5 has been processed but we do not have $N_{i, (k > 5)}$, $\hat{c}_{1, i, (k > 5)}$, $\hat{c}_{2, i, (k > 5)}$, and $\hat{c}_{3, i, (k > 5)}$ available, then

it is not a problem; we are simply unable to give the user an estimate for the query result until we process a key value for which we *do* have the necessary statistics.

Thus, when we process the first run during the sort phase, we choose an acceptably small and yet comprehensive set of key values for which to compute statistics. In our implementation of the SMS join, this is done in the following manner. First, we compute an estimate in seconds s for the amount of time that the merge phase will take to complete (a rough estimate for this is the time that is required for both EMP and SALES to be read in their entirety from disk). Then, to see whether we will compute and store statistics for a given key value k , we hash k and mod the result by s . If the result of the mod operation is zero, then statistics are computed and stored for that key. Assuming no repeated key values, this means that in a very limited amount of main memory we will store approximately enough statistics to update the estimate for the eventual result of the query every second.

6.4 Total Main Memory Requirements

Given this optimization, the main memory required by the SMS join (above and beyond the memory required by a classical sort-merge join) is described as follows. During the sort phase, we store \hat{Q}_i , $\hat{c}_{1,i}$, $\hat{c}_{2,i}$, and $\hat{c}_{3,i}$ (amounting to around 32 bytes total) for each individual ripple join. Note that the number of individual ripple joins is analogous to the number of runs in a classical sort-merge join. Since even for the largest joins, this number will probably never exceed a few thousand, this memory requirement is generally bounded by a few kilobytes.

The sort phase also collects enough data so that we can update $N_{(k > \mathbf{K})}$ (and associated confidence bounds) every few seconds during the shrinking phase of the join. As described in Section 6.2, this requires that we have access to $N_{i,(k > \mathbf{K})}$, $\hat{c}_{1,i,(k > \mathbf{K})}$, $\hat{c}_{2,i,(k > \mathbf{K})}$ and $\hat{c}_{3,i,(k > \mathbf{K})}$ for every value of i (that is, for every individual ripple join) and for enough \mathbf{K} values that we can provide frequent estimate updates. For example, consider a join over two database tables containing 10TB of data in all. If this data were stored on 100 disks, it would take approximately three hours of continuous, sequential I/O to complete the join (at a sequential I/O rate of 50MB/sec per disk). Of this time, approximately one hour would be devoted to the shrinking phase. 1000 sets of statistics would then be enough for one update every 3.6 seconds during the shrinking phase. If approximately 10GB of RAM were devoted to this join (resulting in 1000 individual ripple joins) we might expect on the order of 32MB of main memory in total to store the required statistics: (32 bytes per set of stats) \times (1000 ripple joins) \times (1000 updates) = 32MB.

Though this is relatively little memory for a modern data warehousing system archiving terabytes of data, we do note that multiple sets of values may be required if many estimates are produced simultaneously by a single SMS join (for example, if the task to estimate an aggregate over many groups in a GROUP BY query). This may increase the memory requirement significantly. Continuing with our previous example, for a GROUP BY query with 1,000 groups we would require 32GB of storage total, as opposed to 32MB. This *is* a significant memory requirement.

However, in this case the requirement could be reduced arbitrarily by flushing this data to disk and reading it back as needed during the merge phase. Note that the statistics are both accessed and produced in an orderly fashion, based upon a linear ordering of \mathbf{K} values. Thus, reading these values from disk would cause little thrashing and would lead to only a marginal increase in the total I/O cost. In our example, an additional 32GB of I/O is negligible compared to the 30TB of I/O that would be required to evaluate such a large join using a scalable join algorithm.

7 Reducing the Performance Hit

Aside from the additional memory required by the SMS join compared to the sort-merge join, there is another concern. Might the additional computation render the SMS join significantly slower than the sort-merge join? Note that as described thus far, the following steps are required for the each sort-phase ripple join:

- (1) Read in a run of data from both EMP and SALES.
- (2) Join them using a hashed ripple join.
- (3) Sort both runs in reverse order with respect to the join key.
- (4) Rejoin the data in reverse order using a second hashed ripple join in order to collect the required $N_{i, (k > \mathbf{K})}$, $\hat{c}_{1, i, (k > \mathbf{K})}$, $\hat{c}_{2, i, (k > \mathbf{K})}$ and $\hat{c}_{3, i, (k > \mathbf{K})}$ values.
- (5) Flip the order of both runs so they are sorted in the correct order.
- (6) Write the sorted runs back to disk.

Given all of these steps, significant extra computation may be required compared to the sort phase of a classical sort-merge join. Specifically, the SMS join requires steps (2), (4), and (5) above, while the classic sort-merge join does not. The result is that a naive or straightforward implementation of the SMS join may be somewhat slow compared to typical implementation of the sort-merge join, as we will demonstrate in the next section of the paper.

Fortunately, through careful implementation it turns out that it is possible to speed up the SMS join considerably by effectively eliminating most of the additional computation that is required. To accomplish this, we note that just as no particular sort order is required by a classical sort-merge join, there is nothing particularly special about the ordering used by the reverse ripple join described in Section 6.2. In fact, any lexicographic ordering is sufficient as long as that ordering is the reverse of the ordering used by the merge phase of the SMS join. We can use this observation to our advantage in order to reduce the performance hit associated with the SMS join. In the following, we outline an implementation of the SMS join that will have very little performance penalty compared to a classical sort-merge join. There are three important considerations.

Hashing During the Sort-Phase Ripple Join

The key to achieving the desired efficiency is to carefully choose how tuples are partitioned into buckets during each sort-phase ripple join, so that we can use this bucketing to reduce the amount of computation that is required later on. To accomplish this, we make use of a randomization function H that provides a random mapping from the set of all possible join key values to the set of all possible join

key values when running each sort-phase ripple join. In other words, if the join key is a 32-bit integer, we make sure to use a hash function that is a random bijection (invertible function) from the set of all 32-bit integers to the set of all 32-bit integers.

Then, given a domain size of k for the join key, in order to determine the identifier of the bucket that tuple r is hashed into, we use the function $H(r.key) \text{ div } (k \text{ div } b)$, where b is the desired number of buckets. The result is that after the runs from both input relations have been read into memory and joined using the hashed ripple join, for each tuple pair r_1 and r_2 , if $H(r_1.key) < H(r_2.key)$, then the identifier of the bucket that r_1 is contained in is always less than the identifier of the bucket that r_2 is in. This will facilitate the following two optimizations.

“Eliminating” the Reverse Ripple Join

This ordering is important because we can effectively run our reverse ripple join without sorting (or even looking at) the tuples in a bucket, by simply looking at each bucket as a whole, in order, from back to front. Strictly speaking, this does not “eliminate” the requirement for a reverse ripple join, but it almost eliminates the additional computation associated with it. If, during the “forward” ripple join, each bucket B_i is annotated with the summations $\sum_{\{e,s\} \in B_i} f(e,s)$,

$\sum_{e \in B_i} (\sum_{s \in B_i} f(e,s))^2$, $\sum_{s \in B_i} (\sum_{e \in B_i} f(e,s))^2$, and $\sum_{\{e,s\} \in B_i} f^2(e,s)$, and

each \mathbf{K} is chosen so that \mathbf{K} is the smallest $H(r.key)$ value in some bucket m , then we can collect the statistics we need from the j th reverse ripple join by simply using the formulas:

$$N_{j,(k > \mathbf{K})} = \frac{1}{\alpha\beta} \sum_{i=m}^b \sum_{\{e,s\} \in B_i} f(e,s)$$

$$\hat{c}_{1,j,(k > \mathbf{K})} = \frac{1}{\alpha} \sum_{i=m}^b \sum_{e \in B_i} \left[\left(\frac{1}{\beta} \sum_{s \in B_i} f(e,s) \right)^2 + \left(1 - \frac{1}{\beta} \right) \frac{1}{\beta} \sum_{s \in B_i} f^2(e,s) \right]$$

$$\hat{c}_{2,j,(k > \mathbf{K})} = \frac{1}{\beta} \sum_{i=m}^b \sum_{s \in B_i} \left[\left(\frac{1}{\alpha} \sum_{e \in B_i} f(e,s) \right)^2 + \left(1 - \frac{1}{\alpha} \right) \frac{1}{\alpha} \sum_{e \in B_i} f^2(e,s) \right]$$

$$\hat{c}_{3,j,(k > \mathbf{K})} = \frac{1}{\alpha\beta} \sum_{i=m}^b \sum_{\{e,s\} \in B_i} f^2(e,s)$$

Note that these values for each and every \mathbf{K} can be obtained in a simple, linear scan of the buckets from back to front by simply maintaining a current running total for each of the four statistics.

Facilitating a Linear Time Sort

Finally, there is another important advantage to using the [bucketing](#) described above. Specifically, since each $H(r.key)$ value in a “smaller” bucket is strictly less than each $H(r.key)$ value in a “larger” bucket, if we simply internally sort the contents of each bucket, we can obtain a sorted, lexicographic ordering of all of the tuples with respect to the function H by scanning the hash buckets in order. If the

total number of tuples n is $\Omega(b)$, then there will be $O(1)$ tuples in each hash bucket and after the hash ripple join has completed, obtaining a list of tuples sorted on $H(r.key)$ is then an $O(n)$ operation (requiring only that we sort the $O(1)$ tuples in each bucket). This means that after running the sort-phase ripple join, we can obtain a sorted order that can be used by the merge phase of the join without actually sorting the tuples.

The net result of this is that aside from the fact that we will have some extra computation to produce the running estimates and confidence bounds as the data are loaded, with a careful implementation we may expect little or no performance hit in the SMS join compared to a sort-merge join. As we demonstrate experimentally in the next section, these optimizations typically cut the additional cost associated with the SMS join compared to the sort-merge join by a factor of two.

8 Benchmarking

This section details the results of a set of experiments aimed at benchmarking the SMS join. All of the data used in the experiments can be accessed at <http://www.cise.ufl.edu/~cjermain/SMS>. The goal of the experiments was to answer two basic sets of questions:

- *In reality, what are the statistical properties of the SMS join?* Specifically, does the variance of the SMS join's statistical estimator match the rather complicated formulas derived in the appendix of the paper? What is the shape of the estimator's distribution? Is it safe to assume a normal or bell-curved shape when computing confidence bounds?
- *In practice, how does the join perform?* Specifically, how do the confidence bounds given by the SMS join shrink as a function of time? How does the time to completion of the SMS join compare to its close cousin, the sort-merge join?

8.1 Statistical Properties of the SMS Join

Due to space constraints, details related to first set of experiments are included in Appendix D of the paper (not included in the print version). The key results of this set of experiments are that (1) the variance formulas in the paper do in fact seem to be accurate, and (2) that it is safe to assume normality when using those variance formulas to provide confidence bounds on the SMS join's estimates.

8.2 The SMS Join Over Large, Disk-Based Relations

8.2.1 Experimental Setup

The second set of experiments was designed to demonstrate the performance characteristics of the SMS join when it is used over a relatively large, disk-based data set, and to test the ability of the SMS join to shrink its estimator's confidence interval over time. To perform these tests, we created a series of synthetic data sets. Each data set consisted of a pair of synthetic relations, EMP and SALES, and all queries tested were of the form:

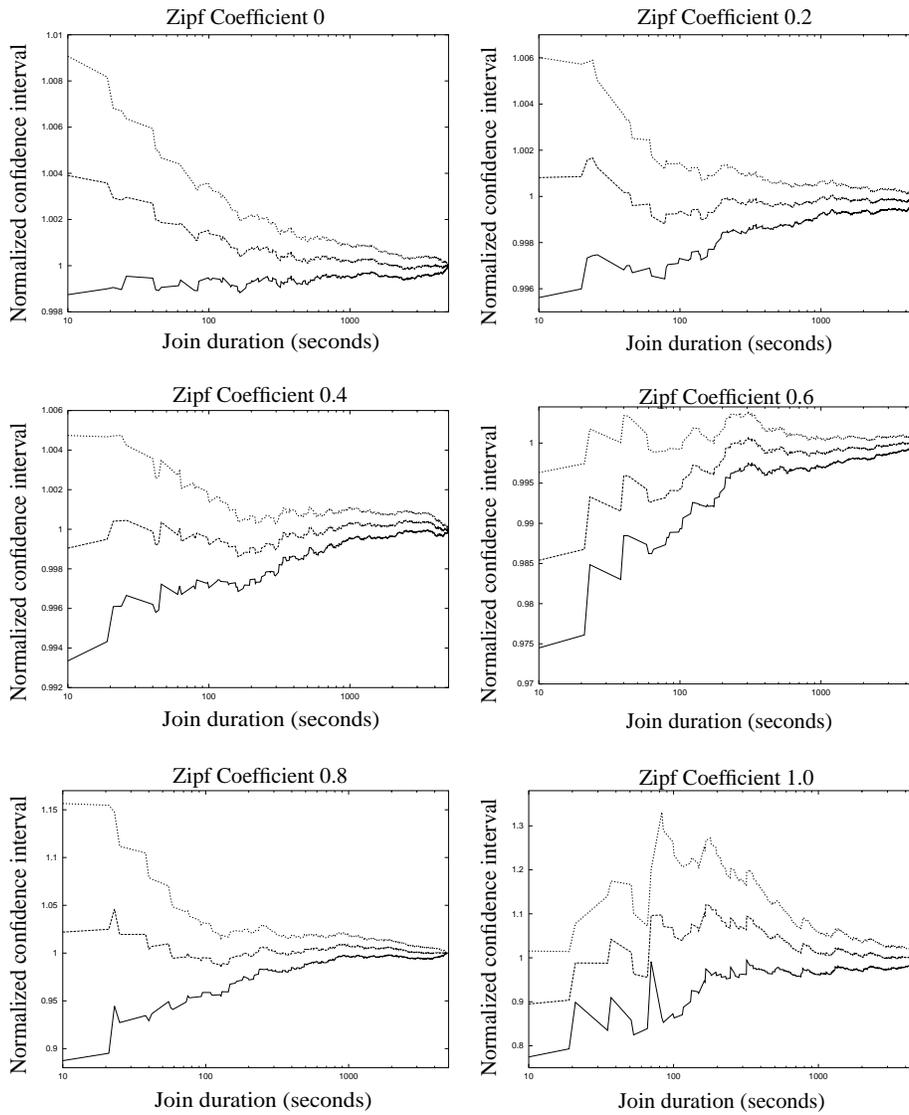


Figure 7: Confidence intervals as a function of time for six joins over synthetic data sets with varying Zipfian skew. For each join, three lines are plotted: the upper and lower limits of the confidence interval, as well as the current estimate over time. The plots begin ten seconds into the SMS join, which is approximately the time that the first ripple join runs out of memory.

```

SELECT SUM (EMP.VAL)
FROM EMP, SALES
WHERE EMP.KEY = SALES.KEY

```

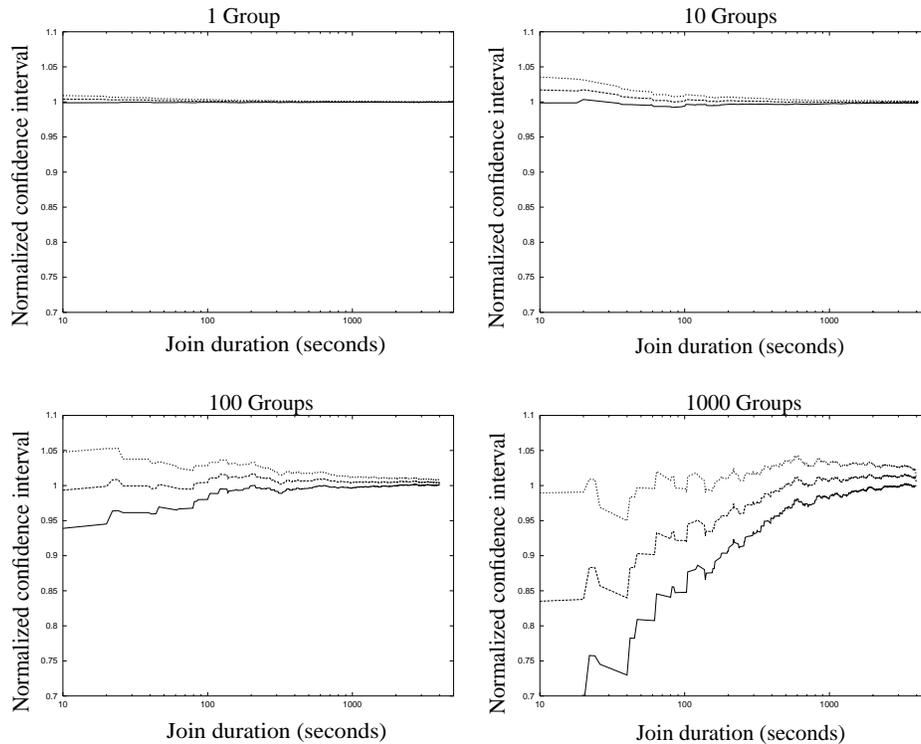


Figure 8: Confidence intervals as a function of time for a zero-skew join, for a single group in a GROUP BY query. For each join, three lines are plotted: the upper and lower limits of the confidence interval, as well as the current estimate over time. The plots begin ten seconds into the SMS join, which is approximately the time that the first ripple join runs out of memory.

Each relation tested had a physical size of 20GB, and consisted of 200 million, 100 byte tuples. The tuples in each data set shared 100,000 distinct key values, and the assignment of key values to tuples was performed according to a Zipfian distribution with a parameter C . Specifically, the most frequent key value in each relation had a frequency proportional to $1/(1^C)$, the second most frequent key value had a frequency proportional to $1/(2^C)$, the third most frequent key value had a frequency proportional to $1/(3^C)$, and so on. Frequencies of key values in both EMP and SALES were independent, so the rank of a given key value in EMP is randomly chosen and is independent of the key value's rank in SALES. The attribute value EMP . VAL for each tuple was generated using a well-behaved normal random variable with a standard deviation that was just about 3% of the variable's mean.

All data sets were stored on a single, 80GB SCSI disk. Experiments were performed on a Pentium 4 machine running the Linux operating system, with 1GB of RAM. In the experiments, six sets of EMP and SALES relations were created, using Zipf parameter C in the set $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. After creating each data set, an SMS join (without the optimizations described in Section 7) was used

to run the corresponding query to completion, and the current estimate and confidence interval were tracked throughout execution.

For comparison, a sort-merge join was also used to evaluate each query, to determine what sort of total-execution-time hit one might expect were an SMS join used instead of a more traditional join algorithm. While the classical sort-merge join is not the most advanced join algorithm available, it is still a good choice for comparison because it is ubiquitous in real database systems, its performance characteristics are very widely understood, and it forms the basis for the SMS join. We also implemented a second, optimized version of the SMS join which makes use of the ideas described in Section 7 of the paper. For each of the six data sets having Zipf parameters $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, we re-ran each join (the sort-merge join, the SMS join, and the optimized SMS join) to completion three separate times with a cold disk buffer, and compute the total running time for each of the (data set, join algorithm) combinations.

We also run an additional experiment to test the effect of other data characteristics on the utility of the SMS join. In this experiment, we keep the amount of skew fixed at 0, but add a `GROUP BY` clause to our query, which then becomes:

```
SELECT SUM (EMP.VAL) , EMP.DEPARTMENT
FROM EMP, SALES
WHERE EMP.KEY = SALES.KEY
GROUP BY EMP.DEPARTMENT
```

We vary the number of groups and test four different numbers: $\{1, 10, 100, 1,000\}$. Note that with a single group, this query is equivalent to the previous set of experiments with a 0 Zipfian skew.

The results of these experiments are depicted in Figure 7, Figure 8, and Figure 9. Figure 7 shows the estimate and confidence interval of the SMS join (at a 95% confidence level) as a function of time for each of the six tests. The confidence intervals were computed using an assumption of normality. The plot begins 10 seconds into the join, which is approximately the time at which the first ripple join runs out of memory. Figure 8 gives similar results for the `GROUP BY` experiment, though only the error bounds for the result of the query applied to the first group are plotted. Figure 9 shows the time-to-completion for the SMS join, the optimized SMS join, and for the sort-merge join over the various joins depicted in Figure 7.

8.2.2 Discussion

These experiments show two main results. First, they demonstrate that the SMS join is effective in continuing to shrink confidence bounds long after main memory has been exhausted. Each of the six plots in Figure 7 shows that, compared to the point at which a standard ripple join would run out of memory (around 10 seconds in our experiments), the SMS join is able to again cut the confidence interval width in half to a third after around 100 seconds total, and then cut it in half to a third again after 1,000 seconds total. After 1,000 seconds (approximately 30% of the SMS join's total execution time for our experiments), the confidence interval has typically shrunk to around 10-12% of the width that was achieved after only 10

Zipf parameter	SMS Join	Opt. SMS Join	Sort-Merge Join	Ratio
$C = 0$	252:52	231:25	212:14	1.19 / 1.09 / 1.0
$C = 0.2$	249:26	229:17	212:29	1.17 / 1.08 / 1.0
$C = 0.4$	250:34	225:40	206:01	1.21 / 1.09 / 1.0
$C = 0.6$	244:57	231:42	205:56	1.19 / 1.12 / 1.0
$C = 0.8$	243:07	227:25	203:57	1.19 / 1.11 / 1.0
$C = 1.0$	228:06	215:59	204:09	1.11 / 1.05 / 1.0

Figure 9: Comparison of the total time (mmm:ss) required to use the SMS join, the optimized SMS join, and a traditional, sort-merge join, to perform three separate joins over two, 20GB relations.

seconds. We point out that such a reduction in confidence interval width is fairly significant. For example, consider the case of the 0.0 Zipf coefficient. The SMS join was able to reduce the $\pm 0.5\%$ error interval at the time that main memory was exhausted to around a $\pm 0.1\%$ error interval after 1,000 seconds. This effect was fairly uniform over all six of the queries tested, with the main difference among the queries being the width of the interval at the time that main memory was exhausted. The width at the time that main memory was exhausted ranged from a confidence interval of approximately $\pm 0.5\%$ for a Zipf coefficient of 0.0 to an interval of $\pm 12\%$ for a Zipf coefficient of 0.8 to even larger for a Zipf coefficient of 1.0 (though it should be noted that the confidence interval early on in the case of the 1.0 Zipf coefficient is somewhat unreliable, due to the high variance of the confidence interval estimator itself).

Similar results were observed for the GROUP BY query, as demonstrated in Figure 8. As the number of groups increases, so does the relative inaccuracy associated with each and every group. In order to better demonstrate this, the scale of the vertical axis is held constant in each plot. For the case of a single group (which is equivalent to a query with no GROUP BY at all), the confidence intervals are very narrow already at the time that memory is exhausted. For just 100 groups, the bounds widen considerably, even with no skew in the data, the error is $\pm 5\%$ at 10 seconds. For 1,000 groups, even with no skew in the data, so little information is available at the time that the first ripple join completes that the confidence interval width is more than $\pm 20\%$ of the current estimate for the final summation over the group at 10 seconds. This clearly demonstrates that a single ripple join will be highly inaccurate in such a situation, and points out the potential utility of the SMS join.

Another point worth mentioning is that for all of the thousands of individual estimates plotted in Figure 7, more than 97% percent were “correct” in the sense that the interval did in fact contain the exact answer (despite the fact that 95% confi-

dence bounds were specified). Since these estimates were computed using the assumption that the SMS join’s estimator is normally distributed, this would seem to be yet another data point that strongly argues for the safety (and perhaps even for the conservativeness) of an assumption of normality.

We note that according to Figure 9, we observed around a 20% hit in performance compared to the sort-merge join if the SMS join is used. This performance hit can be cut in half, to around 10%, if one makes use of the optimization described in Section 7 of the paper. Another, less obvious potential benefit of the improvement described in Section 7 is demonstrated in Figure 8. Note the way that the estimate for the GROUP BY query suddenly becomes very accurate at the end of the join after approximately 3000 seconds (particularly in the case of 100 and 1,000 groups) and does not converge smoothly. The reason for this is that the selection predicate implied by the GROUP BY clause in this query was correlated with the join key due to the way that our data were generated. Such correlations may be present in real data as well. This is problematic, because the merge phase of the SMS join traverses the join keys in order in our naive implementation. Thus, the join is unable to achieve much additional accuracy during the merge phase because it does not happen to process any of the keys associated with the first group until the very end of the join. When the merge phase of the join *does* happen to process these keys, the estimate becomes very accurate, very quickly. The SMS implementation described in Section 7 of the paper will not have this problem, because the optimized SMS implementation processes the keys in a randomized order. As a result, we are guaranteed that a correlation between the order in which tuples are processed and how important they are to the answer to the join cannot exist, and the answer will converge smoothly.

Overall, two main conclusions can be drawn from these results. First, the SMS join offers clear benefits when online aggregation is a viable alternative to more traditional algorithms, as it provides a very effective means of continuing to shrink confidence bounds long after main memory is exhausted. Second, with a careful implementation the performance of the SMS join can be made very competitive with the performance of a classical sort-merge join. This means that the SMS join really does give accuracy guarantees essentially “for free” compared to the traditional sort-merge join.

9 Related Work

Online aggregation was first proposed by Hellerstein et al. [10]. Their work has its roots in a large body of earlier work concerned with using sampling to answer aggregate queries; the most well-known early papers are by Hou et al. [11][12]. Hellerstein et al.’s initial work on online aggregation eventually grew into the UC Berkeley *Control* project [9], resulting in the development of the *ripple join* [6], which serves as the basis for the sort phase estimators used by the SMS join.

[In a follow-up work to the ripple join](#), Luo, Ellman, Haas and Naughton developed a parallel join processing algorithm that maintains statistical guarantees as long as the data that are aggregated over can be buffered in main memory [32]. However, unlike the SMS join, the scalable algorithm of Luo et al. does not main-

tain any statistical guarantees once enough data have been processed that they cannot be buffered in main memory. The only existing, disk-based join algorithm that could conceivably be used to maintain statistical guarantees throughout execution over a large database as the SMS join does is the nested-loops version of the ripple join [6]. However, like other nested-loops joins, the nested-loops ripple join is not scalable. If the data set to be aggregated is n times larger than main memory, then the join will require a quadratic number of disk I/Os with respect to n .

This paper is an extended and revised version of an earlier paper that appeared in the 2005 SIGMOD Conference [22] and introduced the SMS join. The most significant technical differences between that paper and this one are related to how the accuracy guarantees associated with the join's estimate are computed. First, the SIGMOD paper's variance estimate makes use of Haas and Hellerstein's infinite population variance analysis for a single ripple join [6]. This may be a questionable decision, because the general rule-of-thumb is that infinite and finite analyses are interchangeable until 10% of a population has been considered. In the case of the SMS join, all of the database data will be considered during the computation, which results in a violation of the 10% rule. The version of the SMS join described in this paper rectifies this problem by making use of an exact, finite population analysis that is valid at all times during the SMS algorithm.

Second, the SIGMOD paper's algorithm estimates the SMS join's variance and weights the various ripple join estimators by associating an individual and independent variance estimate with each sort-phase ripple join. Unfortunately, we found this to be problematic for less well-behaved data. As the data become more skewed, it becomes likely that the variance estimates associated with one or two individual ripple join may be highly inaccurate. As a result, the weighting of the individual ripple joins may be skewed towards those joins are inaccurately classified as having lower variance. Weighting such joins heavily can then result in wild fluctuations in the estimated query answer, as well as relatively inaccurate confidence bounds. The variance estimator described in this paper makes use of all available information at all times to perform the weighting and avoid this problem, and tends to result in smoother and more reliable estimation behavior.

Finally, the SIGMOD paper's variance estimate also ignores the correlation between the various ripple joins. The SIGMOD paper argues that simply ignoring this correlation is acceptable, since the covariance among ripple join estimates tends to be negative. Thus, by ignoring the correlation the resulting accuracy bounds will actually tend to be more conservative than if the correlation had been taken into account. Unfortunately, a proof that the covariance among ripple join estimates is *always* negative has alluded us to date, and so there is a chance that data sets exist where ignoring the covariance may result in confidence bounds that are too aggressive.

There are a few other, more minor differences between the SIGMOD paper and this one. For example, this paper considers implementation issues more carefully, including how to reduce the performance hit associated with the additional computation necessitated by the SMS join compared to traditional join algorithms, and benchmarks the resulting algorithmic improvements. This paper demonstrates

experimentally the accuracy of the new variance calculations, and shows experimentally that the estimator resulting from the SMS join is approximately normal.

Aside from the work on online aggregation, the other existing work most closely related to this paper can be divided into three categories. The first category is a set of papers that describes how to implement disk-based join algorithms that are able to give early results. The algorithm of this type that is most closely related to the SMS join is a variant of the sort-merge join proposed by Dittrich et al., called the Progressive Merge Join [4]. The SMS join may be viewed as an extension to the basic ideas used to develop the Progressive Merge Join. Like the Progressive Merge Join, the SMS join reads runs from both input relations concurrently, and during the sort phase both the SMS join and the Progressive Merge Join search runs for early output results while their constituent tuples are still in memory. Aside from a few technical differences (such as the fact that the SMS join uses a hashed ripple join algorithm rather than a plane sweep to produce its early results) the primary difference between the Progressive Merge Join and the SMS Join is the addition of considerable machinery required to produce the estimate for the final query result as well as the statistical guarantees as the join progresses.

In later work, Dittrich et al. [3] consider a generalization of the Progressive Merge Join where computational resources are targeted towards specific computations in order to directly control the balance between early results and total query evaluation time. The main relevance of their generalization to the SMS join is that given the close relationship between the SMS join and the Progressive Merge Join, it may be possible to extend these ideas to the SMS join. For example, consider Figure 4 and Figure 5. Rather than continuing with the third and fourth ripple joins (as in Figure 5(e)), we could instead combine the first and second from Figure 4(d), and draw new estimates from this combination before we continue with the additional ripple joins. This may increase accuracy more quickly at the cost of a longer total query evaluation time, just as Dittrich et al. suggest. Unfortunately, an extension to the SMS join of this sort is not straightforward, and thus we leave this as an issue for future work. The difficulty is that it is not immediately clear how to combine two individual ripple joins that do not cover both input relations, so as to give statistically valid accuracy guarantees at the same time. In the SMS join it is possible to combine *all* of the sort-phase ripple join estimates online because we can treat the aggregate value associated as the totally-joined **portion** of the data space (the growing black region of Figure 5) as a constant. **If instead we attempt to merge two ripple joins together that do not cover an entire input relation, no portion of the data space is fully joined; thus, we would still need to estimate a total value for this portion of the data space and take into account the statistical properties of the process.**

Aside from the Progressive Merge Join, other “fast-first” join algorithms have been proposed. Various hash-based schemes have also been proposed for producing initial result tuples fast: the Double-Pipelined Hash Join [28], the X-Join [29], and the Hash-Merge Join [30]. In addition, another recent paper considered how to design a non-blocking algorithm for spatial joins [31].

The next two categories of related work are also closely related to the contents of this paper. The first of these two categories is a set of papers primarily focused on

producing statistical guarantees for various aggregate functions approximated by sampling over joins. The second category is a set of papers that describes algorithms using methods besides sampling to perform the same task.

We begin by briefly considering the first set of papers: those dealing with statistical guarantees for sampling over joins. The papers most closely related our work are a series of papers analyzing joins over samples, where more than one relation is sampled and the samples are joined. This series begins with an analysis of this type of join that is applicable to COUNT queries (the paper in question is due to Haas, Naughton, Seshadri and Swami [8]). The analysis was then extended to more general queries by Haas [7], and ultimately led to the bounds presented in the ripple join paper [10]. However, the key difference between the analysis presented in those papers and the analysis presented in the Appendix was outlined in Section 4.2.1: the exact, finite population, multiple-partition analysis presented here is quite different from the analysis given in those papers.

There is a larger body of work that deals with applying more traditional statistical methods to providing guarantees on sampling over joins. For example, we point out the earlier work on adaptive sampling over joins for selectivity estimation performed by Lipton and Naughton [20][16], the *join synopses* of Acharya et al. [18], and the early work by Hou and Özsoyoglu [19][11][12]. However, this work is only tangentially related to our own work by the fact that these papers consider accuracy bounds over joins. The work is not relevant in the case (as in the SMS join) where multiple input relations are sampled concurrently and on-the-fly, and the result of the join over the samples is used to estimate the eventual query result. These papers either consider sampling the output space of the join or sampling one relation and joining the sample exactly with the other. While such algorithms are amenable to traditional methods of analysis, the statistical properties of the sort of estimators used by the ripple join and the SMS join are far more difficult to reason about (hence the complexity of the analysis in the appendices of this paper). Specifically, there is no i.i.d. structure to the concurrent sampling performed by the ripple join and the SMS join, since the estimates produced by two tuples sampled from the same relation are not independent and identically distributed if they are computed by joining them with samples from another relation. This difficulty corresponds precisely to the observation made by Olken in his Ph.D. thesis [17] and later reiterated by Chaudhuri et al. [1]: sampling and the join operation do not commute. A join of two sampled relations is *not* a sample from the result of a join.

In addition to the work on sampling for estimation over joins, a final category of related work contains a small set of papers that describes algorithms computing statistical estimators with quality guarantees, where the underlying estimator is not based on sampling. For example, several recent papers advocate the use of *sketches*, an idea pioneered by Alon et al. [21][23], and later extended to more complex select-project-join queries by Dobra et al. [24]. Sketches have been the focus of even more recent work, such as a paper applying sketches to spatial joins [25]. In addition to sketches, it is also possible to use histograms to estimate the value of an aggregate function over a join, and indeed histograms have been used for many years as a way to estimate the answer to COUNT queries over joins [27] (though most work of this kind is targeted to performing selectivity estimation).

However, until very recently there has not been a serious attempt to characterize the accuracy of such estimates in a statistically rigorous fashion [26]. Though such methods can be used to develop useful statistical estimators over joins, the fundamental difference between these methods and the samples used by the SMS join is that samples can be used to produce *online estimates*. In other words, if a sample of size n is used to produce an estimate and the user does not like the accuracy of the result, then it is easy to incorporate new samples into the estimate to increase accuracy. This is why the SMS join is able to continuously increase the accuracy of its estimate throughout execution. Sketches and histograms have no such structure; they are constructed in fixed space in a single pass over the data, and so they always return an estimate with fixed precision.

10 Conclusions and Future Work

In this paper, we have introduced the SMS join, which is a join algorithm suitable for online aggregation. The SMS join provides an online estimate and associated statistical guarantees on the accuracy of the estimate from start-up through completion. The key innovation of the SMS join is that it is the first join algorithm to achieve such statistical guarantees with little or no loss in performance compared with off-line alternatives (like the sort-merge join or the hybrid hash join), even if the input tables are too large to fit in memory. Given that the SMS join combines statistical guarantees with efficiency in a disk-based environment, it would be an excellent choice for computing the value of an aggregate function over the result of a relational join.

Several main directions for future work became obvious to us during this project. Even in today's largest data warehouses (the most likely application domain for the SMS join), it is unlikely that more than a handful of relations will be too large to fit into memory. This is because in most data warehouses, having more than one or two central fact tables will be a rare occurrence. However, if there *are* more than two very large, disk-resident fact tables to join, the algorithms described in this paper are not easily applicable. It is an interesting problem to extend the SMS join to handle the multi-table joins that may occur if there are many large fact tables present in the database. We are currently considering how to handle this. An often-observed but almost always ignored problem when performing any sort of statistical estimation is that confidence bounds are themselves estimates, and can be inaccurate when only a small amount of information is used during their computation. In particular, this is the reason for the slightly strange confidence bounds depicted in Figure 7 in the case of a 1.0 Zipf coefficient. Such inaccuracy is dangerous in online aggregation since it can give the user a false sense of security, and cause the user to abort the join early. Research aimed at addressing this danger by deriving safer and less aggressive variance estimates early on is an important future direction. One intriguing idea for allowing even faster reduction in confidence bound width would be to identify during the sort phase those important tuples which join with many other tuples in the data set, and to buffer those tuples in memory and never write them back to disk. This would result in a method that resembles an online version of the bifocal sampling algo-

rithm of Ganguly et al. [5]. Finally, we mention once again the discussion from Section 9 on incorporating the ideas of Dittrich et al. [3] regarding the use of multiple merge phases to trade-off accuracy and join completing time. Working out the technical problems associated with describing the accuracy of such an approach would make the SMS join a much more flexible tool.

References

- [1] S. Chaudhuri, R. Motwani, V.R. Narasayya: On Random Sampling over Joins. SIGMOD Conference 1999: 263-274
- [2] W. Cochran: Sampling Techniques. Wiley and Sons, 1977
- [3] J.-P. Dittrich, B. Seeger, D.S. Taylor, P. Widmayer: On producing join results early. PODS 2003: 134-142
- [4] J.-P. Dittrich, B. Seeger, D.S. Taylor, P. Widmayer: Progressive Merge Join: A Generic and Non-blocking Sort-based Join Algorithm. VLDB 2002: 299-310
- [5] S. Ganguly, P.B. Gibbons, Y. Matias, A. Silberschatz: Bifocal Sampling for Skew-Resistant Join Size Estimation. SIGMOD 1996: 271-281
- [6] P.J. Haas, J.M. Hellerstein: Ripple Joins for Online Aggregation. SIGMOD 1999: 287-298
- [7] P.J. Haas: Large-Sample and Deterministic Confidence Intervals for Online Aggregation. SSDBM 1997: 51-63
- [8] P.J. Haas, J. F. Naughton, S. Seshadri, A. N. Swami: Selectivity and Cost Estimation for Joins Based on Random Sampling. J. Com. Syst. Sci. 52(3): 550-569 (1996)
- [9] J.M. Hellerstein, R. Avnur, A. Chou, C. Hidber, C. Olston, V. Raman, T. Roth, P.J. Haas: Interactive Data Analysis: The Control Project. IEEE Computer 32(8): 51-59 (1999)
- [10] J.M. Hellerstein, P.J. Haas, H.J. Wang: Online Aggregation. SIGMOD 1997: 171-182
- [11] W.-C. Hou, G. Özsoyoglu, B.K. Taneja: Statistical Estimators for Relational Algebra Expressions. PODS 1988: 276-287
- [12] W.-C. Hou, G. Özsoyoglu, B.K. Taneja: Processing Aggregate Relational Queries with Hard Time Constraints. SIGMOD Conference 1989: 68-77
- [13] G. Luo, C. Ellmann, P.J. Haas, J.F. Naughton: A scalable hash ripple join algorithm. SIGMOD Conference 2002: 252-262
- [14] L.D. Shapiro: Join Processing in Database Systems with Large Main Memories. ACM Trans. Database Syst. 11(3): 239-264 (1986)
- [15] J. Shao: Mathematical Statistics. Springer-Verlag, 1999.
- [16] R.J. Lipton, J.F. Naughton: Query Size Estimation by Adaptive Sampling. PODS 1990: 40-46
- [17] F. Olken. Random Sampling from Databases. PhD thesis, University of California, Berkeley, 1993
- [18] S. Acharya, P.B. Gibbons, V. Poosala, S. Ramaswamy: Join Synopses for Approximate Query Answering. SIGMOD Conference 1999: 275-286

- [19] W.-C. Hou, G. Özsoyoglu: Statistical Estimators for Aggregate Relational Algebra Queries. *ACM Trans. Database Syst.* 16(4): 600-654 (1991)
- [20] R.J. Lipton, J.F. Naughton, D.A. Schneider: Practical Selectivity Estimation through Adaptive Sampling. *SIGMOD Conference 1990*: 1-11
- [21] N. Alon, P.B. Gibbons, Y. Matias, M. Szegedy: Tracking Join and Self-Join Sizes in Limited Storage. *PODS 1999*: 10-20
- [22] C. Jermaine, A. Dobra, S. Arumugam, A. Pol, S. Joshi: A Disk-Based Join With Probabilistic Guarantees. *SIGMOD Conference 2005*: 587-598.
- [23] N. Alon, P.B. Gibbons, Y. Matias, M. Szegedy: Tracking Join and Self-Join Sizes in Limited Storage. *J. Comput. Syst. Sci.* 64(3): 719-747 (2002)
- [24] A. Dobra, M.N. Garofalakis, J. Gehrke, R. Rastogi: Processing complex aggregate queries over data streams. *SIGMOD Conference 2002*: 61-72
- [25] A. Das, J. Gehrke, M. Riedewald: Approximation Techniques for Spatial Data. *SIGMOD Conference 2004*: 695-706
- [26] A. Dobra: Histograms Revisited: When are histograms the best approximation method for aggregates over joins? *PODS 2005*: 228-237
- [27] R.P. Kooi: The optimization of queries in relational databases, PhD thesis, CWR University, 1980.
- [28] Z.G. Ives, D. Florescu, M. Friedman, A.Y. Levy, D.S. Weld: An Adaptive Query Execution System for Data Integration. *SIGMOD Conference 1999*: 299-310
- [29] T. Urhan, M.J. Franklin: XJoin: A Reactively-Scheduled Pipelined Join Operator. *IEEE Data Eng. Bull.* 23(2): 27-33 (2000)
- [30] M.F. Mokbel, M. Lu, W.G. Aref: Hash-Merge Join: A Non-blocking Join Algorithm for Producing Fast and Early Join Results. *ICDE 2004*: 251-263
- [31] G. Luo, J.F. Naughton, C. Ellmann: A Non-Blocking Parallel Spatial Join Algorithm. *ICDE 2002*: 697-705
- [32] G. Luo, C. Ellmann, P.J. Haas, J.F. Naughton: A scalable hash ripple join algorithm. *SIGMOD Conference 2002*: 252-262
- [33] S. Chen, A. Ailamaki, P.B. Gibbons, T.C. Mowry: Improving Hash Join Performance through Prefetching. *ICDE 2004*: 116-127

Appendix A: Variance and Covariance Analysis

In this appendix, we compute formulas for the weighted sum of variances of individual estimators and the covariance terms discussed in Section 4.2. Note that our analysis methodology differs somewhat from the analogous analysis given by Haas and Hellerstein [6], in that we use an exact, finite-population analysis, rather than a large-sample, infinite population analysis.

We consider the general problem of computing the following aggregate over the cross-product of the two input relations EMP and SALES:

$$Q = \sum_{(e,s) \in \text{EMP} \times \text{SALES}} f(e,s) \quad (2)$$

where $f(e, s)$ is the value of the aggregate function applied to tuple (e, s) . As described in Section 3.1, this is the most general aggregate of the type SUM we can compute over the cross-product, and it can easily be used to implement any join by incorporating the join and selection predicates in the computation of $f(e, s)$, i.e., we can define $f(e, s)$ to be 0 for tuples not in the join. While this formula is compact and matches the intuition, it is not very useful in the analysis. Using \mathbb{E} as shorthand for EMP and \mathbb{S} as shorthand for SALES, Q can be rewritten as:

$$Q = \sum_{(e,s) \in \mathbb{E} \times \mathbb{S}} f(e, s) = \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f(e, s) \quad (3)$$

by observing that, to get the tuples in the cross-product, we combine every tuple e in \mathbb{E} with every tuple s in \mathbb{S} .

To model the weighted combination of the ripple join estimators described in the paper, let $\mathbb{E}_1, \dots, \mathbb{E}_n$ and $\mathbb{S}_1, \dots, \mathbb{S}_n$ be randomly selected, non-overlapping subsets of \mathbb{E} and \mathbb{S} , respectively. To model the random partitioning of \mathbb{E} and \mathbb{S} , we introduce two families of random variables: $X_{e,i}$ and $Y_{s,i}$. The random variable $X_{e,i}$ takes the value 1 if $e \in \mathbb{E}_i$ and 0 otherwise. $Y_{s,i}$ is similarly defined. Since the partitionings of \mathbb{E} and \mathbb{S} are performed independently, the random variables $X_{e,i}$ and $Y_{s,i}$ are always independent. With this, N_i , the i th ripple join estimator, is:

$$\begin{aligned} N_i &= \frac{|\mathbb{E}||\mathbb{S}|}{|\mathbb{E}_i||\mathbb{S}_i|} \sum_{e \in \mathbb{E}_i} \sum_{s \in \mathbb{S}_i} f(e, s) \\ &= \frac{|\mathbb{E}||\mathbb{S}|}{|\mathbb{E}_i||\mathbb{S}_i|} \sum_{(e,s) \in \mathbb{E}_i \times \mathbb{S}_i} f(e, s) \\ &= \frac{|\mathbb{E}||\mathbb{S}|}{|\mathbb{E}_i||\mathbb{S}_i|} \sum_{e \in \mathbb{E}_i} \sum_{s \in \mathbb{S}_i} X_{e,i} Y_{s,i} f(e, s) \end{aligned} \quad (4)$$

where $\frac{|\mathbb{E}||\mathbb{S}|}{|\mathbb{E}_i||\mathbb{S}_i|}$ makes the correction for the size of \mathbb{E}_i and \mathbb{S}_i with respect to the size of the complete relations.

The overall estimate N is simply the weighted sum over all N_i :

$$N = \sum_{i=1}^n w_i N_i \quad (5)$$

subject to the constraint that $\sum_{i=1}^n w_i = 1$.

We now show N that is an unbiased estimator of Q , the aggregate over the join, and we analyze the variance of N . The following properties of each variable $X_{e,i}$ are useful for the theoretical development (similar properties hold for $Y_{s,i}$):

$$\begin{aligned}
E[X_{e,i}] &= \frac{|\mathbb{E}_i|}{|\mathbb{E}|} \\
E[X_{e,i} X_{e',j}] &= \begin{cases} \frac{|\mathbb{E}_i|}{|\mathbb{E}|} & \text{if } e = e', i = j \\ 0 & \text{if } e = e', i \neq j \\ \frac{|\mathbb{E}_i|(|\mathbb{E}_i| - 1)}{|\mathbb{E}|(|\mathbb{E}| - 1)} & \text{if } e \neq e', i = j \\ \frac{|\mathbb{E}_i||\mathbb{E}_j|}{|\mathbb{E}|(|\mathbb{E}| - 1)} & \text{if } e \neq e', i \neq j \end{cases} \\
&= \frac{|\mathbb{E}_i|}{|\mathbb{E}|} \delta_{e,e'} \delta_{i,j} + \frac{|\mathbb{E}_i|}{|\mathbb{E}|(|\mathbb{E}| - 1)} (1 - \delta_{e,e'}) [(|\mathbb{E}| - 1) \delta_{i,j} + |\mathbb{E}| (1 - \delta_{i,j})] \\
&\quad \text{where } \delta_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \quad (6)
\end{aligned}$$

Note that we expressed the cases using the Kronecker delta symbol that takes the value 1 if both arguments are the same, 0 if they are not. The most important property of $\delta_{i,j}$ is the fact that:

$$\sum_i \sum_j \delta_{i,j} g(i,j) = \sum_i g(i,i) \quad (7)$$

To see why Equation 6 holds, we use the fact that the expectation of the 0, 1 random variables is the probability that the condition on which they depend holds. Furthermore, we notice that:

- (a) when $e = e'$ and $i = j$, the two random variables are identical and we are computing the variable's square. Since the variable always takes the value zero or one, the expected value of its square is its expected value;
- (b) when $e = e'$ but $i \neq j$, the same tuple has to be in two different partitions, but this never happens;
- (c) when $e \neq e'$ and $i = j$, the result is the probability that two particular tuples are in the partition i ; this is simply the probability that e is in \mathbb{E}_i multiplied by the probability that e' is in \mathbb{E}_i , given that e is in \mathbb{E}_i ; since a position in both \mathbb{E}_i and \mathbb{E} is decided, the conditional probability is $(|\mathbb{E}_i| - 1) / (|\mathbb{E}| - 1)$;

(d) when $e \neq e'$ and $i \neq j$, we have a similar situation as in the previous case, but we do not occupy a position in \mathbb{E}_i .

When $i = j$, Equation 6 takes the particular form:

$$\begin{aligned} E[X_{e,i} X_{e',i}] &= \frac{|\mathbb{E}_i|}{|\mathbb{E}|} \left(\delta_{e,e'} + \frac{|\mathbb{E}_i| - 1}{|\mathbb{E}| - 1} (1 - \delta_{e,e'}) \right) \\ &= \frac{|\mathbb{E}_i|}{|\mathbb{E}|} \left(\frac{|\mathbb{E}| - |\mathbb{E}_i|}{|\mathbb{E}| - 1} \delta_{e,e'} + \frac{|\mathbb{E}_i| - 1}{|\mathbb{E}| - 1} \right) \end{aligned} \quad (8)$$

When $i \neq j$, Equation 6 takes the particular form:

$$E[X_{e,i} X_{e',j}] = \frac{|\mathbb{E}_i| |\mathbb{E}_j|}{|\mathbb{E}| (|\mathbb{E}| - 1)} (1 - \delta_{e,e'}) \quad (9)$$

Now we can prove:

$$\begin{aligned} E[N_i] &= \frac{|\mathbb{E}| |\mathbb{S}|}{|\mathbb{E}_i| |\mathbb{S}_i|} \sum_{e \in \mathbb{E}_i} \sum_{s \in \mathbb{S}_i} E[X_{e,i} Y_{s,i}] f(e, s) \\ &= \sum_{e \in \mathbb{E}_i} \sum_{s \in \mathbb{S}_i} f(e, s) = Q \end{aligned} \quad (10)$$

which follows directly from linearity of expectation, Equation 5 and independence of X and Y . Since N is a linear combination of N_i 's, it is unbiased as well.

The variance of N can be expressed in terms of the variance and covariance of N_i 's:

$$\text{Var}(N) = \sum_i^n w_i^2 \text{Var}(N_i) + \sum_i^n \sum_{j \neq i}^n w_i w_j \text{Cov}(N_i, N_j) = Q_V + Q_C \quad (11)$$

where Q_V and Q_C simply refer to the first and second terms, respectively. Now we are ready to introduce the central theoretical result of the paper:

Theorem A.1: The sum of the variances.

$$\begin{aligned} Q_V &= \sum_i^n w_i^2 \text{Var}(N_i) = \\ &= \frac{|\mathbb{E}| |\mathbb{S}| \sum_i^n w_i^2}{(|\mathbb{E}| - 1)(|\mathbb{S}| - 1)} \left[Q^2 \sum_i^n \left(\frac{(|\mathbb{E}_i| - 1)(|\mathbb{S}_i| - 1)}{|\mathbb{E}_i| |\mathbb{S}_i|} - \frac{(|\mathbb{E}| - 1)(|\mathbb{S}| - 1)}{|\mathbb{E}| |\mathbb{S}|} \right) + \right. \\ &+ c_1 \sum_i^n \left(\frac{(|\mathbb{E}| - |\mathbb{E}_i|)(|\mathbb{S}_i| - 1)}{|\mathbb{E}_i| |\mathbb{S}_i|} \right) + c_2 \sum_i^n \left(\frac{(|\mathbb{E}_i| - 1)(|\mathbb{S}| - |\mathbb{S}_i|)}{|\mathbb{E}_i| |\mathbb{S}_i|} \right) + \\ &\left. c_3 \sum_i^n \left(\frac{(|\mathbb{E}| - |\mathbb{E}_i|)(|\mathbb{S}| - |\mathbb{S}_i|)}{|\mathbb{E}_i| |\mathbb{S}_i|} \right) \right] \end{aligned}$$

where:

$$\begin{aligned} c_1 &= \sum_{e \in \mathbb{E}} \left(\sum_{s \in \mathbb{S}} f(e, s) \right)^2 \\ c_2 &= \sum_{s \in \mathbb{S}} \left(\sum_{e \in \mathbb{E}} f(e, s) \right)^2 \\ c_3 &= \sum_{(e, s) \in \mathbb{E} \times \mathbb{S}} f^2(e, s) \end{aligned}$$

Proof. We begin with the formula:

$$\begin{aligned} Q_V &= \sum_i^n w_i^2 \text{Var}(N_i) = \sum_i^n w_i^2 [E[N_i^2] - E[N_i]^2] \quad (12) \\ &= \sum_i^n w_i^2 \left(-Q^2 + \frac{|\mathbb{E}|^2 |\mathbb{S}|^2}{|\mathbb{E}_i|^2 |\mathbb{S}_i|^2} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \right. \\ &\quad \left. E[X_{e,i} X_{e',i}] E[Y_{s,i} Y_{s',i}] f(e, s) f(e', s') \right) \end{aligned}$$

Using Equation 8 applied to both relations, the inner sum becomes:

$$\begin{aligned} &\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} E[X_{e,i} X_{e',i}] E[Y_{s,i} Y_{s',i}] f(e, s) f(e', s') \\ &= \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \frac{|\mathbb{E}_i|}{|\mathbb{E}|} \left(\frac{|\mathbb{E}| - |\mathbb{E}_i|}{|\mathbb{E}| - 1} \delta_{e, e'} + \frac{|\mathbb{E}_i| - 1}{|\mathbb{E}| - 1} \right) \frac{|\mathbb{S}_i|}{|\mathbb{S}|} \left(\frac{|\mathbb{S}| - |\mathbb{S}_i|}{|\mathbb{S}| - 1} \delta_{s, s'} + \frac{|\mathbb{S}_i| - 1}{|\mathbb{S}| - 1} \right) \times \\ &\quad f(e, s) f(e', s') \\ &= \frac{|\mathbb{E}_i| |\mathbb{S}_i|}{|\mathbb{E}| |\mathbb{S}| (|\mathbb{E}| - 1) (|\mathbb{S}| - 1)} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} [(|\mathbb{E}| - |\mathbb{E}_i|) (|\mathbb{S}| - |\mathbb{S}_i|) \delta_{e, e'} \delta_{s, s'} + \\ &\quad (|\mathbb{E}| - |\mathbb{E}_i|) (|\mathbb{S}_i| - 1) \delta_{e, e'} + (|\mathbb{S}| - |\mathbb{S}_i|) (|\mathbb{E}_i| - 1) \delta_{s, s'} + (|\mathbb{E}_i| - 1) (|\mathbb{S}_i| - 1)] \times \\ &\quad f(e, s) f(e', s') \\ &= \frac{|\mathbb{E}_i| |\mathbb{S}_i|}{|\mathbb{E}| |\mathbb{S}| (|\mathbb{E}| - 1) (|\mathbb{S}| - 1)} [(|\mathbb{E}_i| - 1) (|\mathbb{S}_i| - 1) Q^2 + (|\mathbb{E}| - |\mathbb{E}_i|) (|\mathbb{S}_i| - 1) c_1 + \\ &\quad (|\mathbb{E}_i| - 1) (|\mathbb{S}| - |\mathbb{S}_i|) c_2 + (|\mathbb{E}| - |\mathbb{E}_i|) (|\mathbb{S}| - |\mathbb{S}_i|) c_3] \end{aligned}$$

To perform this simplification, we used the rule in Equation 7 to get the following identities:

$$\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \delta_{e,e'} f(e, s) f(e', s') = \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{s' \in \mathbb{S}} f(e, s) f(e, s') = c_1$$

$$\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \delta_{s,s'} f(e, s) f(e', s') = \sum_{s \in \mathbb{S}} \sum_{e \in \mathbb{E}} \sum_{e' \in \mathbb{E}} f(e, s) f(e', s) = c_2$$

$$\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \delta_{e,e'} \delta_{s,s'} f(e, s) f(e', s') = \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f^2(e, s) = c_3$$

$$\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} f(e, s) f(e', s') = \left(\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f(e, s) \right)^2 = Q^2$$

Replacing the formula we have obtained for the quadruple sum in equation 12 and performing straightforward simplifications completes the derivation. \square

A simple corollary of Theorem A.1 is the following:

Corollary A.2: Exact variance of an individual ripple join. *After an individual ripple join has processed a fraction $1/n$ of \mathbb{E} and $1/n$ of \mathbb{S} , the variance of its associated estimator is:*

$$\text{Var}(N_i) = \frac{(n-1)}{(|\mathbb{E}|-1)(|\mathbb{S}|-1)} \times \\ [(n+1-|\mathbb{E}|-|\mathbb{S}|)Q^2 + |\mathbb{E}|(|\mathbb{S}|-n)c_1 + |\mathbb{S}|(|\mathbb{E}|-n)c_2 + (n-1)|\mathbb{E}||\mathbb{S}|c_3]$$

\square

Now we consider the covariances.

Theorem A.3: The sum of the covariances.

$$Q_C = \sum_{i=1}^n \sum_{j \neq i}^n w_i w_j \text{Cov}(N_i, N_j) = (1 - \sum_i^n w_i^2) Q_U$$

where:

$$Q_U = \frac{|\mathbb{E}||\mathbb{S}|}{(|\mathbb{E}|-1)(|\mathbb{S}|-1)} \left[\frac{(|\mathbb{E}|-|\mathbb{S}|-1)}{|\mathbb{E}||\mathbb{S}|} Q^2 - (c_1 + c_2 - c_3) \right]$$

Proof: Again, we begin with the formula:

$$Q_C = \sum_i^n \sum_{j \neq i}^n w_i w_j \text{Cov}(N_i, N_j) = \sum_i^n \sum_{j \neq i}^n w_i w_j [E[N_i N_j] - E[N_i] E[N_j]]$$

$$\begin{aligned}
&= \sum_i^n \sum_{j \neq i}^n w_i w_j \left(-Q^2 + \frac{|\mathbb{E}|^2 |\mathbb{S}|^2}{|\mathbb{E}_i| |\mathbb{E}_j| |\mathbb{S}_i| |\mathbb{S}_j|} \times \right. \\
&\quad \left. \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} E[X_{e,i} X_{e',j}] E[Y_{s,i} Y_{s',j}] f(e, s) f(e', s') \right)
\end{aligned}$$

Just as for Theorem A.1, we first derive an expression for the quadruple sum in the formula above. Using Equation 9 (which is the relevant formula), we have:

$$\begin{aligned}
&\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} E[X_{e,i} X_{e',j}] E[Y_{s,i} Y_{s',j}] f(e, s) f(e', s') \\
&= \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} \frac{|\mathbb{E}_i| |\mathbb{E}_j|}{|\mathbb{E}| (|\mathbb{E}| - 1)} (1 - \delta_{e, e'}) \frac{|\mathbb{S}_i| |\mathbb{S}_j|}{|\mathbb{S}| (|\mathbb{S}| - 1)} (1 - \delta_{s, s'}) f(e, s) f(e', s') \\
&= \frac{|\mathbb{E}_i| |\mathbb{E}_j|}{|\mathbb{E}| (|\mathbb{E}| - 1)} \frac{|\mathbb{S}_i| |\mathbb{S}_j|}{|\mathbb{S}| (|\mathbb{S}| - 1)} \times \\
&\quad \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{e' \in \mathbb{E}} \sum_{s' \in \mathbb{S}} (1 - \delta_{e, e'} - \delta_{s, s'} + \delta_{e, e'} \delta_{s, s'}) f(e, s) f(e', s') \\
&= \frac{|\mathbb{E}_i| |\mathbb{E}_j|}{|\mathbb{E}| (|\mathbb{E}| - 1)} \frac{|\mathbb{S}_i| |\mathbb{S}_j|}{|\mathbb{S}| (|\mathbb{S}| - 1)} (Q^2 - c_1 - c_2 + c_3)
\end{aligned}$$

Again, we use the same simplification rules derived during the computation of Q_V . Replacing this formula in the expression for Q_C and using the fact that $\sum_i \sum_{j \neq i} w_i w_j = 1 - \sum_i w_i^2$ (since $(\sum_i w_i)^2 = 1$) completes the derivation. \square

An interesting corollary of this is the following:

Corollary A.4: Effect of ripple join sizes on covariance. *The total covariance among all partitions is independent of the partition sizes that are chosen.*

Proof. The proof follows directly from the Theorem A.3. \square

Appendix B: Analysis of the Natural Estimators for $\hat{c}_{1,i}$ and $\hat{c}_{2,i}$

The final mathematical issue we formally explore in the Appendix is an analysis of the natural estimator for c_1 (an almost identical analysis applies to c_2). As in the

main body of the paper, we let $\alpha_i = \frac{|\mathbb{E}_i|}{|\mathbb{E}|}$, $\beta_i = \frac{|\mathbb{S}_i|}{|\mathbb{S}|}$. Just as in the previous set of proofs, we define a zero/one random variable $X_{e,i}$ that is one if and only if the tuple e from EMP has been processed by the i th ripple join (the variable $Y_{s,i}$ is

defined similarly for SALES). With this, using the definition of $\hat{c}'_{1,i}$ from Section 4.2.4 we have:

$$\begin{aligned}
E[\hat{c}'_{1,i}] &= E\left[\frac{1}{\alpha_i} \sum_{e \in \mathbb{E}} X_{e,i} \left(\frac{1}{\beta_i} \sum_{s \in \mathbb{S}} Y_{s,i} f(e,s)\right)^2\right] \\
&= \frac{1}{\alpha_i \beta_i^2} \sum_{e \in \mathbb{E}} E[X_{e,i}] E\left[\left(\sum_{s \in \mathbb{S}} Y_{s,i} f(e,s)\right)^2\right] \\
&= \frac{1}{\beta_i^2} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{s' \in \mathbb{S}} E[Y_{s,i} Y_{s',i}] f(e,s) f(e,s')
\end{aligned}$$

To obtain the first simplification, we use the fact that the random variables $X_{e,i}$ and $Y_{s,i}$ are independent. The second simplification uses the fact that $E[X_{e,i}] = \alpha_i$ and linearity of expectation. Now, using the formula for $E[Y_{s,i} Y_{s',i}]$ given by Equation 8 (applied to the SALES relation):

$$\begin{aligned}
E[\hat{c}'_{1,i}] &= \frac{1}{\beta_i^2} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{s' \in \mathbb{S}} \beta_i \left(\frac{(|\mathbb{S}| - |S_i|)}{|\mathbb{S}| - 1} \delta_{s,s'} + \frac{|S_i| - 1}{|\mathbb{S}| - 1} \right) f(e,s) f(e,s') \\
&= \frac{1}{\beta_i} \frac{(|\mathbb{S}| - |S_i|)}{|\mathbb{S}| - 1} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f^2(e,s) + \frac{1}{\beta_i} \frac{(|S_i| - 1)}{|\mathbb{S}| - 1} c_{1,i}
\end{aligned}$$

where we use the following simplifications making use of Equation 8:

$$\begin{aligned}
\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{s' \in \mathbb{S}} f(e,s) f(e,s') &= c_{1,i} \\
\sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} \sum_{s' \in \mathbb{S}} \delta_{s,s'} f(e,s) f(e,s') &= \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f^2(e,s)
\end{aligned}$$

If we observe that $|\mathbb{S}| \gg 1$ and $|S_i| \gg 1$ (so that $\frac{(|\mathbb{S}| - |S_i|)}{|\mathbb{S}| - 1} \approx (1 - \beta_i)$), then:

$$E[\hat{c}'_{1,i}] \approx \frac{1 - \beta_i}{\beta_i} \sum_{e \in \mathbb{E}} \sum_{s \in \mathbb{S}} f^2(e,s) + c_{1,i}$$

Appendix C: What If EMP and/or SALES Are Not Clustered Randomly?

All of the material in the paper thus far (including all of the material in Appendix A and Appendix B) assumes that a randomized ordering of EMP and SALES is available. Even if such an ordering is *not* available, then the SMS join can still be applied after a few additional considerations have been taken into account.

The most obvious way to allow for an SMS join over non-randomly ordered data would be to randomly select tuples from EMP and SALES one-at-a-time, and use the resulting randomly ordered tuples as input into the SMS join. This would allow the SMS join algorithm as described thus far to be used with no additional modification. For example, one could first randomly permute the numbers from 1 to $|\text{EMP}|$ to obtain a list L . Then, the tuples from EMP would be accessed in the order mandated by L . The first tuple from EMP to be processed would be $L[1]$, the second would be $L[2]$, the third $L[3]$, and so on, up to $L[|\text{EMP}|]$. A similar process could be used for SALES. However, this solution is probably not a good one, since it would introduce several additional costs, not the least of which would be the tremendous additional I/O delay incurred due to the access of all of the records from EMP (and SALES) randomly one-at-a-time. This would likely render the SMS join many orders-of-magnitude slower than a scalable alternative such as the sort-merge join or the hybrid hash join.

A more suitable application of this basic technique would be to logically partition EMP and SALES into a large number of contiguous runs of blocks. The resulting relations would then be viewed as two new relations EMP' and SALES' of size $|\text{EMP}'|$ and $|\text{SALES}'|$, respectively. [A run of blocks in EMP is treated as a single "tuple" in EMP'. The SALES relation is treated similarly.](#) Then, one could run the SMS join over the resulting "relations" EMP' and SALES' using exactly the technique described in the previous paragraph, with each run accessed in a randomized fashion. The only additional change to the SMS join's algorithms and analysis is that since each $e \in \text{EMP}'$ and $s \in \text{SALES}'$ are actually runs of blocks rather than single tuples, in order to compute $f(e, s)$ one must use $f(e, s) = \sum_{\tilde{e} \in e} \sum_{\tilde{s} \in s} f(\tilde{e}, \tilde{s})$ where \tilde{e} and \tilde{s} are tuples from EMP and SALES, respectively.

The reason that this solution is preferable to random access of individual tuples is that while accessing tuples randomly is very costly, accessing random runs of blocks is much less so because the randomization is amortized over all of the tuples in the run. For example, if one assumes a 5ms random access time and a 50MB/sec sequential read rate, then accessing a large relation as a series 10MB random sequential reads incurs an overhead of less than 3% due to the random I/O, compared to a single sequential scan.

Still, whenever possible, the SMS join should be used over randomly ordered input relations. The reason for this is that if there is correlation in the ordering of result tuples which tends to result in large $f^2(e, s)$ values for random (e, s) pairs where $e \in \text{EMP}'$ and $s \in \text{SALES}'$, then using a solution such as the one detailed in this appendix may result in confidence bounds that shrink much more slowly compared to the randomly-permuted case.

Appendix D: Experiments: Statistical Properties of the SMS Join

D.1 Experimental Setup

[The first set of questions posed in the introduction to Section 8 was addressed using a simple set of experiments.](#) The SMS join algorithm was used to repeatedly

% query completed	Query #1		Query #2		Query #3		Query #4	
	Pred	Observed	Predicted	Observed	Pred	Obs	Predicted	Observed
10%	9549.0	9558.3	8.8244e8	8.8252e8	37583	37738	4.176e12	4.191e12
20%	4362.3	4364.4	4.0302e8	4.0286e8	16659	16600	1.885e12	1.878e12
30%	2633.5	2625.6	2.4321e8	2.4243e8	9684.4	9677.6	1.122e12	1.120e12
40%	1769.0	1772.1	1.6330e8	1.6358e8	6197.0	6187.9	7.402e11	7.387e11
50%	1250.4	1253.5	1.1536e8	1.1567e8	4104.6	4088.7	5.111e11	5.097e11
60%	1038.5	1042.0	9.5418e7	9.5777e7	2789.9	2787.3	3.576e11	3.573e11
70%	758.92	758.35	6.9156e7	6.9122e7	1825.6	1828.0	2.422e11	2.424e11
80%	527.32	524.56	4.7882e7	4.7629e7	1226.9	1231.6	1.584e11	1.589e11
90%	271.64	271.37	2.4252e7	2.4225e7	490.37	490.40	7.056e10	7.057e10

Figure 10: Observed versus theoretically predicted estimator variances for each of the four test queries over real data sets. The variances for the queries are shown at various stages of query execution.

answer four SQL queries over a set of small, real-life relations. Each SMS join was repeated enough times that it was possible to empirically estimate the variance and probability density function associated with the SMS join's estimator. In this way, the formulas for the estimator's variance could be checked for correctness, and the SMS join's estimator could be checked to see if an assumption of normality is reasonable when attaching confidence bounds to the estimate. The queries that were run all involved the following four relations:

- **PEAKS**(NAME, HEIGHT). This relation is made of 1,817 tuples, each of which gives the height of a mountain peak in the state of California.
- **CLIMBED**(NAME, DATE, CLIMBER). This relation describes 852 mountain climbs around the world that were completed by members of the Sierra Peaks Section (SPS) of the Angeles Chapter of the Sierra Club. CLIMBED.NAME is a foreign key into PEAKS.NAME.
- **FCAT**(ZIP, SCHOOL, SCORE). This relation stores the average third-grade FCAT reading scores for 1,446 elementary schools in Florida, as well as the zip code for each school. The FCAT is a scholastic achievement test administered to all public-school students in Florida at various grade levels.
- **AGRICULTURE**(ZIP, OUTPUT). This relation describes the total agricultural output (in dollars) for each of the 812 zip codes in the state of Florida. FCAT.ZIP is a foreign key into AGRICULTURE.ZIP.

The first two queries tested are the SUM and COUNT SQL queries needed to answer the question: “What is the average height of all of the California peaks climbed by SPS members since January 1, 2000?”¹ The SQL query for this is:

```
SELECT AVG(PEAKS.HEIGHT)
FROM PEAKS, CLIMBED
WHERE PEAKS.NAME = CLIMBED.NAME AND
CLIMBED.DATE > 'Jan 1, 2000'
```

These two queries are subsequently referred to as *Query #1* and *Query #2*, respectively. The next two queries tested were the SUM and the COUNT queries needed to answer the question: “What is the average FCAT reading score for schools located in a zip code with more than \$500,000 in annual agricultural production?” The SQL query for this is:

```
SELECT AVG(FCAT.SCORE)
FROM FCAT, AGRICULTURE AS A
WHERE FCAT.ZIP = A.ZIP AND A.OUTPUT > 500000
```

These two queries are subsequently referred to as *Query #3* and *Query #4*, respectively. To accurately measure the statistical properties of the SMS join as it is used to answer these queries, it was necessary to re-run each query a large number of times. For each query, the following three steps were repeated one million times:

- (1) First, the tuples in each input relation were randomly re-ordered. This is achieved by using the C library function `lrand48()` to associate a 32 bit random number with each tuple in the relation; the tuples are then sorted based upon this random number. At the beginning of the one million trials, `lrand48()` was seeded with the current wall clock time.
- (2) Next, an SMS join over the input relations was performed, using five individual sort phase ripple joins.
- (3) As the SMS join was executed, the current estimate was output nine times. The current estimate was first output after one sort phase ripple join had completed, then again after two had completed, then again after three, four, and five had completed, and again after the sort phase ripple joins had shrunk by 20%, 40%, 60% and 80% during the merge and shrink phases of the SMS join. This process resulted in nine different estimates per join, which roughly correspond to the estimates produced when the join is 10% complete, 20% complete, 30% complete, and so on up to 90% complete.

One thing to note is that our experiments were performed over a set of rather small input relations; certainly the relations were much smaller than the relations one would expect to typically use with an SMS join (experiments over much larger input relations are described in the next subsection). We chose to perform the tests

1. Note that the SQL AVERAGE aggregate function actually computes the ratio of two SQL queries: a SUM query and a COUNT query.

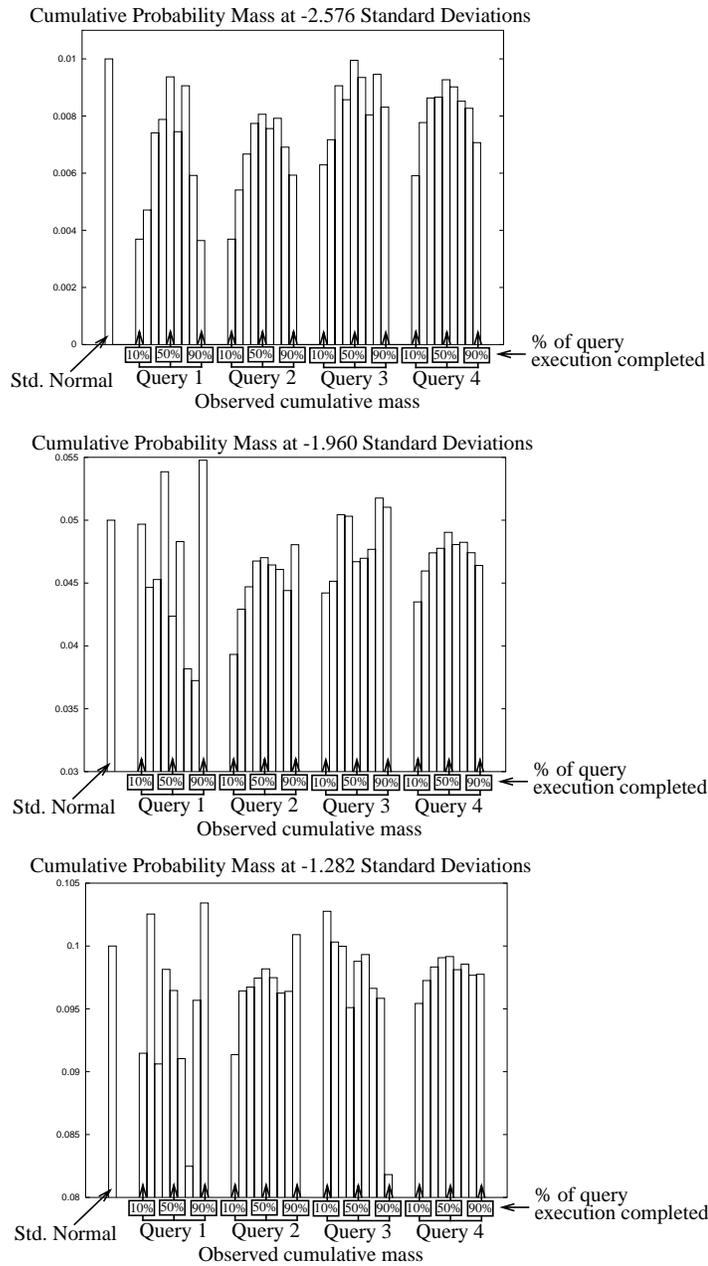


Figure 11: Observed distribution for the SMS join's estimator compared to the standard normal distribution, over four real queries, at various stages of query completion (10%, 20%, ..., 90%).

over such small relations for two primary reasons. First and most obvious, running millions of queries over a large, disk-based database (as is necessary to measure the statistical properties of the join) is simply not possible.

Second and equally important, a fundamental question that we wanted to answer with this set of experiments is: Is the estimator produced by the SMS join normally distributed, particularly with respect to the tail of the distribution? The reason that it makes more sense to test for normality on a small data set is that due to the central limit theorem, normality typically increases with sample size. So if the SMS join is shown to produce relatively normal results even at the end of the shrinking phase over a data set having only a few thousand tuples, it would strongly imply that the SMS join does in fact produce a normally distributed estimator for nearly any database size. Results showing that the SMS join produces a normally distributed estimate for very large sample sizes would be a much weaker observation, indicating only that the SMS join's estimator is normally distributed for very large data sets.

D.2 Experimental Results

The results of this first set of experiments are given in Figure 10, Figure 11, and Figure 12. Figure 10 shows the observed variance over the one million query repetitions for each of the four queries, as well as the theoretically expected variance, calculated using the formulas given in Appendix A of the paper. These values are given at each of the stages of query completion described above.

Figure 11 and Figure 12 compare the observed distribution for the SMS join's estimator to the standard normal distribution. These Figures depict six individual plots. The first plot compares the cumulative probability mass of the standard normal distribution at -2.576 standard deviations from the mean to the same statistic for the SMS join's estimator. This comparison is shown for each of the four queries, at each of the nine stages of query completion (10%, 20%, 30%, and so on). -2.576 standard deviations was chosen because this is the point at which the cumulative probability mass for the standard normal distribution is exactly 1%. The next five plots show similar results for five different numbers of standard deviations from the mean: -1.960 standard deviations, -1.282 standard deviations, 1.282 standard deviations, 1.960 standard deviations, and 2.576 standard deviations. These numbers were chosen because they correspond to the points where the cumulative mass of the standard normal distribution is 5%, 10%, 90%, 95%, and 99%, respectively.

D.3 Discussion

Figure 10 gives very strong evidence that our analytic derivations for the variance of the SMS join's estimator are in fact correct. Nearly all of the 36 observed variance values given in this table agree with the predicted values in at least the first two decimal places, and most of the values are at least close in the third decimal place.

In addition, Figure 11 and Figure 12 also seem to show that it is very safe to assume that the SMS join's estimator is normally distributed for the purpose of giving the user confidence bounds. To understand why such a claim is justified, it is important to consider exactly what evidence would indicate that an assumption

of normality is in fact *not* safe. While considering this question, it will be useful to consider Figure 11 and Figure 12 separately.

Imagine that some of the plots in Figure 11 showed a significant number of bars that were higher than the bar for the standard normal distribution. This would indicate that much of the probability mass for the SMS join's estimator lies in the left or lower tail of the estimator's distribution. This, in turn, might indicate that confidence bounds based on an assumption of normality could be dangerous, since such bounds would tend to understate the probability that the SMS join's estimator would give out a result that is too small. However, very few of the bars depicted in Figure 11 are higher than those of the standard normal distribution. This is a strong indication that normality is a safe assumption in the left tail of the distribution of the SMS join's estimator. The results show that an assumption of normality is reasonable even in the case when the SMS join is only 10% complete and when it is fully 90% complete. These are expectedly the cases where the central limit theorem is least applicable, and thus if an assumption of normality is invalid anywhere, it should be invalid in either of these two cases. The 10% case is difficult because it makes use of only a single sort phase ripple join, and the 90% case is difficult because only a few dozen tuples remain in each of the ripple joins.

Now consider Figure 12. Since all of the plots depicted in Figure 12 consider the cumulative probability mass at points greater than the mean, the problematic case here would be any plots which show bars that are *shorter* than the bar for the standard normal distribution. This would indicate some amount of probability mass that lies far out in the right or upper tail of the distribution of the SMS join's estimator, which again could cause confidence bound problems. In fact, it turns out that most of the SMS join's bars *are* lower than the standard normal distribution's at 1.960 and 2.576 standard deviations. However, we argue that the difference is not very significant. For example, at 2.576 standard deviations, all of the SMS join's bars show that the SMS join's estimator always has at least a cumulative probability mass of at least 0.982 (compared to 0.990 for the standard normal distribution) and 30 out of 36 of the bars show a cumulative probability mass of at least 0.986. This translates to only a very, very small difference between the SMS join's estimator and the standard normal distribution in the right tail, even at the relatively small sample sizes tested. For the larger sample sizes used in more realistic joins, this disparity would likely disappear altogether. Thus, we argue that the results show that an assumption of normality is safe in both the left and right tails of the SMS join's estimator.