

# Offset-value coding in database query processing

Goetz Graefe  
 Google Inc  
 GoetzG@Google.com

Thanh Do\*  
 Celonis Inc  
 T.Do@Celonis.com

## ABSTRACT

Recent work [8] shows how offset-value coding speeds up database query execution, not only sorting but also duplicate removal and grouping (aggregation) in sorted streams, order-preserving exchange (shuffle), merge join, and more. It already saves thousands of CPUs in Google’s Napa and F1 Query systems, e.g., in grouping algorithms and in log-structured merge-forests. In order to realize the full benefit of interesting orderings, however, query execution algorithms must not only consume and exploit offset-value codes but also produce offset-value codes for the next operator in the pipeline. Our research has sought ways to produce offset-value codes without comparing successive output rows one-by-one, column-by-column. This short paper introduces a new theorem and, based on its proof and a simple corollary, describes in detail how order-preserving algorithms (from filter to merge join and even shuffle) can compute offset-value codes for their outputs. These computations are surprisingly simple and very efficient.

## 1 INTRODUCTION

Tree-of-losers priority queues [11, 22] and offset-value coding [6, 20] enable the most efficient sort algorithms. When sorting a database table with  $N$  rows, the provable lower bound for row comparisons is  $\log_2(N!) \approx N \log_2(N/e)$  with  $e \approx 2.718$ . External merge sort using tree-of-losers priority queues for both run generation and merging exceeds this lower bound by only 1-2%. Offset-value coding truncates shared row prefixes and turns prefix sizes into order-preserving surrogate keys. These surrogate keys limit the count of column value comparisons for  $N$  rows with  $K$  key columns to  $N \times K$  [8]. Importantly, there is no  $\log(N)$  factor. Offset-value codes alone decide the remainder of the  $\log_2(N!)$  row comparisons.

Recent work shows how offset-value coding complements “interesting orderings” [27] to speed up database query processing, not only sorting but also duplicate removal and grouping (aggregation) in sorted streams, order-preserving exchange (shuffle), merge join, and more [8, 9]. For the full benefit of interesting orderings, all execution algorithms relying on sort order must also exploit offset-value coding in their comparisons of rows and columns. Moreover, order-preserving query execution algorithms must not only consume but also produce offset-value codes, to be consumed and exploited by the next operator in the pipeline.

Nevertheless, computing offset-value codes in order-preserving query execution algorithms has not received any attention, and the only method known to-date – comparing an operator’s output row-by-row, column-by-column – is so expensive that it renders producing offset-value across database operators much less attractive or even worthless. In contrast, new, simple techniques

efficiently compute offset-value codes for order-preserving algorithms in database query execution. The new techniques do not require comparisons row-by-row, column-by-column; instead, offset-value codes in the output depend only on offset-value codes in the inputs. There is no need for additional column value comparisons beyond those required in the operation itself, e.g., column value comparisons in the merge logic of merge join. The new techniques are implemented in Google’s Napa [1] and F1 Query [26, 28] systems.

## 2 RELATED PRIOR WORK

The context of our work are Google’s Napa [1] and F1 Query [26, 28] systems. Napa is a data warehouse that maintains thousands of materialized views in log-structured merge-forests [25]. F1 Query is a federated query processing platform that executes SQL queries over tables in various Google storage systems such as Spanner [7], BigTable [5], and Napa. Both Napa and F1 Query employ sort order for efficient data access and data manipulation.

Pioneering work on sorting established the benefits of tree-of-losers priority queues and of offset-value coding, both for run generation and external merge steps [6, 11, 20, 22]. Pioneering work in the database field established the value of sort-based query execution, notably merge join but also duplicate removal and grouping [3, 10, 27]. Surveys on database query evaluation [13, 15] cover sorting and offset-value coding but fail to recognize offset-value coding as a significant opportunity throughout sort-based query execution. Recent work [8, 9] fills that gap but fails to offer efficient derivation of offset-value codes for output rows. The present short paper fills this remaining gap.

## 3 BACKGROUND: OFFSET-VALUE CODING AND TREE-OF-LOSERS PRIORITY QUEUES

A tree-of-losers priority queue [11, 22], also known as a tournament tree, embeds a balanced binary tree in an array, with the tree’s unary root in array slot 0. It is efficient due to leaf-to-root passes with one comparison per tree level; root-to-leaf passes with two comparisons per tree level are not required. As in a sports tournament where each round of competition eliminates one contestant, the principal rules are that (i) two candidate key values compete at each binary node in the tree and (ii) after a comparison of two candidates, the loser remains in the node and the winner becomes a candidate in the next tree level. Thus, in a priority queue with  $N$  entries, a new overall winner reaches the root after  $\log_2(N)$  comparisons.

When merging runs, a fixed pair of runs competes at each leaf node. Run generation merges “sorted” runs of a single row each. Run generation by replacement selection can try to extract longer sorted runs from the unsorted input: one additional comparison per input row doubles the expected run size, halves the run count, and saves one comparison per row during merging.


Figure 1, adapted from [22], shows a tree-of-losers priority queue merging  $F = 12$  sorted runs. The left half with runs 0-7 is cropped from Figure 1. The dashed boxes represent the current candidates from runs to be merged; the solid boxes are tree nodes

\*Work done at Google Inc.

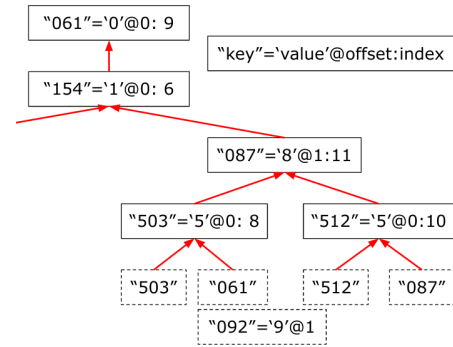
© 2023 Copyright held by the owner/author(s). Published in Proceedings of the 26th International Conference on Extending Database Technology (EDBT), 28th March-31st March, 2023, ISBN 978-3-89318-092-9 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

**Table 1: Offset-value codes in a sorted file or stream.**

Rows and their column values	Descending OVC			Ascending OVC		
	offset	domain - value	OVC	arity - offset	value	OVC
5 7 3 9	0	95	95	4	5	405
5 7 3 12	3	88	388	1	12	112
5 8 4 6	1	92	192	3	8	308
5 9 2 7	1	91	191	3	9	309
5 9 2 7	4	-	400	0	-	0
5 9 3 4	2	97	297	2	3	203
5 9 3 7	3	93	393	1	7	107



**Figure 1: Strings in a tree-of-losers priority queue.**



**Figure 2: Offsets and values in a tree-of-losers priority queue.**

in a tree-of-losers priority queue. The example node in the top-right corner explains the numbers in each node.

The overall smallest key value, “061”, is in the tree’s root node. Interestingly, key value “154” is above key value “087”. After “154” emerged as the winner from the left subtree and “061” from the right subtree, “154” was the loser in the “final match” of this “tournament” and was left behind. The runner-up of the right subtree, “087”, had to remain within that subtree.

The next step in the merge logic moves key value “061” from the tree root to the merge output. The following step gets a new key value from input 9, the origin of key value “061”. This successor key starts a new leaf-to-root pass at the the same leaf. This leaf-to-root pass bubbles the next lowest key value to the root in  $\log_2(F)$  steps: If that next key value is “092”, it wins over “503” but loses to “087” and is left behind; then “087” wins over “154” and reaches the root.

Offset-value coding [6] captures one row’s key value relative to another key that is earlier in the sort sequence. Offset-value codes are set after comparisons. A loser’s new offset is the position where the keys first differ, e.g., a column index; the value is the loser’s data value at that offset. Equivalently, the offset is the size of the shared prefix. For example, a duplicate key shares the entire key and hence has an offset equal to the key size.

Table 1 shows the derivation of descending and ascending offset-value codes in a stream of rows in ascending order on all columns. Each key is encoded relative to its predecessor. With four sort columns, the arity of the sort key is 4; the domain of each column is 1...99. Descending offset-value codes take the actual offset but the negative of the column value. Ascending offset-value codes take the negative offset but the actual column value. Table 1 ignores that small key domains permit encoding multiple key columns together.

If two rows and their key values  $A$  and  $B$  are encoded relative to the same key  $C$ , and if the offsets of  $A$  and  $B$  differ, then the one

with the higher offset is earlier in the sort sequence. Otherwise, if the two data values at that offset differ, then these data values decide the comparison. Otherwise, additional data values in  $A$  and  $B$  must be compared. With offsets and values combined as shown in Table 1, often a single integer comparison can sort two keys  $A$  and  $B$  encoded relative to the same key  $C$ .

If  $A$  sorts lower (earlier) than  $B$ , then  $B$  is the loser in this comparison and can be encoded relative to the winner  $A$ . If offset-value codes decided the comparison, the offset-value code of  $B$  relative to  $C$  is also the offset-value code of  $B$  relative to  $A$ . Otherwise, the offset for  $B$  increases by the count of column comparisons required to determine winner and loser, and the value is taken at that new offset. Similar rules apply if  $A$  is the loser in the comparison.

In a tree-of-losers priority queue with offset-value coding, each tree node was a loser to the local winner and the local offset-value code is relative to the local winner. Along the overall winner’s leaf-to-root path, all key values are coded relative to its key value.

Figure 2 shows the same key values as Figure 1 but adds offset-value codes. In the figure, each node indicates an offset and a character at that offset. Although an offset-value code for an ascending sort order requires negating either the offset or the value, this is not shown in Figure 2. In an implementation, each node in the tree-of-losers priority queue contains only an offset-value code and an index; strings remain in the input buffers reachable because the index values are run identifiers.

Recall that input runs are encoded with prefixes truncated. For example, “092” is encoded relative to its predecessor “061”. After “061” moves from the tree root to the merge output, “092” starts a leaf-to-root pass. Keys “503”, “087”, and “154” on this path are all encoded relative to the prior overall winner “061”. Offset 1 is earlier in the sort order than offset 0; therefore, “092” is

earlier than “503”. “503” stays in place and “092” moves up. There, offset 1 equals offset 1, but data value ‘8’ is less than data value ‘9’; therefore, “092” remains behind as “087” moves up. Finally, “087” wins over “154” and reaches the root.

In this leaf-to-root pass, offset-value codes decide all three comparisons, two by offsets and one by data values encoded within offset-value codes. Not a single string comparison is required and not a single offset-value code needs re-calculation.

Merging sorted runs replaces a winner with its successor in the same merge input. With merge inputs’ fixed assignment to leaf nodes, a successor retraces the leaf-to-root path of the prior overall winner. Since this successor as well as all keys on its leaf-to-root path are encoded relative to the prior overall winner, offset-value coding applies to all comparisons in a tree-of-losers priority queue.

Offset-value codes decide many comparisons in a tree-of-losers priority queue. Column value comparisons are required only if two rows have equal offset-value codes, with column comparisons starting at the offset. After such a row comparison is decided, the count of column value comparisons is added to the loser’s offset.

With  $K$  sort columns, the sum of all offset increments is limited to  $K$  in each row; in an input with  $N$  rows, the sum of all increments and thus the count of all column value comparisons are limited to  $N \times K$ . Importantly, there is no  $\log(N)$  multiplier. Thus, tree-of-losers priority queue and offset-value coding guarantee that the effort for column value comparisons is linear in the count of rows and in the count of sort columns, quite like the effort for computing hash values in hash-based query execution.

Comparisons of offset-value codes are free if they are subsumed in other algorithm activities and data structures. In quick-sort, for example, the inner-most loop not only compares key values but also looping indexes: when the loops from the left and the right meet, the current partitioning step is complete. In priority queues, the inner-most loop compares key values only after testing whether there even are key values. During queue construction, some entries might not have been filled yet; after the end of some merge inputs, some queue entries no longer have valid keys; and during run generation by merging many single-row runs, there practically is only queue build-up and tear-down.

A tree-of-losers priority queue for run generation by replacement selection may hold key values  $-\infty$  and  $+\infty$  as early and late fences, valid keys assigned to the current run, and valid keys for the next run. These cases need some indicator field in each queue entry, but they require only 2 bits. Even with multiple early and late fences, e.g., one for each merge input [15], 30 or 62 bits remain and can hold an offset-value code. A single comparison instruction can test whether the two relevant keys are valid and compare their offset-value codes: if two rows have the same 32- or 64-bit value, then they are both valid (neither early nor late fences), they go to the same output run, they differ from their shared base row (an earlier winner) at the same offset, they have the same value at that offset, and the next step must compare further columns. As the comparison of offset-value codes is already complete when the indicator fields have been inspected for administration of the inner-most loop, offset-value code comparisons are free.

Offset-value coding can speed up not only merge sort but also order-preserving exchange (merging shuffle), merge join, set operations such as intersection, and more. For example, in a query like “select..., count (distinct ...) group by...”, the sort can

detect duplicate rows by offsets equal to the column count and, after the sort, in-stream aggregation can detect group boundaries by offsets smaller than the grouping key.

#### 4 OFFSET-VALUE CODING IN RELATIONAL QUERY EXECUTION OPERATORS

What has been missing to-date are simple and efficient techniques for deriving offset-value codes in the output of database operators other than merge sort. Surprisingly, a new theorem and its corollary enable the wanted solution with no additional column value comparisons beyond those required for computing output rows.

**Definitions:** For key values  $A$  and  $B$ , let  $pre(A, B) \geq 0$  be the length of their maximal shared prefix; let  $val(A, i)$  with  $i \geq 0$  be the data at offset  $i$  within key value  $A$ ,” let  $val(A, B) = val(B, pre(A, B))$  be the first difference in key value  $B$  relative to key value  $A$ ; let  $A < B$  mean that  $A$  sorts lower (earlier) than  $B$ ; and let  $ovc(A, B)$  with  $A < B$  be the offset-value code of key value  $B$  relative to key value  $A$ , computed from  $pre(A, B)$  and  $val(A, B)$  as in Table 1.

**Proposition:** For key values  $A < B < C$  with  $A \neq B$  or  $B \neq C$ ,  $ovc(A, B) \neq ovc(B, C)$ .

Proof (by contradiction):  $ovc(A, B) = ovc(B, C)$  would imply that  $C$  has the same data value as  $B$  at the same offset, e.g., column index, but this would violate the definition of offset-value codes, which requires the maximal shared prefix.

Examples: Table 1 shows no successive equal offset-value codes.

**Theorem:** For key values  $A < B < C$ , in ascending offset-value coding  $ovc(A, C) = \max(ovc(A, B), ovc(B, C))$  and in descending offset-value coding  $ovc(A, C) = \min(ovc(A, B), ovc(B, C))$ .

Proof (for ascending offset-value codes in an ascending sort, in three cases by the lengths of maximal shared prefixes): (i) If  $pre(A, B) > pre(B, C)$ , then  $pre(A, C) = pre(B, C)$ ,  $val(A, C) = val(B, C)$ , and  $ovc(A, C) = ovc(B, C)$ . With  $ovc(A, B) < ovc(B, C)$ , the theorem holds. (ii) Otherwise, if  $pre(A, B) < pre(B, C)$ , then  $pre(A, C) = pre(A, B)$ ,  $val(A, C) = val(A, B)$ , and  $ovc(A, C) = ovc(A, B)$ . With  $ovc(A, B) > ovc(B, C)$ , the theorem holds. (iii) Otherwise,  $pre(A, B) = pre(B, C)$  and, by the lengths of maximal shared prefixes,  $val(A, B) < val(B, C)$ . Thus,  $pre(A, C) = pre(B, C)$ ,  $val(A, C) = val(B, C)$ , and  $ovc(A, C) = ovc(B, C)$ . With  $ovc(A, B) < ovc(B, C)$ , the theorem holds.

Examples: Case (i) in the proof applies to the first three rows in Table 1. If the second row were removed, then the offset-value codes of the third row would change in neither ascending nor descending offset-value coding. As an example of case (ii), if the second-to-last row were removed in Table 1, the offset-value codes of the last row would be those of the removed row. As an example of case (iii), if the third row were removed in Table 1, the offset-value codes of the fourth row would remain unchanged.

**Corollary** (the new “filter theorem”): Our new theorem above extends to multiple intermediate keys: for a sorted list of key values  $X_0 < X_1 < \dots < X_{n-1} < X_n$  and ascending offset-value coding,  $ovc(X_0, X_n) = \max_{i=1 \dots n} ovc(X_{i-1}, X_i)$ .

Proof (sketch): By repeated application of the theorem.

Implication: When a filter drops rows from a sorted stream, simple and efficient integer calculations can derive offset-value codes for the output from offset-value codes of the input.

**Table 2: Offset-value codes after a filter.**

Rows and their column values				Ascending OVC		
				<i>arity - offset</i>	<i>value</i>	OVC
5	7	3	9	4	5	405
5	9	3	7	3	9	309

#### 4.1 Filter

A filter with a predicate computes offset-value codes for its output by directly applying the “filter theorem” above: an output row’s offset-value code is (in ascending encoding) the maximum of its offset-value code in the input and of the offset-value codes of rows that failed the filter predicate since the prior output row. The same applies *mutatis mutandis* for descending offset-value coding and for offset-value coding using byte offsets within normalized keys. Table 2 illustrates the calculation for ascending offset-value codes with the data of Table 1, assuming that only the first input row and the last input row satisfy the filter predicate.

#### 4.2 Projection

Projection, i.e., removal of input columns as well as calculation of new columns from existing columns (all within a single row), typically does not change the sort order. “Relationally pure” projection, however, includes removal of duplicate rows, which of course might change the sequence of rows (see Section 4.4).

If all columns in the sort key survive the projection, offset-value codes in the output are the same as in the input. If not, the offset must be limited to the prefix (column count) that survives.

#### 4.3 Segmented sorting

Segmented query execution means that a single stream is divided into segments, one after another, that can be processed one at a time. A typical example is a stream sorted on (A, B) but required sorted on (A, C) – one can either sort the entire stream on (A, C) or one can segment the input on distinct values of (A) and sort each segment only on (C). In this example, A, B, and C can be individual columns or lists of columns.

To segment a sorted stream with offset-value codes, inspection of these code values suffices: an offset smaller than the segmentation key indicates a segment boundary. All other offsets must be cut to the size of the segmentation key, to be extended again by the sort within each segment. In the example above, all offsets within a segment are cut to the size of (A); the associated value is the first column value of (C). Sorting a segment on (C) refines these offset-value codes in the usual way to reflect the sort order on (A, C). For a more concrete example using Table 1, segmenting on the first two columns does not require any column value comparisons; an offset of less than two (and corresponding ranges of offset-value codes) indicates a boundary between segments.

#### 4.4 Duplicate removal

In a sorted stream with offset-value codes, duplicate removal suppresses input rows with offsets equal to the arity (count of columns). In Table 1, for example, an offset of 4 (and corresponding offset-value codes) indicates a duplicate row. All other rows, i.e., the output rows, retain their offset-value codes from the input. In the duplicate-free output, no row has an offset equal to the arity. The same applies *mutatis mutandis* if the sort key is reduced based on functional dependencies [29].

#### 4.5 Grouping and aggregation

In a stream with offset-value codes sorted on a “group by” list, grouping aggregates input rows with offsets equal to or larger than the “group by” list. In the aggregation output, no row has an offset equal to or larger than the “group by” list. The output rows retain the offset-value codes of the first row in each group of input rows. In Table 1, for example, grouping on the first two columns can use offset-value codes similarly to segmentation (see Section 4.3).

#### 4.6 Pivoting

Pivoting turns rows into columns, e.g., from (year, month, sales) to (year, january\_sales... december\_sales). In many aspects, including the set of useful algorithms, pivoting is like grouping and aggregation. This applies in particular to the benefit of offset-value codes in the input and the calculation of offset-value codes in the output.

#### 4.7 Merge join

The logic of merge join is similar to an external merge sort; hence, it can exploit offset-value codes in its two sorted inputs. Minor variations of merge inner join can provide all join types as well as set operations such as intersection. The merge logic supports cross-table equality predicates, e.g., a primary key and a foreign key; a subsequent filter can enforce other cross-table predicates (see Section 4.1, with caveats for semi joins and outer joins).

Semi joins (SQL “exists” sub-queries) and anti semi joins (SQL “not exists” sub-queries) select rows from one input based on a join predicate rather than an filter predicate within a row (see Section 4.1). Nonetheless, the rule for setting offset-value codes in the output is the same as given in the “filter theorem” above, just like the derivation of Table 2 from Table 1 as discussed in Section 4.1.

Inner joins are similar: the offset-value codes of one input are preserved in the output, rows without match affect the offset-value code of the next row with a match, and rows with duplicate matches must have offset-value codes for duplicate keys.

In summary, the merge join logic is similar to a merge step in an external merge sort: it might require comparisons of column values but it can compute offset-value codes for the output without additional column values comparisons.

#### 4.8 Nested-loops join, look-up join

In the “filter theorem” above, it does not matter whether input rows fail a single-table predicate in a filter, a two-table predicate in a semi join, or a many-table predicate in nested iteration. Thus, the “filter theorem” above applies here just as much as in Sections 4.1 and 4.7. Nested-loops join or lookup join can be order-preserving. Note that there is no requirement that the join predicate is an equality predicate. We ignore here right semi join, right anti semi join, right outer join, and full outer join, which leaves left semi join, left anti semi join, inner join, and left outer join; we assume that the left input is the outer input and the right input is the inner input.

#### 4.9 Order-preserving shuffle

An order-preserving one-to-many “splitting” shuffle resembles a filter with respect to each output partition (see Section 4.1), because each output stream is a selection from the overall input stream. An order-preserving many-to-one “merging” shuffle requires the standard merge logic, very similar to a merge step in

an external merge sort. A tree-of-losers priority queue can exploit offset-value codes in the input and produce new offset-value codes in the output.

#### 4.10 Ordered scans

Data access is a source of offset-value codes as important as sorting. All sorted scans can produce offset-value codes.

Column storage is often sorted with the leading key columns compressed by run-length encoding. Fortunately, as described in recent work [8], such scans can produce row-by-row offset-value codes without sorting and even without any column value accesses or column value comparisons. Thus, these scans can provide offset-value codes practically for free.

Traditional b-trees readily support sorted scans. Page-wide prefix compression [2] gives offset-value coding a head start; compression within index leaves by next-neighbor difference, e.g., in Shore [4], provides offset-value codes practically for free.

Today's ubiquitous log-structured merge-forests and stepped-merge trees [21, 25] support a sorted scan per partition, which can include offset-value codes; a merge of such scans benefits from offset-value codes and merge logic using a tree-of-losers priority queue readily produces offset-value codes for the merge output.

In non-unique secondary indexes, lists of row identifiers are usually sorted and compressed using one of the above compression techniques and thus can deliver such lists with offset-value codes. Range queries need to merge lists of row identifiers; again, the merge logic consumes, benefits from, and produces offset-value codes. Multi-dimensional b-tree access, e.g., MDAM [23], similarly merges sorted lists of row identifiers. Sorted lists of row identifiers are similarly useful for index intersection and index join, i.e., "covering" a query in "index-only retrieval" with multiple secondary indexes of the same table [18].

#### 4.11 Summary of new techniques

In the sort-based and order-preserving query execution algorithms considered above, calculation of new offset-value codes in the operations' output can be simple and efficient. Far from requiring row-by-row column-by-column data comparisons, offset-value codes in the output depend only on offset-value codes in the inputs, as proven by a new theorem and a corollary. There is no need for additional column value comparisons beyond those performed in the operation itself, e.g., column value comparisons in the merge logic of merge join. Ordered storage structures, e.g., b-trees, can preserve the effort for comparisons spent during index creation -- they can do so by storing offset-value codes explicitly, by prefix truncation (encoding each record relative to its immediate predecessor), or by run-length encoding of leading sort columns. Scans over either format can readily produce offset-value codes and, with the new techniques of this paper, sort-based operations can easily carry them forward, i.e., produce offset-value codes in their output derived from offset-value codes in their input.

### 5 IMPLEMENTATION EXPERIENCE

This section briefly summarizes how offset-value coding is used in F1 Query [26, 28]. At a high level, F1 Query employs offset-value coding in its sorting logic and takes advantage of offset-value codes in intermediate results (e.g., spilled files) and across relational operators to further improve query performance.

The F1 sort operator uses external merge sort with tree-of-losers priority queues and offset-value coding for both run generation and merging. Each tree entry encodes offset and value bits in an unsigned 64-bit integer offset-value code. Invalid key values, e.g., after a merge input is exhausted, are also folded into this integer. In a sense, each comparison of offset-value codes is folded into a test whether or not a run is exhausted; thus, the comparison of offset-value codes is practically free.

Offset-value codes for rows in sorted runs are a byproduct of run generation. These offset-value codes later improve the efficiency of merging. For example, the merge logic inspects the offset-value code of the next row from the winner's input. If this next key value has an offset equal to the number of key columns, the row goes directly to the output buffer, bypassing the merge logic entirely.

Since many sort-based operators (e.g., in-stream aggregation, merge join, analytic functions, and order-preserving exchange) can leverage offset-value codes in their inputs, F1 Query supports carrying offset-value codes between operators. An artificial column for offset-value codes is introduced during query planning for order-producing physical operators. Order-producing relational operators use the logic described in this paper to produce offset-value codes in its output. The sorting techniques described in this paper are also employed in Napa [1].

### 6 PERFORMANCE EVALUATION

This section tests two claims or hypotheses about offset-value coding, sort-based query processing, and interesting orderings:

- (1) Offset-value coding can speed up external merge sort and also its consumers, e.g., in-stream aggregation, merge join, segmentation, and order-preserving (merging) exchange.
- (2) Offset-value coding together with interesting orderings can be more efficient than hash-based query plans.

Some readers may think these claims are obviously true, but we nonetheless wish to dispel any conceivable remaining doubt. These hypotheses do not claim universally superior performance, merely new performance advantages for sort-based query processing due to offset-value coding in contexts beyond merge sort.

In order to focus on order-based operators and offset-value coding, all experiments use a single execution thread. The hardware is an engineering workstation; more detailed specifications are omitted on purpose. Each experiment starts with a warm cache, i.e., input data pre-fetched into memory. The measurements below come from Google's public benchmark library [12]. Test data are synthetic yet similar to the actual data in our daily production web analysis with many rows and many key columns. Each key column is an 8-byte integer with only a few distinct values.

Figure 3, in a test of hypothesis 1, shows the effort of offset-value coding in in-stream aggregation. For fast detection of group boundaries, the operation exploits the offset-value codes from a preceding sort operation. This is compared to using full comparisons of multiple key columns. The count of input rows is 1,000,000. The ratio of input rows and output rows varies. For example, a ratio of 1 indicates all input rows are distinct (i.e., all groups have size 1) and a ratio of 100 indicates that on average 100 input rows contribute to each output row. All queries of the type "select ..., count (distinct ...) group by..." require a two-step process. Figure 3 shows that, within the sorted output, testing the offset against the count of grouping columns is much faster






Figure 3: Group boundaries from offset-value codes.




Figure 5: Performance of “intersect distinct” query plans.




Figure 4: Query plans for an “intersect distinct” query.

than full comparisons of multiple key columns. Thus, offset-value codes benefit not only sorting but also other query execution algorithms.

Figure 4 shows two query evaluation plans, one hash-based and the other one sort-based, for a simple SQL query computing the intersection of two tables, e.g., “select B from T1 intersect select B from T2”. If column B is not a primary key in tables T1 and T2, correct execution requires duplicate removal plus a join algorithm. The query plans for “intersect all” would be practically the same, counting duplicates rather than removing them; and similarly for set differences using “except” and “except all” syntax. Index intersection for “and” predicates uses the same kinds of query plans, as does index join [18] when emulating columnar storage in data warehouses with traditional indexes (in the ideal case, compressed).

In the hash-based plan, there are three blocking operators: two hash aggregation operators for duplicate removal and a hash join for set intersection. In contrast, the sort-based plan has only two blocking operators: both are in-sort aggregation operators for duplicate removal. The merge join computing the intersection exploits not only interesting orderings but also offset-value codes in the output of in-sort aggregation. Conceivably, advanced techniques could reduce the counts of blocking operators in both sort- and hash-based query plans, e.g., integrated hashing [14, 18], groupjoin [24], or co-group-by instead of join [9].

Figure 5 shows the performance of the plans in Figure 4. Each input table has 100,000,000 rows; each blocking operator’s memory holds 10,000,000 rows. In a hash-based plan, duplicate removal and join spill to temporary storage such that many rows are spilled twice. In contrast, the sort-based plan spills each input row only once. Thus, interesting orderings cut the spilling effort in half. Moreover, offset-value codes from the in-sort aggregation operators speed up row comparisons in the merge join.

In summary, while the queries in these experiments are simple and common, the measurements support our two claims about the benefits of offset-value coding in database query processing.

## 7 SUMMARY AND CONCLUSIONS

Computing offset-value codes in order-preserving query execution algorithms has not received any attention to-date for two reasons. First, new research has only recently widened the scope of offset-value coding from merge sort to all sort-based query execution algorithms and query plans, including complex join and grouping queries. Second, the only method known to-date – comparing an operator’s output row-by-row, column-by-column – is so expensive that it thwarts all benefit of offset-value codes in query execution.

A new theorem enables an alternative method for computing or propagating offset-value codes. Beyond its use here, this new theorem already enables efficient maintenance of offset-value codes in near-optimal sorting with tree-of-losers priority queues [17] and in b-trees with prefix truncation (during key deletion) [19].

With this solid foundation, calculating offset-value codes for operator output in database query execution is surprisingly simple and very efficient. Importantly, there is no need for any column value comparisons beyond those required to compute output rows. Offset-value coding throughout sort-based query execution takes “interesting orderings” to their full potential, not only in query planning but also in plan execution. In binary operations, e.g., merge join and intersection, offset-value codes decide many or most row comparisons, whereas in unary operations, e.g., duplicate removal and grouping, offset-value codes decide all row comparisons.

Offset-value coding already saves thousands of CPUs in Google’s Napa and F1 Query systems, because ingestion (run generation), compaction (merging), and query processing in log-structured merge-forests rely heavily on sorting and merging. We anticipate ever more savings in the future as sorting and merging are essential in many data processing pipelines and in many storage formats.

In conclusion, if storage structures are kept sorted on column values, if query optimization considers interesting orderings, and if query execution fully employs offset-value coding plus some techniques from earlier work [8, 16], then sort-based query processing can consistently be at least as efficient as hash-based query processing. Hash-partitioning remains recommended for parallel and distributed query execution until future innovation ensures efficient, well-balanced, and failure-tolerant range-partitioning.

## ACKNOWLEDGMENTS

We thank Herald Kllapi and the reviewers for their helpful suggestions.

## REFERENCES

- [1] Ankur Agiwal, Kevin Lai, Gokul Nath Babu Manoharan, Indrajit Roy, Jagan Sankaranarayanan, Hao Zhang, Tao Zou, Min Chen, Zongchang (Jim) Chen, Ming Dai, Thanh Do, Haoyu Gao, Haoyan Geng, Raman Grover, Bo Huang, Yanlai Huang, Zhi (Adam) Li, Jianyi Liang, Tao Lin, Li Liu, Yao Liu, Xi Mao, Yalan (Maya) Meng, Prashant Mishra, Jay Patel, Rajesh S. R., Vijayshankar Raman, Sourashis Roy, Mayank Singh Shishodia, Tianhang Sun, Ye (Justin) Tang, Junichi Tatemura, Sagar Trehan, Ramkumar Vadali, Prasanna Venkatasubramanian, Gensheng Zhang, Kefei Zhang, Yupu Zhang, Zeleng Zhuang, Goetz Graefe, Divyakant Agrawal, Jeff Naughton, Sujata Kosalge, and Hakan Hacigümüş. 2021. Napa: Powering Scalable Data Warehousing with Robust Query Performance at Google. *Proc. VLDB Endow.* 14, 12 (jul 2021), 2986–2997. <https://doi.org/10.14778/3476311.3476377>
- [2] Rudolf Bayer and Karl Unterauer. 1977. Prefix B-Trees. *ACM Trans. Database Syst.* 2, 1 (1977), 11–26. <https://doi.org/10.1145/320521.320530>
- [3] Mike W. Blasgen and Kapali P. Eswaran. 1977. Storage and Access in Relational Data Bases. *IBM Syst. J.* 16, 4 (1977), 362–377. <https://doi.org/10.1147/sj.164.0363>
- [4] Michael J. Carey, David J. DeWitt, Michael J. Franklin, Nancy E. Hall, Mark L. McAuliffe, Jeffrey F. Naughton, Daniel T. Schuh, Marvin H. Solomon, C. K. Tan, Odysseas G. Tsatalos, Seth J. White, and Michael J. Zwilling. 1994. Shoring Up Persistent Applications. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data, Minneapolis, Minnesota, USA, May 24-27, 1994*, Richard T. Snodgrass and Marianne Winslett (Eds.). ACM Press, 383–394. <https://doi.org/10.1145/191839.191915>
- [5] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. 2008. Bigtable: A Distributed Storage System for Structured Data. *ACM Trans. Comput. Syst.* 26, 2, Article 4 (jun 2008), 26 pages. <https://doi.org/10.1145/1365815.1365816>
- [6] W. M. Conner. 1977. Offset-value coding. In *IBM Technical Disclosure Bulletin*.
- [7] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, J. J. Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaure, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, and Dale Woodford. 2013. Spanner: Google’s Globally Distributed Database. *ACM Trans. Comput. Syst.* 31, 3, Article 8 (aug 2013), 22 pages. <https://doi.org/10.1145/2491245>
- [8] Thanh Do and Goetz Graefe. 2022. Robust and Efficient Sorting with Offset-Value Coding. *ACM Trans. Database Syst.* (2022). <https://doi.org/10.1145/3570956>
- [9] Thanh Do, Goetz Graefe, and Jeffrey Naughton. 2022. Efficient sorting, duplicate removal, grouping, and aggregation. *ACM Trans. Database Syst.* 47, 4 (2022). <https://doi.org/10.1145/3568027>
- [10] Robert Epstein. 1979. Techniques for processing of aggregates in relational database systems. In *Univ. of California at Berkeley, UCB/ERL Memorandum M79/8*.
- [11] Martin A. Goetz. 1963. Internal and Tape Sorting Using the Replacement-Selection Technique. *Commun. ACM* 6, 5 (May 1963), 201–206. <https://doi.org/10.1145/366552.366556>
- [12] Google. [n.d.]. A microbenchmark support library. <https://github.com/google/benchmark>
- [13] Goetz Graefe. 1993. Query Evaluation Techniques for Large Databases. *ACM Comput. Surv.* 25, 2 (1993), 73–170. <https://doi.org/10.1145/152610.152611>
- [14] Goetz Graefe. 1994. Volcano— An Extensible and Parallel Query Evaluation System. *IEEE Trans. on Knowl. and Data Eng.* 6, 1 (1994), 120–135. <https://doi.org/10.1109/69.273032>
- [15] Goetz Graefe. 2006. Implementing sorting in database systems. *ACM Comput. Surv.* 38, 3 (2006), 10. <https://doi.org/10.1145/1132960.1132964>
- [16] Goetz Graefe. 2011. A Generalized Join Algorithm. In *Datenbanksysteme für Business, Technologie und Web (BTW), 14. Fachtagung des GI-Fachbereichs "Datenbanken und Informationssysteme" (DBIS), 2.-4.3.2011 in Kaiserslautern, Germany (LNI)*, Theo Härder, Wolfgang Lehner, Bernhard Mitschang, Harald Schöning, and Holger Schwarz (Eds.), Vol. P-180. GI, 267–286. <https://dl.gi.de/20.500.12116/19583>
- [17] Goetz Graefe. 2023. Priority queues for database query processing. In *Datenbanksysteme für Business, Technologie und Web (GI-Fachtagung BTW 2023)*.
- [18] Goetz Graefe, Ross Bunker, and Shaun Cooper. 1998. Hash Joins and Hash Teams in Microsoft SQL Server. In *VLDB’98, Proceedings of 24rd International Conference on Very Large Data Bases, August 24-27, 1998, New York City, New York, USA*, Ashish Gupta, Oded Shmueli, and Jennifer Widom (Eds.). Morgan Kaufmann, 86–97. <http://www.vldb.org/conf/1998/p086.pdf>
- [19] Goetz Graefe and Thanh Do. 2023. Storage and access with offset-value coding. In preparation.
- [20] Bala R. Iyer. 2005. Hardware Assisted Sorting in IBM’s DB2 DBMS. In *International Conference on Management of Data (COMAD)*.
- [21] H. V. Jagadish, P. P. S. Narayan, S. Seshadri, S. Sudarshan, and Rama Kanneganti. 1997. Incremental Organization for Data Recording and Warehousing. In *VLDB’97, Proceedings of 23rd International Conference on Very Large Data Bases, August 25-29, 1997, Athens, Greece*, Matthias Jarke, Michael J. Carey, Klaus R. Dittrich, Frederick H. Lochovsky, Pericles Loucopoulos, and Manfred A. Jeusfeld (Eds.). Morgan Kaufmann, 16–25. <http://www.vldb.org/conf/1997/P016.PDF>
- [22] Donald E. Knuth. 1998. *The Art of Computer Programming, Volume 3: Sorting and Searching (2nd. ed.)*. Addison Wesley Longman Publishing Co., Inc., USA.
- [23] Harry Leslie, Rohit Jain, Dave Birdsall, and Hediéh Yaghmai. 1995. Efficient Search of Multi-Dimensional B-Trees. In *VLDB’95, Proceedings of 21th International Conference on Very Large Data Bases, September 11-15, 1995, Zurich, Switzerland*, Umeshwar Dayal, Peter M. D. Gray, and Shojiro Nishio (Eds.). Morgan Kaufmann, 710–719. <http://www.vldb.org/conf/1995/P710.PDF>
- [24] Guido Moerkotte and Thomas Neumann. 2011. Accelerating Queries with Group-By and Join by Groupjoin. *Proc. VLDB Endow.* 4, 11 (2011), 843–851. <http://www.vldb.org/pvldb/vol4/p843-moerkotte.pdf>
- [25] Patrick E. O’Neil, Edward Cheng, Dieter Gawlick, and Elizabeth J. O’Neil. 1996. The Log-Structured Merge-Tree (LSM-Tree). *Acta Informatica* 33, 4 (1996), 351–385. <https://doi.org/10.1007/s002360050048>
- [26] Bart Samwel, John Cieslewicz, Ben Handy, Jason Govig, Petros Venetis, Chanjun Yang, Keith Peters, Jeff Shute, Daniel Tenedorio, Himani Apte, Felix Weigel, David Wilhite, Jiacheng Yang, Jun Xu, Jiexing Li, Zhan Yuan, Craig Chasseur, Qiang Zeng, Ian Rae, Anurag Biyani, Andrew Harn, Yang Xia, Andrey Gubichev, Amr El-Helw, Orri Erling, Zhepeng Yan, Mohan Yang, Yiqun Wei, Thanh Do, Colin Zheng, Goetz Graefe, Somayeh Sardashti, Ahmed M. Aly, Divy Agrawal, Ashish Gupta, and Shiv Venkataraman. 2018. F1 Query: Declarative Querying at Scale. *Proc. VLDB Endow.* 11, 12 (Aug. 2018), 1835–1848. <https://doi.org/10.14778/3229863.3229871>
- [27] P. Griffiths Selinger, M. M. Astrahan, D. D. Chamberlin, R. A. Lorie, and T. G. Price. 1979. Access Path Selection in a Relational Database Management System. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data (SIGMOD ’79)*. ACM, New York, NY, USA, 23–34. <https://doi.org/10.1145/582095.582099>
- [28] Jeff Shute, Radek Vingralek, Bart Samwel, Ben Handy, Chad Whipkey, Eric Rollins, Mircea Oancea, Kyle Littlefield, David Menestrina, Stephan Ellner, John Cieslewicz, Ian Rae, Traian Stancescu, and Himani Apte. 2013. F1: A Distributed SQL Database That Scales. *Proc. VLDB Endow.* 6, 11 (Aug. 2013), 1068–1079. <https://doi.org/10.14778/2536222.2536232>
- [29] David E. Simmen, Eugene J. Shekita, and Timothy Malkemus. 1996. Fundamental Techniques for Order Optimization. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, Montreal, Quebec, Canada, June 4-6, 1996*, H. V. Jagadish and Inderpal Singh Mumick (Eds.). ACM Press, 57–67. <https://doi.org/10.1145/233269.233320>