



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Incremental View Maintenance with Triple Lock Factorization Benefits

Citation for published version:

Nikolic, M & Olteanu, D 2018, Incremental View Maintenance with Triple Lock Factorization Benefits. in *Proceedings of the 2018 International Conference on Management of Data*. SIGMOD '18, ACM, New York, NY, USA, pp. 365-380, 2018 ACM SIGMOD/PODS International Conference on Management of Data, Houston, Texas, United States, 10/06/18. <https://doi.org/10.1145/3183713.3183758>

Digital Object Identifier (DOI):

[10.1145/3183713.3183758](https://doi.org/10.1145/3183713.3183758)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 2018 International Conference on Management of Data

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Incremental View Maintenance with Triple Lock Factorization Benefits

Milos Nikolic and Dan Olteanu

Department of Computer Science, University of Oxford, UK

ABSTRACT

We introduce F-IVM, a unified incremental view maintenance (IVM) approach for a variety of tasks, including gradient computation for learning linear regression models over joins, matrix chain multiplication, and factorized evaluation of conjunctive queries.

F-IVM is a higher-order IVM algorithm that reduces the maintenance of the given task to the maintenance of a hierarchy of increasingly simpler views. The views are functions mapping keys, which are tuples of input data values, to payloads, which are elements from a task-specific ring. Whereas the computation over the keys is the same for all tasks, the computation over the payloads depends on the task. F-IVM achieves efficiency by factorizing the computation of the keys, payloads, and updates.

We implemented F-IVM as an extension of DBToaster. We show in a range of scenarios that it can outperform classical first-order IVM, DBToaster’s fully recursive higher-order IVM, and plain re-computation by orders of magnitude while using less memory.

1 INTRODUCTION

Supporting modern applications that rely on the availability of accurate and real-time analytics computed over large and continuously evolving databases is a challenging data management problem [6]. Special cases are the classical problems of incremental view maintenance (IVM) [12, 25] and stream query processing [4, 29].

Recent efforts studied the problem of computing machine learning (ML) models over static databases, with the predominant approach loosely integrating the database systems with the statistical packages [20, 27, 30, 39, 40]: First, the database system computes the input to the statistical package by joining the database relations. It then exports the join result, which is ingested by the statistical package and used for training ML models. This approach precludes real-time analytics due to the expensive export/import steps.

Real-time analytics would benefit from a much tighter integration of the two systems. This is provided by Morpheus and F that learn ML models over normalized data. Morpheus decomposes the task of learning generalized linear models into subtasks that are pushed down past the joins to the input database tuples [28]. F goes further and factorizes the join computation [23, 42]. This factorization may significantly lower the complexity of joins by avoiding the computation of their constituent Cartesian products [7, 38]. This performance improvement carries over to learning classification and regression models over factorized joins.

In this paper, we introduce F-IVM, a unified IVM approach for analytics over normalized data. Analytical tasks are expressed as views on joins with group-by aggregates over relations that map keys to payloads. We exemplify the power of F-IVM for matrix chain multiplication, factorized evaluation of conjunctive queries,

and gradient computation used for learning linear regression models. Although these applications achieve different outcomes, they only differ in the specification of the sum and product operations over payloads. These payload operations can be: arithmetic addition and multiplication for queries with joins and group-by aggregates; relational union and join for listing and factorized representation of conjunctive query results; matrix addition and multiplication for gradient computation; in general, sum and product in an appropriate ring. The mechanisms for maintenance and computation over keys stay the same. F-IVM is thus highly extensible: efficient maintenance for new analytics over normalized data is readily available as long as they come with appropriate sum and product operations.

F-IVM has two ingredients. First, it leverages higher-order IVM to reduce the maintenance of the input view to the maintenance of a tree of simpler views. In contrast to classical (first-order) IVM, which does not use extra views and computes changes in the query result on the fly, F-IVM can tremendously speed up the maintenance task and lower its complexity by using carefully chosen extra views. Nevertheless, F-IVM may use substantially fewer and cheaper views than fully-recursive IVM, which is the approach taken by the state-of-the-art IVM system DBToaster [25].

Second, it is the first approach to employ factorized computation for three aspects of incremental maintenance for queries with aggregates and joins: (1) it exploits insights from query evaluation algorithms with best known complexity and optimizations that push aggregates past joins [5, 7, 38]; (2) it can process bulk updates expressed as low-rank decompositions [26, 44]; and (3) it can maintain compressed representation of query results.

F-IVM is implemented on top of DBToaster’s open-source backend, which can generate optimized C++ code from high-level update trigger specifications. In a range of applications, F-IVM outperforms DBToaster and classical first-order IVM by up to two orders of magnitude in both time and space.

Example 1.1. Consider the following SQL query over a database \mathcal{D} with relations $R(A, B)$, $S(A, C, E)$, and $T(C, D)$:

```
Q := SELECT S.A, S.C, SUM(R.B * T.D * S.E)
      FROM R NATURAL JOIN S NATURAL JOIN T
      GROUP BY S.A, S.C;
```

Let us consider two different evaluation strategies for this query. A naïve approach first computes the join result and then the aggregate. This can take time cubic in the size of \mathcal{D} . An alternative strategy exploits the distributivity of the SUM operator over multiplication to partially push the aggregate past joins and later combine these partial aggregates to produce the query result. For instance, one such partial sum over S can be expressed as the view V_S :

```
V_S := SELECT A, C, SUM(E) AS S_E FROM S GROUP BY A, C;
```

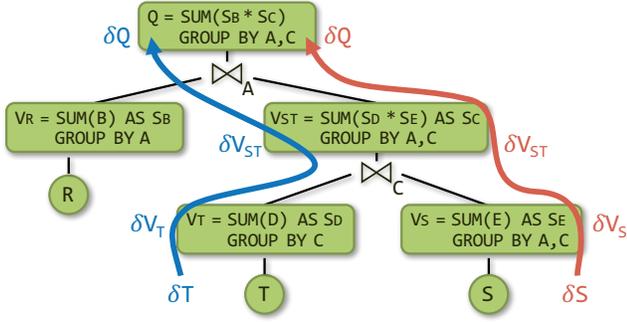


Figure 1: View tree for the query in Example 1.1. Delta propagations for updates to S (right red) and to T (left blue).

In the view V_S , we distinguish keys, which are tuples over (A, C) , and payloads, which are aggregate values S_E . Similarly, we can compute partial sums over R and T as views V_R and V_T . These views are joined as depicted by the *view tree* in Figure 1, which is akin to a query plan with aggregates pushed past joins. This view tree computes the query result is time linear in the size of D .

Consider the problem of learning a linear function f with parameters θ_0 , θ_D and θ_E that predicts the label B given the features D and E , where the training dataset is the natural join of our relations:

$$f(D, E) = \theta_0 + \theta_D \cdot D + \theta_E \cdot E$$

Our insight is that the same above view tree can also compute the gradient vector used for learning f ! The only needed adjustment is the replacement of the SQL SUM and $*$ operators with appropriate new sum and product operators. We next explain the case of one model f for each pair of values (A, C) ; the case of one model for the entire dataset is similar.

As detailed in Section 6.2, the gradient of the square loss objective function requires the computation of three types of aggregates: the scalar c that is the count aggregate SUM(1); the vector s of linear aggregates SUM(B), SUM(D), and SUM(E); and the matrix Q of quadratic aggregates SUM($B \cdot D$), SUM($B \cdot E$), SUM($D \cdot D$), SUM($D \cdot E$), and SUM($E \cdot E$). These aggregates represent sufficient statistics to capture the correlation between the features and the label. If we compute them over each (A, C) group, then we learn one model for each such group; if we compute them over the entire dataset, we then learn one model only.

We treat these aggregates as one compound aggregate (c, s, Q) so we can share computation across them. This compound aggregate can be partially pushed past joins similarly to the SUM aggregate discussed before. Its values are carried in the payloads of keys of the views in the view tree from Figure 1. For instance, the partial compound aggregate (c_T, s_T, Q_T) at the view V_T computes for each C -value the count, sum, and sum of squares of the D -values in T . Similarly, the partial aggregate (c_S, s_S, Q_S) at the view V_S computes for each pair (A, C) the count, sum, and sum of squares of E -values in S . In the view V_{ST} that is the join of V_T and V_S , each key (a, c) is associated with the multiplication of the payloads for the keys c in V_T and (a, c) in V_S . This multiplication is however different from SQL's $*$ as it works on compound aggregates: The scalar c_{ST} is the arithmetic multiplication of c_T and c_S ; the vector of linear aggregates s_{ST} is the sum of the scalar-vector products $c_T s_S$ and

$s_T c_S$; finally, the matrix Q_{ST} of quadratic aggregates is the sum of the scalar-matrix products $c_T Q_S$ and $Q_T c_S$, and of the outer products of the vectors s_T and the transpose of s_S and also of s_S and the transpose of s_T . Our approach significantly shares the computation across the aggregates: The scalar aggregates are used to scale up the linear and quadratic ones, while the linear aggregates are used to compute the quadratic ones.

We now turn to incremental view maintenance. F-IVM operates over view trees. Whereas for non-incremental computation we only materialize the top view in the tree and the input relations, for incremental computation we may materialize additional views to speed up the maintenance task. Our approach is an instance of higher-order delta-based IVM, since an update to one relation may trigger the maintenance for several views.

Figure 1 shows the leaf-to-root paths taken to maintain the view result under changes to S and to T . For updates δS to S , each delta view δV_S , δV_{ST} , and δQ , is computed using delta rules:

$$\begin{aligned} \delta V_S &:= \text{SELECT } A, C, \text{SUM}(E) \text{ AS } S_E \text{ FROM } \delta S \text{ GROUP BY } A, C; \\ \delta V_{ST} &:= \text{SELECT } A, C, \text{SUM}(S_D * S_E) \text{ AS } S_C \\ &\quad \text{FROM } V_T \text{ NATURAL JOIN } \delta V_S \text{ GROUP BY } A, C; \\ \delta Q &:= \text{SELECT } A, C, \text{SUM}(S_B * S_C) \\ &\quad \text{FROM } V_R \text{ NATURAL JOIN } \delta V_{ST} \text{ GROUP BY } A, C; \end{aligned}$$

The update δS may consist of both inserts and deletes, which are encoded as keys with positive and respectively negative payloads. In our example, a negative payload is -1 for the SUM aggregate and $(-1, \mathbf{0}_{5 \times 1}, \mathbf{0}_{5 \times 5})$ for the compound aggregate, where $\mathbf{0}_{n \times m}$ is the n -by- m matrix with all elements 0.

F-IVM materializes and maintains views depending on the update workload. For updates to all input relations, it materializes the view at each node in the view tree. For updates to R only, it only materializes V_{ST} ; for updates to S only, it materializes V_R and V_T ; for updates to T only, it materializes V_R and V_S . F-IVM takes constant time for updates to S and linear time for updates to R and T ; these complexities are in the number of distinct keys in the views. In contrast, the first-order version of F-IVM does not create extra views and takes linear time for updates to any of the three relations. The fully-recursive version of F-IVM would further materialize the join of V_R and V_S as it may consider more than one view tree.

The above analysis holds for our query with one SUM aggregate. For the learning example with the nine SUM aggregates, F-IVM still needs the same views. Classical first-order IVM algorithms would however need to compute a distinct delta query for each of the nine aggregates – nine in total for updates to any of the three relations. DBToaster, which is the state-of-the-art fully recursive IVM, would need to compute 28 views, nine top views and 19 auxiliary ones. Whereas F-IVM shares the computation across the nine aggregates, the classical IVM and DBToaster do not. This significantly widens the performance gap between F-IVM and its competitors.

2 DATA MODEL AND QUERY LANGUAGE

Data Model. A schema \mathcal{S} is a set of variables (or attributes). For a variable $X \in \mathcal{S}$, let $\text{Dom}(X)$ denote its domain. A tuple \mathbf{t} over schema \mathcal{S} has the domain $\text{Dom}(\mathcal{S}) = \prod_{X \in \mathcal{S}} \text{Dom}(X)$. The empty tuple $()$ is the tuple over the empty schema.

A ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$ is a set with two binary operations, $+$ and $*$, which generalize the arithmetic operations of addition and multiplication, the additive identity $\mathbf{0}$, the multiplicative identity $\mathbf{1}$, and an additive inverse for each element of \mathbf{D} (cf. Appendix A for definition). Examples of rings are $\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{R}^2, \mathbb{R}^3$, and matrix ring.

Let $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$ be a ring. A *relation* R over schema \mathcal{S} and the ring \mathbf{D} is a function $R : \text{Dom}(\mathcal{S}) \rightarrow \mathbf{D}$ mapping tuples over schema \mathcal{S} to values in \mathbf{D} such that $R[\mathbf{t}] \neq \mathbf{0}$ for finitely many tuples \mathbf{t} . The tuple \mathbf{t} is called a *key*, while its mapping $R[\mathbf{t}]$ is the *payload* of \mathbf{t} in R . We use $\text{sch}(R)$ to denote the schema of R . The statement $\mathbf{t} \in R$ tests if $R[\mathbf{t}] \neq \mathbf{0}$. The size $|R|$ of R is the size of the set $\{\mathbf{t} \mid \mathbf{t} \in R\}$, which consists of all keys with non- $\mathbf{0}$ payloads. We materialize R as a hash map or a multi-dimensional array. A database \mathcal{D} is a collection of relations over the same ring. Its size $|\mathcal{D}|$ is the sum of the sizes of its relations. This data model is in line with prior work on K -relations over provenance semirings [17], generalized multiset relations [24], and factors over semirings [5].

Query Language. In this paper we consider incremental maintenance for queries with joins and group-by aggregates:

```
SELECT X1, ..., Xf, SUM(gf+1(Xf+1) * ... * gm(Xm))
FROM R1 NATURAL JOIN ... NATURAL JOIN Rn
GROUP BY X1, ..., Xf
```

The group-by variables X_1, \dots, X_f are also called *free*, while the other variables X_{f+1}, \dots, X_m are *bound*. The values of the SUM aggregate are from a ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$. The lifting functions $g_k : \text{Dom}(X_k) \rightarrow \mathbf{D}$, for $f < k \leq m$, map (lift) variable values to elements in \mathbf{D} . The SUM operator uses the addition $+$ from \mathbf{D} . More complex aggregates can be expressed using the sum and product operations from the ring.

Instead of the SQL notation, we use the following encoding:

$$Q[X_1, \dots, X_f] = \bigoplus_{X_{f+1}} \dots \bigoplus_{X_m} \bigotimes_{i \in [n]} R_i[S_i],$$

where \otimes is the join operator, $\bigoplus_{X_{f+1}}$ is the aggregation operator that marginalizes over the variable X_{f+1} , and each relation R_i maps keys over schema S_i to payloads in \mathbf{D} . We also need a union operator \uplus to express updates (insert/delete) to relations.

Given a ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$, relations R and S over schema S_1 and relation T over schema S_2 , a variable $X \in S_1$, and a lifting function $g_X : \text{Dom}(X) \rightarrow \mathbf{D}$, we define the operators as follows:

union:

$$\forall \mathbf{t} \in D_1: (R \uplus S)[\mathbf{t}] = R[\mathbf{t}] + S[\mathbf{t}]$$

join:

$$\forall \mathbf{t} \in D_2: (S \otimes T)[\mathbf{t}] = S[\pi_{S_1}(\mathbf{t})] * T[\pi_{S_2}(\mathbf{t})]$$

aggregation by marginalization:

$$\forall \mathbf{t} \in D_3: (\bigoplus_X R)[\mathbf{t}] = \sum \{ R[\mathbf{t}_1] * g_X(\pi_X(\mathbf{t}_1)) \mid \mathbf{t}_1 \in D_1, \mathbf{t} = \pi_{S_1 \setminus \{X\}}(\mathbf{t}_1) \}$$

where $D_1 = \text{Dom}(S_1)$; $D_2 = \text{Dom}(S_1 \cup S_2)$; $D_3 = \text{Dom}(S_1 \setminus \{X\})$; $\pi_X(\mathbf{t}_1)$ and $\pi_{S_1 \setminus \{X\}}(\mathbf{t}_1)$ are tuples representing the projection of tuple \mathbf{t}_1 on X and respectively on $S_1 \setminus \{X\}$.

Example 2.1. Consider the relations over a ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$:

| $A \ B \rightarrow R[A, B]$ | $A \ B \rightarrow S[A, B]$ | $B \ C \rightarrow T[B, C]$ |
|-----------------------------|-----------------------------|-----------------------------|
| $a_1 \ b_1 \rightarrow r_1$ | $a_2 \ b_1 \rightarrow s_1$ | $b_1 \ c_1 \rightarrow t_1$ |
| $a_2 \ b_1 \rightarrow r_2$ | $a_3 \ b_2 \rightarrow s_2$ | $b_2 \ c_2 \rightarrow t_2$ |

The values $r_1, r_2, s_1, s_2, t_1, t_2$ are non- $\mathbf{0}$ values from \mathbf{D} . The operators \uplus, \otimes, \oplus are akin to classical union, join, and aggregation:

| $A \ B \rightarrow (R \uplus S)[A, B]$ | $A \ B \ C \rightarrow ((R \uplus S) \otimes T)[A, B, C]$ |
|--|---|
| $a_1 \ b_1 \rightarrow r_1$ | $a_1 \ b_1 \ c_1 \rightarrow r_1 * t_1$ |
| $a_2 \ b_1 \rightarrow r_2 + s_1$ | $a_2 \ b_1 \ c_1 \rightarrow (r_2 + s_1) * t_1$ |
| $a_3 \ b_2 \rightarrow s_2$ | $a_3 \ b_2 \ c_2 \rightarrow s_2 * t_2$ |

$$\frac{B \ C \rightarrow (\bigoplus_A (R \uplus S) \otimes T)[B, C]}{b_1 \ c_1 \rightarrow r_1 * t_1 * g_A(a_1) + (r_2 + s_1) * t_1 * g_A(a_2)}$$

$$b_2 \ c_2 \rightarrow s_2 * t_2 * g_A(a_3)$$

where $g_A : \text{Dom}(A) \rightarrow \mathbf{D}$ is the given lifting function for A . \square

Example 2.2. We show how to encode the following SQL query over tables $R(A, B)$, $S(A, C, E)$, and $T(C, D)$ into our formalism:

```
Q = SELECT SUM(1) FROM R NATURAL JOIN S NATURAL JOIN T;
```

Assuming the ring \mathbb{Z} , we encode the table R as a relation $R : \text{Dom}(A) \times \text{Dom}(B) \rightarrow \mathbb{Z}$ that maps tuples (a, b) to their multiplicity in R ; similarly, we encode the tables S and T as relations S and T . Our SQL query is then translated into:

$$Q[] = \bigoplus_A \bigoplus_B \bigoplus_C \bigoplus_D \bigoplus_E R[A, B] \otimes S[A, C, E] \otimes T[C, D]$$

where the lifting functions used in marginalization map all values to 1. Remember that by definition R, S , and T are finite relations. The relation Q maps the empty tuple $()$ to the count. \square

Example 2.3. Let us consider the SQL query from Example 1.1, which computes $\text{SUM}(R.B * T.D * S.E)$ and assume that B, C , and D take values from \mathbb{Z} . We model the tables R, S , and T as relations mapping tuples to their multiplicity, as in Example 2.2. The variables A and C are free while B, D , and E are bound. When marginalizing over the bound variables, we apply the same lifting function to all variables: $\forall x \in \mathbb{Z} : g_B(x) = g_D(x) = g_E(x) = x$. Now, the SQL query can be expressed in our formalism as follows:

$$Q[A, C] = \bigoplus_B \bigoplus_D \bigoplus_E R[A, B] \otimes S[A, C, E] \otimes T[C, D]$$

The computation of the aggregate $\text{SUM}(R.B * T.D * S.E)$ now happens over payloads. \square

One key benefit of using relations over rings is avoiding the intricacies of incremental computation under classical multiset semantics caused by non-commutativity of inserts and deletes. Our data model simplifies delta processing by representing both inserts and deletes as tuples, with the distinction that they map to positive and respectively negative ring values. This uniform treatment allows for simple delta rules for our three operators.

3 FACTORIZED RING COMPUTATION

In this section, we introduce a static query evaluation framework based on factorized computation and data rings. In the next section, we adapt it to incremental maintenance.

Variable Orders. Classical query evaluation uses query plans that dictate the order in which the relations are joined. We use slightly different plans, which we call variable orders, that dictate the order in which we solve each join variable. They may require to join several relations at the same time if these relations have the same variable. Our choice is motivated by the complexity of join evaluation: standard (relation-at-a-time) query plans are provably suboptimal, whereas variable orders can be optimal [31].

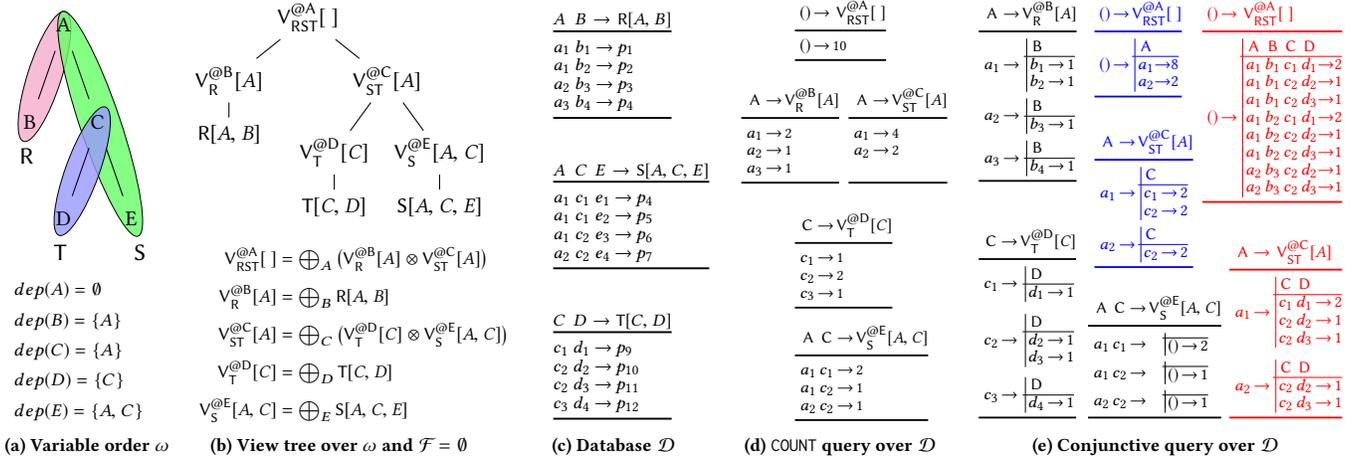


Figure 2: (a) Variable order ω of the natural join of R, S, T . (b) View tree τ over ω and without free variables. (c) Database \mathcal{D} over a ring \mathcal{D} , where $\{p_i\}_{i \in [12]} \subseteq \mathcal{D}$. (d) Computing COUNT using τ and the \mathbb{Z} ring, where $\forall i \in [12] : p_i = 1$. (e) Computing the query from Example 6.5 using τ and the relational ring, where $\forall i \in [12] : p_i = \{() \rightarrow 1\}$. The red views (rightmost column) have payloads storing the listing representation of the intermediate and final query results. The blue views (middle) encode a factorized representation of these results distributed over their payloads. The black views remain the same for both representations.

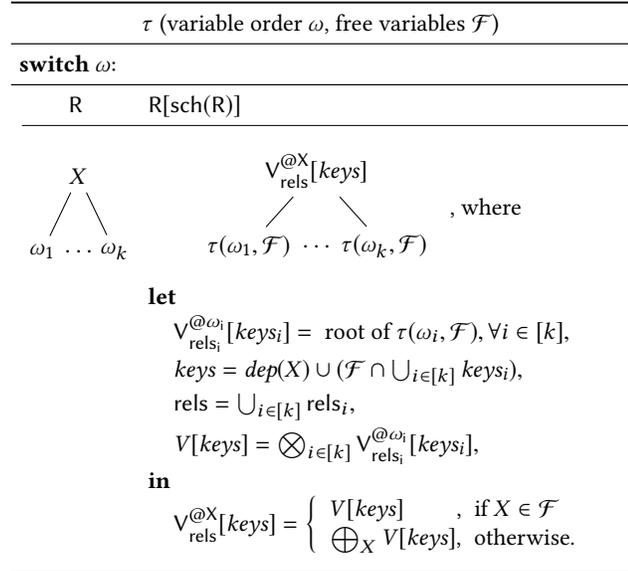


Figure 3: Creating a view tree $\tau(\omega, \mathcal{F})$ for a variable order ω and a set of free variables \mathcal{F} .

Given a join query Q , a variable X depends on a variable Y if both are in the schema of a relation in Q .

Definition 3.1 (adapted from [38]). A variable order ω for a join query Q is a pair (F, dep) , where F is a rooted forest with one node per variable in Q , and dep is a function mapping each variable X to a set of variables in F . It satisfies the following constraints:

- For each relation in Q , its variables lie along the same root-to-leaf path in F .
- For each variable X , $\text{dep}(X)$ is the subset of its ancestors in F on which the variables in the subtree rooted at X depend.

Figure 2a gives a variable order for the natural join of three relations. Variable D has ancestors A and C , yet it only depends on C since C and D appear in the same relation T and D does not occur in any relation together with A . Thus, $\text{dep}(D) = \{C\}$. Given C , the variables D and E are independent of each other. In case Q has free variables, then we prefer variable orders for Q that have all free variables on top of the bound variables.

View Trees. Our framework relies on a variable order ω for the input query Q to describe the structure of computation and indicate which variable marginalizations are pushed past joins. At each variable in ω we define a view that is a query over its children. The view at the root variable corresponds to the entire query Q . The tree of these views is called a *view tree*, which plays the role of a query plan in our framework.

Figure 3 gives an algorithm that constructs a view tree τ for a variable order ω of the input query Q and the set of free variables \mathcal{F} of Q . The input variable order is extended with relations at leaves under their lowest variable (as children or further below). The views in the view tree are defined over the input relations or the views at children. We use the notation $V_{\text{rels}}^{\otimes X}[\text{keys}]$ to state that the view V is recursively defined over the input relations rels , has free variables keys , and is at the variable X in ω ; in case of a view for an input relation R , we use the simplified notation $R[\text{sch}(R)]$.

The base case (leaf in the extended variable order) is that of an input relation: We construct a view that is the relation itself. At a variable X (inner node), we distinguish two cases: If X is a bound variable, we construct a view that marginalizes out X in the natural join of the child views; we thus first join on X , then apply the lifting function for X on its values, and aggregate X away. If X is a free variable, however, we retain it in the view schema without applying the lifting function to its values. The schema of the view consists of those ancestor variables of X in ω on which it depends and the free variables already present at its children.

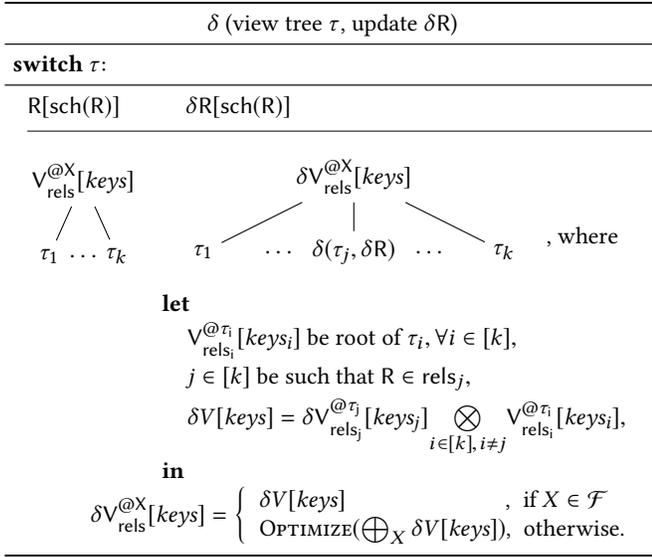


Figure 4: Creating a delta view tree $\delta(\tau, \delta R)$ for a view tree τ to accommodate an update δR to a relation R .

Example 3.2. Figure 2 gives the view tree constructed by our algorithm for the given variable order and the empty set of free variables. The figure also shows the contents of the views computing the COUNT query over the database \mathcal{D} . \square

By default, the algorithm constructs one view per variable in the variable order ω . A wide relation (with many variables) leads to long branches in ω with variables that are only local to this relation. This is, for instance, the case of our retailer dataset used in Section 7. Such long branches create long chains of views, where each view marginalizes one bound variable over its child view in the chain. For practical reasons, we compose such long chains into a single view that marginalizes several variables at a time.

4 FACTORIZED HIGHER-ORDER IVM

We introduce incremental view maintenance in our factorized ring computation framework. In contrast to evaluation, the incremental maintenance of the query result may require the materialization and maintenance of views. An update to a relation R triggers changes in all views from the leaf R to the root of the view tree.

Updates. The insertion (deletion) of a tuple \mathbf{t} into (from) a relation R is expressed as a delta relation δR that maps \mathbf{t} to 1 (and respectively -1). In general, δR can be a relation, thus a collection of tuples mapped to payloads. The updated relation is then the union of the old relation and the delta relation: $R := R \uplus \delta R$.

Delta Views. For each view V affected by an update, a *delta view* δV defines the change in the view contents. In case the view V represents a relation R , then $\delta V = \delta R$ if there are updates to R and $\delta V = \emptyset$ otherwise. In case the view is defined using operators on other views, δV is derived using the following delta rules:

$$\begin{aligned} \delta(V_1 \uplus V_2) &= \delta V_1 \uplus \delta V_2 \\ \delta(V_1 \otimes V_2) &= (\delta V_1 \otimes V_2) \uplus (V_1 \otimes \delta V_2) \uplus (\delta V_1 \otimes \delta V_2) \\ \delta(\bigoplus_X V) &= \bigoplus_X \delta V \end{aligned}$$

The correctness of the rules follows from the associativity of \uplus and the distributivity of \otimes over \uplus ; \bigoplus_X is equivalent to the repeated application of \uplus for the possible values of X . The derived delta views are subject to standard simplifications: If V is not defined over the updated relation R , then its delta view δV is empty, and then we propagate this information using the identities $\emptyset \uplus V = V \uplus \emptyset = V$ and $\emptyset \otimes V = V \otimes \emptyset = \emptyset$.

Delta Trees. Under updates to one relation, a view tree becomes a delta tree where the affected views become delta views. The algorithm in Figure 4 traverses the view tree τ top-down along the path from the root to the updated relation and replaces the views on that path with delta views. The OPTIMIZE method rewrites delta view expressions to exploit factorized updates by avoiding the materialization of Cartesian products and pushing marginalization past joins; we explain this optimization in Section 5.

Example 4.1. Consider the query from Example 2.2 and its view tree from Figure 2b. The update δT triggers delta computation at each view from the leaf T to the root of the view tree:

$$\begin{aligned} \delta V_T^{\otimes D}[C] &= \bigoplus_D \delta T[C, D] \\ \delta V_{ST}^{\otimes C}[A] &= \bigoplus_C \delta V_T^{\otimes D}[C] \otimes V_S^{\otimes E}[A, C] \\ \delta V_{RST}^{\otimes A}[\] &= \bigoplus_A V_R^{\otimes B}[A] \otimes \delta V_{ST}^{\otimes C}[A] \end{aligned}$$

Let us consider the ring \mathbb{Z} and the lifting functions that map all values to 1, and let $\delta T[C, D] = \{(c_1, d_1) \rightarrow -1, (c_2, d_2) \rightarrow 3\}$. Given the contents of $V_S^{\otimes E}$ and $V_R^{\otimes B}$ from Figure 2, we now have:

$$\begin{array}{ccc} C \rightarrow \delta V_T^{\otimes D}[C] & A \rightarrow \delta V_{ST}^{\otimes C}[A] & () \rightarrow \delta V_{RST}^{\otimes A} \\ c_1 \rightarrow -1 & a_1 \rightarrow 1 & () \rightarrow 5 \\ c_2 \rightarrow 3 & a_2 \rightarrow 3 & \end{array}$$

A single-tuple update to T fixes the values for C and D and computing $\delta V_T^{\otimes D}$ takes constant time. The delta view $\delta V_{ST}^{\otimes C}$ iterates over all possible A -values for a fixed C -value, which takes linear time; $\delta V_{RST}^{\otimes A}$ incurs the same linear-time cost. A single-tuple update to either R or S , however, fixes all variables on a leaf-to-root path in the delta view tree, giving a constant view maintenance cost. \square

In contrast to classical (first-order) IVM that only requires maintenance of the query result [12], our approach is higher-order IVM as updates may trigger maintenance of several interrelated views. The fully-recursive IVM scheme of DBToaster [24, 25] creates one materialization hierarchy per relation in the query, whereas we use one view tree for all relations. This view tree relies on variable orders to decompose the query into views and factorize its computation and maintenance.

Which Views to Materialize and Maintain? The answer to this question depends on which relations may change. Let \mathcal{U} be the (multi)set of the updatable relations; if a relation occurs several times in the query, then it has several instances in \mathcal{U} .

Propagating changes along a leaf-to-root path is computationally most effective if each delta view joins with sibling views that are already materialized. Figure 5 depicts an algorithm that reflects this idea: Given a view tree τ and a set of relations \mathcal{U} , the algorithm builds a *materialization tree* $\mu(\tau, \mathcal{U})$ that has the same structure as τ and where each node is a Boolean value indicating whether the corresponding view from τ should be materialized. The root view

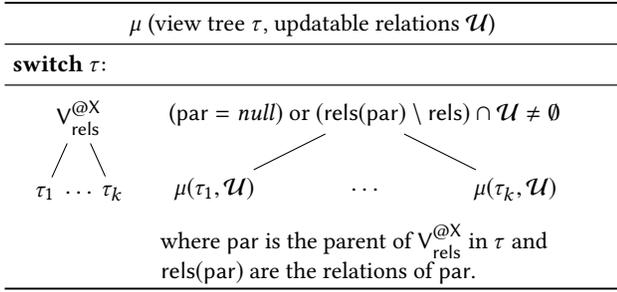


Figure 5: Which views in a view tree τ need to be materialized to accommodate updates to a multiset of relations \mathcal{U} .

is always stored as it represents the query result, that is, it has no parent: $par = null$; every other view V is stored only if it is used to compute the delta of its parent for updates to a relation over which V is not defined, that is, there are updatable relations for the parent and not for V itself: $(rels(par) \setminus rels) \cap \mathcal{U} \neq \emptyset$.

Example 4.2. We continue with our query from Example 4.1. For updates to T only, that is, $\mathcal{U} = \{T\}$, we store the root $V_{RST}^{\otimes A}$ and the views $V_S^{\otimes E}$ and $V_R^{\otimes B}$ used to compute the deltas $\delta V_{ST}^{\otimes C}$ and $\delta V_{RST}^{\otimes A}$. Only the root view is affected by these changes and maintained as:

$$V_{RST}^{\otimes A}[\] = V_{RST}^{\otimes A}[\] \uplus \delta V_{RST}^{\otimes A}[\]$$

It is not necessary to maintain other views. If we would like to also support updates to both R and S , then we would also need to materialize $V_{ST}^{\otimes C}$ and $V_T^{\otimes D}$.

If no updates are supported, then only the root view is stored. \square

For queries with free variables, several views in their (delta) view trees may be identical: This can happen when all variables in their keys are free and thus cannot be marginalized. For instance, a variable order ω for the query from Example 2.3 may have the variables A and C above all other variables, in which case their corresponding views are the same in the view tree for ω . We then only store the top view out of these identical views.

IVM Triggers. For each updatable relation R , our framework constructs a trigger procedure that takes as input an update δR and implements the maintenance schema of the corresponding delta view tree. This procedure also maintains all materialized views needed for the given update workload.

A bulk of updates to distinct relations is handled as a sequence of updates, one per relation. Update sequences can also happen when updating a relation R that occurs several times in the query. The instances representing the same relation are at different leaves in the delta tree and lead to changes along multiple leaf-to-root paths.

5 FACTORIZABLE UPDATES

Our focus so far has been on supporting updates represented by delta relations. We next consider an alternative approach that decomposes a delta relation into a union of factorizable relations. The cumulative size of the decomposition relations can be much less than the size of the original delta relation. Also, the complexity of propagating a factorized update can be much lower than that of its unfactorized (listing) representation, since the factorization makes explicit the independence between query variables and enables

optimizations of delta propagation such as pushing marginalization past joins. Besides the factorized view computation, this is the second instance where our IVM approach exploits factorization.

Factorizable updates arise in many domains such as linear algebra and machine learning. Section 6 demonstrates how our framework can be used for the incremental evaluation of matrix chain multiplication, subsuming prior work on this [33]. Matrix chain computation can be phrased in our language of joins and aggregates, where matrices are binary relations. Changes to one row/column in an input matrix may be expressed as a product of two vectors. In general, an arbitrary update matrix can be decomposed into a sum of rank-1 matrices, each of them expressible as products of vectors, using low-rank tensor decomposition methods [26, 44].

Example 5.1. Arbitrary relations can be decomposed into a union of factorizable relations. The relation $R[A, B] = \{(a_i, b_j) \rightarrow 1 \mid i \in [n], j \in [m]\}$ can be decomposed as $R_1[A] \otimes R_2[B]$, where $R_1[A] = \{(a_i) \rightarrow 1 \mid i \in [n]\}$ and $R_2[B] = \{(b_j) \rightarrow 1 \mid j \in [m]\}$. We thus reduced a relation of size nm to two relations of cumulative size $n + m$. If R were a delta relation, the delta views on top of it would now be expressed over $R_1[A] \otimes R_2[B]$ and their computation can be factorized as done for queries in Section 3. Product decomposition of relations can be done in linearithmic time in both the number of variables and the size of the relation [35].

Consider now the relation $R'[A, B] = R[A, B] \uplus \{(a_{n+1}, b_j) \rightarrow 1 \mid j \in [m-1]\}$ with R as above. We can decompose each of the two terms in R' similarly to R , yielding an overall $n + 2m$ values instead of $nm + m - 1$. A different decomposition with $n + m + 3$ values is given by a factorizable over-approximation of R' compensated by a small product with negative payload: $\{(a_i) \rightarrow 1 \mid i \in [n+1]\} \otimes \{(b_j) \rightarrow 1 \mid j \in [m]\} \uplus \{(a_{n+1}) \rightarrow 1\} \otimes \{(b_m) \rightarrow -1\}$. \square

The OPTIMIZE method in the algorithm from Figure 4 exploits the distributivity of \otimes over \bigoplus_X to push the latter over those views with variable X . This optimization is reminiscent of pushing aggregates past joins in databases and variable elimination in probabilistic graphical models. It has been recently revisited for functional aggregate queries [5]. In case the delta views express Cartesian products, then they are not materialized and instead kept factorized.

Example 5.2. Consider our query from Example 4.1 under updates to relation S . Using the delta view tree derived for updates to S , the top-level delta is computed as follows:

$$\begin{aligned}
 \delta V_{RST}^{\otimes A}[\] &= \bigoplus_A V_R^{\otimes B}[A] \otimes \left(\bigoplus_C V_T^{\otimes D}[C] \otimes \underbrace{\left(\bigoplus_E \delta S[A, C, E] \right)}_{\delta V_S^{\otimes E}[A, C]} \right) \\
 &\qquad\qquad\qquad \underbrace{\hspace{10em}}_{\delta V_{ST}^{\otimes C}[A]}
 \end{aligned}$$

A single-tuple update δS binds variables A , C , and E , and computing $\delta V_{RST}^{\otimes A}$ requires $O(1)$ lookups in $V_T^{\otimes D}$ and $V_R^{\otimes B}$. An arbitrary-sized update δS can then be processed in $O(|\delta S|)$ time.

Assume now that δS is factorizable as $\delta S[A, C, E] = \delta S_A[A] \otimes \delta S_C[C] \otimes \delta S_E[E]$. In the construction of the delta view tree, the OPTIMIZE method exploits this factorization to push the marginalization past joins at each variable; for example, the delta at E becomes:

$$\begin{aligned}
 \delta V_S^{\otimes E}[A, C] &= \bigoplus_E \delta S_A[A] \otimes \delta S_C[C] \otimes \delta S_E[E] \\
 &= \delta S_A[A] \otimes \delta S_C[C] \otimes \bigoplus_E \delta S_E[E]
 \end{aligned}$$

We also transform the top-level delta into a product of three views:

$$\delta V_{RST}^{\textcircled{A}}[\] = \left(\bigoplus_A V_R^{\textcircled{B}}[A] \otimes \delta S_A[A] \right) \otimes \left(\bigoplus_C V_T^{\textcircled{D}}[C] \otimes \delta S_C[C] \right) \otimes \left(\bigoplus_E \delta S_E[E] \right)$$

Computing this delta takes time proportional to the sizes of the three views: $O(\min(|V_R^{\textcircled{B}}|, |\delta S_A|) + \min(|V_T^{\textcircled{D}}|, |\delta S_C|) + |\delta S_E|)$. \square

6 APPLICATIONS

Our IVM framework supports a wide range of application scenarios. In this section, we highlight two scenarios, matrix chain computation and gradient computation for learning linear regression models over joins, in which the payloads have sizes independent of the input relation sizes. We also highlight two distinct scenarios, where the payloads may have arbitrarily large sizes: they can be entire relations under either listing or factorized representations. *All these scenarios behave identically in the key space of the views and are treated uniformly using delta view trees. They differ however in the rings used to define the view payloads.*

6.1 Matrix Chain Multiplication

Consider the problem of computing a product of a series of matrices A_1, \dots, A_n over some ring \mathbf{D} , where matrix $A_i[x_i, x_{i+1}]$ has the size of $p_i \times p_{i+1}$, $i \in [n]$. The product $A = A_1 \cdots A_n$ is a matrix of size $p_1 \times p_{n+1}$ and can be formulated as follows:

$$A[x_1, x_{n+1}] = \sum_{x_2 \in [p_2]} \cdots \sum_{x_n \in [p_n]} \prod_{i \in [n]} A_i[x_i, x_{i+1}]$$

We model a matrix A_i as a relation $A_i[X_i, X_{i+1}]$ with the payload carrying matrix values. The query that computes the matrix A is:

$$A[X_1, X_{n+1}] = \bigoplus_{X_2} \cdots \bigoplus_{X_n} \bigotimes_{i \in [n]} A_i[X_i, X_{i+1}]$$

where each of the lifting functions $\{g_{X_j}\}_{j \in [2, n]}$ maps any key value to payload $\mathbf{1} \in \mathbf{D}$.

Different variable orders lead to different evaluation plans for matrix chain multiplication. The optimal variable order corresponds to the optimal sequence of matrix multiplications that minimizes the overall multiplication cost, which is the textbook Matrix Chain Multiplication problem [13].

Example 6.1. Consider a multiplication chain of 4 matrices of equal size $p \times p$ represented as relations $A_i[X_i, X_{i+1}]$. Let $\mathcal{F} = \{X_1, X_5\}$ be the set of free variables and ω be the variable order $X_1 - X_5 - X_3 - \{X_2, X_4\}$, with the matrix relations placed below the leaf variables in ω . Then, the view tree $\tau(\omega, \mathcal{F})$ has the following views (from bottom to top):

$$\begin{aligned} V_{A_1 A_2}^{\textcircled{X_2}}[X_1, X_3] &= \bigoplus_{X_2} A_1[X_1, X_2] \otimes A_2[X_2, X_3] \\ V_{A_3 A_4}^{\textcircled{X_4}}[X_3, X_5] &= \bigoplus_{X_4} A_3[X_3, X_4] \otimes A_4[X_4, X_5] \\ V_{A_1 A_2 A_3 A_4}^{\textcircled{X_3}}[X_1, X_5] &= \bigoplus_{X_3} V_{A_1 A_2}^{\textcircled{X_2}}[X_1, X_3] \otimes V_{A_3 A_4}^{\textcircled{X_4}}[X_3, X_5] \end{aligned}$$

The views at X_5 and X_1 are equivalent to the view at X_3 , and recomputing them from scratch takes $O(p^3)$ time. A single-value change in any input matrix causes changes in one row or column of the parent view, and propagating them to compute the final delta view takes $O(p^2)$ time. Updates to A_2 and A_3 change every value in A . In case of a longer matrix chain, further propagating δA would require $O(p^3)$ matrix multiplications, same as re-computation.

We exploit factorization to contain the computational effect of such changes. For instance, if δA_2 is a factorizable update (see Section 5) expressible as $\delta A_2[X_2, X_3] = u[X_2] \otimes v[X_3]$, then we can propagate deltas more efficiently, as products of subexpressions:

$$\delta V_{A_1 A_2}^{\textcircled{X_2}}[X_1, X_3] = \underbrace{\left(\bigoplus_{X_2} A_1[X_1, X_2] \otimes u[X_2] \right)}_{u_2[X_1]} \otimes v[X_3]$$

$$\delta V_{A_1 A_2 A_3 A_4}^{\textcircled{X_3}}[X_1, X_5] = u_2[X_1] \otimes \left(\bigoplus_{X_3} v[X_3] \otimes V_{A_3 A_4}^{\textcircled{X_4}}[X_3, X_5] \right)$$

Using such factorizable updates enables IVM in $O(p^2)$ time. The final delta is also in factorized form, suitable for further propagation.

In general, for a chain of k matrices of size $p \times p$, using a binary view tree of the lowest depth, incremental maintenance with factorizable updates takes $O(p^2 \log k)$ time, while re-evaluation takes $O(p^3 k)$ time. The space needed in both cases is $O(p^2 k)$. \square

The above example recovers the main idea of LINVIEW [33]: exploit factorization for incremental computation of linear algebra programs when matrix changes are expressed as vector outer products, $\delta A = uv^T$. This so-called rank-1 updates can capture many practical update patterns such as perturbations of one complete row or column in a matrix, or even changes of the whole matrix when the same vector is added to every row or column. Our framework generalizes this idea to arbitrary join-aggregate queries.

6.2 Gradient Computation

We introduce a ring that captures the data-dependent computation of the gradient for training linear regression models over joins.

Consider a training dataset that consists of k examples with $(X_i)_{i \in [m-1]}$ features/variables and a label X_m arranged into a design matrix \mathbf{M} of size $k \times m$; in our setting, this design matrix is the result of a join query. The goal of linear regression is to learn the model parameters $\theta = [\theta_1 \dots \theta_m]^T$ of a linear function¹ $f(X_1, \dots, X_{m-1}) = \sum_{i \in [m]} \theta_i X_i$ best satisfying $\mathbf{M}\theta \approx \mathbf{0}_{k \times 1}$, where $\mathbf{0}_{k \times 1}$ is the k -by-1 matrix with all elements 0.

We can solve this optimization problem using batch gradient descent. This method iteratively updates the model parameters in the direction of the gradient to decrease the squared error loss and eventually converge to the optimal value. Each convergence step iterates over the entire training dataset to update the parameters, $\theta := \theta - \alpha \mathbf{M}^T \mathbf{M} \theta$, where α is an adjustable step size. The complexity of each step is $O(mk)$. The cofactor matrix $\mathbf{M}^T \mathbf{M}$ quantifies the degree of correlation for each pair of features (or label). Its computation is dependent on the data and can be executed once for all convergence steps [42]. This is crucial for performance in case $m \ll k$ as each iteration step now avoids processing the entire training dataset and takes time $O(m^2)$.

We next show how to incrementally maintain the cofactor matrix. We can factorize its computation over training datasets defined by results to arbitrary join queries [42]. The idea is to compute a triple of regression aggregates (c, s, Q) , where c is the number of tuples in the training dataset (size k of the design matrix), s is an $m \times 1$ matrix (or vector) with one sum of values per variable, and Q is

¹We consider wlog: θ_1 is the bias parameter and then $X_1 = 1$ for all tuples in the input data; θ_m remains fixed to -1 and corresponds to the label/response X_m in the data.

an $m \times m$ matrix of sums of dot products of values for any two variables. Their incremental computation can be captured by a ring.

Definition 6.2. For a fixed $m \in \mathbb{N}$, let \mathbf{D} denote the set of triples $(\mathbb{Z}, \mathbb{R}^m, \mathbb{R}^{m \times m})$, $\mathbf{0} = (0, \mathbf{0}_{m \times 1}, \mathbf{0}_{m \times m})$ and $\mathbf{1} = (1, \mathbf{0}_{m \times 1}, \mathbf{0}_{m \times m})$. For $a = (c_a, s_a, Q_a) \in \mathbf{D}$ and $b = (c_b, s_b, Q_b) \in \mathbf{D}$, define the operations $+^{\mathbf{D}}$ and $*^{\mathbf{D}}$ over \mathbf{D} as:

$$\begin{aligned} a +^{\mathbf{D}} b &= (c_a + c_b, s_a + s_b, Q_a + Q_b) \\ a *^{\mathbf{D}} b &= (c_a c_b, c_b s_a + c_a s_b, c_b Q_a + c_a Q_b + s_a s_b^T + s_b s_a^T) \end{aligned}$$

The structure $(\mathbf{D}, +^{\mathbf{D}}, *^{\mathbf{D}}, \mathbf{0}, \mathbf{1})$ forms a *degree- m matrix ring*.

We next show how to use this ring to compute the cofactor matrix over a training dataset defined by a join query with relations $(R_i)_{i \in [n]}$ over variables $(X_j)_{j \in [m]}$. The payload of each tuple in a relation is the identity $\mathbf{1}$ from the degree- m matrix ring. The query expressing the computation of the cofactor matrix is:

$$Q = \bigoplus_{X_1} \cdots \bigoplus_{X_m} \bigotimes_{i \in [n]} R_i[\text{sch}(R_i)]$$

For each X_j -value x , the lifting function is $g_{X_j}(x) = (1, s, Q)$, where s is an $m \times 1$ vector with all zeros except the value of x at position j , i.e., $s_j = x$, and Q is an $m \times m$ matrix with all zeros except the value x^2 at position (j, j) : $Q_{(j,j)} = x^2$.

Example 6.3. We show how to compute the cofactor matrix over the join, database, and view tree from Figure 2b. We assume alphabetical order over the five variables. The leaves R, S, and T are the input relations that map tuples to $\mathbf{1}$ from the degree-5 matrix ring.

In the view V_T^{D} each D -value d is lifted to a triple $(1, s, Q)$, where s is a 5×1 vector with one non-zero element $s_4 = d$, and Q is a (5×5) matrix with one non-zero element $Q_{(4,4)} = d^2$. Then, those regression triples with the same key c are summed up, yielding $V_T^{\text{D}}[c_1] = (1, s_4 = d_1, Q_{(4,4)} = d_1^2)$, $V_T^{\text{D}}[c_2] = (2, s_4 = d_2 + d_3, Q_{(4,4)} = d_2^2 + d_3^2)$, and $V_T^{\text{D}}[c_3] = (1, s_4 = d_4, Q_{(4,4)} = d_4^2)$. The views V_R^{B} and V_S^{E} are computed similarly.

The view V_{ST}^{C} joins the views V_T^{D} and V_S^{E} and then marginalizes C . For instance, the payload for the key $V_{ST}^{\text{C}}[a_2]$ is as follows:

$$\begin{aligned} V_{ST}^{\text{C}}[a_2] &= V_T^{\text{D}}[c_2] * V_S^{\text{E}}[a_2, c_2] * g_C(c_2) \\ &= \left(2, \begin{bmatrix} 0 \\ 0 \\ 0 \\ d_2 + d_3 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & d_2^2 + d_3^2 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) * \left(1, \begin{bmatrix} 0 \\ 0 \\ 0 \\ e_4 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & e_4^2 \end{bmatrix} \right) * \left(1, \begin{bmatrix} 0 \\ 0 \\ c_2 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ c_2^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) \\ &= \left(2, \begin{bmatrix} 0 \\ 0 \\ 2c_2 \\ d_2 + d_3 \\ 2e_4 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 2c_2^2 & c_2(d_2 + d_3) & 2c_2 e_4 \\ 0 & c_2(d_2 + d_3) & d_2^2 + d_3^2 & (d_2 + d_3)e_4 \\ 0 & 2c_2 e_4 & (d_2 + d_3)e_4 & 2e_4^2 \end{bmatrix} \right) \end{aligned}$$

The root view V_{RST}^{A} maps the empty tuple to the ring element $\sum_{l \in [2]} V_R^{\text{B}}[a_l] * V_{ST}^{\text{C}}[a_l] * g_A(a_l)$. This payload carries aggregates for the entire join result: the count of tuples in the result, the vector with one sum of values per variable, and the cofactor matrix. \square

For performance reasons, in practice we only store as payloads blocks of matrices with non-zero values and assemble larger matrices as the computation progresses towards the root. We can further

exploit the symmetry of the cofactor matrix to compute only the entries above and including the diagonal.

6.3 Factorized Representation of Query Results

Our framework can also support scenarios where the view payloads are themselves relations representing results of conjunctive queries, or even their factorized representations. Factorized representations can be arbitrarily smaller than the listing representation of a query result [38], with orders of magnitude size gaps reported in practice [42]. They nevertheless remain lossless and support constant-delay enumeration of the tuples in the query result as well as subsequent aggregate processing in one pass. Besides the factorized view computation and the factorizable updates, this is the third instance where our framework exploits factorization.

We first introduce the relational data ring that allows us to store entire relations as payloads. We then show how to encode factorized representation of relations in payloads. By using a relational data ring, we can create payloads that hold relations. When marginalizing a variable, we move its values from the key space to the payload space. The payloads of the tuples of a view are now relations over the same schema. These relations have themselves payloads in the \mathbb{Z} ring used to maintain the multiplicities of their tuples.

Definition 6.4. Let $\mathbb{F}[\mathbb{Z}]$ denote the set of relations over the \mathbb{Z} ring, the zero $\mathbf{0}$ in $\mathbb{F}[\mathbb{Z}]$ is a relation that maps every tuple to $0 \in \mathbb{Z}$, while the identity $\mathbf{1}$ is a relation that maps the empty tuple to $1 \in \mathbb{Z}$ and all other tuples to $0 \in \mathbb{Z}$, denoted as $\{() \rightarrow 1\}$. The structure $(\mathbb{F}[\mathbb{Z}], \oplus, \otimes, \mathbf{0}, \mathbf{1})$ forms a ring called the relational data ring.²

We model conjunctive queries as count queries that marginalize every variable but use different lifting functions for the free and bound variables. For a variable X and any of its values x , $g_X(x) = \{(x) \rightarrow 1\}$ if X is a free variable and $g_X(x) = \mathbf{1} = \{() \rightarrow 1\}$ if X is bound; here, (x) is a singleton relation over schema $\{X\}$. We have relational operations occurring at two levels: for keys, we join views from different branches and marginalize variables as before; for payloads, we interpret multiplication and addition of payloads as join and respectively union of relations.

Example 6.5. Consider the conjunctive query

$$Q(A, B, C, D) = R(A, B), S(A, C, E), T(C, D)$$

over the three relations from Figure 2, where each tuple gets the identity payload $\{() \rightarrow 1\} \in \mathbb{F}[\mathbb{Z}]$. The corresponding view is:

$$Q[] = \bigoplus_A \bigoplus_B \bigoplus_C \bigoplus_D \bigoplus_E R[A, B] \otimes S[A, C, E] \otimes T[C, D]$$

The lifting function for E maps each value to $\{() \rightarrow 1\}$, while the lifting functions for all other variables map value x to $\{(x) \rightarrow 1\}$.

Figure 2 shows a view tree for this query and the contents of its views with relational data payloads (in black and red) for the given database. The view keys gradually move to payloads as the computation progresses towards the root. The view definitions are identical to those of the COUNT query (but under a different ring!). The view V_T^{D} lifts each D -value d from T to the relation $\{(d) \rightarrow 1\}$ over schema $\{D\}$, multiplies (joins) it with the payload $\mathbf{1}$ of each

² To form a proper ring, we would need a generalization [24] of our relations and join and union operators, where: tuples have their own schemas; union may apply to tuples with possibly different schemas; join accounts for multiple derivations of output tuples. For our practical needs this generalization is not necessary.

tuple, and sums up (union) all payloads with the same c -value. The views at $V_R^{\textcircled{B}}$ and $V_S^{\textcircled{E}}$ are computed similarly, except the latter lifts e -values to $\{() \rightarrow 1\}$ since E is a bound variable. The view $V_{ST}^{\textcircled{C}}$ assigns to each A -value a payload that is a union of Cartesian products of the payloads of its children and the lifted C -value. The root view $V_{RST}^{\textcircled{A}}$ similarly computes the payload of the empty tuple, which represents the query result (both views are at the right). \square

We next show how to construct a factorized representation of the query result. In contrast to the scenarios discussed above, this representation is *not* available as one payload at the root view, but *distributed* over the payloads of all views. This hierarchy of payloads, linked via the keys of the views, becomes the factorized representation. A further difference lies with the multiplication operation. For the listing representation, the multiplication is the Cartesian product. For a given view, it is used to concatenate payloads from its child views. For the factorized representation, we further project away values for all but the marginalized variable. More precisely, for each view $V_{\text{rels}}^{\textcircled{X}}[S]$ and each of its keys a_S , let $P[\mathcal{T}] = V_{\text{rels}}^{\textcircled{X}}[a_S]$ be the corresponding payload relation. Then, instead of computing this payload, we compute $\bigoplus_{Y \in \mathcal{T} - \{X\}} P[\mathcal{T}]$ by marginalizing the variables in $\mathcal{T} - \{X\}$ and summing up the multiplicities of the tuples in $P[\mathcal{T}]$ with the same X -value.

Example 6.6. We continue Example 6.5. Figure 2e shows the contents of the views with factorized payloads (in black and blue). Each view stores relational payloads that have the schema of the marginalized variable. Together, these payloads form a factorized representation over the variable order ω used to define the view tree in Figure 2. At the top of the factorization, we have a union of two A -values: a_1 and a_2 . This is stored in the payloads of (middle) $V_{RST}^{\textcircled{A}}$. The payloads of (middle) $V_{ST}^{\textcircled{C}}[A]$ store a union of C -values c_1 and c_2 under a_1 , and a singleton union of c_2 under a_2 . The payloads of $V_R^{\textcircled{B}}[A]$ store a union of B -values b_1 and b_2 under a_1 and a singleton union of b_3 under a_2 . Note the (conditional) independence of the variables B and C given a value for A . This is key to succinctness of factorization. In contrast, the listing representation explicitly materializes all pairings of B and C -values for each A -value, as shown in the payload of (right) $V_{RST}^{\textcircled{A}}$. Furthermore, the variable D is independent of the other variables *given* C . This is a further source of succinctness in the factorization: Even though c_2 occurs under both a_1 and a_2 , the relations under c_2 , in this case the union of d_2 and d_3 , is only stored once in $V_T^{\textcircled{D}}[C]$. Each value in the factorization keeps a multiplicity, that is, the number of its derivations from the input data. This is necessary for maintenance.

This factorization is over a variable order that can be used for all queries with same body and different free variables: As long as their free variables sit on top of the bound variables, the variable order is valid and so is the factorization over it. For instance, if the variable D were not free, then the factorization for the new query would be the same except that we now discard the unions of D -values. \square

7 EXPERIMENTS

We compare F-IVM (factorized IVM) against 1-IVM (first-order IVM) and DBT (DBToaster’s fully recursive higher-order IVM). Our experimental results can be summarized as follows:

- Factorized updates lead to two orders of magnitude speedup for F-IVM over competitors for matrix chain multiplication by propagating factorized deltas and avoiding matrix multiplication.
- For cofactor matrices in regression models, F-IVM exhibits the lowest memory utilization and up to two orders of magnitude better performance than 1-IVM and DBT.
- For conjunctive query evaluation, factorized payloads can both speed up view maintenance and reduce memory by up to two orders of magnitude compared to using listing representation.

Due to lack of space, details on the workload and experimental setup, as well as further experiments, are deferred to Appendix C.

Runtime. 1-IVM and DBT are supported by DBToaster [25], a system that compiles a given SQL query into code that maintains the query result under updates to input relations. The generated code represents an in-memory stream processor that is standalone and independent of any database system. DBToaster’s performance on decision support and financial workloads can be several orders of magnitude better than state-of-the-art commercial databases and stream processing systems [25]. We implemented F-IVM as a program that maintains a set of materialized views for a given variable order and a set of updatable relations. We use the intermediate language of DBToaster to encode this program and then feed it into DBToaster’s code generator. We modified the backend of DBToaster v2.2 to enable arbitrary ring payloads and limit the amount of memory over-provisioning to at most one million records. Unless stated otherwise, all the benchmarked approaches use the same runtime environment and materialize views as multi-indexed maps with memory-pooled records. The algorithms and record types used in these approaches, however, can differ greatly.

Workload. We use a real-world dataset *Retailer* used for business decision support and forecasting user demands, a synthetic dataset *Housing* modeling a house price market [42], and dense matrices with random double-precision values from $(-1, 1)$. We run the systems over data streams synthesized from these datasets by interleaving updates to the input relations in a round-robin fashion and grouping them into batches of fixed size.

Matrix Chain Multiplication with Factorized Updates. We consider the problem of maintaining the multiplication $A = A_1 A_2 A_3$ of three $(n \times n)$ matrices under changes to A_2 . We compare F-IVM with factorized updates, 1-IVM that re-computes the delta $\delta A = A_1 \delta A_2 A_3$ from scratch (same as DBT in this particular example), and RE-EVAL that re-computes the entire product from scratch on every update. We consider two different implementations of these maintenance strategies: The first uses DBToaster’s hash maps to store matrices, while the second uses Octave, a numerical tool that stores matrices in dense arrays and offers highly-optimized BLAS routines for matrix multiplication [45]. In both cases, matrix-matrix multiplication takes $O(n^\alpha)$ for $\alpha > 2$; for instance, $\alpha = 2.8074$ for Strassen’s algorithm.

We first consider updates to one row in A_2 . For 1-IVM, the delta $\delta A_{12} = A_1 \delta A_2$ might contain non-zero changes to all n^2 matrix entries, thus computing $\delta A = \delta A_{12} A_3$ requires full matrix-matrix multiplication. RE-EVAL updates A_2 first before computing two matrix-matrix multiplications. F-IVM factorizes δA_2 into a product of two vectors $\delta A_2 = uv^T$, which are used to compute $\delta A_{12} = (A_1 u) v^T = u_1 v^T$ and $\delta A = u_1 (v^T A_3) = u_1 v_1$. Both deltas involve

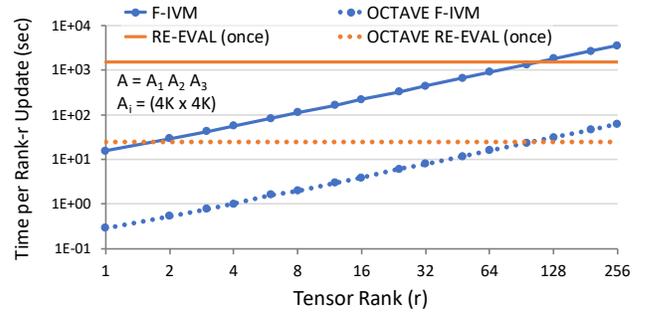
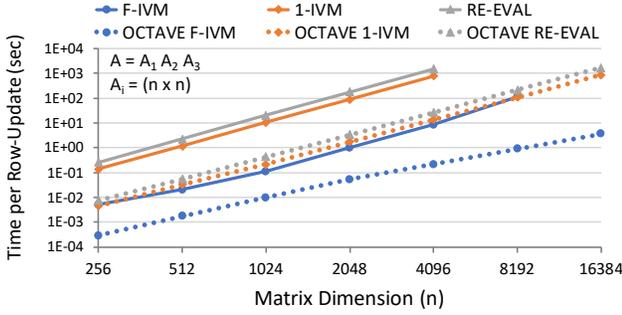


Figure 6: Incremental maintenance and re-evaluation of the product of three $(n \times n)$ matrices, $A = A_1 A_2 A_3$: (left) one-row updates in A_2 ; (right) rank- r updates in A_2 for $n = 4,096$ using the DBToaster and Octave runtimes.

only matrix-vector multiplications computed in $O(n^2)$ time. Figure 6 (left) shows the average time needed to process an update to one randomly selected row in A_2 for different matrix sizes. RE-EVAL performs two matrix-matrix multiplications, while 1-IVM performs only one. In the hash-based implementation, the gap between F-IVM and 1-IVM grows from 28x for $n = 256$ to 92x for $n = 4,096$; similarly, in the Octave implementation, the same gap grows from 16x for $n = 256$ to 236x for $n = 16,384$. This confirms the difference in the asymptotic complexity of these strategies.

Our next experiment considers rank- r updates to A_2 , which can be decomposed into a sum of r rank-1 tensors, $\delta A_2 = \sum_{i \in [r]} u_i v_i^T$. F-IVM processes δA_2 as a sequence of r rank-1 updates in $O(rn^2)$ time, while both RE-EVAL and 1-IVM take as input one full matrix δA_2 and maintain the product in $O(n^3)$ time per each rank- r update (1-IVM is omitted from the plot). Figure 6 (right) shows that the average time F-IVM takes to process a rank- r update for different r values and the matrix size 4,096 is linear in the tensor rank r . Under both implementations in DBToaster and Octave, incremental computation is faster than re-evaluation for updates with rank $r \leq 96$. With larger matrix sizes, the gap between re-evaluation and incremental computation increases, which enables incremental maintenance for updates of higher ranks.

Cofactor Matrix Computation. We benchmark the performance of maintaining a cofactor matrix for learning regression models over a natural join. We compute the cofactor matrix over all variables of the join query (i.e., over all attributes of the input database), which suffices to learn linear regression models over *any label and set of features* that is a subset of the set of variables [36]. This is achieved by specializing the convergence step to the relevant restriction of the cofactor matrix. In end-to-end learning of regression models over factorized joins in the *Retailer* and *Housing* datasets, the convergence step takes orders of magnitude less time compared to the data-dependent cofactor matrix computation [42].

In addition to the three incremental strategies from before, we now also benchmark DBT-RING, DBToaster’s recursive IVM strategy with payloads from the degree- m ring (cf. Section 6.2) instead of scalars, and SQL-OPT, an optimized SQL encoding of cofactor matrix computation. The latter arranges regression aggregates – recall there are quadratically many such aggregates in the number of query variables – into a *single* aggregate column indexed by the degree of each query variable. This encoding takes as input a variable order and constructs one SQL query that intertwines join and

aggregate computation by pushing (partial) regression aggregates (counts, sums, and cofactors) past joins [37].

We consider updates to all relations in the *Retailer* and *Housing* datasets. In the *Retailer* schema, F-IVM and SQL-OPT rely on a given variable order. These two strategies store 9 views each: five views over the input relations, three intermediate views, and the top-level view; DBT-RING stores four additional views, 13 in total. These views are identical to those used for maintaining a sum aggregate (see Appendix C) but have different payloads. DBToaster’s recursive higher-order IVM and first-order IVM use scalar payloads and fail to effectively share the computation of regression aggregates, materializing linearly many views in the size of the cofactor matrix: DBT and 1-IVM use 3,814 and respectively 995 views to maintain 990 aggregates. In the *Housing* schema, where all relations join on one variable, F-IVM and SQL-OPT materialize one view per relation and the root view, 7 in total, while DBT and 1-IVM use 702 and 412 views to maintain 406 aggregates. F-IVM and DBT-RING use identical strategies for the *Housing* dataset.

Figure 7 shows the throughput of these techniques as they process an increasing fraction of the stream of tuple inserts. The *Retailer* stream consists of inserts into the largest relation mostly, and since the variables of this relation form a root-to-leaf path in the variable order, processing a single-tuple update takes $O(1)$ time for F-IVM and SQL-OPT. The former outperforms the latter due to efficient encoding of triples of aggregates (c, s, Q) as payloads containing vectors and matrices. DBT-RING’s additional views cause non-constant update times to the largest relation, which means 8.7x lower average throughput than F-IVM. The two approaches with scalar payloads, DBT and 1-IVM, need to maintain too many views and fail to process the entire stream within a one-hour limit.

The query for *Housing* is a star join with all relations joining on the common variable, which is the root in our variable order. Thus, F-IVM and SQL-OPT can process a single tuple in $O(1)$ time. DBT-RING and F-IVM use the same strategy in this case. DBT exploits the conditional independence in the derived deltas to materialize each input relation separately such that all non-join variables are aggregated away. Although each materialized view has $O(1)$ maintenance cost per update tuple, the large number of such views in DBT is the main reason for its poor performance. In contrast, 1-IVM stores entire tuples of the input relations including non-join variables. On each update, 1-IVM recomputes an aggregate on top of the join of these input relations and the update. Since an update

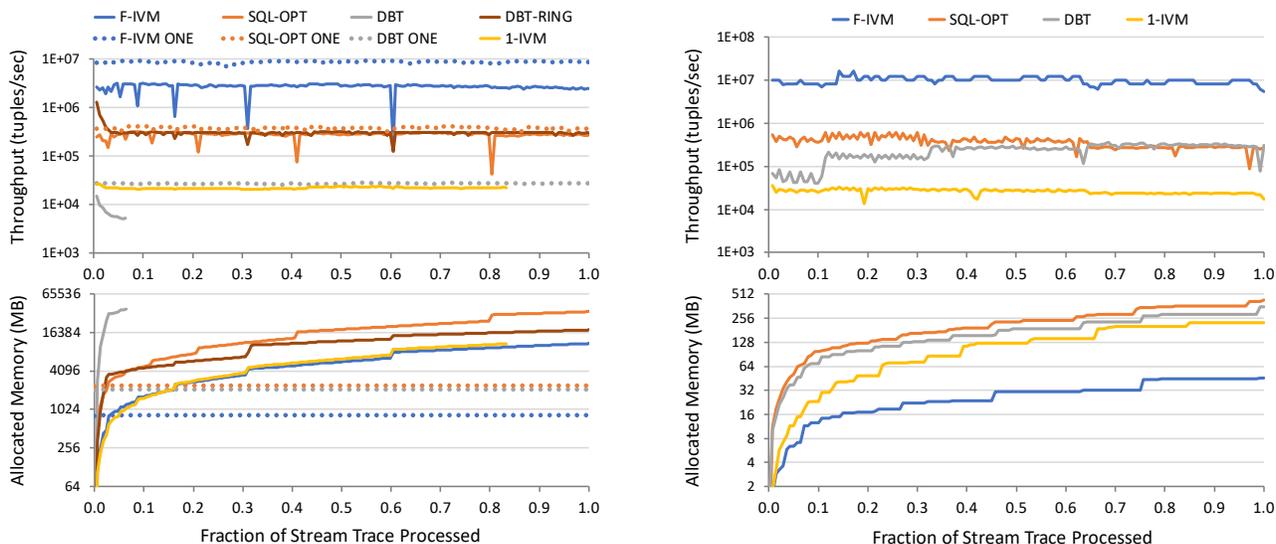


Figure 7: Incremental maintenance of the cofactor matrix over the *Retailer* dataset (left) and *Housing* dataset (right) under updates of size 1,000 to all relations with a one-hour timeout. The ONE plots consider updates to the largest relation only.

tuple binds the value of the common join variable, the delta query consists of disconnected components. DBToaster optimizes such a delta query by placing an aggregate around each component, which means that re-computing a delta means on-the-fly pre-aggregation over each relation followed by a join. Thus, 1-IVM takes linear time, which explains its poor performance.

Memory Consumption. Figure 7 shows that F-IVM achieves the lowest memory utilization on both datasets while providing orders of magnitude better performance than its competitors! The reason behind the memory efficiency of our approach is twofold. First, it uses complex aggregates and factorization structures to express the cofactor matrix computation over a smaller set of views compared to DBT-RING and, even more, to DBT and 1-IVM. Second, it encodes regression aggregates implicitly using vectors and matrices rather than explicitly using variable degrees, like in SQL-OPT. The occasional throughput hiccups in the plot are due to expansion of the underlying data structures used for storing views.

The Effect of Update Workload. Our next experiment studies the effect of different update workloads on performance. We consider the *Retailer* dataset and two possible update scenarios: (1) all relations can change, in which case every view in the view tree needs to be materialized; (2) only the largest relation changes, while all others are static (denoted as ONE in Figure 7). In the latter scenario, we can precompute the views that are unaffected by changes and avoid materialization of those views that do not directly join with the updated relation. Thus, restricting updates to only one relation leads to materializing fewer views, which in turn reduces the maintenance overhead. Figure 7 shows the throughput of processing updates for the incremental maintenance of the cofactor matrix in these two scenarios. If we restrict updates only to a relation, we can avoid materializing all the views on the leaf-to-root path covered by that relation. This corresponds to a streaming scenario where we compute a continuous query and do not store the stream. Restricting updates to only one relation improves the

average throughput, 3.2x in F-IVM and 1.3x in SQL-OPT, and also decreases memory consumption (note the log y -axis). The latter also reflects in smoother throughput curves for the ONE variants. In DBT, restricting updates brings constant time maintenance per view, yet the number of views is still large. F-IVM and DBT-RING use identical materialization strategies in this scenario.

Factorized Computation of Conjunctive Queries. We analyze F-IVM on queries whose results are stored as keys with integer multiplicities using listing representation (List keys) and as relational payloads using factorized and listing representations (Fact payloads and List payloads). Figure 8 (left) considers the natural join of *Retailer* under updates to the largest relation. The factorized payloads reduce the memory consumption by 4.4x, from 34GB to 7.8GB, and improve the average throughput by 2.8x and 3.7x compared to using the two listing encodings. Figure 8 (right) considers the natural join of *Housing* under updates to all input relations. The number of tuples in the dataset varies from 150,000 (scale 1) to 1,400,000 (scale 20), while the size of the listing (factorized) representation of natural join grows cubically (linearly) with the scale factor. The two listing encodings blow up the memory consumption and computation time for large scales. Storing tuples in the listing representation using payloads instead of keys avoids the need for hashing wide keys, which makes the joins slightly cheaper. For *Housing* and factorized representation, the root view stores 25,000 values of the join variable regardless of the scale. The root’s children map these values to relational payloads for each relation. For the largest scale, Fact payloads is 481x faster and takes 548x less memory than List payloads (410ms vs. 197s, 195MB vs. 104GB), and List keys exceeds the available memory.

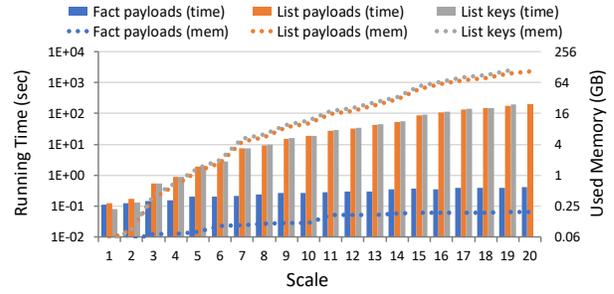
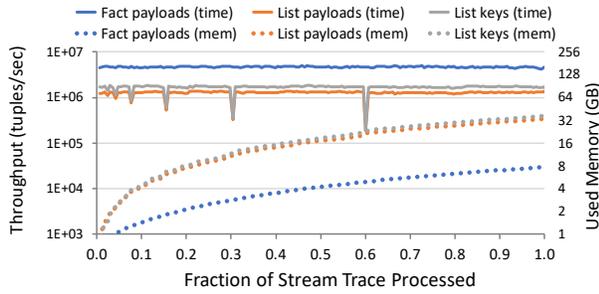


Figure 8: Incremental maintenance using relational and factorized payloads for the natural joins of the *Retailer* (left) and of the *Housing* (right) datasets under updates of size 1,000 to the largest relation (*Retailer*) and all input relations (*Housing*).

8 RELATED WORK

To the best of our knowledge, ours is the first approach to propose factorized IVM for a range of distinct applications. It extends non-trivially two lines of prior work: higher-order delta-based IVM and factorized computation of in-database analytics.

Our view language is modeled after functional aggregate queries over semirings [5] and generalized multiset relations over rings [25]; the latter allowed us to adapt DBToaster to factorized IVM.

IVM. IVM is a well-studied area spanning more than three decades [12]. Prior work extensively studied IVM for various query languages and showed that the time complexity of IVM is lower than of re-computation. We go beyond prior work on higher-order IVM for queries with joins and aggregates, as realized in DBToaster [25], and propose a unified approach for factorized computation of aggregates over joins [7], factorized incremental computation of linear algebra [33], and learning regression models over factorized joins [42]. DBToaster uses one materialization hierarchy per relation in the query, whereas we use one view tree for all relations. DBToaster can thus have much larger space requirements and update times. DBToaster does not target the maintenance of many complex aggregates that share computation (e.g., cofactor matrices), which we observe experimentally. IVM over array data [47] targets scientific workloads but without exploiting data factorization.

Our approach with the relational payload ring strictly subsumes previous work on factorized IVM for acyclic joins [21] as it can support cyclic joins (see Appendix B). The so-called q -hierarchical join queries (such as the Housing query in our experiments) are exactly those self-join-free conjunctive queries that admit constant time update [8]. Recent work on in-database maintenance of linear regression models shows how to compute such models using previously computed models over distinct sets of features [19]. Its contribution is complementary to ours and shares a similar goal with prior work on reusing gradient computation to efficiently explore the space of possible regression models [36]. Exploiting key attributes to enable succinct delta representations and accelerate maintenance complements our approach [22]. Our framework generalizes the main idea of the LINVIEW approach [33] to maintenance of matrix computation over arbitrary join queries.

Most commercial databases, e.g., Oracle [3] and SQLServer [1], support IVM for restricted classes of queries. LogicBlox supports higher-order IVM for Datalog (meta)programs [6, 18]. Trill is a streaming engine that supports incremental processing of relational-style queries but no complex aggregates like cofactor matrices [10].

Static In-DB analytics. The emerging area of in-database analytics has been recently overviewed in two tutorials [27, 39]. Several systems support complex analytics over normalized data via a tight integration of databases and machine learning [20, 27, 30, 39, 40]. Others integrate with R to enable in-situ data processing using domain-specialized routines [9, 46]. The closest in spirit to our approach is work on learning models over factorized joins [23, 36, 41, 42], pushing ML tasks past joins [16] and on in-database linear algebra [11, 14], yet they do not consider incremental maintenance.

Learning. There is a wealth of work in the ML community on incremental or online learning over *arbitrary* relations [43]. Our approach learns over *joins* and crucially exploits the join dependencies in the underlying training dataset to improve the performance.

9 CONCLUSION

This paper introduces a unified IVM approach for analytics over normalized data and shows its applicability for three seemingly disparate analytical asks: matrix chain multiplication, query evaluation with listing/factorized result representation, and gradient computation used for learning linear regression models. These tasks use the same computation paradigm that factorizes the representation and the computation of the keys, the payloads, and the updates. Their differences are factored out in the definition of the sum and product operations in a suitable ring. This approach has been implemented as an extension of DBToaster, a state-of-the-art system for incremental maintenance, and shown to outperform competitors by up to two orders of magnitude in both time and space. An extended version of this paper discusses the computational complexities of this approach and its competitors and explains analytically the reason for the performance gap [34].

Going forward, we would like to apply this approach to further tasks such as inference in probabilistic graphical models and more complex machine learning tasks.

Acknowledgments. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 682588. The authors also acknowledge awards from Microsoft Azure via The Alan Turing Institute and from Fondation Wiener Anspach.

If You Liked It, Then You Should Put A Ring On It.
– Beyoncé.

REFERENCES

- [1] Create Indexed Views. <http://msdn.microsoft.com/en-us/library/ms191432.aspx>.
- [2] Higgs Twitter Dataset. <https://snap.stanford.edu/data/higgs-twitter.html>.
- [3] Materialized View Concepts and Architecture. http://docs.oracle.com/cd/B28359_01/server.111/b28326/repview.htm.
- [4] D. J. Abadi, Y. Ahmad, M. Balazinska, et al. The Design of the Borealis Stream Processing Engine. In *CIDR*, volume 5, pages 277–289, 2005.
- [5] M. Abo Khamis, H. Q. Ngo, and A. Rudra. FAQ: Questions Asked Frequently. In *PODS*, pages 13–28, 2016.
- [6] M. Aref, B. ten Cate, T. J. Green, B. Kimelfeld, et al. Design and Implementation of the LogicBlox System. In *SIGMOD*, pages 1371–1382, 2015.
- [7] N. Bakibayev, T. Kociský, D. Olteanu, and J. Závodný. Aggregation and Ordering in Factorised Databases. *PVLDB*, 6(14):1990–2001, 2013.
- [8] C. Berkholz, J. Keppeler, and N. Schweikardt. Answering Conjunctive Queries under Updates. In *PODS*, pages 303–318, 2017.
- [9] P. G. Brown. Overview of SciDB: Large Scale Array Storage, Processing and Analysis. In *SIGMOD*, pages 963–968, 2010.
- [10] B. Chandramouli, J. Goldstein, et al. Trill: A high-performance Incremental Query Processor for Diverse Analytics. *PVLDB*, 8(4):401–412, 2014.
- [11] L. Chen, A. Kumar, J. F. Naughton, and J. M. Patel. Towards linear algebra over normalized data. *CoRR*, abs/1612.07448, 2016. To appear in *PVLDB* 2017.
- [12] R. Chirkova and J. Yang. Materialized Views. *Found. & Trends in DB*, 4(4):295–405, 2012.
- [13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 2009.
- [14] A. Elgohary, M. Boehm, P. J. Haas, F. R. Reiss, and B. Reinwald. Compressed linear algebra for large-scale machine learning. *PVLDB*, 9(12):960–971, 2016.
- [15] R. Fagin, A. O. Mendelzon, and J. D. Ullman. A simplified universal relation assumption and its properties. *ACM Trans. Database Syst.*, 7(3):343–360, 1982.
- [16] X. Feng, A. Kumar, B. Recht, and C. Ré. Towards a Unified Architecture for In-RDBMS Analytics. In *SIGMOD*, pages 325–336, 2012.
- [17] T. J. Green, G. Karvounarakis, and V. Tannen. Provenance Semirings. In *PODS*, pages 31–40, 2007.
- [18] T. J. Green, D. Olteanu, and G. Washburn. Live programming in the LogicBlox system: A MetaLogiQL approach. *PVLDB*, 8(12):1782–1791, 2015.
- [19] P. Gupta, N. Koudas, E. Shang, R. Johnson, and C. Zuzarte. Processing Analytical Workloads Incrementally. *CoRR*, abs/1509.05066, 2015.
- [20] J. M. Hellerstein, C. Ré, F. Schoppmann, D. Z. Wang, E. Fratkin, et al. The MADlib Analytics Library or MAD Skills, the SQL. *PVLDB*, 5(12):1700–1711, 2012.
- [21] M. Idris, M. Ugarte, and S. Vansummeren. The Dynamic Yannakakis Algorithm: Compact and Efficient Query Processing Under Updates. In *SIGMOD*, pages 1259–1274, 2017.
- [22] Y. Katsis, K. W. Ong, Y. Papakonstantinou, and K. K. Zhao. Utilizing IDs to Accelerate Incremental View Maintenance. In *SIGMOD*, pages 1985–2000, 2015.
- [23] M. A. Khamis, H. Ngo, X. Nguyen, D. Olteanu, and M. Schleich. In-Database Learning with Sparse Tensors. In *PODS*, 2018.
- [24] C. Koch. Incremental Query Evaluation in a Ring of Databases. In *PODS*, pages 87–98, 2010.
- [25] C. Koch, Y. Ahmad, O. Kennedy, M. Nikolic, A. Nötzli, D. Lupei, et al. DBtoaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views. *Vldb J.*, 23(2):253–278, 2014.
- [26] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, 2009.
- [27] A. Kumar, M. Boehm, and J. Yang. Data management in machine learning: Challenges, techniques, and systems. In *SIGMOD*, pages 1717–1722, 2017.
- [28] A. Kumar, J. F. Naughton, and J. M. Patel. Learning Generalized Linear Models Over Normalized Data. In *SIGMOD*, pages 1969–1984, 2015.
- [29] S. R. Madden et al. TinyDB: An Acquisitional Query Processing System for Sensor Networks. *TODS*, 30(1):122–173, 2005.
- [30] X. Meng et al. MLlib: Machine Learning in Apache Spark. *J. Mach. Learn. Res.*, 17(1):1235–1241, 2016.
- [31] H. Q. Ngo, C. Ré, and A. Rudra. Skew Strikes Back: New Developments in the Theory of Join Algorithms. *SIGMOD Record*, 42(4):5–16, 2013.
- [32] M. Nikolic, M. Dashti, and C. Koch. How to win a hot dog eating contest: Distributed incremental view maintenance with batch updates. In *SIGMOD*, pages 511–526, 2016.
- [33] M. Nikolic, M. Elseidy, and C. Koch. LINVIEW: Incremental View Maintenance for Complex Analytical Queries. In *SIGMOD*, pages 253–264, 2014.
- [34] M. Nikolic and D. Olteanu. Incremental view maintenance with triple lock factorization benefits. *CoRR*, abs/1703.07484, 2017.
- [35] D. Olteanu, C. Koch, and L. Antova. World-set decompositions: Expressiveness and efficient algorithms. *Theor. Comput. Sci.*, 403(2-3):265–284, 2008.
- [36] D. Olteanu and M. Schleich. F: Regression Models over Factorized Views. *PVLDB*, 9(13):1573–1576, 2016.
- [37] D. Olteanu and M. Schleich. Factorized Databases. *SIGMOD Rec.*, 45(2):5–16, 2016.
- [38] D. Olteanu and J. Závodný. Size Bounds for Factorised Representations of Query Results. *TODS*, 40(1):2:1–2:44, 2015.
- [39] N. Polyzotis, S. Roy, S. E. Whang, and M. Zinkevich. Data management challenges in production machine learning. In *SIGMOD*, pages 1723–1726, 2017.
- [40] C. Qin and F. Rusu. Speculative Approximations for Terascale Distributed Gradient Descent Optimization. In *DanaC*, pages 1–10, 2015.
- [41] S. Rendle. Scaling Factorization Machines to Relational Data. *PVLDB*, 6(5):337–348, 2013.
- [42] M. Schleich, D. Olteanu, and R. Ciucanu. Learning Linear Regression Models over Factorized Joins. In *SIGMOD*, pages 3–18, 2016.
- [43] S. Shalev-Shwartz et al. Online Learning and Online Convex Optimization. *Found. & Trends in ML*, 4(2):107–194, 2012.
- [44] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos. Tensor decomposition for signal processing and machine learning. *Trans. Sig. Proc.*, 65(13):3551–3582, 2017.
- [45] C. Whaley and J. Dongarra. Automatically tuned linear algebra software. In *PPSC*, 1999.
- [46] M. Zaharia et al. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *NSDI*, pages 15–28, 2012.
- [47] W. Zhao, F. Rusu, B. Dong, K. Wu, and P. Nugent. Incremental View Maintenance over Array Data. In *SIGMOD*, pages 139–154, 2017.

A RINGS

We recall the definition of rings.

Definition A.1. A ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$ is a set \mathbf{D} with closed binary operations $+$ and $*$, the additive identity $\mathbf{0}$, and the multiplicative identity $\mathbf{1}$ satisfying the axioms $(\forall a, b, c \in \mathbf{D})$:

- (1) $a + b = b + a$.
- (2) $(a + b) + c = a + (b + c)$.
- (3) $\mathbf{0} + a = a + \mathbf{0} = a$.
- (4) $\exists -a \in \mathbf{D} : a + (-a) = (-a) + a = \mathbf{0}$.
- (5) $(a * b) * c = a * (b * c)$.
- (6) $a * \mathbf{1} = \mathbf{1} * a = a$.
- (7) $a * (b + c) = a * b + a * c$ and $(a + b) * c = a * c + b * c$.

A semiring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$ satisfies all of the above properties except the additive inverse property and adds the axiom that $\mathbf{0} * a = a * \mathbf{0} = \mathbf{0}$. A (semi)ring for which $a * b = b * a$ is commutative.

Example A.2. The number sets \mathbb{Z} , \mathbb{Q} , \mathbb{R} , and \mathbb{C} with arithmetic operations $+$ and \cdot and numbers 0 and 1 form commutative rings. The set \mathcal{M} of $(n \times n)$ matrices forms a non-commutative ring $(\mathcal{M}, +, \cdot, 0_{n,n}, I_n)$, where $0_{n,n}$ and I_n are the zero matrix and the identity matrix of size $(n \times n)$. The set \mathbb{N} of natural numbers is a commutative semiring but not a ring because it has no additive inverse. Further examples are the max-product semiring $(\mathbb{R}_+, \max, \times, 0, 1)$, the Boolean semiring $(\{\text{true}, \text{false}\}, \vee, \wedge, \text{false}, \text{true})$, and the set semiring $(2^U, \cup, \cap, \emptyset, U)$ of all possible subsets of a given set U . \square

B IVM VARIANT FOR CYCLIC QUERIES

Our framework supports arbitrary conjunctive queries. Whereas for (α) acyclic join queries the size of each view is asymptotically upper bounded by the size of the factorized join, for a cyclic join views may be larger in size than the (factorized) join result. This increase in space may however enable faster incremental view maintenance.

Example B.1. We consider the triangle query over the ring \mathbb{Z} :

$$Q_{\Delta}[] = \bigoplus_A \bigoplus_B \bigoplus_C R[A, B] \otimes S[B, C] \otimes T[C, A]$$

Figure 9 shows the hypergraph of Q_{Δ} and the view tree constructed for the variable order $A - B - C$ by placing each relation directly under its lowest variable. We assume all relations are of size $O(N)$.

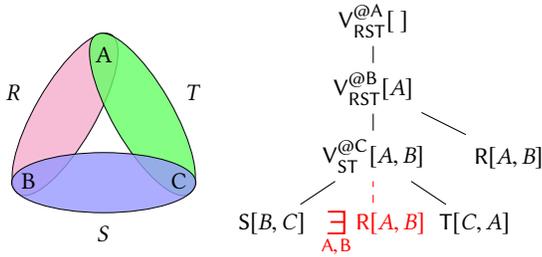


Figure 9: (left) Hypergraph of the triangle query Q_Δ ; (right) View tree for the variable order $\{A-B-C\}$ with an indicator projection $\exists_{A,B}R$.

Computing the triangle query from scratch using a worst-case optimal join algorithm takes $O(N^{3/2})$ time [31].

In the given view tree (without the view in red), we first join S and T and then marginalize out C in the join result. This view at node C may contain $O(N^2)$ pairs of (A, B) values, which is larger than the worst-case size $O(N^{3/2})$. However, by materializing the view at C , we enable single-tuple updates to R in constant time; single-tuple updates to other relations take $O(N)$ time.

To avoid this large intermediate result, we can change the view tree by placing the relation R under variable C . Then, joining all three relations at node C takes $O(N^{3/2})$ time. Updates to any relation now cause re-computation of a 3-way join, like in first-order IVM. For single-tuple updates, recomputing deltas takes $O(N)$ as only two of the three variables are bound to constants. In contrast, the first approach trades off space for time: We need $O(N^2)$ space but then support $O(1)$ updates to one of the three relations. \square

The above example demonstrates how placing a relation under a different node in a view tree can create a cycle of relations and constrain the size of a view. This strategy, however, might not be always feasible or efficient: One relation might form multiple cycles of relations in different parts of a view tree – for example, in a loop-4 cyclic query Q_{\square} with a chord, the chord relation is part of two triangle subqueries. Since this relation cannot be duplicated in multiple subtrees (for correctness reasons so as to avoid multiplying the same payload several times instead of using it once), one would have to evaluate these subqueries in sequence, which yields a view tree that is higher and more expensive to maintain.

Indicator Projections. Instead of moving relations in a view tree, we extend the tree with indicator projections that identify the active domains of these relations [5]. Such projections have no effect on the query result but can constrain view definitions (e.g., create cycles) and bring asymptotic savings in both space and time.

We define a new unary operation $\exists_{\mathcal{A}}R$ that, given a relation R over schema \mathcal{S} with payloads from a ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$, and a set of attributes $\mathcal{A} \subseteq \mathcal{S}$, projects tuples from R with non- $\mathbf{0}$ payload on \mathcal{A} and assigns to these tuples the payload $\mathbf{1}$. We define $\exists_{\mathcal{A}}R$ as:

$$\forall \mathbf{t} \in \text{Dom}(\mathcal{A}) : (\exists_{\mathcal{A}}R)[\mathbf{t}] = \begin{cases} \mathbf{1} & \exists \mathbf{s} \in \text{Dom}(\mathcal{S}), \mathbf{s} \in R, \mathbf{t} = \pi_{\mathcal{A}}(\mathbf{s}) \\ \mathbf{0} & \text{otherwise} \end{cases}$$

Indicator projections may change with updates to input relations. For instance, adding a tuple with a unique \mathcal{A} -value to R enlarges the result of $\exists_{\mathcal{A}}R$; similarly, deleting the last tuple with the given

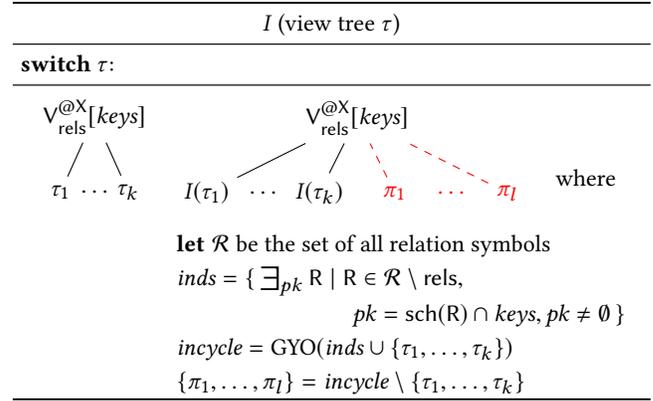


Figure 10: Adding indicator projections to a view tree τ .

\mathcal{A} -value reduces the result. Notice that one change in the input may cause at most one change in the output, that is, $|\delta(\exists_{\mathcal{A}}R)| \leq |\delta R|$.

To facilitate the computation of $\delta(\exists_{\mathcal{A}}R)$, we keep track of how many tuples with non- $\mathbf{0}$ payloads project on each \mathcal{A} -value. For updating the payload of a tuple in R from $\mathbf{0}$ to non- $\mathbf{0}$ (or vice versa), we increase (decrease) the count corresponding to the given \mathcal{A} -value. If this count changes from 0 to 1 (meaning the \mathcal{A} -value is unique) or from 1 to 0 (meaning there are no more tuples with the \mathcal{A} -value), then $\delta(\exists_{\mathcal{A}}R)$ contains a tuple of \mathcal{A} -values with the payload of 1 or -1 , respectively; otherwise, the delta is empty.

Example B.2. Consider a relation R over the schema $\{A, B\}$ with payloads from a ring $(\mathbf{D}, +, *, \mathbf{0}, \mathbf{1})$. We want to maintain the result of the query $Q[A] = \exists_A R[A, B]$. To compute $\delta Q[A]$ for updates to R efficiently, we count tuples from R with non- $\mathbf{0}$ payloads for each A -value, denoted by $\text{CNT}_Q[A]$. For example:

| R | A | B | CNT _Q | A | Q | A |
|---|-------|-------|------------------|-------|-----------------|-------|
| | a_1 | b_1 | | a_1 | $\rightarrow 2$ | a_1 |
| | a_1 | b_2 | | a_2 | $\rightarrow 1$ | a_2 |
| | a_2 | b_3 | | | | |

where r_1, r_2 , and r_3 are non- $\mathbf{0}$ payloads from \mathbf{D} . An update $\delta R = \{(a_1, b_2) \rightarrow -r_2\}$ removes the tuple (a_1, b_2) from R , which in turn decreases $\text{CNT}_Q[a_1]$ by 1. Since there is still a tuple in R that projects on a_1 , the result of Q remains unchanged. A subsequent update $\{(a_1, b_1) \rightarrow -r_1\}$ from R drops the count for a_1 to 0, which triggers a change in the output, $\delta Q = \{(a_1) \rightarrow -1\}$. \square

View Trees with Indicator Projections. Figure 10 gives an algorithm that traverses a given view tree bottom-up and extends each view definition with indicator projections. At each view, the algorithm computes a set of relations inds that could be used as indicator projections, restricting to only those relations that share common variables with that view and that do not appear in its definition. From this set of candidates, only those relations that form a cycle with the children of the given view are used as indicator projections. The algorithm uses the GYO reduction (Fagin et al variant) [15] to determine this set of relations, denoted by incycle , and extends the view definition with the indicator projections of the candidate relations from this set.

In a view tree with indicator projections, changes in one relation may propagate along multiple leaf-to-root paths. We propagate

them in sequence, i.e., updates to one relation are followed by a sequence of updates to its indicator projections.

Example B.3. The algorithm from Figure 10 extends the view tree of the triangle query with an indicator projection $\exists_{A,B} R[A, B]$ placed below the view $V_{ST}^{\text{@C}}$. This view at C is now a cyclic join of the three relations, which can be computed in $O(N^{3/2})$ time. The indicator projection also reduces the size of this view to $O(N)$.

Single-tuple updates to S and T still take linear time; however, bulk updates of size $O(N)$ can now be processed in $O(N^{3/2})$ time, same as re-evaluation. Updates to R might affect the indicator projection: If a single-tuple update δR causes no change in the projection, then incremental maintenance takes constant time; otherwise, joining a tuple $\delta(\exists_{A,B} R)$ with S and T at node C takes linear time. Bulk updates δR of size $O(N)$ can also be processed in $O(N^{3/2})$ time. Thus, using indicator projections in this query takes the best of both approaches from Example B.1, namely faster incremental maintenance and more succinct view representation. \square

C EXPERIMENTS

C.1 Experimental Setup

Workload. We run experiments over three datasets:

- *Retailer* is a real-world dataset from an industrial collaborator and used by a retailer for business decision support and forecasting user demands. The dataset has a snowflake schema with one fact relation *Inventory* with 84M records, storing information about the inventory units for products in a location, at a given date. The *Inventory* relation joins along three dimension hierarchies: *Item* (on product id), *Weather* (on location and date), and *Location* (on location) with its lookup relation *Census* (on zip code). The natural join of these five relations is acyclic and has 43 attributes. We consider a view tree in which the variables of each relation form a distinct root-to-leaf path, and the partial order on join variables is: location - { date - { product id }, zip }.
- *Housing* is a synthetic dataset modeling a house price market [42]. It consists of six relations: *House*, *Shop*, *Institution*, *Restaurant*, *Demographics*, and *Transport*, arranged into a star schema and with 1.4M tuples in total (scale factor 20). The natural join of all relations is on the common attribute (postcode) and has 27 attributes. We consider an optimal view tree that has each root-to-leaf path consisting of query variables for one relation.
- *Higgs Twitter* represents friends/followers relationships among users who were active on Twitter during the discovery of Higgs boson [2]. We split the first 3M records from the dataset into three equally-sized relations, $R(A, B)$, $S(B, C)$, and $T(C, A)$, and consider the triangle query over them and the variable order A-B-C.

We run the systems over data streams synthesized from the above datasets by interleaving insertions to the input relations in a round-robin fashion. We group insertions into batches of different sizes and place no restriction on the order of records in input relations. In all experiments, we use payloads defined over rings with additive inverse, thus processing deletions is similar to that of insertions.

Queries. We next detail the queries used in the experiments.

Matrix Chain Multiplication: The query in standard SQL is over relations $A_1(X, Y, P_1)$, $A_2(Y, Z, P_2)$, $A_3(Z, U, P_3)$:

| | F-IVM | DBT | 1-IVM | F-RE | DBT-RE |
|----------|--------------|--------------|-------------|---------|---------|
| Retailer | 2, 955, 045 | 1, 250, 262 | 2, 925, 828 | 3, 785* | 3, 491* |
| Housing | 22, 857, 143 | 17, 834, 395 | 2, 403, 433 | 79, 226 | 364* |

Figure 11: The average throughput (tuples/sec) of re-evaluation and incremental maintenance of a sum aggregate under updates of size 1, 000 to all relations of the *Retailer* and *Housing* datasets with a one-hour timeout (*).

```
SELECT A1.X, A3.U, SUM(A1.P1 * A2.P2 * A3.P3)
FROM A1 NATURAL JOIN A2 NATURAL JOIN A3
GROUP BY A1.X, A3.U;
```

In our formalism, each relation maps pairs of indices to matrix values, all lifting functions map values to 1, and the query is: $Q[X, U] = \bigoplus_Y \bigoplus_Z A_1[X, Y] \otimes A_2[Y, Z] \otimes A_3[Z, U]$.

Cofactor Matrix Computation: For the *Retailer* schema, the query has one regression aggregate over the natural join of its relations:

```
SELECT SUM(g1(X1) * ... * g43(X43))
FROM Inv NATURAL JOIN It NATURAL JOIN W
NATURAL JOIN L NATURAL JOIN C;
```

where $\{X_i\}_{i \in [43]}$ are all the variables from the *Retailer* schema, the SUM operator uses + and * from the degree-43 matrix ring, and each function g_i maps a value x to $g_i(x) = (c_i = 1, s_i = x, Q_{(i,i)} = x^2)$ (see Example 6.3). Similarly, the queries over *Housing* and *Twitter* use the degree-27 and respectively degree-3 matrix ring. The experiments consider all variables to be continuous; categorical variables can be treated using group-by queries as explained in related work [23].

Factorized Computation of Conjunctive Queries: We consider the natural join queries of all relations in the *Retailer* and respectively *Housing* databases.

Experimental setup. We run all experiments on a Microsoft Azure DS14 instance, Intel(R) Xeon(R) CPU E5-2673 v3 @ 2.40GHz, 112GB RAM, with Ubuntu Server 14.04. We use DBToaster v2.2 for the IVM competitors and code generation in our approach. The generated C++ code is single-threaded and compiled using g++ 6.3.0 with the -O3 flag. We set a one-hour timeout on query execution and report wall-clock times by averaging three best results out of four runs. We profile memory utilization using gperftools, not counting the memory used for storing input streams.

C.2 Further Experiments

Maintenance of sum aggregates. We analyze different strategies for maintaining a sum of one variable on top of a natural join. We measure the average throughput of re-evaluation and incremental maintenance under updates of size 1, 000 to all the relations of *Retailer* and *Housing*. For the former dataset, we sum the inventory units for products in *Inventory*; for the latter, we sum over the common join variable. We also benchmark two re-evaluation strategies that recompute the results from scratch on every update: F-RE denotes reevaluation using variable orders and DBT-RE denotes re-evaluation using DBToaster. Table 11 summarizes the results.

F-IVM achieves the highest average throughput in both cases. For *Retailer*, the maintenance cost is dominated by the update on *Inventory*. DBT’s recursive delta compilation materializes 13 views

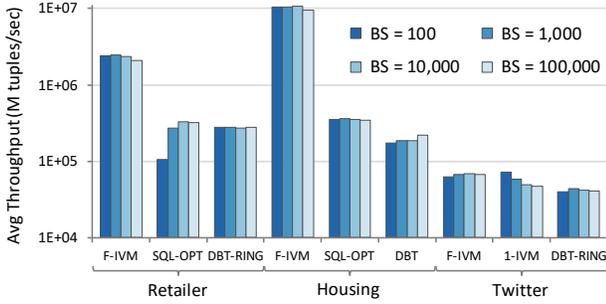


Figure 12: Incremental maintenance of the cofactor matrix under batch updates of different sizes to all input relations.

representing connected sub-queries: five group-by aggregates over the input relations, *Inv*, *It*, *W*, *L*, and *C*; one group-by aggregate joining *L* and *C*; six views joining *Inv* with subsets of the others, namely $\{It\}$, $\{It, W\}$, $\{It, W, L\}$, $\{W\}$, $\{W, L\}$, and $\{W, L, C\}$; and the final aggregate. The two views joining *Inv* with $\{W, L\}$ and $\{It, W, L\}$ require linear maintenance for a single-tuple change in *Inventory*. 1-IVM recomputes deltas from scratch on each update using only the input relations with no aggregates on top of them. Updates to *Inventory* are efficient due to small sizes of the other relations. F-IVM uses the given variable order to materialize 9 views, four of them over *Inventory*, $\{Inv\}$, $\{Inv, It\}$, $\{Inv, It, W\}$, and the final sum, but each with constant maintenance for single-tuple updates to this relation. In contrast to 1-IVM, our approach materializes pre-computed views in which all non-join variables are aggregated away. In the *Housing* schema, both F-IVM and DBT benefit from this pre-aggregation, and since the query is a star join, both materialize the same views. DBT computes $SUM(1)$ and $SUM(postcode)$ for each postcode in the delta for *Inventory*, although only the count suffices. Figure 11 also shows that the re-evaluation strategies significantly underperform the incremental approaches.

The effect of batch size on IVM. This experiment evaluates the performance of maintaining a cofactor matrix for batch updates of different sizes. Figure 12 shows the throughput of batched incremental processing for batch sizes varying from 100 to 100,000 on the *Retailer*, *Housing*, and *Twitter* datasets for updates to all relations. We show only the best three approaches for each dataset.

We observe that using very large or small batch sizes can have negative performance effects: Iterating over large batches invalidates previously cached data resulting in future cache misses, whereas using small batches cannot offset the overhead associated with processing each batch. Using batches with 1,000 – 10,000 tuples delivers best performance in most cases, except when needed to incrementally maintain a large number of views. This conclusion about cofactor matrix computation is in line with similar findings on batched delta processing in decision support workloads [32].

Batched incremental processing is also beneficial for one-off computation of the entire cofactor matrix. Using medium-sized updates can bring better performance, cf. Figure 12, but can also lower memory requirements and improve cache locality during query processing. For instance, incrementally processing the *Retailer* dataset in chunks of 1,000 tuples can bring up to 2.45x better performance compared to processing the entire dataset at once.

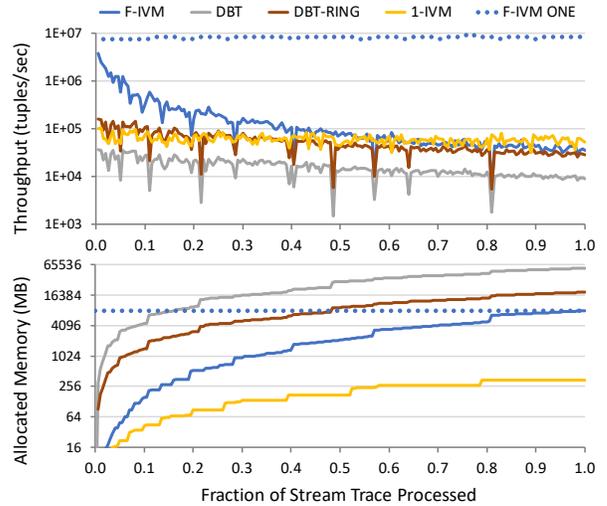


Figure 13: The performance of incremental maintenance of the cofactor matrix on top of the triangle query over the *Twitter* dataset for updates of size 1,000 to all input relations.

Cofactor matrix computation over the triangle query. We analyze the cofactor matrix computation over the triangle query on the *Higgs Twitter* dataset and updates of size 1,000 to all the relations. F-IVM uses the view tree from Figure 9 (right) without the indicator projection and materializes the join of *S* and *T* of size $O(N^2)$. Its time complexity for a single-tuple update to *R* is $O(1)$, but updating the join of *S* and *T* takes $O(N)$. DBT-RING uses pairwise joins, each requiring linear time maintenance. DBT uses scalar payloads and materializes 21 views (to maintain 6 aggregates), out of which 12 views are over two relations. Its time complexity for processing single-tuple updates to either of the three relations is also $O(N)$. The 1-IVM strategy maintains just the input relations and recomputes the delta upon each update in linear time.

The throughput rate of the strategies that materialize views of quadratic size declines sharply as the input stream progresses. DBT exhibits the highest processing and memory overheads caused by storing 12 auxiliary views of quadratic size. DBT-RING underperforms F-IVM due to maintaining two extra views of quadratic size, which contribute to 2.3x higher peak memory utilization. 1-IVM exhibits a 42% decline in performance after processing the entire trace due to its linear time maintenance. The extent of this decrease is much lower compared to the other approaches with the quadratic space complexity. For updates to *R* only, F-IVM-ONE and DBT-RING require one lookup in the materialized join of *S* and *T* per update. This strategy has two orders of magnitude higher throughput than 1-IVM at the cost of using 23x more memory.

Clique queries like triangles provide no factorization opportunities. Materializing auxiliary views to speed up incremental view maintenance increases memory and processing overheads. However, F-IVM can exploit indicator projections to bound the size of such materialized views, as described in Section B.