

Enzyme: Incremental View Maintenance for Data Engineering

Ritwik Yadav*
Supun Abeysinghe*
Min Yang*
Jeffrey Helt*
Manuel Ung*
Yuhong Chen*
Databricks, Inc.
San Francisco, CA, USA

Melody Hu
William Wei
Yiming Yang
Tom van Bussel
Sourav Chatterji
Indrajit Roy
Databricks, Inc.
San Francisco, CA, USA

Paul Lappas†
Yannis
Papakonstantinou†
Google, Inc.
Mountain View, CA, USA

Tahir Fayyaz
Bilal Aslam
Ross Bunker
Michael Armbrust
Shrikanth Shankar*
Databricks, Inc.
San Francisco, CA, USA
enzyme-
paper@databricks.com

Abstract

Materialized views are a core construct in database systems, used to accelerate analytical queries and optimize batch pipelines for extract–transform–load (ETL) workflows. Maintaining view consistency as underlying data evolves is a fundamental challenge, especially in high-throughput and real-time settings. Incremental view maintenance (IVM) has been studied for decades and continues to attract significant investment from major database vendors. However, most industrial systems either offer limited SQL-operator coverage or require users to hand-tune refresh strategies.

This paper presents Enzyme, an IVM engine built at Databricks to power its declarative data pipelines platform. It provides a built-in, end-to-end approach to incremental pipelines, utilizing materialized views as first-class building blocks. By automating refresh planning, Enzyme reduces total cost of ownership and allows users to focus on business logic rather than MV mechanics. Validation across thousands of large-scale production pipelines spanning diverse application domains has demonstrated substantial computational efficiency gains, yielding a cumulative daily compute reduction of billions of CPU seconds. Built atop Apache Spark primitives, Enzyme adds a cost-based optimization layer that selects refresh strategies for collections of materialized views organized into pipelines.

Enzyme’s modular architecture is designed to generalize across data sources and query engines. We present key design decisions for incremental refresh planning and execution, including optimizations that exploit batching opportunities across materialized view sources. Experimental results on standard benchmarks demonstrate significant performance improvements at scale.

CCS Concepts

• **Information systems** → **Online analytical processing engines.**

*Coauthored this paper.

†Work done while at Databricks.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

SIGMOD Companion '26, Bengaluru, India

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2450-3/2026/05

<https://doi.org/10.1145/3788853.3803098>

Keywords

Incremental view maintenance; materialized views; ETL

ACM Reference Format:

Ritwik Yadav, Supun Abeysinghe, Min Yang, Jeffrey Helt, Manuel Ung, Yuhong Chen, Melody Hu, William Wei, Yiming Yang, Tom van Bussel, Sourav Chatterji, Indrajit Roy, Paul Lappas, Yannis Papakonstantinou, Tahir Fayyaz, Bilal Aslam, Ross Bunker, Michael Armbrust, and Shrikanth Shankar. 2026. Enzyme: Incremental View Maintenance for Data Engineering. In *Companion of the International Conference on Management of Data (SIGMOD Companion '26)*, May 31–June 05, 2026, Bengaluru, India. ACM, Bengaluru, India, 14 pages. <https://doi.org/10.1145/3788853.3803098>

1 Introduction

Materialized views (MVs) are a fundamental building block offered by most database management systems [8, 10, 21, 35, 45]. They are used primarily for reducing query execution costs by amortizing expensive, shared computation across multiple queries as well as improving the latency of individual queries. MVs have evolved from their traditional role as query acceleration constructs to become fundamental building blocks in modern data flow architectures. Along with data streams and primitives designed to process changes to datasets, MVs enable efficient extraction, transformation, and loading (ETL) of data at scale.

Modern data architectures, such as the medallion architecture (Figure 1) commonly used in data lakehouses, leverage this capability to organize and transform data across multiple layers. The quality of the data improves as it flows through subsequent layers in this architecture. The bronze layer lands the data as-is along with additional metadata. The silver layer matches, merges, conforms and cleans the data in the bronze layer to enable self-service analytics for ad-hoc reporting, advanced analytics, and machine learning. Finally, the gold layer performs transformations for reporting, utilizing more denormalized and read-optimized data models with fewer joins. MVs primarily feature in the silver and gold layers to apply these transformations.

Efficiently maintaining view consistency as source data evolves has been extensively studied in database theory over the last few decades [26]. However, with the advent of the AI era, the scale of data at organizations has grown exponentially, making incremental refresh of MVs a necessity rather than treating it as an optional optimization. This shift has exposed several critical gaps in existing IVM systems. Modern streaming engines offer IVM as a stream processing construct, updating MVs with every change to

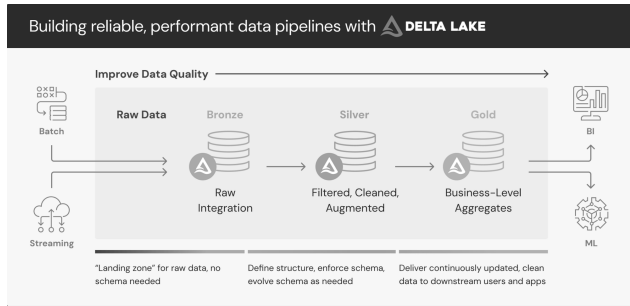


Figure 1: Medallion architecture organizing data into progressively refined layers for analytics.

their sources. While this approach minimizes data latency, batching changes can amortize infrastructure costs, reduce scheduling overhead, and improve resource utilization across large-scale deployments. Additionally, most prior IVM research relies on efficient secondary indexes for incremental query plans to significantly outperform full recomputation. Modern lakehouses, however, lack traditional secondary indexes and instead rely on metadata-based indexing at the file and column level. Research to reduce reliance on secondary indexes [30] is relatively recent. Furthermore, comprehensive coverage for arbitrary compositions of common SQL operators—such as aggregations and window functions—remains rare in industrial-grade IVM engines.

Building a production-ready IVM engine extends far beyond constructing efficient incremental query plans. It requires complex supporting infrastructure, including:

- A query fingerprinter capable of detecting changes to MV definitions across multiple languages (SQL, Python) and in the presence of user-defined functions. The diversity in languages makes it a particularly difficult problem and inaccuracies in this component can lead to incorrect results (false negative) or full recomputes on petabyte scale datasets (false positive).
- A cost model that estimates computational complexity of logically equivalent plans, incorporates feedback from historical refreshes, and optimizes across a DAG of MVs while providing explainable output.
- Specialized query mutators to support incrementalization of operators such as `current_timestamp()` and AI functions. These operators are not addressed in traditional IVM literature but are extremely popular among data engineers and practitioners building ETL pipelines.
- Comprehensive correctness testing frameworks capable of non-deterministically generating test cases spanning the entire domain of valid MV definitions. The level of test coverage obtained using these frameworks cannot be matched by traditional unit or integration tests.

This paper introduces Enzyme, an out-of-the-box IVM engine developed at Databricks that addresses these challenges through automated, cost-based optimization of MV refreshes. Enzyme provides comprehensive operator coverage, supports both scheduled batch and continuous refresh patterns, and automatically selects between

```
CREATE MATERIALIZED VIEW region_avg_sales
SCHEDULE EVERY 1 HOUR
AS
SELECT
  c.region,
  AVG(o.amount) AS avg_order_amount
FROM
  Customers c
  JOIN Orders o ON c.customer_id = o.customer_id
GROUP BY
  c.region
HAVING
  c.region IN ('us-east', 'us-west', 'asia');
```

Figure 2: A standalone MV that computes the average order amount for three selected regions. The MV is configured to refresh every hour. We use this as a running example throughout the paper.

incremental and full refresh strategies based on workload characteristics and historical performance. We begin by enumerating the key primitives and features that Enzyme relies upon (Section 2), followed by a formalization of the incremental view maintenance problem (Section 3). Section 4 presents Enzyme’s design and architecture, while Section 5 discusses lessons learned from its implementation. We validate Enzyme on both production workloads and standard ETL benchmarks, with quantitative performance results presented in Section 6. Section 7 compares Enzyme with related systems before we conclude in Section 8.

2 Background

In this section, we discuss how MVs fit into the Databricks platform. We then introduce the necessary concepts and features for understanding the Enzyme incremental view maintenance system.

2.1 Materialized Views at Databricks

Databricks’s materialized views (MV) physically store query results in Unity Catalog-managed tables [13], in contrast to standard views that are evaluated on demand. Users can create MVs either within Spark Declarative Pipelines or through Databricks SQL (DBSQL) warehouses. When defined inside a pipeline, an MV is refreshed whenever the pipeline itself is refreshed. Standalone MVs, on the other hand, can be refreshed on a user-specified schedule (e.g., every hour) or triggered automatically based on changes to their input tables. Each refresh logically updates the MV to reflect the result of re-running its defining query over the latest versions of its inputs when the refresh starts. Between refreshes, the MV can lag behind its source tables, eventually becoming consistent as subsequent refreshes incorporate newer changes. Delta Lake’s ACID transactions ensure readers always see the last fully committed MV state; refreshes are either entirely visible or not at all.

Figure 2 shows an MV that computes the average order for three selected regions and refreshes hourly. We use this MV as a running example throughout the paper.

Simplified Incrementalization and Orchestration. MVs remove the complexity of incrementalization and orchestration from the user. For instance, for the MV in Figure 2, the underlying engine automatically tracks changes to the input tables (Customers and Orders) and *incrementally* updates the MV accordingly. When MVs are used within a pipeline, the engine also identifies dependencies among pipeline entities and orchestrates their refresh order according to the underlying dependency DAG, exploiting parallelism when possible. This ensures that, for dependencies captured within the same pipeline, all downstream MVs are refreshed against a consistent, committed snapshot of their upstream dependencies.

Databricks additionally provides streaming tables with easy-to-use declarative APIs such as AUTO CDC, which implement change data capture (both SCD Type 1 and Type 2) and handle out-of-order records [16]. These streaming tables, when paired with MVs, allow users to build large, complex pipelines in a fully declarative manner.

Enzyme is the incremental execution engine that powers efficient materialized view refreshes in Databricks. On each refresh, Enzyme evaluates multiple maintenance strategies (several incrementalization techniques plus the option to perform a full recompute and overwrite the MV's contents) and selects the most efficient plan. Whenever possible, Enzyme processes only new data and changes in the source tables, eliminating the need for users to implement manual incremental logic and significantly reducing redundant computation.

Internal Representation of MVs. Each MV consists of two components: a *backing table* and a *top-level view*. The backing table stores the user's data plus internal metadata columns. The top-level view projects away the metadata, exposing only user-requested columns. This separation allows Enzyme to augment internal storage while maintaining a consistent external interface. For instance, as discussed later in Section 3.5, Enzyme decomposes AVG into SUM and COUNT [26], storing both components in the backing table but exposing only the original AVG column through the top-level view. Section 4.3 describes this and other such *technique enablers* in detail, including those for supporting complex aggregations and window functions.

2.2 Spark Primitives

Enzyme relies on two Spark primitives to apply computed changes to materialized views: MERGE INTO [17] applies updates, insertions, and deletions from a source table to a target table. It efficiently executes conditional actions by matching source and target rows on a merge condition.

REPLACE WHERE [18], on the other hand, is simpler. Given a predicate and a set of new rows, it atomically replaces rows matching the predicate in the target table with the provided data.

2.3 Table Format Features

Delta Lake provides several features that Enzyme leverages for incremental maintenance:

2.3.1 Row Tracking. Row tracking [15] assigns a unique identifier to each row at insertion time, and the identifier is preserved across row updates. This enables efficient row-level lineage tracking and MV updates [30].

2.3.2 Change Data Feed (CDF). A change data feed (CDF) [19] tracks row-level changes between table commits. Implemented as a table-valued function, a table's CDF returns one row per insertion or deletion, matching the source table's schema with an additional metadata column indicating the change type (as +1 or -1).

An *effectivized changeset* (analogous to *consolidation* in Differential Dataflow [36]) is the most compact representation of a table's CDF between two commits, canceling operations within the same changeset. For example, if a row was inserted and then deleted from a table, the net effect of these operations cancels. An effectivized changeset, however, can still contain multiple operations for a given row, e.g., an update is represented as a delete followed by an insert.

To effectivize a changeset, its rows are grouped by all columns to compute the sum of the change type column in each group. Then only rows with non-zero net changes are kept.

2.3.3 Deletion Vectors. Deletion vectors [20] implement a merge-on-read strategy for Delta Lake [6]: instead of rewriting entire Parquet [41] files on every change (copy-on-write), DELETE, UPDATE, and MERGE operations mark rows as removed without rewriting files. This reduces write amplification at the cost of additional merging during reads.

2.3.4 Time Travel. Each data manipulation language (DML) operation creates a new table version. Time travel [14] allows reading data from any historical version.

3 Incremental Refresh of Materialized Views

In this section, we start by formally describing the incremental view maintenance problem and then proceed to formally define Enzyme's recursive incremental plan constructions used across a variety of operators and query shapes.

3.1 Incremental View Maintenance

Consider an MV that stores the result of a query $Q(T_1, \dots, T_n)$ defined over a set of base relations T_1, T_2, \dots, T_n . When the contents of one or more base tables are modified, Enzyme aims to automatically compute the corresponding change to the MV, denoted by ΔQ , such that the refreshed view satisfies:

$$Q(T'_1, \dots, T'_n) = Q(T_1, \dots, T_n) \oplus \Delta Q,$$

where T'_i denotes the updated state of table T_i , and \oplus represents the application of the incremental update (typically a combination of insertions and deletions) to the existing MV state. In general, for each base relation we define $\Delta T_i = T'_i - T_i$ where "-" denotes the bag (multi-set) difference, and the incremental view maintenance problem is to compute ΔQ given Q , the current base tables T_i , and their corresponding changes ΔT_i . The Δ terms are concretely represented as relations with rows annotated with the extra change type column denoting insertions (+1) and deletions (-1).

The central challenge addressed by Enzyme is to perform this computation *automatically* for arbitrary SQL queries without user intervention. That is, given any query Q , Enzyme derives a correct and efficient incremental maintenance plan to update the MV in response to changes in its inputs, rather than recomputing the entire query from scratch.

A naive strategy for incrementalization is to exploit the identity $\Delta Q = Q(T'_1, \dots, T'_n) - Q(T_1, \dots, T_n)$. That is, one can evaluate Q on




Figure 3: Simplified query plan for the MV query in Figure 2.

both the previous and the updated states of the inputs, and compute their difference to obtain ΔQ . While this approach is conceptually straightforward and always correct for deterministic queries (Section 3.4 discusses why determinism matters here), it requires evaluating Q twice, which is prohibitively expensive in most practical settings—often more costly than simply recomputing the MV from scratch. Nevertheless, this formulation provides the foundation for constructing more efficient incrementalization strategies for complex queries, as we discuss below.

3.2 Operator-Level Delta Plan Construction

To tackle the problem of efficiently incrementalizing queries, Enzyme operates over the individual operators in the logical query plan of an MV rather than treating the query as a monolithic whole. Spark DataFrames and SQL queries can be expressed as trees of relational operators such as Project, Filter, Aggregate, Window, Join, etc., where each operator transforms one or more input relations into an output relation [7]. Figure 3 shows the simplified query plan for our running example query.

Enzyme operates on these primitive operators and constructs *delta-plan fragments* corresponding to each operator, which are then composed to maintain the result of the overall query [24, 25]. This decomposition works because for any operator node ϕ in the query plan, we can treat its children ψ_i as relations that may receive updates $\Delta\psi_i$. Conceptually, each node behaves as if it materializes an intermediate relation. Thus, incrementalization at the operator level reduces to defining a delta transformation $\Delta\phi$ such that

$$\phi(\psi') = \phi(\psi \oplus \Delta\psi) = \phi(\psi) \oplus \Delta\phi.$$

For simplicity, the above formulation assumes that ϕ is a unary operator with one child ψ . For operators with multiple children, the definition generalizes naturally by taking deltas from all child nodes as inputs.

In general, computing $\Delta\phi$ may require access to any combination of ψ' , ψ , and $\Delta\psi$, as illustrated by the operator-level rules discussed later. Although ψ' can be obtained by applying $\Delta\psi$ to the previous state ψ , Enzyme propagates both ψ and ψ' explicitly for efficiency, as reflected in the algebraic transformations described later. Once per-operator delta transformations are defined, they can be composed bottom-up—following the structure of the query plan—to produce the final incremental plan ΔQ , as illustrated in Figure 4.

Below we outline the intuition behind incrementalizing several key operator types. Many of these operator-level formulations have




Figure 4: Illustration of operator-level delta plan composition. The incremental plan is constructed by traversing the query plan from the leaf nodes upward and progressively applying delta rules. At each operator node ψ , three outputs are generated: ψ (the previous output), ψ' (the output over the updated input), and $\Delta\psi$ (the change up to this node). These intermediate deltas are then used by parent operators, following the algebraic formulations in Section 3.2, to compose the final incremental plan ΔQ .

been extensively studied in prior work on incremental view maintenance [24–26, 40].

Projection (π) and Filter (σ_θ). For projection and filter operators, the delta propagates directly through: $\Delta(\pi_{a_1, \dots, a_k}(T)) = \pi_{a_1, \dots, a_k}(\Delta T)$ and $\Delta(\sigma_\theta(T)) = \sigma_\theta(\Delta T)$. Since projection and filter operate on each tuple independently, their incremental update can be obtained by simply applying the same operator to the input delta.

Aggregation ($\mathcal{G}_{k, \text{agg}}$). Aggregations are more complex because each output tuple summarizes a group of input tuples. A naive recomputation would re-aggregate all groups, but given ΔT , only those groups affected by the changes need to be updated. This leads to the following formulation, where π_- and π_+ denote deletions and insertions produced at this operator:

$$\Delta(\mathcal{G}_{k, \text{agg}}(T)) = \pi_-(\mathcal{G}_{k, \text{agg}}(T \bowtie_k \Delta T)) + \pi_+(\mathcal{G}_{k, \text{agg}}(T' \bowtie_k \Delta T)).$$

For example, for $\mathcal{G}_{k, \text{SUM}(f)}(T)$, it suffices to recompute the sums for only the groups (i.e., distinct k values) that appear in ΔT . This principle generalizes to a wide range of aggregation functions, including SUM, AVG, COUNT, and MEDIAN, as long as the computation is deterministic (see Section 3.4 to see how non-deterministic functions are handled). As discussed later in Section 3.5, Enzyme further extends this formulation with specialized techniques for handling top-level and compositional aggregations.

Window. For window functions with a PARTITION BY clause, the maintenance is analogous to that of aggregations above; only the partitions affected by ΔT require recomputation.

Joins. Join operators are among the most challenging to incrementalize. The delta for an inner join $L \bowtie R$ can be expressed as $\Delta(L \bowtie R) = (\Delta L \bowtie R) + (L' \bowtie \Delta R)$ [26].

Outer joins can be maintained by conceptually decomposing them into an inner join and one (for left- or right-outer) or two (for full-outer) anti-joins, yielding a formulation similar to that proposed in prior work [24]. Although this approach is theoretically correct

and applies to all types of outer joins, its performance degrades significantly when multiple joins are composed. To address this, Enzyme employs specialized delta-construction techniques for the most common join patterns, such as equi-joins and primary-foreign key joins. We omit the full details of these constructions due to space constraints.

To illustrate how we combine the individual rules described above, we now present the delta plan construction for our running example in Figure 2. In the above notation, this query is expressed as $\mathcal{G}_{\text{region, AVG(amount)}}(\text{Customers} \bowtie_{\text{customer_id}} \text{Orders})$.

We can now repeatedly apply the operator-level delta rules from above, pushing the Δ terms down to the leaf relations. For brevity, we write C and O for `Customers` and `Orders`, respectively, and omit the predicates, and join and grouping keys.

$$\begin{aligned} \Delta Q &= \Delta(\sigma(\mathcal{G}(C \bowtie O))) \\ &= \sigma(\Delta(\mathcal{G}(C \bowtie O))) \\ &= \sigma(\pi_-(\mathcal{G}(C \bowtie O) \times \Delta(C \bowtie O)) + \pi_+(\mathcal{G}(C' \bowtie O') \times \Delta(C \bowtie O))) \\ &= \sigma(\pi_-(\mathcal{G}(C \bowtie O) \times (C \bowtie \Delta O + \Delta C \bowtie O)) + \\ &\quad \pi_+(\mathcal{G}(C' \bowtie O') \times (C \bowtie \Delta O + \Delta C \bowtie O))) \end{aligned}$$

3.3 Applying Computed Changes

Given a query Q , we have seen how a corresponding ΔQ plan can be derived through the algebraic transformations discussed above. These transformations are implemented in Enzyme as a semantic rewrite layer that generates executable delta plans. To make the MV up to date, however, the computed changes must be *applied* to the MV's backing table.

To support efficient application of updates, Enzyme introduces an internal `row_id` column that uniquely identifies each output tuple at every stage of the query plan. At the leaf level, Enzyme relies on Delta Lake's row tracking (see Section 2.3.1) to provide stable, unique identifiers for each input row. At higher operator levels, Enzyme deterministically constructs new identifiers that are unique to the output of each operator. For example, for a join operator, the `row_id` is derived by combining the row identifiers from the left and right inputs; for an aggregation, the grouping keys serve as the identifier for each aggregated row.

Once each operator output is associated with a unique `row_id`, the resulting ΔQ can be used to determine which rows in the MV should be deleted or inserted. Enzyme first identifies the set of `row_id` values corresponding to deletions and removes them from the MV's backing table. It then inserts the newly produced rows corresponding to insertions. The precise Delta operations used to apply a change set to the MV's backing table are discussed further in Sections 3.5.2 and 4.6.

3.4 Handling Non-Determinism

A largely overlooked aspect in the IVM literature is how the non-deterministic behavior of certain operators or expressions can affect the correctness of incremental plan generation. Most existing incrementalization techniques, including those discussed above, rely on the ability to deterministically reconstruct the previous output state of each operator given its prior inputs.

Consider the example query below, however:

```
SELECT a, rand() AS r FROM T WHERE r > 0.5
```

When a row is deleted from T , a naive application of the rules above would evaluate `rand()` for the deleted tuple to compute the corresponding delta. But since `rand()` is non-deterministic, it will likely produce a different value upon re-evaluation, and as a result, the computed change would no longer reflect the difference between the old and new states. In this example, a tuple that initially satisfied the predicate ($r > 0.5$) might later not during deletion, incorrectly skipping its removal from the MV.

While functions such as `rand()` are explicitly non-deterministic, non-determinism also arises in more subtle ways. Examples include aggregate functions like `collect_set` or `collect_list`, whose output order is not guaranteed; floating-point aggregates like `SUM` or `AVG`, where arithmetic precision introduces non-deterministic variation; and window functions with ordering clauses, where ties may be broken arbitrarily.

Enzyme mitigates these issues through semantic-preserving rewrites whenever possible and when doing so does not introduce significant performance overhead. For example, Enzyme introduces an explicit local sort to enforce a deterministic order for the output of each `collect_set`. Because these sorts are applied locally within each partition, they do not require cross-node data shuffles and thus incur minimal overhead.

In cases where a deterministic rewrite is not possible—such as queries using `current_timestamp()`, whose evaluation depends on the current system time—Enzyme tries to employ specialized incrementalization strategies tailored to the most common patterns observed in production workloads. Some of these are discussed in Section 3.5. If neither a rewrite nor a specialized technique is possible, e.g., when handling a non-deterministic UDF, Enzyme simply falls back to complete recomputation.

3.5 Optimizations & Special Cases

In Section 3.2, we described a general approach for constructing incremental plans for arbitrary queries. This approach performs well across a wide range of query patterns but can be suboptimal for certain query shapes. Moreover, there are patterns that are not supported by the general formulation. In this section, we outline several cases where Enzyme leverages specialized logic to ensure correctness and efficiency.

These specializations are *local* in nature: each one replaces or refines the default delta rule for a specific operator node in the query plan. The overall bottom-up composition procedure described in Section 3.2 remains unchanged. During plan construction, Enzyme checks whether a node matches one of the specialized cases below and, if so, substitutes the corresponding optimized rule in place of the general formulation. The resulting delta fragments compose with those of surrounding operators exactly as before.

3.5.1 Temporal Filters. As mentioned in Section 3.4, queries containing `current_timestamp()` (and related expressions such as `current_date()`) are non-deterministic and cannot be handled directly by the general IVM formulation. But a very common pattern observed in production workloads is the use of these expressions to define a *temporal filter*, so the MV contains only the most recent n days (or hours) of data. As an example, suppose we update our running example MV to compute each region's 30-day

rolling average sales by adding a filter on the Orders relation:
WHERE date >= **current_date()** - **INTERVAL 30 DAYS**.

To handle such cases, Enzyme includes specialized logic to capture the values of time-dependent functions used in filter predicates in the previous and current refreshes. It then uses the previous and current timestamps to calculate the source rows entering and leaving the temporal window.

Formally, consider a filter $f(t)$ that depends on a time variable t and is applied to a base table T : $\sigma_{f(t)}(T)$. Let $f(t = \text{prev})$ and $f(t = \text{curr})$ denote the filter predicate with t set to the previous and current refresh timestamps, respectively. The delta for this temporal filter can be computed as:

$$\begin{aligned} \Delta(\sigma_{f(t)}(T)) = & \pi_{-}(\sigma_{-f(t=\text{curr})} \wedge f(t=\text{prev})(T)) + \\ & \pi_{+}(\sigma_{-f(t=\text{prev})} \wedge f(t=\text{curr})(T)) + \\ & \sigma_{f(t=\text{curr})}(\Delta T) \end{aligned}$$

Intuitively, the first term removes source rows that have fallen out of temporal window; the second term adds rows that have now moved into the window; and the third term applies any changes that remain valid under the current temporal predicate.

This incrementalization formula is actually a special case of handling modifications to an MV's definition. For instance, handling a change in a filter from **WHERE col = 'x'** to **WHERE col = 'y'** can be handled similarly if we capture the parameters used in the previous and current refreshes (values 'x' and 'y' in this case), although this approach may not be efficient in general.

3.5.2 Leveraging Materialized Data. The general incrementalization formulation described earlier relies on the pre-state (T), post-state (T'), and the computed changes (ΔT) of base tables to derive the delta of a query. But for queries whose top-level operator is an aggregation or a window function, we can further optimize the process by exploiting the fact that the previous output of the query is already materialized in the MV itself. This allows Enzyme to avoid recomputing the previous state of the output from scratch.

Consider again our running example MV from Figure 2. In the general formulation, Enzyme would compute the changeset resulting from the join of the Customers and Orders tables (using $\Delta\text{Customers}$ and ΔOrders). It would then use this to identify the distinct region values (i.e., grouping keys) requiring updates and recompute both the previous and current `avg_order_amount` values for those grouping keys to determine ΔQ . However, as the MV already stores the previous aggregate results, Enzyme can instead simply delete all rows whose grouping keys appear in computed join changeset and then append the newly computed aggregate values. This approach eliminates the need to re-evaluate the previous aggregate state entirely and applies analogously to queries with top-level window functions.

For certain aggregation functions, such as SUM and COUNT, Enzyme can push this optimization further by maintaining incremental adjustments to the aggregate value, rather than recomputing it from scratch. To support this, Enzyme maintains additional metadata—such as the count of tuples per group—if it is not already present among the aggregated columns.

Aggregations such as AVG and STDDEV are internally *decomposed* into simpler aggregates (SUM and COUNT) that can be incrementally

updated using these merge-based delta adjustments [26]. As a result, for the running example query, Enzyme runs the aggregation only on the computed join changeset to compute the delta of the aggregate and then uses a `MERGE INTO` operation to update the MV's rows in place.

3.5.3 Partition Overwrite. The general formulation computes changes at the granularity of individual rows; it determines the precise set of tuples to delete and insert. In some cases, however, this fine-grained approach can be suboptimal. When both the input tables and the MV are partitioned on the same columns, it is sometimes more efficient to compute changes at a coarser granularity by refreshing entire partitions rather than individual rows. This *partition-overwrite* strategy is applicable when the query does not contain operations that span multiple partition keys.

For example, consider tables and an MV all partitioned by a date column. If the query does not include aggregations or window functions that operate across multiple date values, the system can safely refresh data at the partition level. In such cases, Enzyme overwrites the affected partitions instead of computing and applying row-level deltas, which can significantly reduce execution overhead associated with computing these deltas and applying changes at the finer granularity. Enzyme's cost model (discussed in Section 4.5) determines when to use partition overwrite versus row-granularity refresh based on the relative cost of two incremental strategies.

4 System Architecture


Building a production-ready incremental view maintenance system presents challenges beyond the theoretical foundations. Although the algorithmic techniques for incrementalizing individual operators are well-established, productionizing these techniques requires addressing practical concerns including query plan stability, compatibility with existing systems, and robustness across diverse workloads. This section describes Enzyme's system architecture and the key components that enable reliable incremental refresh in a production environment.

Enzyme's architecture comprises several interconnected components that work together to provide end-to-end incremental view maintenance. The system processes user queries through six main stages: normalization, fingerprinting, decomposition, incremental plan generation, costing, and execution. Figure 5 shows a high-level diagram of Enzyme's architecture.

4.1 Normalization

The main entry point for Enzyme is a Catalyst logical plan [7] representing the MV definition. This allows Enzyme to support all Spark front-ends, including SQL and the entire DataFrame API (e.g., for PySpark).

The first step of Enzyme's processing is *normalization*, which serves two critical purposes: (1) simplifying query plans to enable simpler incremental plan construction, and (2) serving as the first step to ensure that logically equivalent queries with minor cosmetic differences produce identical fingerprints (discussed in Section 4.2). As shown below, the normalizer produces a simplified plan that lies between the analyzed and fully optimized plans.



Query plans arriving at Enzyme are logical plans and thus often contain redundancies and complexities that complicate the writing of incrementalization transformations. The normalizer applies a comprehensive set of simplification rules, including inlining common table expressions (CTEs), merging and simplifying filter predicates, collapsing unnecessary projections, and eliminating redundant operations.

Enzyme requires a normalized plan (instead of the fully optimized plan) because some of Spark’s Catalyst optimizations [7], while beneficial for execution, obscure or eliminate information essential for incremental processing. For example, timestamp substitution and empty-relation propagation may rewrite or remove expressions that are semantically meaningful for computing incremental updates. Normalizing the logical plan allows Enzyme to preserve the semantic structure necessary for correct incremental refresh, while still producing a canonical representation for fingerprinting and subsequent transformations.

4.2 Query Fingerprinter

To ensure correctness, Enzyme must detect when users modify an MV’s definition and trigger complete recomputation on the next refresh when changes occur. While catalog-based change tracking could address some scenarios, several factors make this approach insufficient for our purposes.

First, Enzyme must handle MVs defined in both SQL and Python, where source code and logical definitions can diverge. A common pattern observed in production is the use of meta-programming in Python to define complex pipelines with repetitive structures. Such cases make static, catalog-based tracking particularly difficult: source code can change without altering an MV’s semantics (e.g., a code refactor) or an MV’s definition can change without any source code modifications (e.g., changes in external parameters). Second, fingerprints derived from the normalized query plan naturally capture transitive changes through dependency chains—for example, when an MV references a view that itself references

another view. Finally, by operating on the normalized plan, the fingerprinter ignores semantics-preserving changes, avoiding unnecessary recomputation.

Fingerprint Construction. Beyond the normalization performed in the previous stage, the fingerprinter applies additional canonicalization steps to eliminate semantics-independent differences. These include standardizing the order of commutative operators, such as joins and unions, and the order of operands within commutative expressions, as well as ensuring consistent assignment of expression identifiers across sub-expressions.

For UDFs implemented in Python, fingerprinting is more involved because their semantics cannot be easily analyzed through static inspection alone. Enzyme addresses this by incorporating UDF signatures into the MV’s fingerprint. The system tracks not only the UDF definition (using the Python bytecode) but also its transitive dependencies, ensuring that changes to imported libraries or nested function calls trigger appropriate recomputation. This comprehensive approach enables safe incremental processing even in environments with complex and evolving user code.

Fingerprint Stability. Fingerprinting the normalized plan, however, also introduces a key challenge: maintaining fingerprint stability across system upgrades and algorithmic improvements. When Enzyme introduces new incremental techniques or optimizes existing ones, the fingerprints of affected MVs may change even though the user’s query does not. Naively, this would trigger costly recomputation for all affected MVs during system upgrades.

Enzyme addresses this through multi-versioning of each MV’s fingerprint. Before deploying changes affecting fingerprints, we introduce a new fingerprint version that maps the updated plan representation to the existing fingerprint, preserving continuity across the upgrade. Old fingerprint versions are retired alongside the corresponding MV data versions once they are no longer needed. This approach minimizes recomputation while ensuring all MVs can leverage improved incremental techniques.

4.3 Decomposition & Technique Enablers

The split architecture of MVs at Databricks, comprising a backing Delta table and a top-level view, provides crucial flexibility for incremental computation. This design allows Enzyme to modify the internal representation of an MV while maintaining a consistent external interface for users.

Enzyme leverages this flexibility through *technique enablers*, which are transformations applied to MV definitions to make incremental computation possible or more efficient. The most fundamental enabler is the decomposition of aggregate functions into incrementally computable components [26]. For example, an MV containing $AVG(x)$ is internally augmented with additional columns storing $SUM(x)$ and $COUNT(*)$, enabling efficient incremental updates when new data arrives. The decomposer also performs targeted expression rewriting, such as converting semantically equivalent expressions like $FIRST$ to MIN , where the ordering guarantees make them equivalent. Another critical enabler is the explicit propagation of aggregate and window keys to the top-level output of the query plan. This transformation exposes grouping and partition keys in the plan output, which can be exploited by our incremental

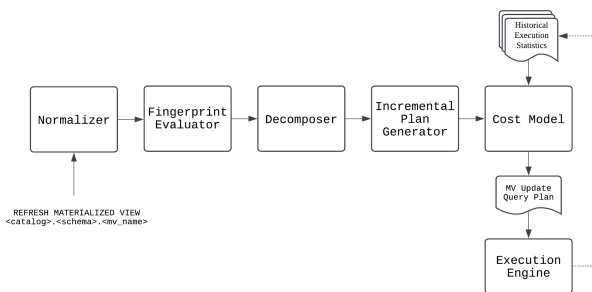


Figure 5: Enzyme transforms a query into an incremental execution in six phases: normalization, fingerprinting, decomposition, plan generation, costing, and execution.

plan generator (discussed below). Finally, the technique enablers handle the construction of the unique row identifiers described in Section 3.3.

4.4 Incremental Plan Generator

The incremental plan generator transforms a normalized and enabled query plan into one that computes the MV’s incremental changes based on the rules defined in Section 3.2. This component employs a recursive visitor pattern that traverses the query plan bottom-up, producing at each node a representation of how to compute changes to that node’s output given changes to its inputs.

Composable Delta Plan Construction. At the core of the generator is the concept of a delta plan, which comprises three pieces: the pre-state (the data before changes), the post-state (the data after changes), and the delta (the actual change representation with insertion/deletion markers). As the visitor traverses each operator in the query plan, it transforms the delta plans from child nodes into a delta plan for the current node by applying the operator-level delta rules described in Section 3.2. Finally, the completed traversal yields an incremental plan to refresh the MV.

Effectivization. Enzyme’s plan generator selectively performs effectivization for changesets consumed by intermediate operators in the plan. This is technically optional but can be useful to reduce downstream computation when the change volume is large and contains redundant pairs.

Further Optimizations. As discussed in Section 3.5.2, for MVs with top-level aggregations or window functions, Enzyme applies an additional optimization. Rather than computing the full changeset with complete before-and-after images of affected rows, the generator recognizes that only the aggregation or partition keys are needed to identify which rows to delete, along with the after-image values for insertion. This eliminates the need to materialize the full pre-state of affected groups, significantly reducing the cost when the aggregation or window is expensive to compute.

4.5 Cost Model

An effective IVM engine requires sophisticated cost modeling. Multiple logically equivalent query plans exist for updating an MV, yet cost models in most online analytical processing (OLAP) databases lack sufficient accuracy to reliably choose between them [34]. Enzyme’s cost model addresses this gap by estimating the end-to-end execution cost of data transformations across layers of the medallion architecture. For each MV, it selects the optimal refresh strategy—incremental update or full recomputation (Figure 6)—by modeling the total refresh cost as the sum of executor CPU times across key physical operators, including joins, aggregates, window functions, shuffles, file scans, and file writes.

Enzyme leverages execution profiles from prior refreshes to ground cost estimates in empirical data. For each operator in a refresh plan under consideration, the system identifies structurally similar historical executions through normalized physical plan matching and derives expected CPU time from observed execution metrics. This approach implicitly captures runtime phenomena such as data skew, spilling behavior, and dynamic pruning effects recorded in prior executions. When no historical information is




Figure 6: Enzyme selects between incremental and full recomputation techniques based on estimated execution cost grounded by historical executions.

available, the model falls back to default parameters derived from large-scale query processing logs.

This costing framework, with its historical execution feedback loop, allows Enzyme to function as an adaptive, data-driven system capable of accurately modeling diverse query structures and workload characteristics while maintaining operational reliability at scale. The cost model serves as a fundamental component that opens opportunities for automatically determining which intermediate MVs would benefit customer workloads by trading additional storage for compute savings [2, 3]. Due to its complexity and significance, a comprehensive treatment of Enzyme’s cost model is reserved for future work.

4.6 Refresh Execution

While Enzyme’s planning components operate on logical query plans, the actual execution of these plans occurs within Apache Spark, with Enzyme as a client. This architectural decision allows Enzyme to benefit from ongoing improvements to the Spark optimizer and execution engine without re-implementing physical planning or execution logic. Enzyme operates on Spark logical plans, specifying transformations and actions that are then optimized and executed by Spark’s Catalyst optimizer [7] and execution engine [9, 44].

Spark Configuration and Optimization. Before executing an incremental plan, Enzyme tunes several Spark configuration parameters to optimize performance based on expected workload characteristics. Because Enzyme constructs the incremental plans itself and has access to historical change statistics, it has additional information that can be used to steer the optimizer toward more informed decisions. For example, when a delta term participates in a join, Enzyme may use historical changeset sizes to hint the optimizer toward appropriate join strategies (e.g., broadcast vs. shuffle). Similarly, Enzyme may adjust broadcast thresholds dynamically based on the changeset sizes observed in previous refreshes.

For complex plans involving multiple aggregations or joins, Enzyme may selectively cache intermediate results that will be reused multiple times within the same refresh operation. This is particularly valuable when computing both the pre- and post-states of inputs for aggregate changesets, where caching the input changeset avoids redundant computation.

Choosing Update Mechanisms. Given an incremental plan produced by the plan generator, Enzyme applies the computed changes to the MV's backing table. It selects between two mechanisms based on the changeset characteristics and query structure.

REPLACE WHERE [18] performs an atomic delete-then-insert operation. Since it semantically applies all deletions before any insertions, the final changeset must be effectivized to avoid incorrect results. Otherwise, a row appearing in both the deletion and insertion sets could introduce duplicate or stale rows. The execution layer is responsible for adding a final effectivization operation to un-effectivized REPLACE WHERE incremental plans.

For certain MVs with top-level aggregates, Enzyme uses an optimized MERGE INTO [17] formulation to update materialized results in place by combining existing values with incremental deltas (see Section 3.5.2).

Transactional Provenance Updates. Enzyme maintains provenance metadata for each MV that records the query fingerprint, source table versions, and other information necessary to validate consistency and detect definition changes. This provenance is stored as part of the MV's metadata and must remain consistent with its data. Enzyme ensures this consistency by committing provenance changes in the same Delta transaction [6] that applies the data changes. Delta's ACID transaction semantics guarantee that either both the data and provenance are updated atomically, or neither update is visible. This prevents the MV from entering an inconsistent state where the data reflects one set of source versions while the provenance metadata references others.

5 Operational Experience & Lessons

Deploying an IVM engine in a production data platform requires rigorous attention to correctness, performance, and reliability. This section describes key operational practices and lessons learned from running Enzyme at scale.

Ensuring Correctness. Enzyme employs multiple complementary strategies to validate correctness. The system leverages a Random Query Generator (RQG) framework to generate random schemas, synthetic data, and SQL queries covering complex combinations of different operators. For each generated query, Enzyme creates an MV, applies randomized changes to source tables, and validates that the incremental refresh produces identical results to complete recomputation. To complement synthetic testing, Enzyme leverages an A/B correctness testing system for production workloads that compares MV contents to their definition queries. This dual approach catches subtle correctness issues that manifest only under specific data distributions or timing conditions not covered by synthetic tests.

Reliability Through Fallback. Beyond correctness validation, ensuring reliability requires graceful degradation in the face of bugs and other issues. Enzyme implements automatic fallback to complete recomputation when planning encounters exceptions or incorrectly generates an invalid incremental plan. The system monitors these exceptions closely, as they often reveal bugs in the underlying query engine. Enzyme also integrates with the Health Mediated Release framework in Databricks, which automatically rolls back to the previously deployed version when a refresh fails after a system

upgrade. The complexity of incremental plans exercises execution paths rarely encountered by typical user queries, making Enzyme an effective stress test for uncovering optimizer and execution bugs in Spark itself.

Performance Optimization Challenges. Operating Enzyme at scale revealed several performance challenges requiring careful attention. While Enzyme leverages Spark's query optimizer when submitting incremental plans for execution, several scenarios required considering alternative plan formulations. For example, when dynamic file pruning optimizations failed to activate, we adjusted Enzyme to use an equivalent incremental formulation that leverages semijoins to provide more explicit pruning. Additionally, some incremental plans caused out-of-memory failures in the driver during optimization due to plan complexity. Enzyme addresses this by adding rules to the normalizer to simplify plans, or by reformulating incremental formulas to be more tractable for the optimizer.

Pipeline-Aware Cost Modeling. A key insight from deploying Enzyme's cost model is that optimization decisions must consider the entire graph of dependent MVs rather than optimizing each view in isolation. Even when an upstream MV would benefit from complete recomputation, performing incremental refresh may be preferable because it generates a smaller change feed that dramatically reduces refresh costs for downstream dependent MVs. This holistic approach can reduce total system cost despite being locally suboptimal for individual MVs.

Fingerprint Stability Challenges. Maintaining fingerprint stability over time required addressing spurious fingerprint changes by leveraging multi-versioning. Changes to the normalizer, for instance, should not change fingerprints and invalidate existing MVs. Similarly, the upgrade from Scala 2.12 to 2.13 potentially introduced many spurious changes due to differences in how Scala objects were hashed. The multi-versioned fingerprinter has been critical for allowing Enzyme to handle such cases without triggering unnecessary recomputation.

6 Experimental Study

We evaluate the effectiveness of Enzyme in performing incremental updates compared to both full recomputation and a state of the art IVM system. Our experiments aim to demonstrate three key outcomes: (1) Enzyme delivers superior performance and efficiency

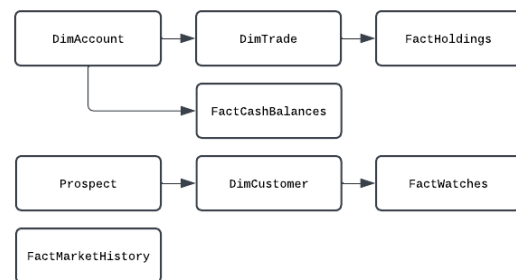


Figure 7: Dependency Graph of MVs in TPC-DI.




Figure 8: Enzyme performance overview on TPC-DI benchmark. Note that the y-axis is log scale.

over full recomputation across scale factors, (2) it outperforms existing commercial alternatives in incremental update workloads, and (3) its cost model reliably identifies when incrementalization is most beneficial.

6.1 Experimental Setup

In the experiments for this paper, we use the TPC-DI benchmark [38] as our workload. TPC-DI is well-suited for evaluating ETL and incremental processing systems because it includes built-in incremental updates, realistic operational domains, and a mix of append-only and CDC-style changes.

6.1.1 Data and Query Workload. The TPC-DI benchmark models end-to-end data warehousing across multiple operational domains—Customer, Account, Trade, Company, Security, Financial, and Market—integrated into a unified analytical schema. It consists of three batches: a large historical load covering over two years of data, followed by two single-day incremental batches. Although TPC-DI does not prescribe an ingestion method, we adopt a streaming ingestion layer, consistent with modern systems and aligned with the Databricks Medallion Architecture. This design preserves the benchmark’s temporal semantics while cleanly separating record acquisition (append or CDC processing, deduplication, surrogate key generation) from analytical derivation.

Operational datasets such as TradeHistory, DailyMarket, and Financial are modeled as append-only streaming tables, while entity datasets, such as Customer, Account, Company, and Security, are represented as CDC streaming tables that maintain current state through merges.

On top of this ingestion layer, analytical datasets—including DimCustomer, DimAccount, DimSecurity, FactHoldings, and FactMarketHistory—are expressed as MVs that perform multi-table joins, aggregations, and enrichment, corresponding to the Silver/Gold layers of the Medallion Architecture.

To ensure parity, we use the same logical definitions for both Enzyme and the compared cloud system. Our study focuses on evaluating the IVM performance of each system, measuring compute efficiency during updates. Datasets that do not receive changes in a given batch are excluded from evaluation.

6.1.2 Enzyme Setup. We create a Spark Declarative Pipeline (SDP) with our ingestion layer defined as Streaming Tables and the analytical datasets defined as MVs. A Databricks Serverless cluster is used as the underlying compute. Data is generated for all 3 batches with the DGen tool provided by TPC-DI and stored into a staging location. Before each batch, we copy the data for the batch to an ingestion location for the pipeline and execute an update of the pipeline.

Data is generated at the 1000, 2000, 5000, and 10,000 scale factors and executed with the same batch sequencing. In all scale factors, we used a serverless autoscaling pipeline backed by AWS *c7g.2xlarge* executor instances (8 CPU cores, 16 GiB memory each). To reduce variance and improve reproducibility, we capped the number of executors at 64, which does not fully reflect production capacity.

6.1.3 CV-IVM Setup. We compare Enzyme against a commercial Cloud Vendor. Due to licensing restrictions, we refer to it simply as CV-IVM. In the CV-IVM setup, we use raw tables for ingestion and defined IVM objects on top of them. For each batch, we copy the data into the raw tables, and trigger a refresh of the IVM objects in topological order to ensure functional equivalence with Enzyme’s automated pipeline updates.

TPC-DI data at the 10,000 scale factor was used to test the CV-IVM performance. The underlying compute used 64 *c7g.2xlarge* executor nodes.

6.2 Experimental Results

To account for differences in hardware and resource configurations, we report relative speedups—expressed as a multiple of the runtime of a full recomputation of the same query. Although the cost models may select full recomputation for certain queries, we report incremental performance results for all queries to present a comprehensive comparison.

These results show that:

- Enzyme successfully supports and incrementalizes **100% of the TPC-DI workloads**, demonstrating broad applicability across diverse dataset types.
- The Enzyme cost model made the **optimal decision in 7 out of 8 cases** with one false negative selection.




Figure 9: Enzyme performance compared to a leading cloud vendor. Note that the y-axis is log scale.

*Speedups were reported as 1 for queries that the vendor did not support incrementalizing.

- When incrementalization is applied, Enzyme achieves **better performance than full recomputation in 6 out of 8 datasets**, confirming the effectiveness of its incremental execution model.
- The observed performance benefits are **consistent across larger scale factors**, indicating that Enzyme’s efficiency scales with data volume.

6.2.1 Comparison with Full Recomputes in Databricks. Figure 8 shows the relative speedup for each scale factor tested in Databricks for each dataset. We summarize our findings as an in-depth breakdown of the overhead and query characteristics is outside the scope of this paper.

Across the benchmark, six out of eight datasets exhibit substantial speedups when executed incrementally compared to full recomputation. The speedups remain roughly constant across scale factors, demonstrating Enzyme’s ability to maintain efficiency and scalability as data volume increases.

FactMarketHistory is a computationally intensive workload that computes 52-week highs and lows for each stock symbol based on roughly two years of historical data. Each refresh requires scanning one year of records per symbol, leading to high memory pressure and increased disk spilling at larger scale factors. This explains the dip in speedup between scale factors 2000 and 5000.

Prospect performs poorly under incremental refresh because a large fraction of the dataset is updated in each batch. The column `SK_RecordDateID` represents the most recent date a prospective customer appears in the input. Since more than 95% of prospects persist across batches, nearly all rows in the dataset are rewritten during each refresh to update this column. In such cases, identifying and applying fine-grained row-level updates incurs more overhead than simply performing a full table refresh.

DimAccount is a lightweight workload that completes in under a minute even at the largest scale. Small datasets like this are more sensitive to transient overheads—including cache contention and scheduling delays—caused by concurrently running queries. Although `DimAccount` itself shows limited speedup, incrementalizing

it remains important because it enables more efficient incremental updates for its downstream consumers.

6.2.2 Comparison with CV-IVM. Figure 9 shows the relative speedup of Enzyme compared to CV-IVM, tested on the 10,000 scale factor. Although the static cost model in the CV-IVM offering did not choose to incrementalize any of the datasets, we overrode the cost model to perform incremental updates where possible.

When setting up the workload in CV-IVM, `DimAccount` and `FactMarketHistory` were not able to be refreshed incrementally due to them containing unsupported operators. Although `DimTrade`, `FactHoldings`, `FactCashBalances` did not have unsupported operators, they were forced to perform full refreshes because their upstream dependencies were not incrementalized.

Of the three datasets that did incrementalize, we observed better performance with `FactWatches`, and the incremental refresh in CV-IVM regressed in `Prospect` and `DimCustomer`. `Prospect` likely suffered for the same reason discussed in section 6.2.1. However it was not apparent why `DimCustomer` regressed with CV-IVM.

6.2.3 Cost model evaluation on TPC-DI. Across the TPC-DI workload, the Enzyme cost model made the right decision for 7 out of the 8 datasets, representing an accuracy of 87.5%.

Prospect, as discussed in Section 6.2.1, is more efficiently handled by a full recompute—consistent with the cost model’s choice.

FactCashBalances was excluded from incrementalization due to operator nesting detected by the model; however, empirical results show that its incremental refresh actually outperforms a full recompute.

DimAccount was chosen to be incremental. Although it did not show significant gains when incrementalized, it was globally optimal as it helped the performance of downstream datasets.

The cost model used by CV-IVM makes static decisions based on the user query, and did not choose to incrementalize any of the datasets in the benchmark.

6.2.4 Autoscaling outcomes. Spark Declarative Pipelines in Databricks support automatic scaling, dynamically adjusting compute resources

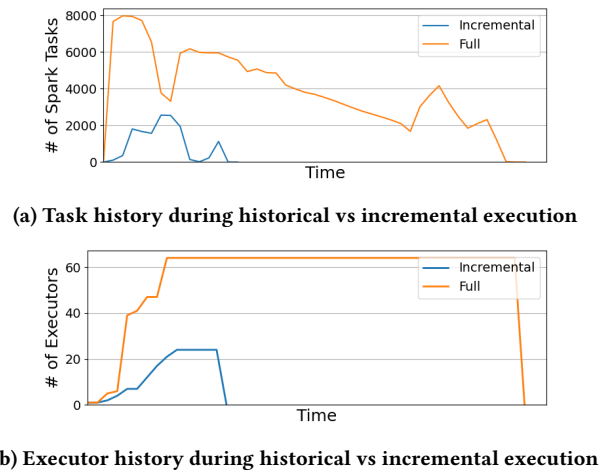


Figure 10: Autoscaling takes advantage of smoother task load during incremental runs to scale down and improve efficiency.

based on workload demand. As shown in Figure 10a, the incremental runs exhibit a curve in the number of submitted Spark tasks with smaller spikes, indicating reduced data volume and more consistent scheduling. Consequently, as illustrated in Figure 10b, the cluster requires fewer executors overall, since the smoother and lower task load eliminates the sharp peaks in parallelism seen during full recomputes. This leads to more efficient resource utilization and lower overall compute cost, as the cluster scales only to the level necessary for the incremental workload.

7 Related Work

Incremental view maintenance (IVM) has been a long-standing problem in databases. Early work [11, 24–26] developed first-order delta rules that propagate inserts, deletes, and updates through relational algebra, with follow-ups covering both set and bag semantics; nested and correlated subqueries remained difficult and only partially addressed by later theory.

Higher-order IVM advanced the state of the art. DBToaster [31, 33] applies recursive finite differencing [32], materializing not only the view but also higher-order deltas, trading additional state for substantially lower update costs on join-heavy queries. Differential Dataflow [36] generalizes incremental computation using partially ordered timestamps and consolidation, while DBSP [12] recasts databases as streams over commutative groups with integration/differentiation operators; the formalisms are closely related.

Several prior works have proposed algorithms and data structures for incrementalizing specific classes of queries. For example, DYN [27] and IDYN [28, 29] target acyclic conjunctive queries (CQs) with equality and inequality join predicates, respectively; AJU [42] focuses on incremental maintenance of acyclic foreign-key joins; and RPAI [1] addresses correlated nested aggregate queries.

A complementary line exploits keys. ID-based IVM [30] represents change sets as identifiers rather than full tuples, often avoiding base-table lookups on the update path and yielding large speedups on multi-join workloads, subject to key/foreign-key assumptions.

Its modular compilation—ID inference, rule instantiation/composition, and algebraic minimization—reduces both work and state.

Production systems span a wide spectrum. Traditional MVs in Redshift [5], BigQuery [22], and Postgres-style [39] warehouses target query acceleration with periodic refresh. Modern streaming MVs (e.g., Materialize [35], RisingWave [43]) are typically deployed as long-running dataflows that continuously process input streams, as opposed to scheduled batch refreshes, enabling low-latency maintenance. Cloud warehouses bridge paradigms: Snowflake Streams and Dynamic Tables [4] pair CDC with scheduled/target-freshness maintenance. Databricks provides MVs refreshed manually or on a schedule; streaming tables which provide append-oriented ingestion with exactly-once processing in managed pipelines. Both surfaces integrate with Delta Lake Change Data Feed for row-level CDC. Streaming tables inherit structured streaming semantics such as watermarks for state bounding and exactly-once processing when using appropriate sources/sinks.

Materialization strategy has been studied from view selection (AND-OR DAGs; multi-query optimization [37]) to recursive maintenance (semi-naive evaluation [23]; DRed and counting variants [26]). These classics underpin today’s systems: DBToaster’s aggressive intermediate caches exemplify the state-vs-latency trade-off, while differential/DBSP frameworks formalize streaming execution and consolidation.

Enzyme delivers a fully automated, built-in approach for Databricks MVs capable of handling arbitrary compositions of complex SQL operators including aggregations, joins, and window functions without user intervention. Its cost-based optimization layer intelligently selects refresh strategies across entire pipelines of MVs, automatically choosing between incremental updates and full recomputation based on historical execution profiles and empirical feedback. On the other hand, streaming tables are designed to seamlessly process and store incoming data streams in real time. Unlike traditional batch tables that require explicit refreshes or data overwrites, streaming tables automatically ingest new data as it arrives, maintaining a continuously updated, up-to-date view. This design makes streaming tables inherently incremental. The combination of streaming tables and MVs creates a unified, powerful framework that elegantly addresses both batch and streaming use cases.

8 Conclusion

Enzyme is a state-of-the-art incremental view maintenance engine that synthesizes decades of foundational academic research into an industrial-grade system, addressing the critical challenge of maintaining MV consistency in modern data lakehouse architectures as source data evolves. Built atop Apache Spark primitives, it achieves exceptional operator coverage and seamlessly handles complex operator compositions including aggregations, joins, and window functions. Among the few production-grade IVM engines, Enzyme demonstrates superior performance, validating both its scalability and practical effectiveness across thousands of real-world pipelines. Successfully bridging the gap between theoretical advances and operational demands, Enzyme is an essential optimization system for data processing in the Lakehouse era.

References

- [1] Supun Abeysinghe, Qiyang He, and Tiark Rompf. 2022. Efficient Incrementalization of Correlated Nested Aggregate Queries using Relative Partial Aggregate Indexes (RPAI). In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '22)*. ACM, 136–149. doi:10.1145/3514221.3517889
- [2] Sanjay Agrawal, Surajit Chaudhuri, and Vivek R. Narasayya. 2000. Automated Selection of Materialized Views and Indexes in SQL Databases. In *Proceedings of the 26th International Conference on Very Large Data Bases (VLDB '00)*. Morgan Kaufmann, Cairo, Egypt, 496–505.
- [3] Rafi Ahmed, Randall Bello, Andrew Witkowski, and Praveen Kumar. 2020. Automated Generation of Materialized Views in Oracle. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3046–3058.
- [4] Tyler Akidau, Paul Barbier, Istvan van Hovell, Adrian Ionescu, Alicja Luszczyk, Michał Świątkowski, Michał Szafranski, Xiao Li, Takuya Ueshin, Mostafa Mokhtar, Peter Boncz, Ali Ghodsi, Sameer Paranjpye, Pieter Senster, Reynold Xin, and Matei Zaharia. 2020. Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3411–3424. doi:10.14778/3415478.3415560
- [5] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. 2015. Spark SQL: Relational Data Processing in Spark. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '15)*. ACM, 1383–1394. doi:10.1145/2723372.2742797
- [6] Nikos Armatzoglou, Sanuj Basu, Naga Bhanoori, Mengchu Cai, Naresh Chainani, Kiran Chinta, Venkatraman Govindaraju, Todd J. Green, Monish Gupta, Sebastian Hillig, Eric Hotinger, Yan Leshinsky, Jintian Liang, Michael McCreedy, Fabian Nagel, Ippokratis Pandis, Panos Parchas, Rahul Pathak, Orestis Polychroniou, Foyzur Rahman, Gaurav Saxena, Gokul Soundararajan, Sriram Subramanian, and Doug Terry. 2022. Amazon Redshift Re-invented. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '22)*. ACM, 2205–2217.
- [7] Alexander Behm, Shoumik Palkar, Utkarsh Agarwal, Timothy Armstrong, David Cashman, Ankur Dave, Todd Greenstein, Shant Hovsepian, Ryan Johnson, Arvind Sai Krishnan, Paul Leventis, Ala Luszczyk, Prashanth Menon, Mostafa Mokhtar, Gene Pang, Sameer Paranjpye, Greg Rahn, Bart Samwel, Tom van Bussel, Herman van Hovell, Maryann Xue, Reynold Xin, and Matei Zaharia. 2022. Photon: A Fast Query Engine for Lakehouse Systems. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '22)*. ACM, 2326–2339. doi:10.1145/3514221.3526054
- [8] Randall G. Bello, Karl Dias, Alan Downing, James J. Feenan, James L. Finnerty, William D. Norcott, Harry Sun, Andrew Witkowski, and Mohamed Ziauddin. 1998. Materialized Views in Oracle. In *Proceedings of the 24th International Conference on Very Large Data Bases (VLDB '98)*. Morgan Kaufmann, New York, NY, USA, 659–664.
- [9] José A. Blakeley, Per-Åke Larson, and Frank Wm. Tompa. 1986. Efficiently Updating Materialized Views. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '86)*. ACM, Washington, DC, USA, 61–71. doi:10.1145/16894.16861
- [10] Mihai Budiu, Tej Chajed, Frank McSherry, Leonid Ryzhyk, and Val Tannen. 2023. DBSP: Automatic Incremental View Maintenance for Rich Query Languages. *Proceedings of the VLDB Endowment* 16, 7 (2023), 1601–1614. doi:10.14778/3587136.3587137
- [11] Ramesh Chandra, Haogang Chen, Ray Matharu, Sarah Cai, Jeff Chen, Priyam Dutta, Bogdan Ghita, Todd Greenstein, Gopal Holla, Peng Huang, Yuchen Huo, Adrian Ionescu, Adriana Ispas, Tim Januschowski, Vihang Karajgaonkar, Stefania Leone, David Lewis, Andrew Li, Nong Li, Cheng Lian, Stephen Link, Qing Lu, Yesheng Ma, Chris Pettitt, Vijayan Prabhakaran, Bogdan Raducanu, Kyle Rong, Paul Roome, Samarth Shetty, Sean Smith, Xiaotong Sun, Yuyuan Tang, Weitao Wen, Lei Xia, Junlin Zeng, Ben Zhang, Reynold Xin, and Matei Zaharia. 2025. Unity Catalog: Open and Universal Governance for the Lakehouse and Beyond. In *Companion of the ACM International Conference on Management of Data (SIGMOD '25)*. ACM, 310–322.
- [12] Databricks. 2019. *Introducing Delta Time Travel for Large Scale Data Lakes*. Retrieved November 1, 2025 from <https://www.databricks.com/blog/2019/02/04/introducing-delta-time-travel-for-large-scale-data-lakes.html>
- [13] Databricks. 2024. *Use Row Tracking for Delta Tables*. Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/delta/row-tracking>
- [14] Databricks. 2025. *The AUTO CDC APIs: Simplify Change Data Capture with Pipelines*. Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/ldp/cdc>
- [15] Databricks. 2025. *MERGE INTO (Delta Lake SQL Reference)*. Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/sql/language-manual/delta-merge-into>
- [16] Databricks. 2025. *Selectively Overwrite Data with Delta Lake*. Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/delta/selective-overwrite>
- [17] Databricks. 2025. *Use Delta Lake Change Data Feed on Databricks*. Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/delta/delta-change-data-feed>
- [18] Databricks. 2025. *What Are Deletion Vectors?* Retrieved November 1, 2025 from <https://docs.databricks.com/aws/en/delta/deletion-vectors>
- [19] Jonathan Goldstein and Per-Åke Larson. 2001. Optimizing Queries Using Materialized Views: A Practical, Scalable Solution. *ACM SIGMOD Record* 30, 2 (2001), 331–342.
- [20] Google. 2025. *Introduction to Materialized Views*. Retrieved November 1, 2025 from <https://cloud.google.com/bigquery/docs/materialized-views-intro>
- [21] Todd J. Green, Shan Shan Huang, Boon Thau Loo, and Wenchao Zhou. 2013. Datalog and Recursive Query Processing. *Foundations and Trends in Databases* 5, 2 (2013), 105–195. doi:10.1561/19000000017
- [22] Timothy Griffin and Bharat Kumar. 1998. Algebraic Change Propagation for Semijoin and Outerjoin Queries. *ACM SIGMOD Record* 27, 3 (1998), 22–27.
- [23] Timothy Griffin and Leonid Libkin. 1995. Incremental Maintenance of Views with Duplicates. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '95)*. ACM, San Jose, CA, USA, 328–339. doi:10.1145/223784.223849
- [24] Ashish Gupta, Inderpal Singh Mumick, and V. S. Subrahmanian. 1993. Maintaining Views Incrementally. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '93)*. ACM, Washington, DC, USA, 157–166. doi:10.1145/170035.170066
- [25] Muhammad Idris, Martin Ugarte, and Stijn Vansummeren. 2017. The Dynamic Yannakakis Algorithm: Compact and Efficient Query Processing Under Updates. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '17)*. ACM, 1259–1274. doi:10.1145/3035918.3064027
- [26] Muhammad Idris, Martin Ugarte, Stijn Vansummeren, Hannes Voigt, and Wolfgang Lehner. 2018. Conjunctive Queries with Inequalities Under Updates. *Proceedings of the VLDB Endowment* 11, 7 (2018), 733–745.
- [27] Muhammad Idris, Martin Ugarte, Stijn Vansummeren, Hannes Voigt, and Wolfgang Lehner. 2019. Efficient Query Processing for Dynamically Changing Datasets. *ACM SIGMOD Record* 48, 1 (2019), 33–40.
- [28] Yannis Katsis, Kian Win Ong, Yannis Papakonstantinou, and Kevin Keliang Zhao. 2015. Utilizing IDs to Accelerate Incremental View Maintenance. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '15)*. ACM, 1985–2000.
- [29] Oliver Kennedy, Yanif Ahmad, and Christoph Koch. 2011. DBToaster: Agile Views for a Dynamic Data Management System. In *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR '11)*. www.cidrdb.org, 284–295.
- [30] Christoph Koch. 2010. Incremental Query Evaluation in a Ring of Databases. In *Proceedings of the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS '10)*. ACM, 87–98. doi:10.1145/1807085.1807100
- [31] Christoph Koch, Yanif Ahmad, Oliver Kennedy, Milos Nikolic, Andres Nötzli, Daniel Lupei, and Amir Shaikhha. 2014. DBToaster: Higher-Order Delta Processing for Dynamic, Frequently Fresh Views. *The VLDB Journal* 23, 2 (2014), 253–278. doi:10.1007/s00778-013-0348-4
- [32] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? *Proceedings of the VLDB Endowment* 9, 3 (2015), 204–215. doi:10.14778/2850583.2850594
- [33] Frank McSherry. 2022. Materialize: A Platform for Building Scalable Event Based Systems. In *Proceedings of the 16th ACM International Conference on Distributed and Event-Based Systems (DEBS '22)*. ACM, 3.
- [34] Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. 2013. Differential Dataflow. In *Proceedings of the 6th Biennial Conference on Innovative Data Systems Research (CIDR '13)*. www.cidrdb.org.
- [35] Hoshi Mistry, Prasan Roy, S. Sudarshan, and Krithi Ramamritham. 2001. Materialized View Selection and Maintenance Using Multi-Query Optimization. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '01)*. ACM, 307–318.
- [36] Meikel Poess, Tilmann Rabl, Hans-Arno Jacobsen, and Brian Caulfield. 2014. TPC-DI: The First Industry Benchmark for Data Integration. *Proceedings of the VLDB Endowment* 7, 13 (2014), 1367–1378.
- [37] PostgreSQL Global Development Group. 2025. *Materialized Views*. Retrieved November 1, 2025 from <https://www.postgresql.org/docs/current/rules-materializedviews.html>
- [38] Dallan Quass. 1996. Maintenance Expressions for Views with Aggregation. In *Proceedings of the Workshop on Materialized Views: Techniques and Applications (VIEWS '96)*. 110–118.
- [39] Deepak Vohra. 2016. Apache Parquet. In *Practical Hadoop Ecosystem: A Definitive Guide to Hadoop-Related Frameworks and Tools*. Springer, 325–335.

- [42] Qichen Wang and Ke Yi. 2020. Maintaining Acyclic Foreign-Key Joins under Updates. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '20)*. ACM, 1225–1239.
- [43] Yanghao Wang and Zhi Liu. 2022. A Sneak Peek at RisingWave: A Cloud-Native Streaming Database. In *Proceedings of the 16th ACM International Conference on Distributed and Event-Based Systems (DEBS '22)*. ACM, 190–193. doi:10.1145/3524860.3543284
- [44] Maryann Xue, Yingyi Bu, Abhishek Somani, Wenchen Fan, Ziqi Liu, Steven Chen, Herman van Hovell, Bart Samwel, Mostafa Mokhtar, Rk Korlapati, Andy Lam, Yunxiao Ma, Vuk Ercegovic, Jiexing Li, Alexander Behm, Yuanjian Li, Xiao Li, Sriram Krishnamurthy, Amit Shukla, Michalis Petropoulos, Sameer Paranjpye, Reynold Xin, and Matei Zaharia. 2024. Adaptive and Robust Query Execution for Lakehouses at Scale. *Proceedings of the VLDB Endowment* 17, 12 (2024), 3947–3959.
- [45] Daniel C. Zilio, Calisto Zuzarte, Sam Lightstone, Wenbin Ma, Guy M. Lohman, Roberta Cochrane, Hamid Pirahesh, Latha S. Colby, Jarek Gryz, Eric Alton, Dongming Liang, and Gary Valentin. 2004. Recommending Materialized Views and Indexes with the IBM DB2 Design Advisor. In *Proceedings of the International Conference on Autonomic Computing (ICAC '04)*. IEEE, 180–187.