

DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views

Christoph Koch · Yanif Ahmad · Oliver Kennedy · Milos Nikolic · Andres Nötzli · Daniel Lupei · Amir Shaikhha

Received: date / Accepted: date

Abstract Applications ranging from algorithmic trading to scientific data analysis require realtime analytics based on views over databases receiving thousands of updates each second. Such views have to be kept fresh at millisecond latencies. At the same time, these views have to support classical SQL, rather than window semantics, to enable applications that combine current with aged or historical data.

In this article, we present the DBToaster system, which keeps materialized views of standard SQL queries continuously fresh as data changes very rapidly. This is achieved by a combination of aggressive compilation techniques and DBToaster’s original recursive finite differencing technique which materializes a query and a set of its higher-order deltas as views. These views support each other’s incremental maintenance, leading to a reduced overall view maintenance cost. This article provides a first description of the complete system, and a thorough experimental evaluation of its performance. DBToaster supports tens of thousands of complete view refreshes a second for a wide range of queries.

Keywords incremental view maintenance, compilation

1 Introduction

Data analytics has been dominated by after-the-fact exploration in classical data warehouses for decades. This is now

Yanif Ahmad
The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, USA. Tel.: (410) 516-6781 E-mail: yanif@jhu.edu

Oliver Kennedy
SUNY Buffalo, 338 Davis Hall, Buffalo, NY 14260, USA. Tel.: (716) 645-1597 E-mail: okennedy@buffalo.edu

Christoph Koch, Milos Nikolic, Andres Nötzli, Daniel Lupei, Amir Shaikhha
École Polytechnique Fédérale de Lausanne; EPFL IC DATA, Station 14, CH-1015 Lausanne, Switzerland. Tel.: (021) 693-7547 E-mail: christoph.koch@epfl.ch

beginning to change: Today, businesses, engineers and scientists are increasingly placing data analytics engines earlier in their workflows to react to signals in fresh data. These dynamic datasets exhibit a wide range of update rates, volumes, anomalies and trends. Responsive analytics is an essential component of computing in finance, telecommunications, intelligence, and critical infrastructure management, and is gaining adoption in operations, logistics, scientific computing, and web and social media analysis.

Developing suitable analytics engines remains challenging. The combination of frequent updates, long-running queries and a large stateful working set precludes the exclusive use of OLAP, OLTP, or stream processors. Furthermore query requirements on updates often do not fall singularly into the functionality and semantics provided by the available technologies, from Complex Event Processing (CEP) engines to triggers, active databases, and database views.

Consider the example of algorithmic trading (see [20] for an application overview). Here, strategy designers want to use analytics — expressible in a declarative language like SQL — on order book data in their algorithms. Order books consist of the orders waiting to be executed at a stock exchange and change very frequently. However, some orders may stay in the order book relatively long before they are executed or revoked, precluding the use of stream engines with window semantics. Applications such as scientific simulations and intelligence analysis also involve entities of interest over widely ranging periods of time, resulting in large stateful and dynamic computation.

The DBToaster project [4, 22, 21] builds and studies data management systems designed for large datasets that evolve rapidly through high-rate update streams. We aim to combine the advantages of DBMS (rich queries over both recent and historical data, without restrictive window semantics) and CEP engines (low latency and high view refresh rates).

The technical focus of this article is on an extreme form of incremental view maintenance that we call *higher-order IVM*. We make use of discrete forward differences (delta queries) recursively, on multiple levels of derivation. That is, we use delta queries (“first-order deltas”) to incrementally maintain the view of the input query, then materialize the delta queries as views too. We then maintain these views using delta queries to the delta queries (“second-order deltas”), use third-order delta queries to incrementally maintain the second-order views, and so on. Our use of higher-order deltas is quite different from earlier work on choosing which query *subexpressions* to materialize and incrementally maintain for best performance [35]. Instead, our technique for constructing higher-order delta views is somewhat reminiscent of discrete wavelets and numerical differentiation methods, and we use a superficial analogy to the Haar wavelet transform as motivation for calling the base technique a *viewlet transform*.

Example 1 Consider a query Q that counts the number of tuples in the product of relations R and S . For now, we only want to maintain the view of Q under insertions. Denote by Δ_R (resp. Δ_S) the change to a view as one tuple is inserted into R (resp. S). Suppose we simultaneously materialize the four views:

- Q (0-th order),
- $\Delta_R Q = \text{count}(S)$ and $\Delta_S Q = \text{count}(R)$ (first order), and
- $\Delta_R \Delta_S Q = \Delta_S \Delta_R Q = 1$ (second order, a “delta of a delta query”).

Then we can simultaneously maintain all these views based on each other, using exclusively summation and avoiding the computation of any products. The fourth view is constant and does not depend on the database. Each of the other views is refreshed when a tuple is inserted by adding the appropriate delta view. For instance, as a tuple is inserted into R , we add $\Delta_R Q$ to Q and $\Delta_R \Delta_S Q$ to $\Delta_S Q$. (No change is required to $\Delta_R Q$, since $\Delta_R \Delta_R Q = 0$.) Suppose R contains 2 tuples and S contains 3 tuples. The table below presents the sequence of states of the materialized views when performing several insertions into R and S . When we add a tuple to S , we increment Q by 2 ($\Delta_S Q$) to obtain 8 and $\Delta_R Q$ by 1 ($\Delta_S \Delta_R Q$) to get 4. If we subsequently insert a tuple into R , we increment Q by 4 ($\Delta_R Q$) to 12 and $\Delta_S Q$ by 1 to 3. A similar process applies for the next two insertions of S tuples.

time point	insert into	$ R $	$ S $	Q	$\Delta_R Q$	$\Delta_S Q$	$\Delta_R \Delta_S Q,$ $\Delta_S \Delta_R Q$
0	—	2	3	6	3	2	1
1	S	2	4	8	4	2	1
2	R	3	4	12	4	3	1
3	S	3	5	15	5	3	1
4	S	3	6	18	6	3	1

Again, the main benefit of using the auxiliary views is that we can avoid computing the product $R \times S$ (or in general,

joins) by simply summing up views. In this example, the view values of the $(k + 1)$ -th row can be computed by just three pairwise additions of values from the k -th row. \square

This is the simplest possible example query for which the viewlet transform includes a second-order delta query, omitting any complex query features (e.g., predicates, self-joins, nested queries). Viewlet transforms can handle general update workloads including deletions and updates, as well as queries with multi-row results.

For a large fragment of SQL, *higher-order IVM avoids join processing*, reducing all the view refreshment work to summation. Joins are only needed in the presence of inequality joins and nested aggregates in view definitions. The viewlet transform repeatedly (recursively) performs delta rewrites. In the absence of nested aggregates, each k -th order delta is structurally simpler than the $(k - 1)$ -th order delta query. The viewlet transform terminates, as for some n , the n -th order delta is guaranteed to be constant, only depending on the update but not on the database. In the above example, the second-order delta is constant, independent of any database relation.

Our higher-order IVM framework, DBToaster, realizes as-incremental-as-possible query evaluation over SQL with a query language extending the bag relational algebra, query compilation and a variety of novel materialization and optimization strategies. DBToaster bears the promise of providing materialized views of long-running SQL queries, without window semantics or other restrictions, at very high view refresh rates. The data may change rapidly, and still part of it may be long-lived. DBToaster can use this functionality as the basis for richer query constructs than those supported by stream engines. DBToaster takes SQL view queries as input, and automatically incrementalizes them into procedural C++ or functional Scala trigger code where all work reduces to fine-grained, low-cost updates of materialized views.

We have developed and released DBToaster as a query compiler toolchain [1]. DBToaster produces dedicated binaries (or source code) that implement long-running query engines for SQL views. The resulting source code or binary can be embedded in client applications, or can operate as standalone systems that consume data from files or sockets. We present a system and architecture overview in [21].

Example 2 Consider the query

```
Q = SELECT SUM(LI.PRICE * O.XCH)
FROM Orders O, Lineitem LI
WHERE O.ORDK = LI.ORDK;
```

on a TPC-H like schema of Orders and Lineitem where line items have prices and orders have currency exchange rates. The query asks for total sales across all orders weighted by exchange rates. We materialize the views for query Q as well

as the first-order views Q_{LI} , representing $\Delta_{LI}Q$, and Q_O , representing Δ_OQ . The second-order deltas are constant with respect to the database and have been inlined in the following insert trigger programs for query Q .

on insert into O values ($ordk, custk, xch$):

```
 $Q$  +=  $xch * Q_O[ordk]$ 
 $Q_{LI}[ordk]$  +=  $xch$ 
```

on insert into LI values ($ordk, ptk, price$):

```
 $Q$  +=  $price * Q_{LI}[ordk]$ 
 $Q_O[ordk]$  +=  $price$ 
```

The query result is again scalar, but the auxiliary views are not, and our language generalizes them from SQL's multi-sets to maps that associate multiplicities with tuples. This is again a very simple example (more involved ones are presented throughout the article), but it illustrates something notable: while classical incremental view maintenance has to evaluate the first-order deltas, which takes linear time in this example (e.g., $\Delta_OQ[ordk]$ is `SELECT SUM(LI.PRICE) FROM Lineitem LI WHERE LI.ORDK=ordk`), we sidestep this by performing IVM on the deltas. This way our triggers can be evaluated in *constant* time for single-tuple inserts in this example. The delete triggers for Q are the same as the insert triggers with `+=` replaced by `-=` everywhere. \square

This example presents single-tuple update triggers. The viewlet transform is not limited to this but supports bulk updates. However, delta queries for single-tuple updates have further optimization potential, which the DBToaster compiler leverages to create very efficient code that refreshes views whenever a new update arrives. By not queueing updates for bulk processing, DBToaster maximizes view availability and minimizes view refresh latency, enabling ultra-low latency monitoring and algorithmic trading applications.

On paper, higher-order IVM clearly dominates classical IVM. If classical IVM is a good idea, then doing it recursively is an even better idea. The same efficiency improvement argument in favor of IVM of the base query also holds for IVM of the delta query. Considering that joins are expensive and this approach eliminates them, higher-order IVM has the potential for excellent query performance.

In practice, how well do our expectations of higher-order IVM translate into real performance gains? A priori, the costs associated with storing and managing auxiliary materialized views for higher-order delta queries might be more considerable than expected. This article presents the lessons learned in an effort to realize higher-order IVM, and to understand its strengths and drawbacks. Our contributions are:

1. We present the concept of higher-order IVM and describe the viewlet transform. This part of the article generalizes and consolidates our earlier work [4, 22].
2. There are cases (inequality joins and certain nesting patterns) when a naive viewlet transform is too aggressive, and certain parts of queries are better re-evaluated than

incrementally maintained. We develop heuristic rules for trading off between materialization and lazy evaluation for the best performance.

3. We have built the DBToaster system that implements higher-order IVM. It combines an optimizing compiler that creates efficient update triggers, based on the techniques discussed above, with a runtime system to keep views continuously fresh as updates stream in at high rates. (The runtime system is currently single-core and main-memory based, but this is not an intrinsic limitation of our method. In fact, our trigger programs are particularly parallel-friendly [22]. See [21] for a more detailed system description.)
4. We present the first set of extensive experimental results on higher-order IVM obtained using DBToaster. Our experiments indicate that frequently, particularly for queries consisting of many joins or nested aggregation subqueries, our compilation approach dominates the state of the art, often by multiple orders of magnitude. On a workload of automated trading and ETL queries, we show that current systems cannot sustain fresh views at the rates required in algorithmic trading and real-time analytics, while higher-order IVM takes a big step to making these applications viable. We have released the DBToaster system as [1].

Most of our benchmark queries contain features like nested subqueries that no commercial IVM implementation supports, while our approach handles them all.

2 Related Work

2.1 A Brief Survey of IVM Techniques

Database view management is a well-studied area with over three decades of supporting literature. A recent survey of the topic can be found in [11]. We focus on the aspects of view materialization most pertinent to the DBToaster project. Our work adds the high-order aspect to IVM, which is orthogonal to all previous work.

Incremental Maintenance Algorithms and Formal Semantics.

Maintaining query answers has been considered under both the set [7, 8] and bag [10, 16] relational algebra. Generally, given a query on N relations $Q(R_1, \dots, R_N)$, classical IVM uses a first-order delta query $\Delta_{R_i}Q = Q(R_1 \cup \Delta R_i, R_2, \dots, R_N) - Q(R_1, \dots, R_N)$ for each input relation R_i in turn. The creation of delta queries has been studied for query languages with aggregation [33] and bag semantics [16], but we know of no work to formally examine delta queries of nested and correlated subqueries. [19] has considered view maintenance in the nested relational algebra (NRA), however this has not been widely adopted in any commercial DBMS. Finally, [42] considered temporal views, and [28]

outer joins and nulls, all for flat SPJAG queries without generalizing to subqueries, the full compositionality of SQL, or the range of standard aggregates.

Materialization and Query Optimization Strategies. Selecting queries to materialize and reuse during processing has spanned fine-grained approaches from subqueries [35] and partial materialization [25,36], to coarse-grained methods as part of multiquery optimization and common subexpressions [18,44]. Picking views from a workload of queries typically uses the AND-OR graph representations from multiquery optimization [18,35], or adopts signature and subsumption methods for common subexpressions [44]. [35] selects sets of subqueries of view definitions for materialization. This is closely related to higher-order IVM, as the delta of a simple query is typically a subquery. Naturally, both approaches select the same (optimal) materialization strategy. However, our delta operation also captures nonlinearities inherent in the deltas of more complex queries (e.g., self-joins), producing different materialization strategies. Our experiments include results for a DBMS that uses a very similar nested subquery materialization strategy. Additionally, our work builds on higher-order IVM, extending it into a full query compilation pipeline.

Physical DB designers [3,45] often use the query optimizer as a subcomponent to manage the search space of equivalent views, reusing its rewriting and pruning mechanisms. For partial materialization methods, ViewCache [36] and DynaMat [25] use materialized view fragments, the former materializing join results by storing pointers back to input tuples, the latter subject to a caching policy based on refresh time and cache space overhead constraints.

Evaluation Strategies. In the interest of efficient maintenance with first-order delta queries, [12,43] studies eager and lazy evaluation to balance query and update workloads, [37] argues for asynchronous view maintenance, and [13] investigates interaction of different view freshness models. In addition, evaluating maintenance queries has been studied extensively in Datalog with semi-naive evaluation (which also uses first-order deltas) and DRed (delete-rederive) [17]. Finally, [15] argues for view maintenance in stream processing, which reinforces our position of IVM as a general-purpose change propagation mechanism for collections, on top of which window and pattern constructs can be defined.

2.2 Update Processing Mechanisms

Triggers and Active Databases. Triggers, active database and event-condition-action (ECA) mechanisms [5] provide general purpose reactive behavior in a DBMS. The literature considers recursive and cascading trigger firings, and restrictions to ensure restricted propagation. Trigger-based approaches require developers to manually convert queries

to delta form, a painful and error-prone process especially in the higher-order setting. Without manual incrementalization, triggers suffer from poor performance and cannot be optimized by a DBMS when written in C or Java.

Data Stream Processing. Data stream processing [2,31] and streaming algorithms combine two aspects of handling updates: (1) shared, incremental processing (e.g. paired vs panned windows, sliding windows), (2) sublinear algorithms (with polylogarithmic space bounds). The latter are approximate processing techniques that are difficult to compose, and have had limited adoption in commercial DBMS. Advanced processing techniques in the streaming community also focus mostly on approximate techniques when processing cannot keep up with stream rates (e.g. load shedding, prioritization [39]), on shared processing (e.g. on-the-fly aggregation [27]), or specialized algorithms and data structures [14]. Our approach to streaming is about generalizing incremental processing to (non-windowed) SQL semantics (including nested subqueries and aggregates). Of course, windows can be expressed in this semantics if desired. Similar principles are discussed in [15].

Automatic Differentiation and Incrementalization, and Applications. Beyond the database literature, the programming language literature has studied automatic incrementalization [29] and automatic differentiation. Automatic incrementalization is by no means a solved challenge, especially when considering general recursion and unbounded iteration. Automatic differentiation considers deltas of functions applied over scalars rather than sets or collections, and lately in higher-order fashion [34]. Bridging these two areas of research would be fruitful for supporting UDFs and general computation on scalars and collections in DBToaster.

3 Queries and Deltas

In this section, we present the internal data model, *generalized multiset relations* (GMRs), and the query language, *AGgregate Calculus* (AGCA), of DBToaster and show how to compute delta queries. The design of the data model and query language avoids complex case distinctions when processing different forms of updates (specifically, deletions) during incremental view maintenance. The language is algebraic in flavor, expressive (it captures most of SQL), and has few syntactic cases, which facilitates the construction of powerful optimizers and compilers.

GMRs generalize multiset relations (as in SQL) to collections of tuples with rational multiplicities (i.e., from \mathbb{Q}). This for one allows us to treat databases and updates uniformly; for instance, a deletion is a relation with negative multiplicities, and applying it to a database means unioning/adding it to the database. It also allows us to use multiplicities to represent aggregate query results (which do not need to be integers). As a consequence, when performing

delta processing on aggregate queries, growing an aggregate means adding to the aggregate value rather than to delete the tuple with the old aggregate value and insert a tuple with the new aggregate value. Maintaining aggregates in the multiplicities allows for simpler and cleaner bookkeeping.

AGCA is a very simple language, essentially constructed from GMRs and infinite interpreted relations (which capture conditions) using just three operations — addition, multiplication, and sum-aggregation. The syntactic simplicity facilitates rich optimizations in the DBToaster compiler. For the purpose of understanding the delta processing framework (and proving it correct), one can view the query language as a polynomial ring over GMRs with an addition operation that at once generalizes multiset union (as known from SQL) and updating, and a multiplication operation that generalizes the natural join operation. This ring-theoretic framework has been initially developed in [22] and has been refined in [23].

The multiplication operation also implements sideways binding passing and enforces range restriction as known in the context of relational calculus. This allows the language to be algebraic without a need for an explicit selection operation. AGCA encodes selection as a multiplication of a query with a condition just like relational calculus does. Multiplication is defined in such a way that query results are guaranteed to be always finite.

3.1 Data Model

It is convenient to model tuples to have their own schema; this way we can use the same definition for variable environments. Formally, we define a tuple \vec{t} as a partial function from a vocabulary of column names to data values. We write \vec{t} as $\langle A : v \mid A \in \text{dom}(\vec{t}) \rangle$, where v is a value from the codomain which we may also identify as $\vec{t}(A)$; $\langle \rangle$ signifies the empty tuple. The set of all tuples is denoted by \mathbb{T} .

A *generalized multiset relation (GMR)* is a function from tuples to rational numbers $R : \mathbb{T} \rightarrow \mathbb{Q}$ such that $R(\vec{t}) \neq 0$ for at most a finite number of tuples \vec{t} . Such a function indicates the *multiplicity* with which each tuple of \mathbb{T} occurs in the GMR. The set of all such functions is denoted by \mathbb{Q}_{Rel} . We require the tuples of a GMR R to have the same schema; the set of column names of R is denoted by $\text{sch}(R)$.

Below, we also use classical singleton relations (without multiplicities) and the natural join operator \bowtie . We write $\{\vec{t}\}$ to construct a singleton relation from tuple \vec{t} with the schema $\text{sch}(\{\vec{t}\}) = \text{dom}(\vec{t})$. For tuples \vec{s}, \vec{t} that are consistent ($\{\vec{s}\} \bowtie \{\vec{t}\} \neq \emptyset$), we can write $\vec{s}\vec{t}$ for the consistent concatenation ($\{\vec{s}\vec{t}\} = \{\vec{s}\} \bowtie \{\vec{t}\}$).

3.2 Query Language

We now formally define AGCA.

Syntax. AGCA expressions are built from constants, variables, relational atoms, conditions, and variable assignments ($:=$), using operations bag union $+$, natural join $*$, and aggregate sums $\text{Sum}_{\vec{A}}$. The abstract syntax is:

$$q ::= q * q \mid q + q \mid -q \mid c \mid x \mid R(\vec{t}) \mid \text{Sum}_{\vec{A}}(q) \mid x \theta 0 \mid x := q$$

Here x denotes variables (which we also call columns), \vec{t} tuples of variables, \vec{A} tuples of group-by variables, R relation names, c constants from \mathbb{Q} , and θ denotes comparison operators ($=, \neq, >, \geq, <, \leq$). We also use $x \theta y$ as syntactic sugar for $(x - y) \theta 0$.

Note that $x := q$ is a special condition essentially equivalent to $x = q$, with one catch. In relational calculus, both variables x and y are safe in $\phi \wedge x = y$ if at least one of them is safe in ϕ (the other variable can be *assigned* the value of the safe variable from ϕ). To simplify our semantics definition, we make the distinction between the case where only one of the variables is safe from the left ($:=$) and the case where both are safe ($=$) clear by syntax.

Informal Semantics. Intuitively, the following steps lead to the definition of AGCA:

1. Take the fragment of relational algebra on multiset relations with just the operations selection σ , natural join \bowtie , union $+$, and (multiplicity-preserving) projection π . Allow queries of this algebra to be nested into selection conditions; projections to nullary relations yield a numerical multiplicity — a tuple count — that can be compared with numerical database fields (e.g., we can write $\sigma_{A < \pi_{\emptyset}(R)}(S)$, when S has a numerical column A).
2. Promote selection conditions to interpreted relations and enforce a range restriction policy to ensure that all queries define finite relations, both as in relational calculus. That is, we may write $R \bowtie (A < B)$ for $\sigma_{A < B}(R)$, and thus have no further need for an explicit selection operation. However, $(A < B)$ by itself is not a valid query because an unbounded number of tuples satisfy the condition. This simplifies the algebra: there are only three effective operations, and we will see that we can treat union and projection alike in delta processing, reducing the operations to two.
3. Generalize this language to GMRs while preserving distributivity of $+$ and \bowtie . This turns GMRs with these two operations into a ring.

There is essentially only one way to do this, as shown in [23], and this solution is the semantics of AGCA. This ring structure makes delta processing for AGCA extremely simple.

With the above intuition in mind, the reader should be able to understand the remainder of the paper and validate that the formal semantics presented below is a solution to the specification just given. Note that in the following the generalized union, join, and projection operations are denoted by $+$, $*$, and $\text{Sum}_{\vec{A}}$.

Semantics. The formal semantics of AGCA is given by an evaluation function $\llbracket \cdot \rrbracket$ that, for a query Q , a database \mathbf{D} , and a *context* — a tuple \vec{b} of “bound variables” — evaluates to an element $\llbracket Q \rrbracket(\mathbf{D}, \vec{b})$ of \mathbb{Q}_{Rel} constructed as follows:

Bag Union

$$\llbracket Q_1 + Q_2 \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \llbracket Q_1 \rrbracket(\mathbf{D}, \vec{b})(\vec{t}) + \llbracket Q_2 \rrbracket(\mathbf{D}, \vec{b})(\vec{t})$$

Natural Join

$$\llbracket Q_1 * Q_2 \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \sum_{\substack{\{\vec{r}\} = \{\vec{r}\} \bowtie \{\vec{s}\} \\ \{\vec{b}\} \bowtie \{\vec{r}\} \neq \emptyset}} \llbracket Q_1 \rrbracket(\mathbf{D}, \vec{b})(\vec{r}) * \llbracket Q_2 \rrbracket(\mathbf{D}, \vec{b})(\vec{s})$$

Group-by Summation

$$\llbracket \text{Sum}_{\vec{A}} Q \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \begin{cases} \sum_{\{\vec{r}\} \bowtie \{\vec{s}\} = \{\vec{s}\}} \llbracket Q \rrbracket(\mathbf{D}, \vec{b})(\vec{s}) & \text{.. } \text{dom}(\vec{t}) = \vec{A}, \\ & \{\vec{b}\} \bowtie \{\vec{r}\} \neq \emptyset \\ 0 & \text{.. otherwise} \end{cases}$$

Constant Value

$$\llbracket c \rrbracket(\cdot, \cdot) := \vec{t} \mapsto \begin{cases} c & \text{.. } \vec{t} = \langle \rangle \\ 0 & \text{.. otherwise} \end{cases}$$

Variable Value

$$\llbracket x \rrbracket(\cdot, \vec{b}) := \vec{t} \mapsto \begin{cases} \text{fail} & \text{.. } x \notin \text{dom}(\vec{b}) \\ \vec{b}(x) & \text{.. otherwise, if } \vec{t} = \langle \rangle \\ 0 & \text{.. otherwise} \end{cases}$$

Variable Comparison

$$\llbracket x \theta 0 \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \begin{cases} 1 & \text{.. } \llbracket x \rrbracket(\mathbf{D}, \vec{b})(\langle \rangle) \theta 0, \vec{t} = \langle \rangle \\ 0 & \text{.. otherwise} \end{cases}$$

Variable Assignment

$$\llbracket x := Q \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \begin{cases} 1 & \text{.. } \vec{t} = \langle x : \llbracket Q \rrbracket(\mathbf{D}, \vec{b})(\langle \rangle) \rangle \\ 0 & \text{.. otherwise} \end{cases}$$

Relation

$$\llbracket R(x_1, \dots, x_k) \rrbracket(\mathbf{D}, \vec{b}) := \vec{t} \mapsto \begin{cases} R^{\mathbf{D}}(\langle A_i : \vec{t}(x_i) \mid A_i \in \text{sch}(R) \rangle) & \text{.. } \{\vec{b}\} \bowtie \{\vec{t}\} \neq \emptyset, \\ & |\text{dom}(\vec{t})| = |\text{sch}(R)| \\ 0 & \text{.. otherwise} \end{cases}$$

AGCA admits sideways information passing. That is, query expressions are evaluated relative to a given *context* \vec{b} — an association of variables and their values — which is provided from the outside. The language — specifically the multiplication operation — dictates how such bindings are to be passed to the right during query evaluation.

The definition of $\llbracket R(x_1, \dots, x_k) \rrbracket$ allows column renaming. The evaluation of variables x using $\llbracket x \rrbracket$ *fails* if they are unbound at evaluation time. We consider a query in which this may happen illegal, and therefore exclude such queries

from AGCA. We define $-Q$ in AGCA as $(-1) * Q$, that is, the product of query Q with the constant -1 . Observe that $R - S = R + (-S)$ does not refer to the difference operation of relational algebra, but to the additive inverse for GMRs: for instance, $\emptyset - R = -R$ in AGCA (\emptyset can be written in AGCA as the constant 0), while the syntactically same expression in relational algebra results in \emptyset . It is more appropriate to think of a GMR $-R$ as a deletion, where deleting “too much” results in a database with *negative tuples*.

Example 3 Let R be a GMR of \mathbb{Q}_{Rel}

$R^{\mathbf{D}}$	A	B	#
	1	2	$\mapsto q_1$
	3	5	$\mapsto q_2$
	4	2	$\mapsto q_3$

where q_1, q_2, q_3 denote rational multiplicities. Then

$\llbracket R(x, y) \rrbracket(\mathbf{D}, \langle x : 3 \rangle)$	x	y	#
	3	5	$\mapsto q_2$

The query renames the columns (A, B) to (x, y) and selects on x since it is a bound variable.

The AGCA version of the query $\sigma_{A < B}(R)$ evaluates to

$\llbracket R(x, y) * (x < y) \rrbracket(\mathbf{D}, \langle \rangle)$	x	y	#
	1	2	$\mapsto q_1$
	3	5	$\mapsto q_2$

For instance,

$$\begin{aligned} & \llbracket R(x, y) * (x < y) \rrbracket(\mathbf{D}, \langle \rangle)(\langle x : 1, y : 2 \rangle) \\ &= \sum_{\{\langle x : 1, y : 2 \rangle\} = \{\vec{r}\} \bowtie \{\vec{s}\}} \llbracket R(x, y) \rrbracket(\mathbf{D}, \langle \rangle)(\vec{r}) * \llbracket x < y \rrbracket(\mathbf{D}, \vec{r})(\vec{s}) \\ &= \llbracket R(x, y) \rrbracket(\mathbf{D}, \langle \rangle)(\langle x : 1, y : 2 \rangle) \\ & *^{\mathbb{Q}} \llbracket x < y \rrbracket(\mathbf{D}, \langle x : 1, y : 2 \rangle)(\langle \rangle) = q_1 \quad \square \end{aligned}$$

Sum-aggregates serve as multiplicity-preserving projections: The result of $\text{Sum}_{\vec{A}} Q$ are the tuples of the projection of Q on \vec{A} , and each tuple’s multiplicity is the sum of the multiplicities of the tuples that were projected down to it. An aggregation $\text{Sum}_{\vec{A}} R$ almost works like the SQL query `SELECT \vec{A} , SUM(1) FROM R GROUP BY \vec{A}` . The only difference is that SQL puts the aggregate values into a new column, while $\text{Sum}_{\vec{A}} R$ puts them into the multiplicity of the group-by tuples. We can express more general aggregate summations using clever arithmetics on multiplicities.

Example 4 The sum-aggregate query $\text{Sum}_{[y]}(R(x, y) * 2 * x)$ generalizes the SQL query `SELECT B, SUM(2 * A) FROM R GROUP BY B` to GMRs. Applied to the GMR of Example 3, it yields

$\llbracket \text{Sum}_{[y]}(R(x, y) * 2 * x) \rrbracket(\mathbf{D}, \langle \rangle)$	y	#
	2	$\mapsto 2q_1 + 8q_3$
	5	$\mapsto 6q_2$

For instance,

$$\begin{aligned}
& \llbracket \text{Sum}_{[y]}(R(x,y) * 2 * x) \rrbracket(\mathbf{D}, \langle \rangle)(\langle y : 2 \rangle) \\
&= \sum_{\vec{r}, \vec{s}, \vec{t}} \llbracket R(x,y) \rrbracket(\mathbf{D}, \langle \rangle)(\vec{r}) * \llbracket 2 \rrbracket(\mathbf{D}, \vec{r})(\vec{s}) * \llbracket x \rrbracket(\mathbf{D}, \vec{r}\vec{s})(\vec{t}) \\
&= q_1 * 2 * \llbracket x \rrbracket(\mathbf{D}, \langle x : 1, y : 2 \rangle)(\langle \rangle) \\
&+ q_3 * 2 * \llbracket x \rrbracket(\mathbf{D}, \langle x : 4, y : 2 \rangle)(\langle \rangle) \\
&= 2 * q_1 + 8 * q_3 \quad \square
\end{aligned}$$

Using the assignment operator, variables can also take on values of non-grouping aggregates, or those that evaluate to a single value for a given set of bindings. That way we can express queries with nested aggregates. Nested aggregates may be correlated with the outside as usual in SQL.

Example 5 Assume relation R has columns (A, B) and relation S has columns (C, D) . The SQL query

```
SELECT * FROM R
WHERE B < (SELECT SUM(D) FROM S WHERE A > C)
```

is equivalent to $\text{Sum}_{[A,B]}(R(A,B) * (z := Q_n) * (B < z))$ with $Q_n = \text{Sum}_{[C,D]}(S(C,D) * (A > C) * D)$. \square

AGCA has no explicit syntax for universal quantification or aggregates other than Sum, but these features can be expressed using (nested) sum-aggregate queries (a popular homework exercise in database courses). Special handling of these features in delta processing and query optimization could yield performance better than what we report in our experiments. However, granting these definable features specialized treatment is beyond the scope of this article. As a consequence, our implementation provides native support for only the fragment presented above and the experiments use only techniques described in the article. This language specification covers all of the core features of SQL with the exception of null values and outer joins.

3.3 Binding Patterns

Query expressions have binding patterns: There are input variables or parameters without which we cannot evaluate these expressions, and output variables, the columns of the query result schema. Each expression Q has input variables or parameters \vec{x}_{in} and a set of output variables \vec{x}_{out} , which form the schema of the query result. We denote such an expression as $Q[\vec{x}_{in}][\vec{x}_{out}]$. The input variables are those that are not range-restricted in a calculus formulation, or equivalently, *parameters* in an SQL query because their values cannot be computed from the database: They have to be provided so that the query can be evaluated.

The most interesting case of input variables occurs in a correlated nested subquery, viewed in isolation. In such a subquery, the correlation variable from the outside is such

an input variable. The subquery can only be computed if a value for the input variable is given. For instance, consider Example 5. Here, all columns of R 's schema are output variables. In the subexpression Q_n , A is an input variable; there are no output variables since the aggregate is non-grouping.

3.4 Computing the Delta of a Query

Next, we show how to construct delta queries. AGCA has the nice property of being *closed under taking deltas*: For each query expression Q , there is an expression ΔQ of the *same language* that expresses how the result of Q changes as the database \mathbf{D} is changed by update workload $\Delta \mathbf{D}$,

$$\Delta Q(\mathbf{D}, \Delta \mathbf{D}) := Q(\mathbf{D} + \Delta \mathbf{D}) - Q(\mathbf{D}).$$

Due to the strong compositionality of the language, we can turn any AGCA expression into its delta by repeatedly applying the following rules syntactically to expressions until we obtain an AGCA expression over GMRs and delta GMRs (updates). We write u to denote an update, and $\Delta_u Q$ for the delta of expression Q with respect to that update. Thus for a GMR R , $\Delta_u R$ is the change to R made in update u .

$$\begin{aligned}
\Delta_u(Q_1 + Q_2) &:= (\Delta_u Q_1) + (\Delta_u Q_2) \\
\Delta_u(Q_1 * Q_2) &:= ((\Delta_u Q_1) * Q_2) + (Q_1 * (\Delta_u Q_2)) \\
&\quad + ((\Delta_u Q_1) * (\Delta_u Q_2)) \\
\Delta_u c &:= 0 \\
\Delta_u x &:= 0 \\
\Delta_u(x \theta 0) &:= 0 \\
\Delta_u(x := Q) &:= (x := (Q + \Delta_u Q)) - (x := Q) \\
\Delta_u(\text{Sum}_{\vec{A}} Q) &:= \text{Sum}_{\vec{A}}(\Delta_u Q)
\end{aligned}$$

The correctness of the rules follows from the fact that the GMRs with $+$ and $*$ form a ring (for example, the delta rule for $*$ is a direct consequence of distributivity) and that $\text{Sum}_{\vec{A}}$ can be thought of as the repeated application of $+$ [23].

The special case of single-tuple updates is interesting since it allows to simplify delta queries further and to generate particularly efficient view refresh code. We write $\pm R(\vec{t})$ to denote the insertion or deletion of a tuple \vec{t} into/from relation R of the database.

$$\begin{aligned}
\Delta_{\pm R(\vec{t})}(R(x_1, \dots, x_{|\text{sch}(R)|})) &:= \pm \prod_{i=1}^{|\text{sch}(R)|} (x_i := t_i) \\
\Delta_{\pm R(\vec{t})}(S(x_1, \dots, x_{|\text{sch}(S)|})) &:= 0 \quad (R \neq S)
\end{aligned}$$

Example 6 Consider the AGCA query

$$Q = \text{Sum}_{[]}(R(A,B) * S(C,D) * (B = C) * A * D)$$

which is equivalent to the query of Example 2. We abbreviate the Sum subexpression as Q_1 . Let us study the insertion/deletion of a single tuple $\langle A : x, B : y \rangle$ to/from R . Since

$$\Delta_{\pm R(x,y)} R(A,B) = \pm(A := x) * (B := y)$$

by the delta rule for $*$,

$$\begin{aligned}\Delta_{\pm R(x,y)}Q_1 &= \pm(A := x) * (B := y) * S(C, D) * (B = C) * A * D \\ &= \pm S(C, D) * (y = C) * x * D\end{aligned}$$

For the main query, we get $\Delta_{\pm R(x,y)}Q = \text{Sum}_{[\]}(\Delta_{\pm R(x,y)}Q_1)$. \square

Note that taking a delta *adds input variables* parameterizing the query with the update. In Example 6, the delta query has input variables x and y to pass in the update.

Example 7 Consider the overall query with a nested aggregate from Example 5. The delta for insertion/deletion of a tuple $\langle C : x, D : y \rangle$ to/from relation S is:

$$\text{Sum}_{[A,B]}(R(A, B) * \Delta_{\pm S(x,y)}(z := Q_n) * (B < z))$$

Following the delta rule for $:=$,

$$\Delta_{\pm S(x,y)}(z := Q_n) = (z := (Q_n \pm \Delta_{\pm S(x,y)}Q_n)) - (z := Q_n)$$

where

$$\begin{aligned}\Delta_{\pm S(x,y)}Q_n &= \text{Sum}_{[\]}((C := x) * (D := y) * (A > C) * D) \\ &= (A > x) * y\end{aligned}$$

Note that A , x , and y are input variables in $\Delta_{\pm S(x,y)}Q_n$, and the expression is equivalent to $\langle \rangle \mapsto (A > x) ? y : 0$. \square

4 The Viewlet Transform

We are now ready for the viewlet transform. In this section, we exclude variable assignments with nested aggregates from the query language, a restriction eliminated in the next section. This query language fragment has the following nice property: ΔQ is structurally strictly simpler than Q when query complexity is measured as follows. For union-free queries, the *degree* $\text{deg}(Q)$ of query Q is the number of relations joined together. Distributivity allows to push unions above joins and thus give a degree to queries with unions: the maximum degree of the union-free subqueries. Queries are strongly analogous to polynomials, and the degree of queries is defined precisely as it is defined for polynomials (where the relation atoms of the query correspond to the variables of the polynomial).

Theorem 1 ([23])

If $\text{deg}(Q) > 0$ then $\text{deg}(\Delta Q) = \text{deg}(Q) - 1$.

The viewlet transform uses the simple fact that a delta query is a query too. Thus it can be incrementally maintained, using the delta query of the delta query, which again can be materialized and incrementally maintained, and so on, recursively. By the above theorem, this recursive query transformation terminates in the $\text{deg}(Q)$ -th recursion level,

when the rewritten query is a “constant” independent of the database, and dependent only on updates.

All queries, whether aggregate or not, map tuples to rational numbers (= define GMRs). Thus it is natural to think of the views as map data structures (dictionaries). In this section, we make no notational difference between queries and (materialized) views, but it will be clear from the context that we are using views when we increment them.

Definition 1 The viewlet transform turns a query into a set of update triggers that together maintain the view of the query and a set of auxiliary views. Assume the query Q has input variables (parameters) \vec{x}_{in} . For each relation R used in the query, the viewlet transform creates a trigger

on insert into R values $(U_R) : T_R$.

where U_R is the update — a generalized multiset relation — to the relation named R and T_R is the trigger body, a set of statements. The trigger bodies are most easily defined by their computation by a recursive imperative procedure VT defined below (initially, the statement lists T_R are empty):

Algorithm 1 $VT(Q, M_Q)$

Require: A query Q to be maintained as M_Q .

Ensure: A list of update statements $T_{\pm R}$ for each update event $\pm R$.

for all Relation Name R used in Q **do**

for all $\pm \in \{+, -\}$ **do**

let \vec{x} = the input/output variables of Q

update $T_{\pm R} = T_{\pm R} :: (\text{foreach } \vec{x} \text{ do } M_Q[\vec{x}] += \Delta_{\pm R}Q[\vec{x}])$

$VT(\Delta_{\pm R}Q, M_{Q,\pm R})$

Here, $::$ on T_R appends a statement to an imperative code block, and $+=$ on GMRs uses the $+$ of Section 3. Only queries occurring in triggers that have degree greater than zero are materialized. Of course, these are exactly the queries that are added to by trigger statements: the queries that need to be incrementally maintained. \square

Example 8 For example, if the schema contains two relations R and S and query Q has degree 2, then $VT(Q, \langle \rangle)$ returns as T_R the code block

```
Q += ΔRQ[UR];
foreach D1 do ΔRQ[D1] += ΔRΔRQ[D1, UR];
foreach D2 do ΔSQ[D2] += ΔRΔSQ[D2, UR];
```

The body for the update trigger of S , T_S , is analogous. Note that the order of the first two statements matters. For correctness, we read the old versions of views in a trigger. \square

The viewlet transform owes its name to a superficial analogy with the Haar wavelet transform, which also materializes a hierarchy of differences.

At runtime, each trigger statement loops over the domain of all attributes of all views appearing in the statement. For

views defined with unsafe attributes, this is a priori astronomically expensive, as the view must be evaluated from scratch. Fortunately, it is often possible to bound the domain of these attributes or optimize such attributes out of the expression. We address such optimizations in Section 5.

Single-tuple updates offer particular optimization potential, and we focus on these in this paper. $+R(\vec{t})$ denotes the insertion of a single tuple \vec{t} into relation R , while $-R(\vec{t})$ denotes the corresponding deletion. We create insert and delete triggers in which the argument is the tuple, rather than a GMR, and thus avoid looping over relation-typed variables.

Example 9 Consider query Q from Example 6, with single-tuple updates. This query has degree two. Let $\text{sgn}_R, \text{sgn}_S \in \{+, -\}$. Then the second-order deltas are:

$$(\Delta_{\text{sgn}_R R(x,y)} \Delta_{\text{sgn}_S S(z,u)} Q)[x,y,z,u] = \text{sgn}_R \text{sgn}_S (y=z) * x * u$$

which is equivalent to $\langle \rangle \mapsto \text{sgn}_R \text{sgn}_S$ if $y=z$ then $x * u$ else 0. Variables x and y are arguments of the trigger and are bound at runtime, but variables z and u need to be looped over. On the other hand, the right-hand side of the trigger is only non-zero in case that $y=z$. So we can substitute z by y everywhere and eliminate z . Using this simplification, the viewlet transform produces the following trigger for $+R(x,y)$:

```
Q += Δ+R(x,y)Q[x,y];
foreach u do Δ+S(y,u)Q[y,u] += {⟨⟩ ↦ x * u};
foreach u do Δ-S(y,u)Q[y,u] -= {⟨⟩ ↦ x * u};
```

The construction of the remaining triggers happens analogously. The trigger contains an update rule for the (in this case, scalar) view Q for the overall query result, which uses the auxiliary view $\Delta_{\pm R(x,y)}Q$ which is maintained in the update triggers for S , plus update rules for the auxiliary views $\Delta_{\pm S(z,u)}Q$ that are used to update Q in the insertion and deletion triggers on updates to S .

The reason why we omitted deltas $\Delta_{\pm R(\dots)}\Delta_{\pm R(\dots)}Q$ and $\Delta_{\pm S(\dots)}\Delta_{\pm S(\dots)}Q$ is that these are guaranteed to be 0, as the query does not have a self-join.

A further optimization, presented in the next section, exploits distributivity and eliminates the loops on u and leads to the triggers of Example 2. \square

We observe that the structure of the work that needs to be done is extremely regular and (conceptually) simple. Moreover, there are no classical large-granularity operators left, so it makes no sense to give this workload to a classical query optimizer. There are for-loops over many variables, which have the potential to be very expensive. But the work is also perfectly data-parallel, and there are no data dependencies comparable to those present in joins. All this provides justification for making heavy use of compilation.

We refer to the viewlet transform as presented in this section as the naive viewlet transform. The following section presents improvements and optimizations.

5 Higher-Order IVM

The DBToaster compilation algorithm is a practical implementation of the viewlet transform, which we call *Higher-Order IVM*. Like the viewlet transform, Higher-Order IVM transforms a query Q into a *trigger program* — a set of triggers that maintain the materialized view on Q as relations appearing in Q are modified and a set of supplemental materialized views. As before, each trigger consists of *update statements*, each of the form **foreach** \vec{x} **do** $M_Q[\vec{x}] += Q'[\vec{x}]$;

The key observation behind Higher-Order IVM is that full delta queries materialized by the naive viewlet transform may be very expensive, or simply impossible to maintain.

An example of the former situation is a delta query including a Cartesian product (i.e., a product of two subqueries $Q_1 * Q_2$ with no output variables in common). As we will soon see, such queries arise quite frequently in the viewlet transform, and are expensive to maintain.

Full delta queries that are impossible to maintain include (1) Delta queries that contain input variables, and therefore lack finite support, and (2) Deltas of queries with nested subqueries, to which Theorem 1 does not apply.

The insight behind DBToaster that makes Higher-Order IVM possible is that it is unnecessary to materialize the full delta query. When generating update statements for a materialized view Q , DBToaster materializes the delta terms $\Delta_{\pm R}Q$ as one or more subqueries of each $\Delta_{\pm R}Q$. When the corresponding update statement is executed, $\Delta_{\pm R}Q$ is reconstituted from the materialized subqueries.

Although materializing subqueries instead of the original delta query increases the cost of evaluating trigger statements, by carefully selecting an appropriate set of subqueries, the increased *execution cost* is offset by a substantial reduction in view *maintenance costs*.

We now formally define Higher-Order IVM (HO-IVM). Recall that the viewlet transform produces a sequence of statements, each of the form: $Q[\vec{x}] += \Delta_{\pm R}Q[\vec{x}]$. Unlike the viewlet transform, HO-IVM does not materialize $\Delta_{\pm R}Q$ as a single materialized view. Instead, HO-IVM identifies a set of subqueries $\vec{M}_{\Delta_{\pm R}Q}$ to materialize and rewrites the statement into an equivalent statement $Q[\vec{x}] += \Delta_{\pm R}Q'[\vec{x}]$, evaluated over these materialized views. Then, instead of recurring on $\Delta_{\pm R}Q$ as in the viewlet transform, HO-IVM recurs individually on each $M_i \in \vec{M}_{\Delta_{\pm R}Q}$. We refer to the rewritten query and the set of materialized subqueries as a *materialization decision* for $\Delta_{\pm R}Q$, denoted $\langle \Delta_{\pm R}Q', \vec{M}_{\Delta_{\pm R}Q} \rangle$.

Example 10 Consider the following query

$$Q[] = \text{Sum}_{[]} (R(A,B) * S(B,C) * T(C,D))$$

The insertion trigger for $+S(b,c)$ includes the statement

$$Q[] += \text{Sum}_{[]} (R(A,b) * T(c,D));$$

Under the viewlet transform, we materialize and maintain the expression $\text{Sum}_{[b,c]}(R(A,b) * T(c,D))$. DBToaster materializes that expression in terms of two sub-expressions:

$$M_1[b] := \text{Sum}_{[b]}(R(A,b)) \quad M_2[c] := \text{Sum}_{[c]}(T(c,D))$$

The insertion trigger then includes the statement:

$$Q[] += \text{Sum}_{[]} (M_1[b] * M_2[c]); \quad \square$$

Algorithm 2 summarizes HO-IVM and Sections 5.1–5.3 discuss heuristics for obtaining a materialization decision.

Algorithm 2 HO-IVM(Q, M_Q)

Require: A query Q to be maintained as M_Q .

Ensure: A list of update statements $T_{\pm R}$ for each update event $\pm R$.

for all Relation Name R used in Q **do**

for all $\pm \in \{+, -\}$ **do**

let \vec{x} = the input/output variables of $\Delta_{\pm R}Q$

let $\langle Q', \{M_i := Q_i\} \rangle = \text{materialize}(\Delta_{\pm R}Q)$

update $T_{\pm R} = T_{\pm R} :: (\text{foreach } \vec{x} \text{ do } M_Q[\vec{x}] += Q'[\vec{x}])$

for all i **do** HO-IVM(Q_i, M_i)

Algorithm 3 Generalized HO-IVM(Q)

Require: A query Q to be maintained.

Ensure: A query Q' , equivalent to Q .

Ensure: A list of update statements $T_{\pm R}$ for each update event $\pm R$.

let $\langle Q', \{M_i := Q_i\} \rangle = \text{materialize}(Q)$

for all i **do** HO-IVM(Q_i, M_i)

Note that the use of materialization decisions need not be limited to *delta* queries. At the user's request, DBToaster can materialize the user-provided (i.e., top-level) query piecewise as opposed to maintaining a single view with the full result that the user is interested in. Although doing so creates a computational overhead when the view contents are accessed, it can substantially reduce maintenance costs. Computing averages is a common example where piecewise materialization is beneficial, as it involves maintaining two separate, simpler aggregates: the count and sum of the input. The average value can be easily reconstituted from the sum and count when it is necessary. This generalized form of Higher-Order IVM is made explicit in Algorithm 3.

5.1 Heuristic Optimization

We present our approach to selecting a materialization decision as a set of independent heuristic rewrite rules that are repeatedly applied to the naive materialization decision $\langle (M_{Q,1}), \{M_{Q,1} := Q\} \rangle$, up to a fixed point.

For clarity, we present these rules in terms of a materialization operator \mathcal{M} . For example, one possible materialization decision for the query $Q := Q_1 * Q_2$ is

$$\mathcal{M}(Q_1) * \mathcal{M}(Q_2) \equiv \langle (M_{Q,1} * M_{Q,2}), \{M_{Q,i} := Q_i\} \rangle.$$

Query Decomposition

$$\mathcal{M}(\text{Sum}_{\vec{A}\vec{B}}(Q_1 * Q_2)) \Rightarrow \mathcal{M}(\text{Sum}_{\vec{A}}(Q_1)) * \mathcal{M}(\text{Sum}_{\vec{B}}(Q_2)) \quad (1)$$

\vec{A} and \vec{B} are any disjoint sets of variables.

Factorization and Polynomial Expansion

$$\mathcal{M}(\text{Sum}_{\vec{A}}(Q_L * (Q_1 + Q_2 + \dots) * Q_R)) \Leftrightarrow$$

$$\mathcal{M}(\text{Sum}_{\vec{A}}(Q_L * Q_1 * Q_R)) + \mathcal{M}(\text{Sum}_{\vec{A}}(Q_L * Q_2 * Q_R)) + \dots \quad (2)$$

Input Variables

$$\mathcal{M}(\text{Sum}_{\vec{A}}(Q * f(\vec{B}\vec{C}))) \Rightarrow \text{Sum}_{\vec{A}}(\mathcal{M}(\text{Sum}_{\vec{A}\vec{B}}Q) * f(\vec{B}\vec{C})) \quad (3)$$

Q is the maximal subquery that contains no input variables.

f is a subquery that contains no relation terms

\vec{A} is any set of variables

\vec{B} is the set of output variables of Q referenced by f

\vec{C} is the set of input variables referenced by f

Nested Aggregates and Decorrelation

$$\mathcal{M}(\text{Sum}_{\vec{A}}(Q_O * (x := Q_N) * f(x, \vec{B}))) \Rightarrow$$

$$\text{Sum}_{\vec{A}}(\mathcal{M}(\text{Sum}_{\vec{A}\vec{B}}(Q_O)) * (x := \mathcal{M}(Q_N)) * f(x, \vec{B})) \quad (4)$$

Q_O is the maximal subquery for which x is not an *input* variable.

Q_N is a subquery containing at least one relation term.

f is a subquery containing no relation terms.

\vec{A} is any set of variables

\vec{B} is the set of output variables of Q_O referenced by f or Q_N

Fig. 1 Rewrite rules for partial materialization. Bidirectional arrows indicate rules that are applied heuristically from left to right while materializing an expression, but applied in reverse to some expressions. Note that for any query Q with output variables \vec{A} , $Q = \text{Sum}_{\vec{A}}(Q)$.

All but the trivial rewrite rules are presented in Figure 1. The full array of heuristic optimizations is discussed in depth below. Figure 2 shows how these rules apply to the experimental workload discussed in Section 9.

Duplicate View Elimination. As the simplest optimization, we observe that the viewlet transform produces many duplicate views. This is primarily because the delta operation typically commutes with itself; $\Delta_R \Delta_S Q = \Delta_S \Delta_R Q$ for any Q without nested aggregates over R or S . Structural equivalence on the materialized view queries is typically sufficient to identify this type of view duplication. View deduplication substantially reduces the number of views created.

Query Decomposition. In most relational optimizers, the generalized distributive law[6] plays an important role in widening the search space for optimal materialization decisions. It is also significant in DBToaster's heuristic optimization strategy. Queries with disconnected join graphs are particularly expensive to materialize. The query decomposition rewrite rule presented in Figure 1.1 exploits the generalized distributive law to break up such queries into smaller components for materialization.

If the join graph of Q includes multiple disconnected components Q_1, Q_2, \dots (i.e., Q is the Cartesian product $Q_1 \times Q_2 \times \dots$), it is better to materialize each component independently as $\mathcal{M}(Q_1) * \mathcal{M}(Q_2) * \dots$

The cost of selecting from (iterating over) $\mathcal{M}(Q)$ is similar to the cost of selecting from $\mathcal{M}(Q_1) * \mathcal{M}(Q_2) * \dots$, as both require an iteration over $|Q_1| \times |Q_2| \times \dots$ elements. Fur-

	Query	Features				Rules			
		T, J	Wc	Gb	Nst.	D	P	I S/C	N R/I
TPC-H	Q1	1	<	✓	-	-	✓	S	I
	Q2	5,=	∧, =	-	1	✓	✓	S	I
	Q3	3,=	∧, <	✓	-	✓	✓	-	-
	Q4	1	∧, <	✓	1	-	✓	S	I
	Q5	6,=	∧, <	✓	-	✓	✓	-	-
	Q6	1	∧, <	-	-	-	-	-	-
	Q7	6,=	∧, ∨, <	✓	1	✓	✓	-	-
	Q8	7,=	∧, =, <	✓	1	✓	✓	S	R
	Q9	6,=	∧, =	✓	1	✓	✓	-	-
	Q10	4,=	∧, =, <	✓	-	✓	✓	-	-
	Q11	2,=	-	✓	-	✓	✓	S	I
	Q11a	2,=	-	✓	-	-	-	-	-
	Q12	2,=	∧, =, <	✓	-	-	✓	-	-
	Q13	2,=	≠	✓	1	-	✓	S	I
	Q14	2,=	∧, <	-	-	-	✓	S	R
	Q15	2,=	∧, <	✓	2	-	✓	S	I
	Q16	2,=	∨, =, ≠	✓	1	✓	✓	S	R
	Q17	2,=	<	-	1	✓	✓	S	I
	Q17a	2,=	<	-	1	✓	✓	S	I
	Q18	3,=	<	✓	2	✓	-	S	R,I
	Q18a	3,=	<	✓	2	✓	✓	S	I
	Q19	2,=	∨, =, <	-	-	-	✓	-	-
Q20	2,=	∧, =, <	-	2	-	✓	S	I	
Q21	4,=	∧, =, <	✓	1	✓	✓	S	I	
Q22	1	=, <	✓	1	-	✓	S	R,I	
Q22a	1	=, <	✓	1	-	✓	S	R,I	
SSB4	7,=	<	✓	-	✓	-	-	-	
Finance	AXF	2,=	∨, <	✓	-	-	✓	S	-
	BSP	2,=	∧, <	✓	-	-	✓	-	-
	BSV	2,=	-	-	-	-	✓	-	-
	MST	2,x	∧, <	✓	1	-	✓	S	R,I
	PSP	2,x	∧, <	-	1	-	✓	S	R,I
	VWAP	1	<	-	1	-	✓	C	R
Sci.	MDDb1	4,=	∧, =	✓	-	✓	✓	-	-
	MDDb2	10,=	∧, ∨, <	-	-	✓	✓	-	-

Fig. 2 Workload features and rewrite rules applied to each query. Features notation: **T, J** = Number of tables, Join type (=: equi, x: cross); **Wc** = Where-clause (\wedge : conjunction, \vee : disjunction, =: equality, \neq : inequality, <: range inequality); **Gb** = GroupBy-clause; **Nst.** = Nesting depth. Rules notation: **D** = query decomposition; **P** = factorization and polynomial expansion; **I** = input variables, with a subquery, **S**, or a view cache, **C**; **N** = nested aggregates and decorrelation, with complete re-evaluation of the nested query, **R**, or incremental evaluation, **I**.

thermore, maintaining each individual Q_i is less computationally expensive: the decomposed materialization stores (and maintains) only $|Q_1| + |Q_2| + \dots$ values, while the combined materialization handles $|Q_1| * |Q_2| * \dots$ values.

Taking a delta of a query with respect to a single-tuple update replaces a relation in the query by a singleton constant tuple, effectively eliminating one hyperedge from the join graph and creating new disconnected components that can be further decomposed. Consequently, this optimization plays a major role in the efficiency of DBToaster, and for ensuring that the number of maps created for any acyclic query is polynomial. Example 10 shows this optimization.

Polynomial Expansion and Factorization. As described, the query decomposition optimization operates exclusively over conjunctive queries (i.e., AGCA expressions without addition). In order to support decomposition across unions, we observe that addition and aggregate summations commute:

$$\text{Sum}_{\vec{A}}(Q_1 + Q_2) = \text{Sum}_{\vec{A}}(Q_1) + \text{Sum}_{\vec{A}}(Q_2)$$

and that the generalized distributive law[6] applies:

$$Q_1 * (Q_2 + Q_3) = (Q_1 * Q_2) + (Q_1 * Q_3)$$

Consequently, any query can be expanded into a sum of multiplicative clauses, where each clause is a conjunctive query (analogous to a query in disjunctive normal form).

As part of the heuristics-based materialization strategy, queries are fully expanded so that each multiplicative clause may be materialized independently. This in turn makes it possible to effectively perform query decomposition. The rewrite rule for this process, which we refer to as polynomial expansion, is presented in Figure 1.2.

Note that this rule can also be applied in reverse. If a common term (Q_L and Q_R in the rewrite rule) appears in multiple multiplicative clauses, the term can be *factorized* out of the sum of these multiplicative clauses to create a smaller, typically cheaper to evaluate, equivalent query expression. It is typically possible (and beneficial) to factorize the rewritten query Q' after obtaining a final materialization decision $\langle Q', \{ \dots \} \rangle$, and the expression is no longer required to be in expanded form.

Input Variables. The delta operation introduces input variables, which in turn makes it possible to create delta queries without finite support. For example, consider the query

$$Q[A, B, C] = R(A, B) * S(C) * (B < C) * A$$

The delta $\Delta_{+R(x,y)}Q[x, y, C] = S(C) * (y < C) * x$ has two input variables (x, y) , making it impossible to fully materialize it.

A trivial solution to this problem is to simply avoid materializing terms that contain input variables. The rewrite rule shown in Figure 1.3 uses the generalized distributive law[6] to do precisely this, by pulling terminals that contain input variables out of the materialization operator. Note that as with the query decomposition operator, this rewrite rule relies on polynomial expansion to simplify the expression into a sum of multiplicative clauses.

In addition to extracting input variables from the materialized query, this rewrite rule also pushes summation into the materialized expression¹. This is analogous to a common optimization in view maintenance and query processing, where aggregation and projection operators are pushed down as far as possible into the evaluation pipeline.

In addition to the trivial solution, several other strategies for dealing with input variables are possible. For queries where the input variables appear in comparison predicates (e.g., $S(C) * (y < C)$), data structure-based solutions such as range indices are possible, but beyond the scope of this work. A third strategy based on caching is shown below.

¹ Note that the operation need not actually have a sum aggregate. An expression Q with output variables \vec{A} is equivalent to the expression $\text{Sum}_{\vec{A}}(Q)$

Deltas of Nested Aggregates. AGCA encodes nested subqueries using the assignment operator ($:=$). Recall that the delta rule for this operator is

$$\Delta_u(x:=Q) := (x:=Q + \Delta_u Q) - (x:=Q)$$

The delta query references the original query (twice), and is clearly not simpler than the original query (as per Theorem 1). On such expressions, the (naive) viewlet transform fails to terminate.

Of course, queries with assignment are not always catastrophic. If $\Delta_u Q = 0$, then

$$(x:=Q + \Delta_u Q) - (x:=Q) = (x:=Q) - (x:=Q) = 0$$

For assignments where the query Q being assigned to x corresponds to a simple arithmetic expression, the delta is always empty. However, if Q contains a relation term $R(\vec{A})$ (i.e., Q represents a nested subquery), then the delta $\Delta_{\pm R}$ must be special-cased.

The rewrite rule presented in Figure 1.4 identifies nested subqueries, and uses the generalized distributive law[6] to extract them for independent materialization. Thus, only expressions without nested subqueries are materialized, and Higher-Order IVM terminates.

As with rules 1 and 3, this rule relies on polynomial expansion to simplify the expression into a sum of multiplicative clauses. Furthermore, just like the input-variable rewrite rule, this rule aggressively pushes aggregates down into the newly created expressions.

Although this rule is necessary to guarantee compiler termination, it can introduce unnecessary overheads into the evaluation of delta queries. When naively used, this rule might separately materialize a nested subquery that does not reference the delta relation. A refinement of this optimization analyzes a given delta query before applying the rewrite rule: Considering the expression from Figure 1.4 and updates $\pm R$, it is only necessary to apply the rewrite rule to Q_N when Q_N references the delta relation R . If it does not, then $\Delta_{\pm R} Q_N = 0$, and the rewrite is unnecessary to ensure termination of Higher-Order IVM.

Example 11 Recall the delta query $\Delta_{\pm S(x,y)} Q$ of Example 7.

$$\text{Sum}_{[A,B]}(R(A,B) * (z := (Q_n \pm (A > x) * y)) * (B < z) - R(A,B) * (z := Q_n) * (B < z))$$

where $Q_n = \text{Sum}_{[C]}(S(C,D) * (A > C) * D)$. The materialization decision for this delta query materializes two subqueries: $M_{Q,1} := R(A,B)$ and $M_{Q,2} := \text{Sum}_{[C]}(S(C,D) * D)$. On every update of relation S , the delta evaluation effectively evaluates the outer query twice: once using the new value $Q_n \pm \Delta Q_n$, and once using the old value of Q_n . Conversely, the delta for updates to R , $\Delta_{\pm R} Q$ always has a lower degree than Q , and is materialized as a single map (the nested-aggregates rewrite rule is ignored). \square

In this example, we see one additional possibility for optimization of nested aggregates. The delta for insertions into S is actually *more* expensive than re-evaluating the entire update (The outer query is evaluated twice in the delta, but just once in the original).

Thus, in some situations DBToaster produces an update statement that replaces the map being maintained, instead of updating it. As a general rule, the incremental approach pays off when the inner query is correlated on an equality, and the delta's arguments bind at least one of these variables; then the delta query only aggregates over a subset of the tuples in the outer query. For instance, if the nested query from Example 7 were to have $(A = C)$ instead of $(A > C)$, then only a subset of the aggregated tuples would have been affected by the delta, leading to the incremental approach as a much better choice. Based on this analysis, the heuristic optimizer decides whether to re-evaluate or incrementally maintain any given delta query.

Note that though Q is being recomputed, we can still accelerate the computation by materializing Q piecewise. Although the expression being materialized is not a delta, we still compute a materialization decision (as in Generalized Higher-Order IVM).

Because we are already materializing the expression Q , care must be taken to avoid creating a self-referential loop in this materialization decision. The default materialization decision $\mathcal{M}(Q)$ is meaningless, as Q defines the view being maintained. We avoid this by first applying the nested-query rewrite heuristic as aggressively as possible to eliminate all nested subqueries in the expressions being materialized. Because recomputation is only appropriate for queries with nested subqueries, the resulting expression is guaranteed to be simpler than Q .

5.2 Specialized Data Structures

Thus far we have considered only straightforward view materialization, where views are stored in map-like data structures. But for some queries, advanced data storage primitives can provide opportunities to materialize more complex expressions — particularly those involving input variables. As many of these opportunities are data-dependent, DBToaster's heuristic optimizations rely on user-input to direct selection of an appropriate data structure. In practice, a cost-based tuning advisor could also be used to automate the selection process with minimal user involvement.

As one example of a specialized data structure, we discuss a data structure capable of materializing expressions with arbitrary input variables: *view caches*. This data structure is analogous to partially materialized views [25,36].

A view cache materializes AGCA expressions with input variables. The cache stores multiple full copies of the

materialized view, each for a different valuation of the input variable(s) that appear in the cache’s defining expression.

When a lookup is performed on the cache, these input variables must be bound to specific values. If the cache contains a materialized view for that particular valuation, the materialized view is returned as a normal map. Otherwise, the cache’s defining query is evaluated as normal and the result is stored in the cache.

Unlike a traditional cache, the contents of a view cache are not invalidated when the underlying data changes. Instead, whenever the data is updated, each materialized view stored in the view cache is updated as normal.

View caches are only beneficial when the size of the active domain of an input variable is small, and so the heuristic optimizer refrains from creating them. The cost-based optimizer, however, includes all available data structures when computing the cost of materializing a given expression.

5.3 Simplifying Delta Expressions

Although the delta operation reduces the expression degree, it makes the expression itself more complicated and harder to evaluate. It introduces input variables into the expression. For self-joins (or equivalently multiple nested aggregates on the same relation), it creates additional additive terms.

Example 12 Consider the following expression:

$$Q[A, B] = R(A) * R(A) * S(B)$$

Applying the delta rules leaves us with the expression

$$\begin{aligned} \Delta_{+R(x)} Q[A, B] = & ((A := x) * R(A) + \\ & R(A) * (A := x) + (A := x) * (A := x)) * S(B) + \\ & R(A) * R(A) * 0 + ((A := x) * R(A) + \\ & R(A) * (A := x) + (A := x) * (A := x)) * 0 \end{aligned}$$

This expression is quite complex, and can be simplified to $\Delta_{+R(x)} Q[x, B] = (2 * R(x) + 1) * S(B)$ \square

This added complexity increases both compilation and evaluation costs. Therefore, as part of Higher-Order IVM, we regularly apply several simplifying transformations to AGCA expressions; some of these correspond to common relational algebra transformations. Like with the heuristic rules, these simplifications are applied repeatedly, up to a fixed point.

Unification. Like unification in relational algebra and first-order logic, we propagate range restrictions through AGCA expressions through a two-step process. First, we transform equality predicates into equivalent assignment expressions. Second, we propagate assignments through the expression, eliminating them if appropriate.

In the first stage, we identify equality comparison terms that can be rewritten into an assignment-compatible form, where a single variable appears on the left-hand side of the expression. Each such equality comparison is commuted left through product terms and out of Sum operators until either (1) the left-hand variable falls out of scope, (2) commuting it further would cause a variable appearing on the right-hand side to fall out of scope. If condition 1 is satisfied, we convert the equality into an assignment.

An equality comparison may have multiple assignment-compatible rewritings. At most one of these rewritings can possibly satisfy condition 1, as a second variable falling out of scope would violate condition 2. When multiple rewritings are available, we continue commuting until condition 1 is satisfied for precisely 1 rewriting, or until condition 2 is violated for all rewritings.

Once all equality comparisons have been converted into assignments, we propagate assignments throughout the expression. This is analogous to beta reduction in lambda calculus, although there are semantic restrictions on AGCA expressions that can prevent us from fully reducing the assignment. There are three such limitations: (1) AGCA forbids range restrictions to be incorporated directly into relation terms, (2) AGCA disallows computationally intensive (i.e. nested aggregates) expressions to be incorporated directly into comparison operations, and (3) If the assignment creates a range restriction on the domain of the query output, the assignment must remain in the expression.

The assignment is propagated as aggressively as possible. If none of the above limitations are violated, the assignment can be safely removed.

Partial Evaluation and Algebraic Identities. Delta derivation frequently produces expressions containing sums of terms that differ only by a constant multiplier. The polynomial factorization heuristic presented in Section 5.1 can be applied (even in the absence of a materialization operator) to group these constant multipliers into a single sum of constants for evaluation.

During this optimization stage, AGCA expressions are partially evaluated by merging constant values that appear in a sum or product together, and by applying the standard algebraic identities $Q + 0 = Q$, $Q * 1 = Q$, and $Q * 0 = 0$. This last identity is especially useful during delta computation, as the delta operation produces many zeroes, as well as expressions of the form $Q - Q$.

Extracting Range Restrictions. Assignments that create a range restriction on the output of a query can sometimes be pulled out of the query. The primary application of this technique is for update trigger statements, where a range restriction on the statement’s loop variables can be applied directly to the map being updated.

The procedure is as follows. After the query has been fully simplified, we identify all assignments where the right-

hand value is a single trigger variable, and which can be commuted up to the left-most position of a product term at the root of the query. These assignments are extracted from the query and used to create a mapping from loop variables to trigger variables, which is applied to both the query and the variables of the map being updated.

For instance, consider the delta query from Example 12. Its simplified version contains terms of the form: $(A := x) * R(A) * S(B)$. Here, we can extract the assignment, and use it to eliminate the loop over variable A in the update statement; the final statement is **foreach** B **do** $Q[x, B] += \Delta_R Q[x, B]$. Note that one of the variables appearing on the left-hand side of the update statement has been bound to a corresponding trigger variable (x).

We incorporate a similar technique into the delta rule for nested queries in order to range-restrict the domain of the generated expression.

Example 13 Consider the following query and its delta.

$$Q[A, B] = (B := R(A))$$

$$\Delta_{+R(x)} Q[A, B] = (B := R(A) + (A := x)) - (B := R(A))$$

Observe that although only a single value of the delta map receives a nonzero update, the range of the delta query is the full domain of R . A more efficient, equivalent query is

$$\Delta_{+R(x)} Q[A, B] = (A := x) * ((B := R(A) + 1) - (B := R(A)))$$

or, simplifying further

$$\Delta_{+R(x)} Q[x, B] = (B := R(x) + 1) - (B := R(x)) \quad \square$$

Recall the delta rule for nested queries

$$\Delta_u(x := Q) = (x := Q + \Delta_u Q) - (x := Q).$$

After computing the nested delta ($\Delta_u Q$), we extract all range restrictions and prepend them to the delta of the full expression. The revised delta rule is thus:

$$\Delta_u(x := Q) = (\Delta_u Q)_{rr} * ((x := Q + (\Delta_u Q)_e) - (x := Q)).$$

Here, $(\Delta_u Q)_{rr}$ represents the extracted range-restricting assignments, and $(\Delta_u Q)_e$ is the nested delta after extraction.

6 Examples

In this section, we provide several examples of Higher-Order IVM. Our goal is to illustrate how the heuristic optimizations interact to produce an efficient view maintenance program, and to highlight interesting behaviors of DBToaster.

on insert into C **values** (ck):

```
01  $Q[] [ck] += Q_C[] [ck]$ 
02 foreach  $OK$  do  $Q_{LI}[] [ck, OK] += Q_{LI,C}[] [ck, OK]$ 
03  $Q_{O1}[] [ck] += 1$ 
```

on insert into O **values** (ck, ok):

```
04  $Q[] [ck] += Q_{O1}[] [ck] * Q_{O2}[] [ok] * (x := Q_{O2}[] [ok]) * (100 < x)$ 
05  $Q_{LI}[] [ck, ok] += Q_{O1}[] [ck]$ 
06  $Q_{LI,C}[] [ck, ok] += 1$ 
07  $Q_C[] [ck] += Q_{O2}[] [ok] * (x := Q_{O2}[] [ok]) * (100 < x)$ 
```

on insert into LI **values** (ok, qty):

```
08 foreach  $CK$  do
    $Q[] [CK] += Q_{LI}[] [CK, ok] * ($ 
      $((Q_{O2}[] [ok] + qty) * (x := Q_{O2}[] [ok] + qty)) -$ 
      $(Q_{O2}[] [ok] * (x := Q_{O2}[] [ok])) * (100 < x)$ 
   )
09 foreach  $CK$  do
    $Q_C[] [CK] += Q_{LI,C}[] [CK, ok] * ($ 
      $((Q_{O2}[] [ok] + qty) * (x := Q_{O2}[] [ok] + qty)) -$ 
      $(Q_{O2}[] [ok] * (x := Q_{O2}[] [ok])) * (100 < x)$ 
   )
10  $Q_{O2}[] [ok] += qty$ 
```

Fig. 3 DBToaster insert trigger program for Q18a.

6.1 Simplified TPC-H Query 18

We explain the DBToaster compilation process on a query that contains an equality-correlated nested aggregate:

```
SELECT C.CK, SUM(LI.QTY) FROM C, O, LI
WHERE C.CK = O.CK AND O.OK = LI.OK AND
      100 < (SELECT SUM(LI1.OK) FROM LI AS LI1
           WHERE LI1.OK = LI1.OK)
GROUP BY C.CK
```

The query is a simplified, but equivalent, version of Q18a from our test workload (see the appendix provided in our technical report [24]). For simplicity, we use the condensed schema $C(CK)$, $O(CK, OK)$, and $LI(OK, QTY)$. Figure 3 shows the generated trigger program. The equivalent ACGA expression $Q[] [CK]$ for the query is:

$$\text{Sum}_{[CK]} (C(CK) * O(CK, OK) * LI(OK, QTY) * QTY * (x := Q_n) * (100 < x))$$

where $Q_n = \text{Sum}_{[]} (LI(OK_1, QTY_1) * (OK = OK_1) * QTY_1)$. First, we simplify the subquery Q_n ; unification eliminates the equality predicate to get an equivalent expression with no input variables: $Q'_n = \text{Sum}_{[OK]} (LI(OK, QTY_1) * QTY_1)$.

Due to space limitations we only show the derivation of insertions into Orders O and Lineitem LI . Insertions into Customer C are a simple extension, while deletions are duals of insertions and are omitted entirely.

Insertions into Orders. The first-order delta of Q for insertion of a single tuple $\langle CK : ck, OK : ok \rangle$ is

$$\Delta_{+O(ck, ok)} Q := \text{Sum}_{[CK]} (C(CK) * LI(OK, QTY) * (OK := ok) * (CK := ck) * QTY * (x := Q'_n) * (100 < x))$$

The delta expression gets simplified after propagating the assignments: every occurrence of OK and CK is replaced with ok and ck , respectively; the new assignments are also safe to delete. The delta expression is the following:

$\text{Sum}_{[ck]}(C(ck) * LI(ok, QTY) * QTY * (x := Q'_n) * (100 < x))$
with OK replaced by ok inside Q'_n .

Query decomposition splits the delta into three parts: $C(ck)$ has no common columns with the rest of the expression, and is materialized as a separate map. The remaining expression can also be divided into two subexpressions that share only the trigger variable ok . Then, since the selection predicate is being applied to a singleton, we can safely materialize only the aggregate in the assignment. Applying these optimizations yields the following materialization decision:

$$\mathcal{M}(\text{Sum}_{[ck]}(C(ck))) * \mathcal{M}(\text{Sum}_{[ok]}(LI(ok, QTY) * QTY) * \text{Sum}_{[]}((x := \mathcal{M}(Q'_n)) * (100 < x)))$$

The trigger statement uses the following set of views (note that Q_{O2} is used twice):

$$Q_{O1}[][CK] := \text{Sum}_{[CK]}(C(CK))$$

$$Q_{O2}[][OK] := \text{Sum}_{[OK]}(LI(OK, QTY) * QTY)$$

$Q_{O1}[][CK]$ is maintained on insertions into C with:

$$\Delta_{+C(ck)}Q_{O1} := \{ \langle CK : ck \rangle \mapsto 1 \}$$

which corresponds to trigger statement 03. $Q_{O2}[][OK]$ is maintained similarly with trigger statement 10.

Insertions into Lineitem. Since the nested subquery is correlated on an equality, DBToaster chooses to compute the first-order delta of Q for insertion of a single tuple $\langle OK : ok, QTY : qty \rangle$. The revised rule for nested subqueries yields:

$$\Delta_{+LI(ok, qty)}(x := Q'_n) := (OK := ok) * ((x := Q'_n + qty) - (x := Q'_n))$$

Following the delta rule for products

$$\Delta_{+LI(ok, qty)}Q := \text{Sum}_{[CK]}(C(CK) * O(CK, OK) * \Delta_{QLI} * QTY * (100 < x))$$

where

$$\Delta_{QLI} := (OK := ok) * ((QTY := qty) * (x := Q'_n) + LI(OK, QTY) * ((x := Q'_n + qty) - (x := Q'_n)) + (QTY := qty) * ((x := Q'_n + qty) - (x := Q'_n)))$$

Polynomial expansion, partial evaluation, and unification result in:

$$\Delta_{+LI(ok, qty)}Q := \text{Sum}_{[CK]}(C(CK) * O(CK, ok) * (LI(ok, QTY) * QTY * (x := Q'_n + qty) - LI(ok, QTY) * QTY * (x := Q'_n) + qty * (x := Q'_n + qty))) * (100 < x))$$

on insert into B values (bv, bp) :

```

01 QB1[][bv] += bp
02 QB2[][ ] += bv
03 QB3[][bv] += 1
04 Q1[][ ] :=
    (Sum1[(v1 := QB2[][ ] * 0.0001) * QB3[][BV'] * (BV' > v1)] *
    Sum1[(v2 := QA2[][ ] * 0.0001) * QA1[][AV'] * (AV' > v2)]) -
    (Sum1[(v3 := QB2[][ ] * 0.0001) * QB1[][BV'] * (BV' > v3)] *
    Sum1[(v4 := QA2[][ ] * 0.0001) * QA3[][AV'] * (AV' > v4)])

```

Fig. 4 The DBToaster trigger program for PSP insertions into B . The deletion trigger for B and the triggers for A are symmetric

Decomposition and polynomial expansion allow to extract $\text{Sum}_{[ok, CK]}(C * O)$ and $\text{Sum}_{[ok]}(LI(ok, QTY) * QTY)$ into separate maps. The rewriting rules for nested aggregates and input variables materialize Q'_n . The final materialization is:

$$\mathcal{M}(\text{Sum}_{[CK, ok]}(C(CK) * O(CK, ok))) * (\mathcal{M}(Q_2) * (x := \mathcal{M}(Q'_n) + qty) - \mathcal{M}(Q_2) * (x := \mathcal{M}(Q'_n)) + qty * (x := \mathcal{M}(Q'_n) + qty)) * (100 < x)$$

where $Q_2 = \text{Sum}_{[ok]}(LI(ok, QTY) * QTY)$

Apart from the outermost materialization (of $C \bowtie O$), the remaining five materializations in this expression are not only equivalent, but identical to Q_{O2} , which is already being maintained. Thus, only one new view, $Q_{LI} := \text{Sum}_{[ok, CK]}(C * O)$, has to be maintained. Rewriting the materialization decision produces trigger statement 08.

Note that this statement requires a loop. We update Q by iterating over $\text{dom}_{CK}(\Delta_{+LI}Q) = \text{dom}_{CK}(Q_{LI})$. For this example, the loop never encounters more than one tuple due to the foreign key dependency from O to C .

Q_{LI} can be maintained in a manner analogous to that of Example 6, resulting in trigger statements 03, 05, and 06.

6.2 The Pricespread Query (PSP)

Next, we look at the query PSP from our test workload that has two nested aggregates:

```

SELECT SUM(A.P - B.P) FROM A, B WHERE B.V >
    (SELECT SUM(B'.P * 0.0001) FROM B, B') AND
    A.V > (SELECT SUM(A'.P * 0.0001) FROM A, A')

```

Again, for simplicity, we use the condensed schema $B(P, V)$ and $A(P, V)$. The ACGA expression Q for the query is:

$$\text{Sum}_{[]} (B(BP, BV) * A(AP, AV) * (AP - BP) * (v1 := \text{Sum}_{[]} (B(BP', BV') * BV' * 0.0001)) * (BV > v1) * (v2 := \text{Sum}_{[]} (A(AP', AV') * AV' * 0.0001)) * (AV > v2))$$

Figure 4 shows the trigger program. As the aggregates have no correlated variables, they can be decorrelated. Because

of this, there is no benefit to using deltas to update the final query result and our compilation heuristics decide on a full recomputation on updates to both A and B . Hence, rather than describing the full compilation process for this example, we focus on the process of materializing the full query.

The join graph of this expression is intriguing. It consists of two mostly disconnected, symmetric components, one for $B(BP, BV)$ and one for $A(AP, AV)$. In fact, the only edge between these two is the term $(AP - BP)$. Our materialization strategy exploits both this, and the fact that integer addition and bag union are identical in AGCA.

Starting with the default materialization strategy $\mathcal{M}(Q)$, we perform polynomial expansion (Rule 2). Because AGCA does not separate integer addition from bag union, this distributes the rest of the expression over the term $(AP - BP)$.

$$\begin{aligned} & \mathcal{M} \left(\text{Sum}_{\square} (B(BP, BV) * A(AP, AV) * AP * \right. \\ & \quad (v1 := \text{Sum}_{\square} (B(BP', BV') * BV' * 0.0001)) * \\ & \quad (v2 := \text{Sum}_{\square} (A(AP', AV') * AV' * 0.0001)) * \\ & \quad \left. (BV > v1) * (AV > v2)) \right) - \\ & \mathcal{M} \left(\text{Sum}_{\square} (B(BP, BV) * A(AP, AV) * BP * \right. \\ & \quad (v1 := \text{Sum}_{\square} (B(AP', BV') * BV' * 0.0001)) * \\ & \quad (v2 := \text{Sum}_{\square} (A(AP', AV') * AV' * 0.0001)) * \\ & \quad \left. (BV > v1) * (AV > v2)) \right) \end{aligned}$$

We can now decorrelate the nested aggregates (Rule 4). This expression contains two identical aggregates, each computing the total volume of B or A . We call these Q_{B2} and Q_{A2} . As only one relation appears in each aggregate, maintenance requires only a single statement each, shown in the trigger program for B as statement 02.

$$\begin{aligned} & \text{Sum}_{\square} \left(\mathcal{M} (\text{Sum}_{[BV, AV]} (B(BP, BV) * A(AP, AV) * AP)) * \right. \\ & \quad (v1 := Q_{B2}[\square] * 0.0001) * (BV > v1) * \\ & \quad \left. (v2 := Q_{A2}[\square] * 0.0001) * (AV > v2)) \right) - \\ & \text{Sum}_{\square} \left(\mathcal{M} (\text{Sum}_{[BV, AV]} (B(BP, BV) * A(AP, AV) * BP)) * \right. \\ & \quad (v1 := Q_{B2}[\square] * 0.0001) * (BV > v1) \\ & \quad \left. (v2 := Q_{A2}[\square] * 0.0001) * (AV > v2)) \right) \end{aligned}$$

After polynomial expansion the expression computes two joins instead of one. The hypergraphs of the simpler joins, however, contain disconnected components. We can apply decomposition (Rule 1) to each.

$$\begin{aligned} & \text{Sum}_{\square} \left(\mathcal{M} (\text{Sum}_{[BV]} (B(BP, BV))) * \right. \\ & \quad \mathcal{M} (\text{Sum}_{[AV]} (A(AP, AV) * AP)) * \\ & \quad (v1 := Q_{B2}[\square] * 0.0001) * (BV > v1) * \\ & \quad \left. (v2 := Q_{A2}[\square] * 0.0001) * (AV > v2)) \right) - \\ & \text{Sum}_{\square} \left(\mathcal{M} (\text{Sum}_{[BV]} (B(BP, BV) * BP)) * \right. \\ & \quad \mathcal{M} (\text{Sum}_{[AV]} (A(AP, AV))) * \\ & \quad (v1 := Q_{B2}[\square] * 0.0001) * (BV > v1) * \\ & \quad \left. (v2 := Q_{A2}[\square] * 0.0001) * (AV > v2)) \right) \end{aligned}$$

Now, no further rules are applicable. We materialize four additional maps: for each volume we maintain both the count and sum of prices of both relations. In the trigger program, maps Q_{B3} and Q_{A3} maintain the counts using statement 03 and its dual in A ; maps Q_{B1} and Q_{A1} maintain the price sums using statements 01 and its dual in A .

Because the total volume of each relation changes with every insertion, we must recompute the price and count totals for the relation that changes. Specialized data structures such as range trees could further reduce the cost of doing so by allowing us to efficiently maintain expressions of the form $\mathcal{M} (\text{Sum}_{[BV]} (B(BP, BV) * (BV > v1)))$.

Nevertheless, by exploiting the connection between addition and bag union, DBToaster is able to evaluate this expression exclusively using scans (as opposed to needing to first compute a Cartesian product, as a traditional database system would do).

7 Query Engine Compilation and System Overview

DBToaster is a compiler toolchain implemented in OCaml. It consists of several compilation stages, their intermediate representations that transform AGCA into an efficient imperative implementation, and a runtime library to provide basic data loading and instrumentation of our generated query engines. This article has primarily focused on frontend compilation, which implements AGCA and its generalized multiset relations from Section 3. In this section, we briefly outline the intermediate representations used in our compiler's backend, including a simple trigger language based on AGCA, as well as a functional language and a procedural language. Compilation applies optimizations in all of these stages prior to code generation.

7.1 Backend Compilation

We have implemented our AGCA queries and its GMR data model as a trigger program language that combines pure AGCA expressions with view maintenance as side effects. We parse AGCA expressions directly from SQL, and produce trigger programs from Higher-Order IVM.

While we can directly interpret trigger programs, for efficient execution, we translate them to a lower-level program. The next step is a functional language, K3 [38], which is inspired by the Collection Programming Language in the Kleisli functional query system [9, 41]. Its main features are its use of a nested collections data model, the incorporation of group-by aggregation to the query language, and a rich set of optimizing program rewrites on collection transformers. Examples of our collection transformers include `map`, `fold`, `groupby`, and `flatten`, as frequently found in functional programming languages such as Scala, Haskell and OCaml.

Our program transformations rely on the definitions of collection transformers through structural recursion, and the subsequent axiomatization of the monad construct realized by structural recursion. By bringing programming language inspired (compile-time) optimizations to query processing, we are able to perform holistic query optimization, which is frequently limited by operator abstractions in query plans. Other recent works have observed a variety of benefits from holistic optimization [26,32]. In the rest of this section, we highlight K3’s compile-time optimizations, relating them to the common methods used in database query optimizers. A full coverage of this topic is outside the scope of this work.

Our transformations start off with normalizing program expressions to a lambda-conditional form, where `if`-statements are lifted to the minimal lambda expressions binding any variables present in the conditional. With conditionals often determining the use of delta queries or initial value queries, this ensures maximal applicability of optimizations inside the two different forms of queries. Next, DBToaster applies constant simplification as well as a conservative form of beta reduction to inline repeated occurrences of scalar values and collection arguments that are used at most once during function application.

Subsequently, we aggressively apply code inlining and simplification through a series of function composition and fusion operations. Along with our collection transformer optimizations, this helps to eliminate large intermediate collections in delta processing and initial value computations, in a similar way to deforestation algorithms [30]. Other notable optimizations include common sub-expression elimination to factor repeated expressions into function application, as well as delayed tuple construction to minimize the flow of packed values throughout expressions. These optimizations are summarized in Table 5. DBToaster’s Scala code generator emits source code directly from our functional representation, making extensive use of Scala’s standard libraries.

Our intermediate procedural language defines a minimal set of imperative control constructs (conditionals, loops, and sequential evaluation), and is designed to facilitate code generation in procedural target languages such as C++. Currently, we perform a limited set of optimizations on our procedural language. As ongoing work, we are working on optimizations that are not captured in our functional layer. This includes low-level loop and peephole optimizations, for example loop unrolling and tiling, and the improved data locality and vector operations exposed therein. Many structural loop optimizations, such as loop fusion and ordering, are already captured by our collection transformer optimizations.

From a procedural representation, the DBToaster code generator supplements source code synthesis with the ability to select and adapt specialized data structures to the query being maintained. We implement our generalized multiset relations as a map data structure (e.g., an STL map) that as-

sociates tuples to their multiplicities. View caches are implemented as two-level nested maps, with the outer-tier corresponding to input variables in our binding patterns, and the inner tier mapping output variables to multiplicities as above. This two-level mapping forms the primary index of our data structure, and is accompanied by a range of secondary indexes based on access patterns in delta queries.

Specifically, we maintain secondary indexes for all binding patterns present in triggers produced by Higher-Order IVM. While our index selection could be refined by physical database design techniques, we found the number of extra indexes to be relatively small throughout our workloads.

DBToaster currently uses the Boost Multi-Index (BMI) library for its C++ data structures. BMI provide in-memory data structures that implement a generic container over composite types (e.g., tuples or C structs), with an ability to compose multiple types of indexes into a single data structure. Thus we can specify secondary indexes over different attributes, for in-memory implementations of hash, tree and sequential data structures. We have developed an equivalent library collection for Scala, and have included it as part of our query runtime library.

7.2 System Overview and Application Usage

The DBToaster compiler produces query processors that are aggressively specialized to a specific query workload, rather than ad-hoc queries. End-users interact with a DBToaster-generated query processor in one of three ways:

1. *Standalone binaries*, where users may run the binary on a file, or specify a listening socket through which data can be sent to the engine. The engine can output a stream of view query results to a file or a network connection.
2. *Shared libraries*, where application developers may link against our library and directly access the data structures representing our views while they are concurrently maintained. We are exploring asynchronous notification methods to support push-based application logic, including callback functions and futures registered with our views.
3. *Source code*, where application developers may adapt and extend our query processing engine as desired, for example to use custom data structures to implement views.

DBToaster produces extensible query engines capable of custom stream pre-processing and workload generation, as well as on-demand querying of views. Our object-oriented design enables users to inherit our engine in their applications, where they may override a pre-processing method invoked on each arriving event. This allows users to perform basic data extraction, transformation, cleaning and logging functionality prior to delta processing.

Users may also direct our compiler to use the generalized form of Higher-Order IVM. The result is a query engine that mixes pull- and push-based processing, providing

Optimization	Description
Condition normalization	Yields a normalized program where <code>if</code> -expressions appear only as the first expression in a function body.
Beta reduction	Inlines function arguments, with a compile-time cost analysis to avoid expensive argument re-evaluation.
Fusion and deforestation	Removes intermediate collection construction by fusing multiple collection transformations.
Partial aggregation	Reduces intermediate collection sizes, especially for nested collections undergoing flattening.
Common subexpressions	Extracts repeated expressions as function application, with consideration of inlining applied by beta reductions.
Effect normalization	Lifts effects to their earliest feasible application to reduce the memory footprint of large values.
Index construction	Rewrites collection transformers applying equality predicates to build and probe index data structures.

Fig. 5 A summary of program transformations and optimizations applied in the K3 language.

users a rich API to retrieve query results. Our API methods pull and compute query results from a set of materialized views that are maintained with higher-order delta queries. This mode of operation produces results at a lower frequency than the application’s update rates. The set of materialized views are those considered by Higher-Order IVM when starting delta rewrites one level down in the query.

Our compiled binaries implement a single-core, single-threaded query executor. Our implementation strategy has focused on novel view maintenance rather than the full range of state-of-the-art query execution mechanisms. Thus, our experiments represent a lower limit on performance and scalability, both of which could be substantially improved with a parallel engine. As ongoing work [21], we are developing a distributed main-memory runtime that exploits aggregate memory and network bandwidth available in large clusters and datacenters. Furthermore, we plan to use K3 for multithreaded and vectorized engines that utilize flexible view data structures as inspired by database cracking.

8 Experiment Setup and Methodology

For the experiments in this paper, we use the DBToaster Public Beta, rev. 2827, released on February 11th, 2013 [1]. We evaluate the experimental performance of DBToaster on Redhat Enterprise Linux on an Intel Xeon E5620 2.4 GHz processor with 16 GB of RAM (on a single-core). Generated C++ code was compiled using `g++ 4.4.6` and linked against the Boost library v1.50; generated Scala code was compiled with version 2.10.0 of the Scala compiler and ran on the Java HotSpot VM (build 23.6-b04). We used `gperftools v2.0` for estimating the memory consumption of our query binaries.

We compare our compilation algorithm with a database management system with incremental view maintenance capabilities (DBX) and a stream processing system (SPY). Because these systems are not optimized for our workload, we also provide a shared-infrastructure comparison by emulating their functionalities — query re-evaluation and IVM — with DBToaster.

Data and Query Workload. Our workload covers algorithmic order book trading (financial), online business decision support scenarios (TPC-H) and scientific queries (MDDDB), which involve computing a variety of statistics. Figure 2 lists

the processing properties of our workload.² Due to limitations of DBToaster, we did several changes to the TPC-H queries: (1) We ignored all the ORDER BY clauses, and, to make the results comparable, also dropped the LIMIT clause; (2) We rewrote MIN, MAX aggregates using equivalent nested subqueries; (3) We replaced the LEFT OUTER join in Q13 with a natural join; (4) Finally, out of convenience, we rewrote HAVING clauses using subqueries, and we inlined INTERVAL expressions into constants.

The financial queries VWAP, MST, AXF, BSP, PSP, and BSV were run on a 2.63 million tuple trace of an order book update stream, representing one day of stock market activity for MSFT. These are updates to a Bids and Asks table with the schema (`timestamp`, `order_id`, `broker_id`, `price`, `volume`). The TPC-H benchmark queries Q1-Q22, and SSB4 were run on a stream of updates adapted from a database generated by DBGEN[40]. We simulate a system that monitors a set of “active” orders by randomly interleaving insertions on all relations and injecting random deletions of Orders and Lineitem rows to keep the Orders and Lineitem tables at around 30000 tuples and 120000 tuples, respectively. All updates preserve the foreign key constraints that exist between the TPC-H tables. Most experiments use a stream synthesized from a scaling factor 0.1 database (100 MB), while our scaling experiments extend these results up to a scaling factor of 10 (10 GB). Finally, the scientific workload was run on a 3.6 million tuple trace (128 MB) of insertions into a table of atom positions during a molecular dynamics simulation.

DBToaster Setup. The DBToaster compiler produces incremental view maintenance code for both C++ and Scala. The compilers for these languages produce binaries with distinct (and surprising) performance characteristics. Our evaluation includes the results for both languages.

DBToaster emulates the behavior of a traditional view maintenance system by terminating recursive delta materialization early. The remaining compiler stages (functional optimization and target-language generation) operate as usual. Our evaluation includes: (1) The HO-IVM algorithm as presented in this paper (DBToaster), (2) A full re-evaluation of the query on every change (REP), (3) The HO-IVM algo-

² The detailed queries can be found in our technical report [24] or on our website <http://www.dbtoaster.org>.

rithm used without recursion (first-order deltas are materialized) to emulate traditional IVM (IVM), and (4) A naive application of the viewlet transform that aggressively materializes as much of each query as possible, creating view caches and employing partial materialization to decorrelate nested subqueries, but ignoring the rules for join graph decomposition and delta expression simplification.

For each compilation method, we measured the memory consumption of the C++ programs. To this end, we produced instrumented binaries for each experiment and processed the same fraction of the stream as without profiling.

DBMS Setup. We compare DBToaster against a commercial DBMS. Due to license restrictions, we refer to it using the anonymized name DBX. In order to measure the rate at which DBX is able to refresh the query results as consistently as possible with other systems, we preload all updates to be performed on all base tables into a single table called Agenda. The Agenda table’s schema is the union of all of the input table schemas, and includes columns identifying the type of update (insert or delete), the table being updated, and the update’s sequence number. Each trial iterates over the updates in Agenda in order, inserting or deleting one tuple and then refreshing the query results, either by re-evaluating the query (DBX-REP), or by using the system’s built-in capability to incrementally maintain materialized views (DBX-IVM). In order to minimize the overheads of the system, we disable log collection as much as possible.

For re-evaluation, we completely re-evaluate the query after each update and store its results in a separate table that gets truncated before each re-evaluation. Because generating materialized views that can be incrementally maintained is non-trivial, has many restrictions, and requires extra update logs, for IVM, we use the provided tuning advisor in order to derive the proper view setup for each of the queries.

In many cases, the tuning advisor suggested views that were not precisely identical to the input queries. We encountered situations in which the advisor added group-by columns or relaxed WHERE clauses by dropping conditions or replacing disjunctions with single expressions, covering a superset of the original condition. We can only speculate that these transformations were meant to allow the generated view to support answering a larger class of queries. For complex queries that could not be maintained as a single view, the advisor generated nested subviews to be incrementally maintained and a top-level view to be re-evaluated on every commit. Out of 36 queries that we experimented with, 20 required up to 5 nested subviews.

SPY Setup. As a second comparison point, we use a commercial stream processor. We refer to the stream processor using the anonymized name SPY due to license restrictions. One major semantic difference between traditional stream processing engines and DBToaster is that stream processing engines are optimized to operate on windows of input

streams, while DBToaster is designed to handle the whole history of streams. We benchmark SPY by reading the same Agenda table used for DBX directly into a stream to minimize event dispatch overheads.

We implemented the queries using the dialect of SQL supported by SPY. Since the queries in our benchmark cannot be efficiently expressed using window semantics, we used auxiliary in-memory tables for all relations. Our implementation of the queries assign a monotonically increasing number to each event and dispatch it to a stream corresponding to the affected relation. This stream updates the in-memory table of the relation by inserting or removing the affected tuple. Then, the query result is re-evaluated and recorded together with the event number and a timestamp. Full recomputation is necessary as the SPY does not support IVM.

Although we attempted to maintain the original query semantics, the SQL dialect employed by SPY imposes some limitations. A severe limitation is that in-memory tables may not be joined together; each in-memory table may only be joined with a stream, requiring manual selection of a join order. Our heuristic for this order was to minimize the size of intermediate streams.

9 Experimental Results

Our results show view refresh rates on stream traces, replayed with a timeout of two hours. Details of the traces are provided in Section 8. These results show that:

- DBToaster consistently outperforms the two commercial systems we tested against, often by multiple orders of magnitude (Figures 6 and 7).
- The performance gap between Higher-Order IVM and Traditional IVM is even greater within the DBToaster runtime. Through aggressive optimization, we believe that DBToaster’s performance can be improved by at least another order of magnitude.
- DBToaster exhibits consistent performance and memory usage over time (Figures 8, 9, and 10).
- These results scale to longer streams (Figure 11).

In all figures, we use the following notation:

- **DBToaster** is the full HO-IVM algorithm.
- **REP** and **IVM** are DBToaster repeatedly re-evaluating queries, and emulating non-recursive IVM, respectively.
- **Naive** is a simplified form of the viewlet transform that aggressively materializes entire delta queries.
- **DBX-REP** and **DBX-IVM** are a commercial database system performing view maintenance by re-evaluation and non-recursive IVM respectively.
- **SPY** is a commercial stream processing engine.

DBToaster results are presented with both C++ and Scala as target languages.




Fig. 6 DBToaster performance overview. Note the log scale on the y-axis. (*) For VWAP, where DBToaster uses view caching, we compare against a strategy that avoids input variables.

9.1 Higher-Order IVM Performance

We now compare the performance of DBToaster with a commercial DBMS (DBX) and a stream processor (SPY).

Comparison with Commercial Systems. Figure 7 shows the performance of DBToaster’s Higher-Order IVM alongside all comparison systems. We summarize our findings, because an in-depth itemized breakdown of overheads is outside the scope of this article.

When recomputing the query results after each update (DBX-REP), DBX experienced view refresh rates between 0.08 and 972.22, with average and median values of 37.03 and 6. When using DBX’s support for IVM, however, view refresh rates dropped to between 0.14 and 2.94. This drop in performance when using IVM is counter-intuitive and prompted us to trace the execution of our program. DBX’s tracing utility revealed that most of the execution time was spent parsing several parametrized system queries used in the bookkeeping. As the amount of useful work to be performed after a single update is quite small, the time spent parsing those system queries ends up dominating the overall running time. Additionally, maintaining catalog information across many tables for high rate updates also substantially impacts latencies and throughput.

The performance gap between SPY and DBToaster is a result of the lack of support for IVM in SPY, and synchronization used to prevent the asynchronous system from producing inconsistent results. Due to the nature of the test

queries, we are unable to make use of SPY’s window semantics and are forced to use in-memory tables instead. Even though we use indexes on the in-memory tables wherever it makes sense, SPY seems to be unable to take full advantage of them in queries with complex predicates, contributing to poor performance, as exemplified in Q19.

Join-free queries. The simplest queries in our workload: Q1 (Figure 8a) and Q6, aggregate TPC-H’s Lineitem relation. As these queries involve only one relation, the first-order delta depends solely on the values being inserted or deleted.

The materialized view of Q6 stores a single aggregate value and has a constant update cost. Thus, the view refresh rates of the DBToaster, Naive, and IVM methods are almost identical. In all these cases, the generated Scala programs outperform the C++ programs. The REP compilation exhibits low refresh rates as it performs a complete scan over Lineitem upon every update. Unlike the other methods for which the memory overhead is negligible, REP requires a bounded amount of memory to store the set of active tuples.

Q1 evaluates multiple group-by aggregates over Lineitem. DBToaster treats these aggregates as separate AGCA expressions and maintains each individually. Since many of these share common subexpressions, duplicate view elimination and polynomial expansion are essential for achieving high view refresh rates. Consequently, although Q1 has substantially more aggregates than Q6 (8 vs. 1), the view refresh rate of the C++ code is only 30% lower. Because the result set contains a fixed number of tuples (based on the limited

Query	REP C++	REP Scala	DBX Rep	DBX IVM	SPY	DBToaster C++	DBToaster Scala	Naive C++	Naive Scala	IVM C++	IVM Scala
Q1	5.39	6.41	17.74	1.58	29.46	14,109.08	11,214.98	5,378.30	1,500.17	9.83	11.16
Q2	1.76	0.75	10.74	1.24	1.01	12,742.67	3,239.07	0.03	0.03	659.94	60.33
Q3	8.19	3.58	18.22	0.7	22.55	15,045.62	8,021.88	3.54	2.02	121.10	14.03
Q4	24.99	8.42	32.3	2.06	252.04	17,604.34	16,911.45	63.26	51.65	11,614.26	7,678.18
Q5	5.46	1.10	9.67	0.32	19.05	34.18	131.68	0.15	0.19	25.22	4.14
Q6	27.76	28.29	44.24	0.43	361.97	20,021.32	25,509.86	19,025.53	26,485.42	21,030.09	26,910.07
Q7	1.70	0.17	3.84	1.35	14.1	2.76	261.01	0.03	0.03	10.49	10.17
Q8	0.39	0.23	34.83	2.45	23.88	19.42	1,122.93	0.09	0.07	0.34	0.25
Q9	0.29	0.07	6.75	1.76	23.74	3,778.52	578.81	0.03	0.03	1.15	0.80
Q10	5.57	1.34	4.6	0.47	38.23	13,697.35	7,317.77	0.12	0.19	79.45	4.92
Q11	0.63	0.23	3.74	2.84	12.32	45.12	149.39	0.26	0.16	0.41	0.12
Q11a	6.77	1.60	1.57	0.86	5.28	28,060.43	17,315.17	15,164.32	3,506.81	29,152.02	8,808.10
Q12	2.77	2.06	24.71	1.62	74.47	17,440.02	9,353.63	20.68	23.17	14,900.63	7,576.58
Q13	0.19	0.12	9.64	2.94	10.9	3.69	22.37	1.26	0.95	0.16	0.10
Q14	3.47	0.99	39.6	1.56	464.88	15,953.58	28,047.48	199.45	152.52	3.60	1.01
Q15	0.13	0.12	2.49	1.93	6.14	2.07	3.10	–	–	0.12	0.08
Q16	3.25	2.01	2.94	1.87	8.82	713.07	843.56	0.20	0.37	3.28	2.00
Q17	8.07	3.40	11.77	2.1	19.64	16,456.54	25,408.05	0.76	0.66	14,918.33	1,373.52
Q17a	6.22	1.99	1.51	1.34	13.06	15,617.53	8,285.00	1.18	0.99	6,190.91	2,060.98
Q18	1.44	2.45	0.08	1.2	11.16	33.51	24.86	0.34	0.27	1.75	2.95
Q18a	3.91	2.09	0.6	1.31	7.42	18,725.53	10,085.95	0.11	0.09	107.89	7.38
Q19	0.20	0.06	23.84	1.42	0.57	83.98	236.54	28.44	81.93	69.82	187.78
Q20	48.28	13.91	7.16	1.19	33.63	1,586.76	5,427.83	0.85	0.91	3,553.37	502.33
Q21	5.17	1.65	8.72	1.33	14.72	3,703.25	3,782.58	0.29	0.55	189.88	8.50
Q22	0.27	0.39	36.05	1.6	58.22	201.72	742.50	6.76	0.47	0.33	0.47
Q22a	0.94	1.08	1.4	1.98	41.68	7,868.03	3,687.80	176.88	68.31	1.19	1.31
SSB4	2.42	1.13	3.43	0.51	16.92	2,877.63	1,039.36	0.03	0.02	64.10	4.15
AXF	3.63	3.66	5.62	1.32	6.91	23,817.13	5,764.69	2,168.17	779.99	15,808.05	5,677.57
BSP	3.31	2.99	6	1.61	5.18	23,040.81	3,673.54	192.16	191.56	1,261.01	703.86
BSV	3.28	2.87	5.23	1.55	10.63	90,116.98	77,797.66	54,810.63	6,921.83	1,284.29	702.49
MST	3.59	2.07	4.37	1.26	3.73	5.81	3.81	5.90	5.88	1.52	2.06
PSP	2.90	2.68	5.93	1.96	5.38	7,319.93	2,658.12	362.82	365.24	2.87	2.67
VWAP	4.42	3.08	7.93	2.12	4.81	2,649.42	2,087.87	2,436.26	2,692.96	4.36	2.75
MDDDB1	5.62	1.54	972.22	1.02	5.96	29,842.28	48,784.54	–	–	9,163.44	230.31
MDDDB2	3.42	1.14	0.31	0.26	2.11	6,093.05	189.56	0.02	0.03	3,656.16	131.68

Fig. 7 Comparison between DBToaster and two commercial query engines (in view refreshes per second). Both the DBMS (DBX) and stream system (SPY) columns show the cost of full refresh on each update. Higher numbers are better.

domain of the group-by columns), DBToaster uses only a fixed amount of memory to store the additional maps.

DBToaster inlines the computation of algebraic aggregates. For instance, DBToaster computes an average aggregate as sum and count aggregates: Because the current incarnation of AGCA supports only one “multiplicity” per tuple, average is expressed as the product of the sum and inverse count. HO-IVM requires two recursive steps to separate out the (linear) count from the (non-linear) inverse count. This accounts for IVM’s poor performance on Q1, as it must fully recompute the inverse count on every change. As future work, we plan to extend AGCA to generalize GMRs to have multiple “multiplicities”. This will allow DBToaster to store multiple aggregate values per tuple, and improve the efficiency of this class of queries.

Equijoins. Q11a, Q12, Q14, and Q19 (Figure 9c) contain two-way joins without nested aggregates. The first level deltas correspond nearly to the base relations. For Q11a and Q12, DBToaster and IVM compilations produce virtually identical results. IVM has a slightly lower refresh rate for

Q19 than DBToaster, because it materializes entire base relations, while DBToaster projects away unused columns. As in Q1, Q14 has to maintain an inverse count, resulting in poor performance for IVM. In Naive, range restrictions are not extracted from deltas of nested aggregate expressions (Section 5.3), necessitating a full scan of each materialized nested aggregate whenever it changes. The effect of this optimization is most evident in these four queries.

Query decomposition also plays an important role in efficiency of DBToaster for queries containing linear joins of 3 or more relations. Decomposition avoids materialization of cross products, improving performance and reducing memory consumption. For instance, the delta of Q10 (a 4-way equijoin) with respect to the Orders relation creates a cross product between Customer and Lineitem (which are only connected through Orders in the original query). In Naive, the entire cross product is materialized, resulting in performance five orders of magnitude worse.

Due to foreign key constraints in the TPC-H schema (of which DBToaster is not made aware) most loops in Q3’s trigger program have only one iteration, and the cost of up-

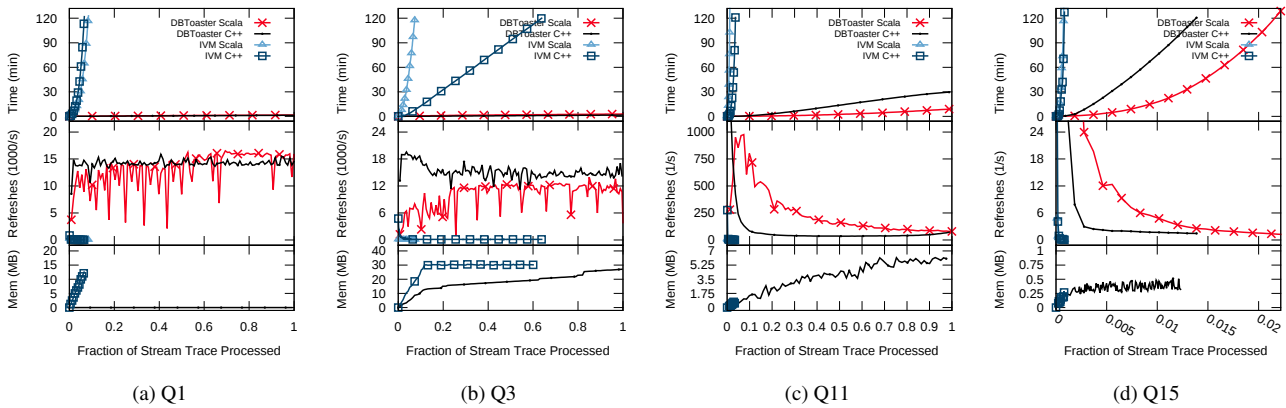


Fig. 8 (a) A join-free query. (b) A 3-way linear join. (c) A 2-way join with an aggregate subquery in the FROM clause and an uncorrelated nested aggregate. (d) A 2-way join with an aggregate subquery in the FROM clause and an inequality-correlated nested aggregate in the EXIST clause. DBToaster completes only a small fraction of the trace since the update cost grows quadratically with the number of distinct suppliers.

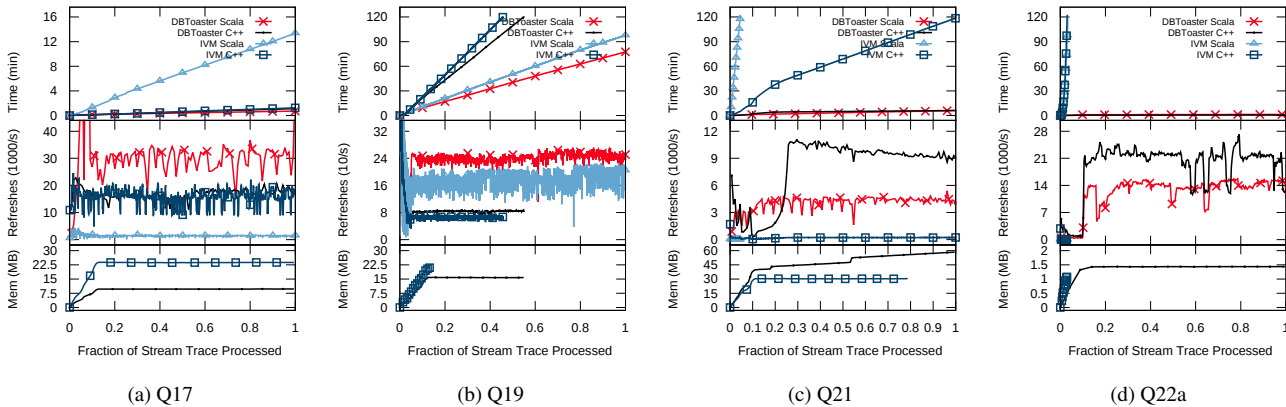


Fig. 9 (a) A 2-way join with an equality-correlated nested aggregate. (b) A 2-way join with three disjunctive clauses. (c) A 4-way join with an equality- and an inequality-correlated subqueries. (d) A single relation with an equality- and an inequality-correlated nested aggregates. Insertions into the Customer relation complete within the first 10% of the stream.

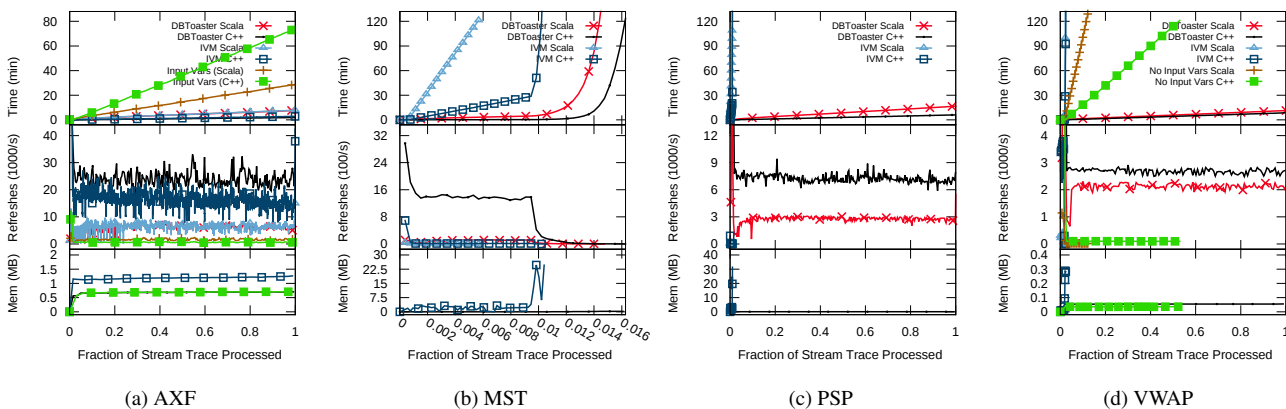


Fig. 10 (a) A 2-way inequality join. DBToaster outperforms view caching due to the large domain of the input variables. (b) A 2-way join with two uncorrelated, and two inequality-correlated nested aggregates. None of the tested engines completed the trace within the 2-hour limit. (c) A 2-way join with two uncorrelated nested aggregates. (d) A single relation with an inequality-correlated and an uncorrelated nested aggregate. DBToaster chooses the view cache method, so we compare against an approach that aggressively avoids input variables.

dating either the Orders or Lineitem relation is constant. For queries with multi-way joins and selection predicates (Q3, Q5, Q10, SSB4, MDDDB1, and MDDDB2), DBToaster further outperforms IVM by pushing predicates into the materialized views and projecting away unused columns.

DBToaster considers the contents of Nation, Region, and all the scientific relations except AtomPositions as static. It loads static relations into memory before processing the streams. It avoids materialization of deltas needed to support updates to these relations, effectively reducing the join width of certain queries (Q5, Q10, SSB4, MDDDB1, and MDDDB2) and eliminating several potentially high maintenance maps.

Nested Aggregates. Q17 (Figure 9a), Q17a, and Q18a³ are multi-way join queries with nested aggregates that are correlated on an equality with the outer query. In these cases, DBToaster’s strong performance comes from decorrelating the nested subquery and range-restricting the domain of the generated delta expressions for updates to the Lineitem relation (on which all nested subqueries are based).

Q22a (Figure 9d) includes two nested aggregates, an uncorrelated aggregate on Customer and an equality-correlated aggregate on Orders. The first subquery causes DBToaster to re-evaluate the top-level query (as per Section 5.1) since the delta with respect to updates to Customer is not simpler than the whole query. This re-evaluation strategy iterates over the whole Customer relation (materialized with only necessary columns) to compute the top-level aggregate for the customers with no orders. In contrast, DBToaster uses an incremental strategy for updates into Orders, since the equality-based correlation between the second subquery and the outer query restricts the domain of the corresponding delta expression, making the update cost constant. Therefore, Q22a has two different maintenance costs for updates into the base relations. This is seen in the performance graph as the query’s slow startup ends once the last customer has been inserted.

VWAP (Figure 10d) has a nested aggregate correlated on an inequality. The small domain of the correlation variable (price) makes this an ideal candidate for view caching. The performance graph shows the benefit of view caching over avoiding the materialization of maps with input variables.

PSP (Figure 10c) includes two uncorrelated nested aggregates. It benefits from top-level query re-evaluation on each update. As in Section 6.2, polynomial expansion and graph decomposition are essential to avoid computation of a cross product between the base relations. DBToaster evaluates the query using six auxiliary materialized views with constant time updates: Two views maintain single aggregate values, while the others are linear in the number of distinct values of the column being compared to the nested aggregate (volume). The finite domain of these values results in nearly constant view refresh rate and memory consumption.

MST (Figure 10b) is fundamentally similar to PSP, but rather than comparing its uncorrelated aggregates against columns from the base relations, they are each compared against another nested aggregate correlated on an inequality. This is a worst case scenario for DBToaster, as it cannot incrementally process this query in better than $O(n^2)$ time without specialized indexes (e.g., aggregate range trees).

Inequijoins. AXF (Figure 10a) and BSP are 2-way joins with inequality join-predicates. The performance graph of AXF shows the inefficiency of view caching in this case. The view caching approach treats both the join variable (price) and one of the aggregate variables (volume) as input variables; together, these input variables have an extremely large domain. In BSP, the join variable (timestamp) also has an unbounded domain. In both cases, DBToaster outperforms view caching by precluding materialized views with input variables. DBToaster also achieves a small speed boost compared to IVM by not materializing the entire base relation.

Queries with EXIST or IN clauses. Q2, Q4, Q16, and Q21 (Figure 9c) contain clauses that check for the existence of the nested subquery results. DBToaster transforms each subquery into a count aggregate, assigns this value to a fresh variable, and adds an additional constraint over that variable according to the semantics of the clause (e.g., $x = 0$ for the NOT EXIST clause). As all the subqueries of the above queries are correlated on an equality, DBToaster decides to incrementally maintain the top-level views for updates to the subquery relations. For queries that are also correlated on an inequality (Q2 and Q21), DBToaster avoids materializing maps with input variables due to the large domain of the correlation variables (supplycost and suppley, respectively). Q21 has constant time updates to Lineitem and Orders, and a linear time update in the number of orders of one supplier. The higher update cost results in a lower view refresh rate within the first 20% of the stream, until insertions into the Supplier relation complete.

Subqueries in FROM clauses. DBToaster maintains separate materialized views for subqueries that appear in the FROM clause (Q7, Q8, and Q9). For Q7 and Q8, we observe that the C++ backend fails to transform DBToaster’s functional representation into efficient procedural code, causing huge memory overheads and poor performance. In contrast, DBToaster derives Scala code directly from its internal representation. The Scala compiler further optimizes the code, resulting in performance better by two orders of magnitude.

Complex queries. The remaining TPC-H queries Q11, Q13, Q15, Q18, Q20, and Q22 combine the above characteristics. Our experiments show that the update costs for these queries coincide with their structural complexity.

Q11 (Figure 8c) has a group-by aggregate in its FROM clause and an uncorrelated nested aggregate that is compared against on the top level using an inequality. DBToaster exploits the fact that both subqueries share the same struc-

³ Q17, Q17a produce incorrect results due to floating point errors.

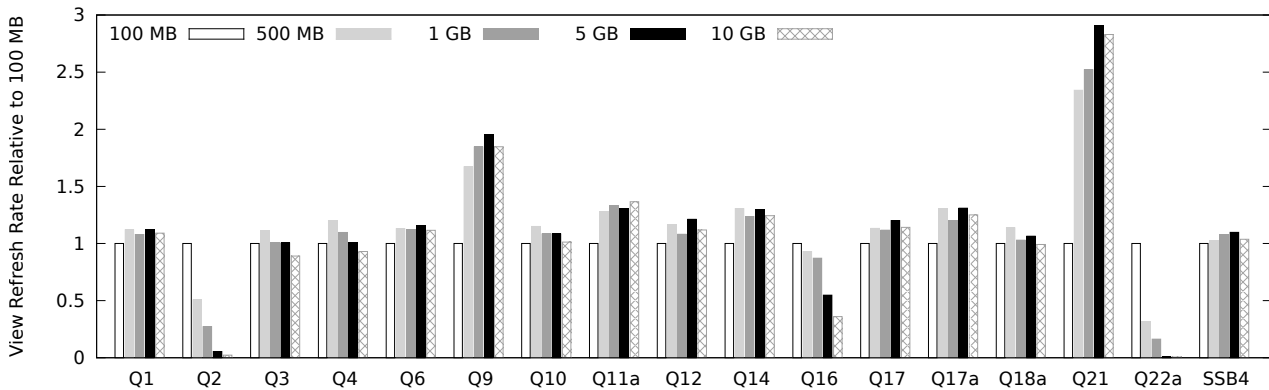


Fig. 11 Performance scaling on a subset of TPC-H queries.

ture to reduce the number of generated maps. Since the update stream contains only insertions to the base relations, the amount of memory used to store additional views grows continuously. The costs of updating Supplier and Partsupp are linear in the number of distinct partkey values. Thus, the view refresh rate levels off as the number of tuples in the materialized views reaches the maximum number of distinct group-by values.

Q15 (Figure 8d) departs from the original TPC-H definition in that it replaces the max aggregate with an equivalent nested subquery and an EXIST clause. Since both subqueries are identical, duplicate view elimination reduces the number of auxiliary views. However, the update cost for this query grows quadratically with the number of distinct supkey values in Lineitem, as shown on the graph. In order to improve performance of MIN, MAX, and theta-joins in general, we plan to extend DBToaster with specialized tree-based data structures.

9.2 Stream Scalability

This section analyzes the scaling behavior of DBToaster for a subset of the TPC-H queries over a larger stream of updates. Our focus is on measuring view refresh rates in terms of the stream length and query complexity, rather than the working set size.

The workload for this experiment was synthesized from databases created by DBGEN at scaling factors 0.5, 1, 5, and 10 (500MB, 1GB, 5GB, and 10GB, respectively). An update stream was built by randomly interleaving tuples from the base relations, while preserving the reference integrity. As before, after inserting 30 thousand Orders tuples, respectively 120 thousand Lineitem tuples, we started to randomly inject deletions into these two relations in order to keep their sizes roughly constant. Tuples of other TPC-H relations are inserted and never deleted.

The length of the update stream increases with larger scaling factors. However, the size of the working set depends on the query structure. Materialized views that refer-

ence Customer, Part, Supplier, or Partsupp might grow with larger scaling factors, while views defined solely over Orders or Lineitem have a bounded working set size.

Figure 11 presents the results of our scaling experiments. For most queries performance stays roughly constant as the stream length grows. Q2 and Q16 select over insert-only relations (Part, Supplier, and Partsupp); thus, the memory overhead of DBToaster grows with the scaling factors. The view refresh rates drop as the maintenance cost for these queries is linear in the size of in-memory data structures. In contrast, Q11a also queries insert-only relations, but exhibits good scaling behavior due to the cardinality constraints between its base relations and use of index data structures. The running time of Q22a is dominated by the first 10% of the stream in all cases, before the Customer relation has been fully inserted. The cost of inserting a new customer is linear in the size of the Customer relation. After all customer tuples have been processed, the refresh rate increases to a constant 8000 tuples per second, regardless of scale.

Q9 and Q21 demonstrate an increase of the view refresh rates for larger stream lengths. The reason for this behavior is as follows. In our workload, the working set sizes of Orders and Lineitem are irrespective of the scale. With larger scaling factors the base relations get larger; thus, we have to place more deletions to maintain the size invariant (As an extreme case, imagine that the working set size of Orders is 1; then we have to double the number of Orders tuples in the stream as every insertion is followed by a deletion). Placing more deletions increases the fraction of Orders and Lineitem tuples in the stream. This in turn affects the view refresh rates of these queries, as both have constant costs with respect to updates to the Lineitem relation.

10 Conclusion

We presented DBToaster, a compiler and optimizer framework for higher-order IVM that uses aggressive simplification of recursive delta queries and a plethora of materialization strategies to make recursive IVM viable. Our compila-

tion method is effective on a wide range of select-project-join-aggregate queries, including those with nested subqueries, which are unsupported by current IVM mechanisms. Our methods provide view refresh rates that often improve on today's tools by several orders of magnitude.

Acknowledgements This work was supported by ERC Grant 279804. We thank Prof. Tom Woolf and Dr. Sarana Nutanong at the Johns Hopkins University Departments of Physiology and Computer Science respectively for their assistance with the MDDB workload.

References

- DBToaster Public Beta revision 2827, Feb. 11, 2013. <http://www.dbtoaster.org/index.php?page=download>
- Abadi, D., Ahmad, Y., Balazinska, M., Cetintemel, U., Cherniack, M., Hwang, J., Lindner, W., Maskey, A., Rasin, A., Ryvkina, E., et al.: The design of the Borealis stream processing engine. In: CIDR, pp. 277–289 (2005)
- Agrawal, S., Chaudhuri, S., Narasayya, V.R.: Automated selection of materialized views and indexes in SQL databases. In: VLDB, pp. 496–505 (2000)
- Ahmad, Y., Koch, C.: DBToaster: A SQL compiler for high-performance delta processing in main-memory databases compiler for high-performance delta processing in main-memory databases. PVLDB **2**(2), 1566–1569 (2009)
- Aiken, A., Hellerstein, J.M., Widom, J.: Static analysis techniques for predicting the behavior of active database rules. ACM TODS **20**(1), 3–41 (1995)
- Aji, S.M., McEliece, R.J.: The generalized distributive law. IEEE Trans. Inf. Theory **46**(2), 325–343 (2000)
- Blakeley, J.A., Larson, P.Å., Tompa, F.W.: Efficiently updating materialized views. In: SIGMOD, pp. 61–71 (1986)
- Buneman, P., Clemons, E.K.: Efficiently monitoring relational databases. ACM TODS **4**(3), 368–382 (1979)
- Buneman, P., Naqvi, S.A., Tannen, V., Wong, L.: Principles of programming with complex objects and collection types. Theor. Comput. Sci. **149**(1), 3–48 (1995)
- Chaudhuri, S., Krishnamurthy, R., Potamianos, S., Shim, K.: Optimizing queries with materialized views. In: ICDE, pp. 190–200 (1995)
- Chirkova, R., Yang, J.: Materialized views. Foundations and Trends in Databases **4**(4), 295–405 (2012)
- Colby, L.S., Griffin, T., Libkin, L., Mumick, I.S., Trickey, H.: Algorithms for deferred view maintenance. In: SIGMOD, pp. 469–480 (1996)
- Colby, L.S., Kawaguchi, A., Lieuwen, D.F., Mumick, I.S., Ross, K.A.: Supporting multiple view maintenance policies. In: SIGMOD, pp. 405–416 (1997)
- Cormode, G., Muthukrishnan, S.: What's hot and what's not: Tracking most frequent items dynamically. ACM TODS **30**(1), 249–278 (2005)
- Ghanem, T.M., Elmagarmid, A.K., Larson, P.Å., Aref, W.G.: Supporting views in data stream management systems. ACM TODS **35**(1), 1–47 (2010)
- Griffin, T., Libkin, L.: Incremental maintenance of views with duplicates. In: SIGMOD, pp. 328–339 (1995)
- Gupta, A., Mumick, I.S., Subrahmanian, V.S.: Maintaining views incrementally. In: SIGMOD, pp. 157–166 (1993)
- Gupta, H., Mumick, I.S.: Selection of views to materialize in a data warehouse. IEEE TKDE **17**(1), 24–43 (2005)
- Kawaguchi, A., Lieuwen, D.F., Mumick, I.S., Ross, K.A.: Implementing incremental view maintenance in nested data models. In: DBPL, pp. 202–221 (1997)
- Kearns, M., Ortiz, L.: The penn-lehman automated trading project. IEEE Intelligent Systems **18**(6), 22–31 (2003)
- Kennedy, O., Ahmad, Y., Koch, C.: DBToaster: Agile views for a dynamic data management system. In: CIDR, pp. 284–295 (2011)
- Koch, C.: Incremental query evaluation in a ring of databases. In: PODS, pp. 87–98 (2010)
- Koch, C.: Incremental query evaluation in a ring of databases (2013). Technical Report EPFL-REPORT-183766, <https://infoscience.epfl.ch/record/183766>
- Koch, C., Ahmad, Y., Kennedy, O., Nikolic, M., Nötzli, A., Lupei, D., Shaikhha, A.: Dbtoaster: Higher-order delta processing for dynamic, frequently fresh views (2013). Technical report EPFL-REPORT-183767, extends this article by an appendix that lists the full query workload as well as experimental parameters and trace figures that did not find space in this article; <http://infoscience.epfl.ch/record/183767>
- Kotidis, Y., Roussopoulos, N.: A case for dynamic view management. ACM TODS **26**(4), 388–423 (2001)
- Krikellas, K., Viglas, S., Cintra, M.: Generating code for holistic query evaluation. In: ICDE (2010)
- Krishnamurthy, S., Wu, C., Franklin, M.J.: On-the-fly sharing for streamed aggregation. In: SIGMOD, pp. 623–634 (2006)
- Larson, P.Å., Zhou, J.: Efficient maintenance of materialized outer-join views. In: ICDE, pp. 56–65 (2007)
- Liu, Y.A., Stoller, S.D., Teitelbaum, T.: Static caching for incremental computation. ACM TOPLAS **20**(3), 546–585 (1998)
- Marlow, S., Wadler, P.: Deforestation for higher-order functions. In: Functional Programming, pp. 154–165 (1992)
- Motwani, R., Widom, J., Arasu, A., Babcock, B., Babu, S., Datar, M., Manku, G.S., Olston, C., Rosenstein, J., Varma, R.: Query processing, approximation, and resource management in a data stream management system. In: CIDR (2003)
- Neumann, T.: Efficiently compiling efficient query plans for modern hardware. PVLDB **4**(9), 539–550 (2011)
- Palpanas, T., Sidle, R., Cochrane, R., Pirahesh, H.: Incremental maintenance for non-distributive aggregate functions. In: VLDB, pp. 802–813 (2002)
- Pearlmuter, B.A., Siskind, J.M.: Lazy multivariate higher-order forward-mode AD. In: POPL, pp. 155–160 (2007)
- Ross, K.A., Srivastava, D., Sudarshan, S.: Materialized view maintenance and integrity constraint checking: Trading space for time. In: SIGMOD, pp. 447–458 (1996)
- Roussopoulos, N.: An incremental access method for ViewCache: Concept, algorithms, and cost analysis. ACM TODS **16**(3), 535–563 (1991)
- Salem, K., Beyer, K.S., Cochrane, R., Lindsay, B.G.: How to roll a join: Asynchronous incremental view maintenance. In: SIGMOD, pp. 129–140 (2000)
- Shyamshankar, P., Palmer, Z., Ahmad, Y.: K3: Language design for building multi-platform, domain-specific runtimes. In: First International Workshop on Cross-model Language Design and Implementation (XLDI) (2012)
- Tatbul, N., Çetintemel, U., Zdonik, S.B., Cherniack, M., Stonebraker, M.: Load shedding in a data stream manager. In: VLDB, pp. 309–320 (2003)
- Transaction Processing Performance Council: TPC-H benchmark specification. <http://www.tpc.org/hspec.html> (2011)
- Wong, L.: Kleisli, a functional query system. J. Funct. Program. **10**(1), 19–56 (2000)
- Yang, J., Widom, J.: Incremental computation and maintenance of temporal aggregates. VLDB J. **12**(3), 262–283 (2003)
- Zhou, J., Larson, P.Å., Elmongui, H.G.: Lazy maintenance of materialized views. In: VLDB, pp. 231–242 (2007)
- Zhou, J., Larson, P.Å., Freytag, J.C., Lehner, W.: Efficient exploitation of similar subexpressions for query processing. In: SIGMOD, pp. 533–544 (2007)
- Zilio, D.C., Zuzarte, C., Lightstone, S., Ma, W., Lohman, G.M., Cochrane, R., Pirahesh, H., Colby, L.S., Gryz, J., Alton, E., Liang, D., Valentin, G.: Recommending materialized views and indexes with IBM DB2 design advisor. In: ICAC, pp. 180–188 (2004)

Appendix A Workload Queries

In the following we list our workload queries, excluding the standard TPC-H queries that we ran with only shallow modifications, as detailed in section 8.

Appendix A.1 TPC-H schema based queries

Q11a SELECT ps.partkey,
SUM(ps.supplycost * ps.availqty) AS query11a
FROM Partsupp ps, Supplier s
WHERE ps.suppkey = s.suppkey
GROUP BY ps.partkey;

Q17a SELECT SUM(l.extendedprice) AS query17a
FROM Lineitem l, Part p
WHERE p.partkey = l.partkey
AND l.quantity < 0.005 *
(SELECT SUM(l2.quantity)
FROM Lineitem l2 WHERE l2.partkey = p.partkey);

Q18a SELECT c.custkey, SUM(l1.quantity) AS query18a
FROM Customer c, Orders o, Lineitem l1
WHERE 1 <=
(SELECT SUM(l1) FROM Lineitem l2
WHERE l1.orderkey = l2.orderkey
AND 100 < (SELECT SUM(l3.quantity) FROM Lineitem l3
WHERE l2.orderkey = l3.orderkey))
AND c.custkey = o.custkey
AND o.orderkey = l1.orderkey
GROUP BY c.custkey;

Q22a SELECT c1.nationkey, SUM(c1.acctbal) AS query22a
FROM Customer c1
WHERE c1.acctbal < (SELECT SUM(c2.acctbal)
FROM Customer c2
WHERE c2.acctbal > 0)
AND 0 = (SELECT SUM(1)
FROM Orders o
WHERE o.custkey = c1.custkey)
GROUP BY c1.nationkey;

SSB4 SELECT sn.regionkey, cn.regionkey, p.type,
SUM(li.quantity)
FROM Customer c, Orders o, Lineitem li,
Part p, Supplier s, Nation cn, Nation sn
WHERE c.custkey = o.custkey
AND o.orderkey = li.orderkey
AND p.partkey = li.partkey
AND s.suppkey = li.suppkey
AND o.orderdate >= DATE('1997-01-01')
AND o.orderdate < DATE('1998-01-01')
AND cn.nationkey = c.nationkey
AND sn.nationkey = s.nationkey
GROUP BY sn.regionkey, cn.regionkey, p.type;

Appendix A.2 Financial workload

AXE SELECT b.broker_id, SUM(a.volume-b.volume)
FROM Bids b, Asks a
WHERE b.broker_id = a.broker_id
AND (a.price-b.price > 1000 OR b.price-a.price > 1000)
GROUP BY b.broker_id;

BSV SELECT x.broker_id, SUM(x.volume*x.price - y.volume*y.price)
FROM Bids x, Bids y
WHERE x.broker_id=y.broker_id AND x.t>y.t
GROUP BY x.broker_id;

BSV SELECT x.broker_id, SUM(x.volume*x.price*y.volume*y.price*0.5)
FROM Bids x, Bids y
WHERE x.broker_id = y.broker_id
GROUP BY x.broker_id;

MST SELECT b.broker_id, SUM(a.price*a.volume - b.price*b.volume)
FROM Bids b, Asks a
WHERE 0.25*(SELECT SUM(a1.volume) FROM Asks a1) >
(SELECT SUM(a2.volume) FROM Asks a2 WHERE a2.price>a.price)
AND 0.25*(SELECT SUM(b1.volume) FROM Bids b1) >
(SELECT SUM(b2.volume) FROM Bids b2 WHERE b2.price>b.price)
GROUP BY b.broker_id;

PSP SELECT SUM(a.price - b.price)
FROM Bids b, Asks a
WHERE b.volume>0.0001*(SELECT SUM(b1.volume) FROM Bids b1)
AND a.volume>0.0001*(SELECT SUM(a1.volume) FROM Asks a1);

VWAP SELECT SUM(b1.price * b1.volume)
FROM Bids b1
WHERE 0.25 * (SELECT SUM(b3.volume) FROM Bids b3) >
(SELECT SUM(b2.volume) FROM Bids b2
WHERE b2.price>b1.price);

Appendix A.3 Scientific workload

MDDBI SELECT p.trj_id, p.t,
AVG(vec_length(p.x-p2.x, p.y-p2.y, p.z-p2.z)) AS rdf
FROM AtomPositions p, AtomMeta m,
AtomPositions p2, AtomMeta m2
WHERE p.trj_id = p2.trj_id
AND p.t = p2.t
AND p.atom_id = m.atom_id
AND p2.atom_id = m2.atom_id
AND m.residue_name = 'LYS'
AND m.atom_name = 'NZ'
AND m2.residue_name = 'TIP3'
AND m2.atom_name = 'OH2'
GROUP BY p.trj_id, p.t;

MDDB2 SELECT p1.trj_id, p1.t,
dihedral_angle(p1.x,p1.y,p1.z,
p2.x,p2.y,p2.z,
p3.x,p3.y,p3.z,
p4.x,p4.y,p4.z)
AS phi_psi,
dm.dim_id
FROM Dihedrals d, Dimensions dm,
AtomPositions p1, AtomPositions p2,
AtomPositions p3, AtomPositions p4,
AtomMeta m1, AtomMeta m2,
AtomMeta m3, AtomMeta m4
WHERE p1.t = p2.t
AND p1.t = p3.t
AND p1.t = p4.t
AND p1.trj_id = p2.trj_id
AND p1.trj_id = p3.trj_id
AND p1.trj_id = p4.trj_id
AND d.atom_id1 = m1.atom_id AND m1.atom_id = p1.atom_id
AND d.atom_id2 = m2.atom_id AND m2.atom_id = p2.atom_id
AND d.atom_id3 = m3.atom_id AND m3.atom_id = p3.atom_id
AND d.atom_id4 = m4.atom_id AND m4.atom_id = p4.atom_id
AND d.atom_id1 = dm.atom_id1 AND d.atom_id2 = dm.atom_id2
AND d.atom_id3 = dm.atom_id3 AND d.atom_id4 = dm.atom_id4
AND ((m1.atom_name = 'N' AND m2.atom_name = 'CA' AND
m3.atom_name = 'C')
OR (m2.atom_name = 'N' AND m3.atom_name = 'CA' AND
m4.atom_name = 'C')));

Appendix B TPC-H Queries

In the following we list the TPC-H queries, as we ran them, with the small changes summarized in section 8.

```

SELECT returnflag, linestatus,
  SUM(quantity) AS sum_qty,
  SUM(extendedprice) AS sum_base_price,
  SUM(extendedprice * (1-discount)) AS sum_disc_price,
  SUM(extendedprice * (1-discount)*(1+tax)) AS sum_charge,
  AVG(quantity) AS avg_qty,
  AVG(extendedprice) AS avg_price,
  AVG(discount) AS avg_disc,
  COUNT(*) AS count_order
FROM Lineitem
WHERE shipdate <= DATE('1997-09-01')
GROUP BY returnflag, linestatus;

SELECT s.acctbal, s.name, n.name, p.partkey, p.mfgr,
  s.address, s.phone, s.comment
FROM Part p, Supplier s, Partsupp ps,
  Nation n, Region r
WHERE p.partkey = ps.partkey
AND s.suppkey = ps.suppkey
AND p.size = 15
AND (p.type LIKE '%BRASS')
AND s.nationkey = n.nationkey
AND n.regionkey = r.regionkey
AND r.name = 'EUROPE'
AND (NOT EXISTS
  (SELECT 1
   FROM Partsupp ps2, Supplier s2,
        Nation n2, Region r2
   WHERE p.partkey = ps2.partkey
        AND s2.suppkey = ps2.suppkey
        AND s2.nationkey = n2.nationkey
        AND n2.regionkey = r2.regionkey
        AND r2.name = 'EUROPE'
        AND ps2.supplycost < ps.supplycost));

SELECT o.orderkey,
  o.orderdate,
  o.shippriority,
  SUM(extendedprice * (1 - discount)) AS query3
FROM Customer c, Orders o, Lineitem l
WHERE c.mktsegment = 'BUILDING'
AND o.custkey = c.custkey
AND l.orderkey = o.orderkey
AND o.orderdate < DATE('1995-03-15')
AND l.shipdate > DATE('1995-03-15')
GROUP BY o.orderkey, o.orderdate, o.shippriority;

SELECT o.orderpriority, COUNT(*) AS order_count
FROM Orders o
WHERE o.orderdate >= DATE('1993-07-01')
AND o.orderdate < DATE('1993-10-01')
AND (EXISTS (
  SELECT * FROM Lineitem l
  WHERE l.orderkey = o.orderkey
  AND l.commitdate < l.receiveptdate
))
GROUP BY o.orderpriority;

SELECT n.name,
  SUM(l.extendedprice * (1 - l.discount)) AS revenue
FROM Customer c, Orders o, Lineitem l, Supplier s,
  Nation n, Region r
WHERE c.custkey = o.custkey
AND l.orderkey = o.orderkey
AND l.suppkey = s.suppkey
AND c.nationkey = s.nationkey
AND s.nationkey = n.nationkey
AND n.regionkey = r.regionkey
AND r.name = 'ASIA'
AND o.orderdate >= DATE('1994-01-01')
AND o.orderdate < DATE('1995-01-01')
GROUP BY n.name;

SELECT SUM(l.extendedprice*l.discount) AS revenue
FROM Lineitem l
WHERE l.shipdate >= DATE('1994-01-01')
AND l.shipdate < DATE('1995-01-01')
AND (l.discount BETWEEN (0.06 - 0.01) AND (0.06 + 0.01))
AND l.quantity < 24;

SELECT supp_nation, cust_nation, l_year,
  SUM(volume) AS revenue
FROM (
  SELECT n1.name AS supp_nation,
        n2.name AS cust_nation,
        EXTRACT(year from l.shipdate) AS l_year,
        l.extendedprice * (1 - l.discount) AS volume
  FROM Supplier s, Lineitem l, Orders o, Customer c,
        Nation n1, Nation n2
  WHERE s.suppkey = l.suppkey
  AND o.orderkey = l.orderkey
  AND c.custkey = o.custkey
  AND s.nationkey = n1.nationkey
  AND c.nationkey = n2.nationkey
  AND ((n1.name = 'FRANCE' AND n2.name = 'GERMANY')
  OR
        (n1.name = 'GERMANY' AND n2.name = 'FRANCE'))
  AND (l.shipdate BETWEEN DATE('1995-01-01')
        AND DATE('1996-12-31'))
) AS shipping
GROUP BY supp_nation, cust_nation, l_year;

SELECT total.o_year,
  (SUM(CASE total.name WHEN 'BRAZIL'
    THEN total.volume
    ELSE 0 END) /
  LISTMAX(1, SUM(total.volume))) AS mkt_share
FROM (
  SELECT n2.name,
    EXTRACT(year from o.orderdate) AS o_year,
    l.extendedprice * (1-l.discount) AS volume
  FROM Part p, Supplier s, Lineitem l, Orders o,
    Customer c, Nation n1, Nation n2, Region r
  WHERE p.partkey = l.partkey
  AND s.suppkey = l.suppkey
  AND l.orderkey = o.orderkey
  AND o.custkey = c.custkey
  AND c.nationkey = n1.nationkey
  AND n1.regionkey = r.regionkey
  AND r.name = 'AMERICA'
  AND s.nationkey = n2.nationkey
  AND (o.orderdate BETWEEN DATE('1995-01-01')
    AND DATE('1996-12-31'))
  AND p.type = 'ECONOMY ANODIZED STEEL'
) total
GROUP BY total.o_year;

SELECT nation, o_year, SUM(amount) AS sum_profit
FROM (
  SELECT n.name AS nation,
    EXTRACT(year from o.orderdate) AS o_year,
    ((l.extendedprice * (1 - l.discount)) -
    (ps.supplycost * l.quantity)) AS amount
  FROM Part p, Supplier s, Lineitem l, Partsupp ps,
    Orders o, Nation n
  WHERE s.suppkey = l.suppkey
  AND ps.suppkey = l.suppkey
  AND ps.partkey = l.partkey
  AND p.partkey = l.partkey
  AND o.orderkey = l.orderkey
  AND s.nationkey = n.nationkey
  AND (p.name LIKE '%green%')
) AS profit
GROUP BY nation, o_year;

```

```

Q10 SELECT c.custkey, c.name,
      c.acctbal,
      n.name,
      c.address,
      c.phone,
      c.comment,
      SUM(l.extendedprice * (1 - l.discount)) AS revenue
FROM Customer c, Orders o, Lineitem l, Nation n
WHERE c.custkey = o.custkey
AND l.orderkey = o.orderkey
AND o.orderdate >= DATE('1993-10-01')
AND o.orderdate < DATE('1994-01-01')
AND l.returnflag = 'R'
AND c.nationkey = n.nationkey
GROUP BY c.custkey, c.name, c.acctbal, c.phone,
        n.name, c.address, c.comment;

Q11 SELECT p.partkey, SUM(p.value) AS QUERY11
FROM
  (
    SELECT ps.partkey,
          SUM(ps.supplycost * ps.availqty) AS value
    FROM Partsupp ps, Supplier s, Nation n
    WHERE ps.suppkey = s.suppkey
    AND s.nationkey = n.nationkey
    AND n.name = 'GERMANY'
    GROUP BY ps.partkey
  ) p
WHERE p.value > (
  SELECT SUM(ps.supplycost * ps.availqty) * 0.001
  FROM Partsupp ps, Supplier s, Nation n
  WHERE ps.suppkey = s.suppkey
  AND s.nationkey = n.nationkey
  AND n.name = 'GERMANY'
)
GROUP BY p.partkey;

Q12 SELECT l.shipmode,
      SUM(CASE WHEN o.orderpriority IN ('1-URGENT',
                                       '2-HIGH')
              THEN 1 ELSE 0 END) AS high_line_count,
      SUM(CASE WHEN o.orderpriority NOT IN ('1-URGENT',
                                           '2-HIGH')
            THEN 1 ELSE 0 END) AS low_line_count
FROM Orders o, Lineitem l
WHERE o.orderkey = l.orderkey
AND (l.shipmode IN ('MAIL', 'SHIP'))
AND l.commitdate < l.receiveptdate
AND l.shipdate < l.commitdate
AND l.receiveptdate >= DATE('1994-01-01')
AND l.receiveptdate < DATE('1995-01-01')
GROUP BY l.shipmode;

Q13 SELECT c_count, COUNT(*) AS custdist
FROM (
  SELECT c.custkey AS c_custkey,
        COUNT(o.orderkey) AS c_count
  FROM Customer c, Orders o
  WHERE c.custkey = o.custkey
  AND (o.comment NOT LIKE '%special%requests%')
  GROUP BY c.custkey
) c_orders
GROUP BY c_count;

Q14 SELECT (100.00 *
          SUM(CASE WHEN (p.type LIKE 'PROMO')
                    THEN l.extendedprice * (1 - l.discount)
                    ELSE 0 END) /
          LISTMAX(1,
                SUM(l.extendedprice * (1 - l.discount)))
        ) AS promo_revenue
FROM Lineitem l, Part p
WHERE l.partkey = p.partkey
AND l.shipdate >= DATE('1995-09-01')
AND l.shipdate < DATE('1995-10-01');

```

```

Q15 SELECT s.suppkey, s.name, s.address, s.phone,
      r1.total_revenue AS total_revenue
FROM Supplier s,
  (SELECT l.suppkey AS supplier_no,
        SUM(l.extendedprice * (1 - l.discount))
        AS total_revenue
   FROM Lineitem l
   WHERE l.shipdate >= DATE('1996-01-01')
   AND l.shipdate < DATE('1996-04-01')
   GROUP BY l.suppkey) r1
WHERE s.suppkey = r1.supplier_no
AND (NOT EXISTS
  (SELECT 1
   FROM (SELECT l.suppkey,
              SUM(l.extendedprice *
                  (1 - l.discount))
              AS total_revenue
        FROM Lineitem l
        WHERE l.shipdate >= DATE('1996-01-01')
        AND l.shipdate < DATE('1996-04-01')
        GROUP BY l.suppkey) AS r2
   WHERE r2.total_revenue > r1.total_revenue) );

Q16 SELECT p.brand,
      p.type,
      p.size,
      COUNT(DISTINCT ps.suppkey) AS supplier_cnt
FROM Partsupp ps, Part p
WHERE p.partkey = ps.partkey
AND p.brand <> 'Brand#45'
AND (p.type NOT LIKE 'MEDIUM POLISHED%')
AND (p.size IN (49, 14, 23, 45, 19, 3, 36, 9))
AND (ps.suppkey NOT IN (
  SELECT s.suppkey
  FROM Supplier s
  WHERE s.comment LIKE '%Customer%Complaints%'
))
GROUP BY p.brand, p.type, p.size;

Q17 SELECT SUM(l.extendedprice) / 7.0 AS avg_yearly
FROM Lineitem l, Part p
WHERE p.partkey = l.partkey
AND p.brand = 'Brand#23'
AND p.container = 'MED BOX'
AND l.quantity < (
  SELECT 0.2 * AVG(l2.quantity)
  FROM Lineitem l2
  WHERE l2.partkey = p.partkey
);

Q18 SELECT c.name, c.custkey, o.orderkey, o.orderdate,
      o.totalprice,
      SUM(l.quantity) AS query18
FROM Customer c, Orders o, Lineitem l
WHERE o.orderkey IN
  ( SELECT 13.orderkey FROM (
    SELECT 12.orderkey, SUM(l2.quantity) AS QTY
    FROM Lineitem l2 GROUP BY 12.orderkey ) l3
    WHERE QTY > 100
  )
AND c.custkey = o.custkey
AND o.orderkey = l.orderkey
GROUP BY c.name, c.custkey, o.orderkey, o.orderdate,
        o.totalprice;

```

```

Q19 | SELECT SUM(l.extendedprice * (1 - l.discount)) AS revenue
    | FROM Lineitem l, Part p
    | WHERE
    | (
    |   p.partkey = l.partkey
    |   AND p.brand = 'Brand#12'
    |   AND ( p.container IN ( 'SM CASE', 'SM BOX',
    |                         'SM PACK', 'SM PKG' ) )
    |   AND l.quantity >= 1 AND l.quantity <= 1 + 10
    |   AND ( p.size BETWEEN 1 AND 5 )
    |   AND ( l.shipmode IN ( 'AIR', 'AIR REG' ) )
    |   AND l.shipinstruct = 'DELIVER IN PERSON'
    | )
    | OR
    | (
    |   p.partkey = l.partkey
    |   AND p.brand = 'Brand#23'
    |   AND ( p.container IN ( 'MED BAG', 'MED BOX',
    |                         'MED PKG', 'MED PACK' ) )
    |   AND l.quantity >= 10 AND l.quantity <= 10 + 10
    |   AND ( p.size BETWEEN 1 AND 10 )
    |   AND ( l.shipmode IN ( 'AIR', 'AIR REG' ) )
    |   AND l.shipinstruct = 'DELIVER IN PERSON'
    | )
    | OR
    | (
    |   p.partkey = l.partkey
    |   AND p.brand = 'Brand#34'
    |   AND ( p.container IN ( 'LG CASE', 'LG BOX',
    |                         'LG PACK', 'LG PKG' ) )
    |   AND l.quantity >= 20 AND l.quantity <= 20 + 10
    |   AND ( p.size BETWEEN 1 AND 15 )
    |   AND ( l.shipmode IN ( 'AIR', 'AIR REG' ) )
    |   AND l.shipinstruct = 'DELIVER IN PERSON'
    | )
    | );

```

```

Q20 | SELECT s.name, s.address
    | FROM Supplier s, Nation n
    | WHERE s.supkey IN
    |   ( SELECT ps.supkey
    |     | FROM Partsupp ps
    |     | WHERE ps.partkey IN
    |     |   ( SELECT p.partkey
    |     |     | FROM Part p
    |     |     | WHERE p.name like 'forest%' )
    |     | AND ps.availqty >
    |     |   ( SELECT 0.5 * SUM(l.quantity)
    |     |     | FROM Lineitem l
    |     |     | WHERE l.partkey = ps.partkey
    |     |     | AND l.supkey = ps.supkey
    |     |     | AND l.shipdate >= DATE('1994-01-01')
    |     |     | AND l.shipdate < DATE('1995-01-01') ) )
    | AND s.nationkey = n.nationkey
    | AND n.name = 'CANADA';

```

```

Q21 | SELECT s.name, COUNT(*) AS numwait
    | FROM Supplier s, Lineitem l1, Orders o, Nation n
    | WHERE s.supkey = l1.supkey
    | AND o.orderkey = l1.orderkey
    | AND o.orderstatus = 'F'
    | AND l1.receiptdate > l1.commitdate
    | AND ( EXISTS (SELECT * FROM Lineitem l2
    |               | WHERE l2.orderkey = l1.orderkey
    |               | AND l2.supkey <> l1.supkey) )
    | AND ( NOT EXISTS
    |       | (SELECT * FROM Lineitem l3
    |         | WHERE l3.orderkey = l1.orderkey
    |         | AND l3.supkey <> l1.supkey
    |         | AND l3.receiptdate > l3.commitdate) )
    | AND s.nationkey = n.nationkey
    | AND n.name = 'SAUDI ARABIA'
    | GROUP BY s.name;

```

```

Q22 | SELECT cntrycode,
    | COUNT(*) AS numcust,
    | SUM(custsale.acctbal) AS totalacctbal
    | FROM (
    |   SELECT SUBSTRING(c.phone, 0, 2) AS cntrycode,
    |          c.acctbal
    |   FROM Customer c
    |   WHERE (SUBSTRING(c.phone, 0, 2) IN
    |         | ('13', '31', '23', '29', '30', '18', '17'))
    |   AND c.acctbal > (
    |     | SELECT AVG(c2.acctbal)
    |     | FROM Customer c2
    |     | WHERE c2.acctbal > 0.00
    |     | AND (SUBSTRING(c2.phone, 0, 2) IN
    |           | ('13', '31', '23', '29',
    |           | '30', '18', '17'))
    |   )
    | AND (NOT EXISTS (SELECT * FROM Orders o
    |                  | WHERE o.custkey = c.custkey))
    | ) custsale
    | GROUP BY cntrycode;

```

Appendix C Compilation flags

DBToaster used the highest optimization level to compile all the queries from our workload, except TPC-H Q11. For queries where we wanted to analyze the effect of view caching, like VWAP and AXF, we used flags to enable materialization of maps with input variables, and also to predefine their sizes (only for the C++ backend). For Q16, we suggested to DBToaster to incrementally maintain the query; this is an exception where DBToaster fails to choose the correct maintenance strategy; we plan to fix this issue in our future release.

The naive recursive compilation materialized the whole delta expressions with input variables, without simplifying them or applying join graph decomposition. Under the naive method, Q15, MDDB1, and MDDB2 were unable to compile within a 4-hour time limit. The C++ backend failed to produce compilable code for Q2, Q4, Q18a and Q20; we fixed these errors by hand.

Depth-1 and Depth-0 used a command line option to limit the compiler's maximum recursive depth. For the C++ backend and queries Q8, SSB4, MDDB1, and MDDB2, we used the WIDE-TUPLE flag to overcome the Boost's limitation that tuples may contain at most 50 attributes.

Appendix D Traces

Figures 13-18 present full traces of view refresh rates and memory footprint for every query in our workload, excepting those previously discussed in section 9.

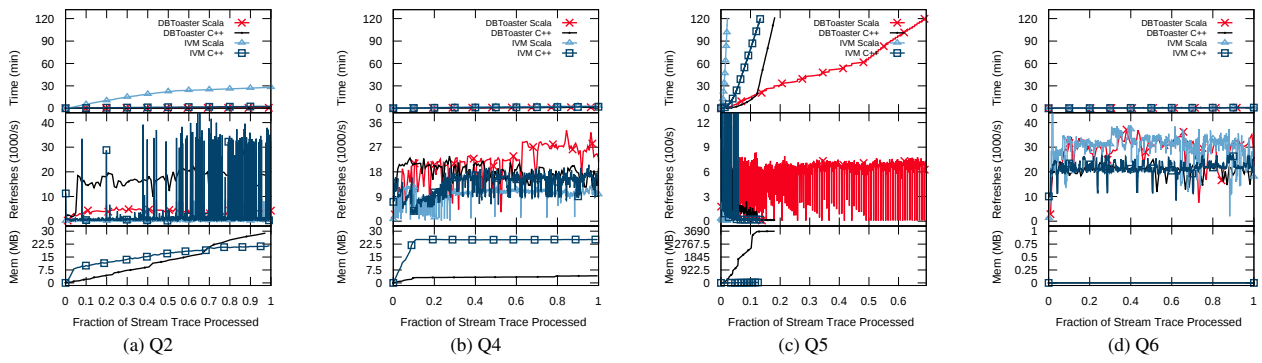


Fig. 13 (a) A 5-way linear join with an equality/inequality correlated nested aggregate in the EXISTS clause. (b) A join-free query with an equality/inequality correlated nested aggregate in the EXISTS clause. (c) A 6-way linear join. (d) A join-free query.

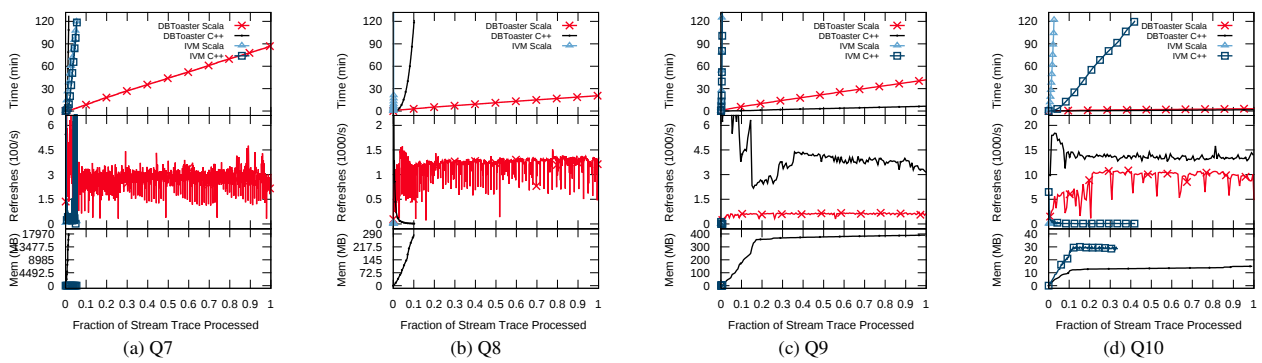


Fig. 14 (a) A 6-way linear join. (b) A 8-way linear join. (c) A 6-way star join. (d) A 4-way linear join.

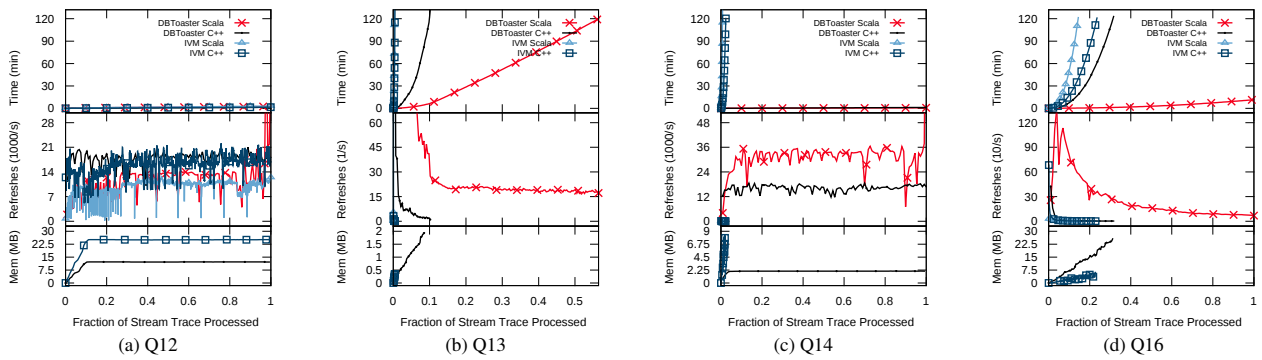


Fig. 15 (a) A 2-way join. (b) A 2-way join with an aggregate subquery in the FROM clause. (c) A 2-way join. (d) A 2-way join with a nested inequality correlated query.

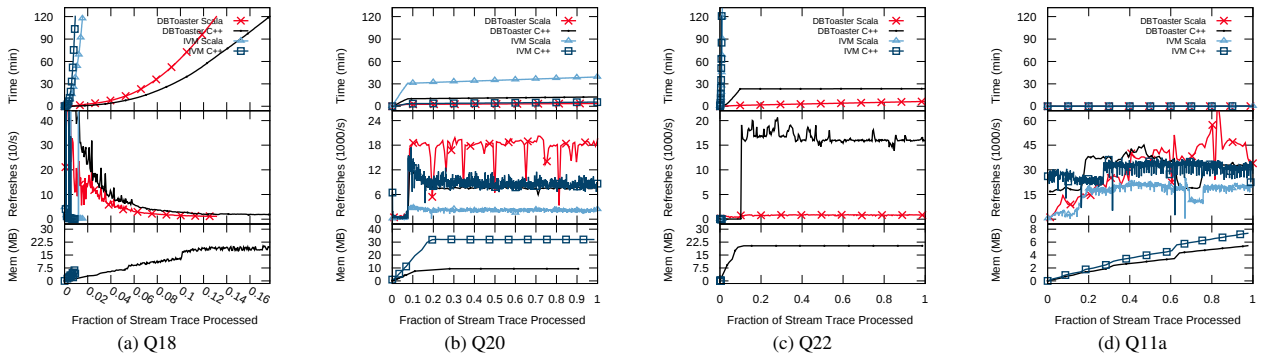


Fig. 16 (a) A 3-way linear join with an equality correlated nested query and an uncorrelated aggregate in the FROM clause. (b) A 2-way join with 3 equality-correlated nested queries. (c) Join-free query with an uncorrelated nested aggregate and a correlated nested aggregate in the EXISTS clause. (d) A 2-way join.

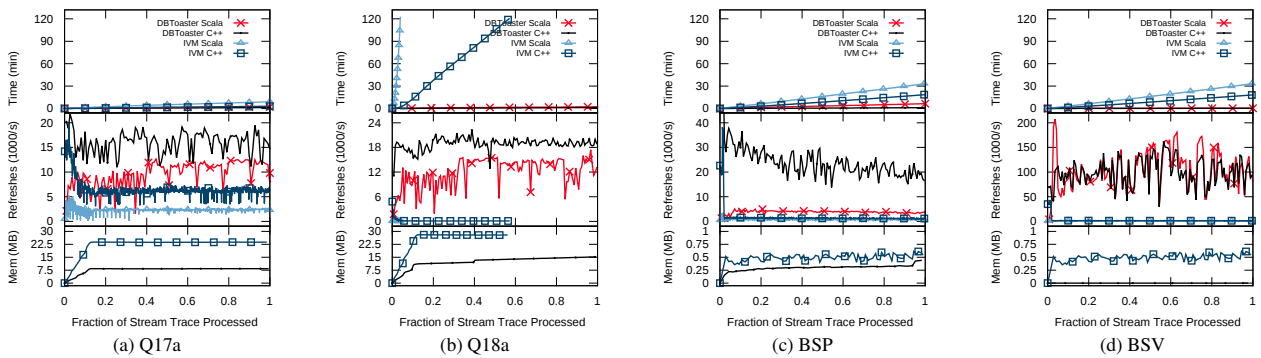


Fig. 17 (a) A 2-way join with an equality correlated nested aggregate. (b) A 3-way linear join with two equality-correlated nested aggregates. (c) A self join with both equality and inequality join conditions. (d) A self join.

	Queries	Opt Level	DBT Flags
DBToaster	Q11	-O2	-d DELETE-ON-ZERO
	Q16	-O3	-d HEURISTICS-ALWAYS-UPDATE
	VWAP w/ ivars	-O3	-d HEURISTICS-ENABLE-INPUTVARS
	AXF w/ ivars	-O3	-d HEURISTICS-ENABLE-INPUTVARS -d HEURISTICS-AGGRESSIVE-INPUTVARS -g -DDEFAULT_MAP_SIZE=10000 (C++)
others	-O3	-	
Naive	all	-O2	-d CALC-NO-OPTIMIZE -d CALC-NO-DECOMPOSITION -d HEURISTICS-ENABLE-INPUTVARS -g -DDEFAULT_MAP_SIZE=10000 (C++)
Depth-0 or 1	Q8		
	SSB4	-O3	-d WIDE-TUPLE (C++)
	Mddb1		--depth [0]1
	Mddb2		
others	-O3	--depth [0]1	

Fig. 12 Compilation flags used by DBToaster to compile queries from our workload. Description of the flags can be found at www.dbtoaster.org.

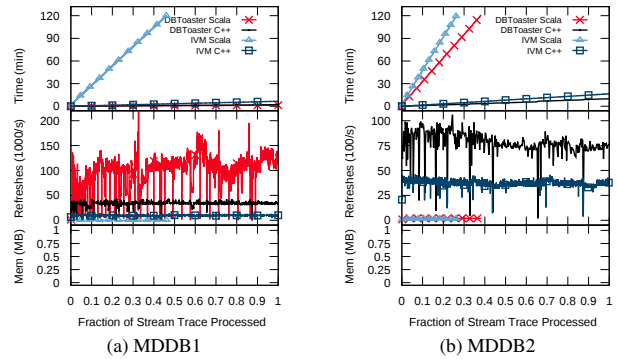


Fig. 18 (a) A 4-way join. (b) A 10-way join.