

SyncSQL: A Language to Express Views over Data Streams

Thanaa M. Ghanem¹

Per-Åke Larson²

Walid G. Aref¹

Ahmed K. Elmagarmid¹

¹Department of Computer Science, Purdue University, ghanemtm,aref,ake@cs.purdue.edu

²Microsoft Research, palarson@microsoft.com

Abstract

Prior work on languages to express continuous queries over streams has defined a stream as a sequence of tuples that represents an infinite append-only relation. In this paper, we show that composition of queries is not possible in the append-only model. Query composition is a fundamental property of any query language - composition makes it possible to build up complex queries from simpler queries. We then propose a query language, termed Synchronized SQL (or SyncSQL), that defines a stream as a sequence of modify operations (i.e., insert, update, and delete) against a relation with a specified schema. Inputs and outputs in any SyncSQL query are interpreted in the same way and, hence, SyncSQL expressions can be composed. Coarser periodic refresh requirements are typically expressed as sliding-window queries. We generalize this approach by introducing the synchronization principle that empowers SyncSQL with a formal mechanism to express queries with arbitrary refresh conditions. After introducing the semantics and syntax, we lay the algebraic foundation for SyncSQL and propose a query matching algorithm for deciding containment of SyncSQL expressions.

1 Introduction

Query languages in the streaming literature (e.g., [2, 7, 8, 11, 25]) define a stream as a sequence of tuples that represents an infinite append-only relation. Languages based on the append-only model are not closed, that is, the result of a query expression is not necessarily an append-only relation. This has the effect that query expressions cannot be freely composed, that is, expressing a query in terms of one or more sub-queries as can be done, for example, with SQL queries in relational databases. Composition is a fundamental property of any query language but it requires that query inputs and outputs are interpreted in the same way. However, in the append-only stream model a continuous query may not be able to produce an append-only output

even when input streams represent append-only relations.

For example, consider an application monitoring a parking lot where two sensors continuously monitor the lot's entrance and exit. The sensors generate two streams of identifiers, say S_1 and S_2 , for cars entering and exiting the lot, respectively. A reasonable query in this environment is Q_1 : "Continuously keep track of the identifiers of all cars inside the parking lot". The answer of Q_1 is a view that, at any time point T , contains the identifiers for cars that are inside the parking lot. S_1 can be modeled as a stream that inserts tuples into an append-only relation, say $\mathcal{R}(S_1)$ and, similarly, S_2 inserts tuples into the append-only relation $\mathcal{R}(S_2)$. Then, Q_1 can be regarded as a *materialized view* that is defined as the set-difference between the two relations $\mathcal{R}(S_1)$ and $\mathcal{R}(S_2)$. As tuples arrive on S_1 and S_2 , the corresponding relations are modified, and the relation representing the result of Q_1 is updated to reflect the changes in the inputs. The result of Q_1 is updated by *inserting* identifiers of cars entering the lot and *deleting* identifiers of cars exiting the lot. Notice that, although the input relations in Q_1 change by only inserting tuples (i.e., append only), the output of Q_1 changes by both insertions and deletions.

The answer to query Q_1 can be output either as (1) a *complete* answer, or (2) an *incremental* answer. In the first case, at any time point T , the issuer of Q_1 sees a state, i.e., a relation containing identifiers of all cars inside the lot at time T . In the second case, the issuer of Q_1 receives a stream that represents the changes (i.e., insertions and deletions) in the state. The output in the incremental case is interpreted in the same way as the inputs, namely, as a stream that represents modifications to an underlying relation. However, Q_1 's incremental answer cannot be produced or consumed by a query in a language that models a stream as an append-only relation. Existing languages may produce output streams from Q_1 but the output streams are interpreted differently from the input streams. For example, the output may be modeled as a stream representing a concatenation of serializations of the complete answer (e.g., RStream in CQL [2], and the output of window queries in TelegraphCQ [8]). As another alternative, CQL divides the

output into two append-only streams such that one stream represents the insertions in the output and the other stream represents the deletions (i.e., IStream and DStream).

The different interpretation and the division of an output stream prevents composition of queries, that is, using the output of a query as the input to another queries or building up complex query expressions from simpler expressions. Composition is a fundamental requirement on any query language and particularly important in streaming environments that are characterized by concurrent, overlapping queries. For example, consider the following query, Q_2 , from the same application: “Group the cars inside the parking lot by type (e.g., trucks, cars, or buses). Continuously keep track of the number of cars in each group”. By analyzing the two queries, Q_1 and Q_2 , it is obvious that Q_2 is an aggregate query over the output of Q_1 . This observation motivates the idea of defining Q_1 as a view, say V_1 and then, expressing both Q_1 and Q_2 in terms of V_1 . However, realizing this requires a language that allows query composition.

In streaming applications with high tuple arrival rates, an issuer of continuous queries may not be interested in refreshing the answer in response to every tuple arrival. Instead, coarser refresh periods may be desired. For example, instead of reporting the count of cars with every change in the parking lot, Q_2 may be interested in updating the count of cars in each group *every five minutes*. This refresh condition is based on time but a powerful language should allow a user to express more general refresh conditions based on time, tuple arrival, events, relation state, and so on.

In addition to preventing query composition, the append-only model limits the applicability of the language because streams may have denotations other than the append-only relation [22]. For example, update streams are used in applications where objects continuously update their values. For example, consider a temperature-monitoring application in which sensors are distributed in rooms and each sensor continuously reports the room temperature. A reasonable query in this environment is, Q_3 : “Continuously keep track of the rooms that have temperature greater than 80”. Neither the input nor the output streams in Q_3 represent append-only relations. The input in Q_3 is an update stream in which, a room identifier is considered a key and an input tuple is an update over the previous tuple with the same key value. The output tuples from Q_3 represent incremental changes in the answer and include insertions and deletions for rooms that switch between satisfying and not satisfying the query predicate.

1.1 Our Approach

We can summarize the limitations of the existing continuous query languages as follows. (a) Cannot express queries over streams other than the append-only relation

representation. (b) Cannot produce incremental answer for queries that do not produce an append-only output. (c) Cannot always compose queries because of the different interpretation and/or division of the output streams. (d) Refresh condition are restricted to be either time or tuple-based.

In this paper, we introduce a continuous query language for data streams, termed Synchronized SQL (SyncSQL for short), that avoids the previous limitations. In contrast to other languages, SyncSQL defines the stream as a sequence of modify operations (i.e., insert, update, and delete) against a relation with a specified schema. Basically, a continuous query in SyncSQL is semantically equivalent to a *materialized view* where the inputs are relations that are modified by streams of modify operations. The answer of the query is another stream of modify operations that represent changes in the result of the view. This is equivalent to incremental maintenance of materialized views [17]. The unified representation of query inputs and outputs enables the composition of SyncSQL expressions, and as a result, gives the ability to express and exploit views over streams.

To cope with the coarser refresh requirement of continuous queries, we introduce the *synchronization principle*. The idea is to formally specify synchronization time points at which the input tuples are processed by the query pipeline. Input tuples that arrive between two consecutive synchronization points are accumulated and reflected in the output at once at the next synchronization point. The synchronization principle makes it possible to (1) express queries with arbitrary refresh conditions, and (2) formally reason about the containment relationship among queries with different refresh periods.

The contributions of this paper are summarized as follows:

- **SyncSQL semantics and syntax:** We define concise semantics and syntax for continuous queries and views over streams.
- **SyncSQL algebra:** We lay the algebraic foundation for SyncSQL by providing data types, operators, algebraic laws and transformation rules that are needed to enumerate query plans.
- **Shared execution using query composition:** Based on the algebraic framework, we propose a query matching algorithm that is used to deduce the containment relationships among query expressions. The containment relationship is used to achieve shared execution using query composition.
- **Execution model:** We present a pipelined and incremental execution model to efficiently realize SyncSQL queries in a data stream management system.

1.2 Paper Outline

The rest of the paper is organized as follows. Section 2 introduces the semantics and syntax of SyncSQL. The synchronization principle is explained in Section 3. In Section 4, we lay the algebraic foundation for SyncSQL. The shared query execution algorithm is given in Section 5. In Section 6, we give an incremental execution model for SyncSQL queries. Section 7 surveys the existing works for continuous queries and contrasts our approach with the other approaches. Finally, Section 8 concludes the paper.

2 SyncSQL Semantics and Syntax

In short, a continuous SyncSQL query is semantically equivalent to a materialized view over one or more relations where the input relations are updated by *streams* of modify operations.

2.1 Stream, Query, and View Semantics

Stream semantics. A data stream is a sequence of tuples that are sent by remote data sources (e.g., sensors) [2, 8, 25]. By considering the wide variety of streaming applications, the same stream may be interpreted differently by the different applications. For example, consider a stream, say `TemperatureSource`, that is sent by sensors in a temperature-monitoring application. Assume that a tuple in `TemperatureSource` reports a certain room’s temperature and is denoted by “`<RoomID, Temperature>Timestamp`”. One application, say `Application1`, may consider `TemperatureSource` as an update stream over the various rooms temperature. In this case, `RoomID` is considered a key and a tuple is considered an update over the previous tuple with the same key value. On the other hand, another application, say `Application2`, may view the `TemperatureSource` stream as just a sequence of temperature readings and ignore the `RoomID` attribute. Notice that, the same continuous query over `TemperatureSource` produces different answers under the different application-dependent semantics.

In order to capture the various application-dependent semantics, SyncSQL distinguishes between two types of streams: *raw* and *tagged*. A *raw* stream is a sequence of tuples that are sent by the remote sources. On the other hand, a *tagged* stream is a stream of modify operations (i.e., insert (+), update(u), and delete(-)) against a relation with a specified schema. A raw stream must be transformed into a tagged stream before being used as input in a query. This is similar to the relational model in traditional databases where the raw data has to be transformed into relations before being used in a query.

The function that transforms a raw input stream to a corresponding tagged stream is application-dependent where the same raw input stream can produce different tagged streams under different transformation functions. For example, in a temperature-monitoring application, a *raw* input stream, say `TemperatureSource`, is sent by remote sensors where an input tuple in the raw stream reports a room temperature. A tuple in the `TemperatureSource` stream consists of two attributes: `RoomID` and `Temperature`. One application, say `Application1`, may consider `TemperatureSource` as an update stream over the various rooms temperature. In this case, `RoomID` is considered a key and a tuple is considered an update over the previous tuple with the same key value. On the other hand, another application, say `Application2`, may view the `TemperatureSource` stream as just a sequence of temperature readings and ignore the `RoomID` attribute.

Assume that an input tuple in `TemperatureSource` is denoted by “`<RoomID, Temperature>Timestamp`”. In `Application1`, `TemperatureSource` represents an update stream over the various room temperatures and the corresponding tagged stream, say `RoomTempStr`, consists of *insert* and *update* operations. Basically the tagging procedure takes an input `TemperatureSource` tuple and produces a corresponding tagged tuple in `RoomTempStr` as follows: the first tuple in `TemperatureSource` with a certain `RoomID` value is transformed into a corresponding *insert* operation “`+<RoomID, Temperature>Timestamp`” in `RoomTempStr`. A subsequent tuple in `TemperatureSource` with the same `RoomID` is transformed into an *update* tuple “`u<RoomID, Temperature>Timestamp`” in `RoomTempStr`. Notice that the tagging function needs to keep a list of the observed key (i.e., `RoomID`) values so far.

On the other hand, in the case of `Application2`, `TemperatureSource` represents an infinite append-only relation and the corresponding tagged stream, say `TempStr`, is a sequence of *insert* operations where each tuple “`<RoomID, Temperature>Timestamp`” in `TemperatureSource` is transformed to a corresponding insert operation “`+<RoomID, Temperature>Timestamp`” in `TempStr`.

In the query processing phase, the transformation (or tagging) function is implemented inside an operator, called `Tagger`, that is placed at the bottom of a query pipeline. In `Application1`, the functionality of the `Tagger` operator is similar to that of the `MERGE` (or `UPSERT`) operator in the SQL:2003 standard [12]. Basically, in `Application1`, `Tagger` needs to keep a list of all the observed key values (i.e., `RoomID`) so far. The size of

the key list has an upper bound that is equal to the maximum number of rooms. However, implementing the tagging function as an operator opens the room for the query optimizer to re-order the pipeline and optimize the memory consumption. For example, the Tagger operator can be pulled above the Select operator so that only qualified rooms are stored in the key list. The details of query processing and optimization is beyond the scope of this paper.

Example 1 This example demonstrates the syntax for defining streams and the mapping from raw to tagged streams. The raw `TemperatureSource` stream is defined in `SyncSQL` by the following statement:

```
REGISTER SOURCE TemperatureSource
(char RoomID, int Temperature) From
port5501
```

where `RoomID` and `Temperature` represent the stream schema and `port5501` is the port at which external sources report tuples. The tagged streams are defined over the source `TemperatureSource` as follows:

```
RoomTempStr: CREATE STREAM
RoomTempStr
                OVER TemperatureSource
                KEY RoomID

TempStr: CREATE STREAM TempStr
          OVER TemperatureSource
          KEY NULL
```

Running example: Assume the following tuples arrived at `TemperatureSource`: $\langle a, 100 \rangle 1$, $\langle b, 75 \rangle 2$, $\langle c, 80 \rangle 3$, $\langle a, 95 \rangle 4$, $\langle b, 85 \rangle 5$.

`RoomTempStr`: The following tuples represent the corresponding tagged `RoomTempStr` stream: $+\langle a, 100 \rangle 1$, $+\langle b, 75 \rangle 2$, $+\langle c, 80 \rangle 3$, $u\langle a, 95 \rangle 4$, $u\langle b, 85 \rangle 5$. Notice that the tuple $\langle a, 100 \rangle 1$ is mapped to $+\langle a, 100 \rangle 1$ while $\langle a, 95 \rangle 4$ is mapped to $u\langle a, 95 \rangle 4$.

`TempStr`: The following tuples represent the corresponding `TempStr` tuples: $+\langle a, 100 \rangle 1$, $+\langle b, 75 \rangle 2$, $+\langle c, 80 \rangle 3$, $+\langle a, 95 \rangle 4$, $+\langle b, 85 \rangle 5$. Notice that all the tuples in `TempStr` are *insert* operations.

The relational view of a tagged stream: The semantics of query operators (e.g., select and join) are defined over relations. However, inputs in a continuous query are tagged streams where each stream represents modifications against a relation. Hence, in order to adopt the well-known semantics of relational operators, `SyncSQL` queries are expressed over the tagged streams' corresponding relations. Notice that, streams of insert and delete tuples are frequently used when addressing continuous query processing [1, 5, 13, 14, 23]. However, to the best of the authors'




Figure 1. Illustrating Time-varying Relations.

knowledge, `SyncSQL` is the first language that addresses query semantics over tagged streams.

Basically, any tagged stream, say S , has a corresponding time-varying relation, termed $\mathcal{R}(S)$, that is continuously modified by S 's tuples. An input tuple in a tagged stream is denoted by "Type<Attributes>Timestamp", where Type can be either +, u, or - and Timestamp indicates the time at which the modification takes place. $\mathcal{R}(S)$'s schema consists of two parts as follows: (1) a set of attributes that corresponds to S 's Attributes, and (2) a timestamp attribute, termed TS, that corresponds to the Timestamp field of S 's tuples. Timestamp is mapped to $\mathcal{R}(S)$ in order to be able to express time-based windows over S as will be shown in Section 2.2. At any time point, say T , $\mathcal{R}(S)$ is denoted by $R[S(T)]$ and is the relation resulting from applying S 's operations with timestamps less than or equal to T in an increasing order of timestamp. According to the underlying application, $\mathcal{R}(S)$ can be modified by either inserting tuples (i.e., append-only), or by general modify operations.

Definition 1. Time-varying relation. A time-varying relation $\mathcal{R}(S)$ is the relational view of a tagged stream S such that $\mathcal{R}(S) = R[S(T)] \forall T$, where T is any point in time.

Example 2 This example demonstrates the mapping from `RoomTempStr` (as defined in Example 1) to a time-varying relation. Figure 1a shows the following input tuples: $+\langle a, 99 \rangle 1$, $+\langle b, 75 \rangle 2$, $+\langle c, 80 \rangle 3$, $u\langle a, 95 \rangle 4$. Figure 1b gives $\mathcal{R}(\text{RoomTempStr})$ with a schema of three attributes: `RoomID`, `Temperature`, and `TS`. Figure 1b shows that, at time 1, $\mathcal{R}(\text{RoomTempStr})$ reflects the insertion of Room "a" with temperature 99. At time 4, $\mathcal{R}(\text{RoomTempStr})$ reflects the update of Room "a" temperature to 95.

Query semantics. A continuous query over n tagged streams, $S_1 \dots S_n$, is semantically equivalent to a *materialized view* that is defined by an SQL expression over the time-varying relations, $\mathcal{R}(S_1) \dots \mathcal{R}(S_n)$. At any time point, T , the query answer reflects the contents of the underlying relations at time T , (i.e., $R[S_1(T)] \dots R[S_n(T)]$). Whenever any of the underlying relations is modified by the arrival of a stream tuple, the modify operation is propagated to produce the corresponding set of




Figure 2. Q_3 Running Example.

modify operations in the query answer in a way similar to incremental maintenance of materialized views [17].

Query outputs. The output of a query can be provided in two forms as follows:

- (1) **COMPLETE** output, where, at any time point, the issuer of the query sees a state or a relation that represents the complete answer of the query. The output relation is continuously modified by the query pipeline whenever any of the input relations is modified.
- (2) **STREAMED** output, where the issuer of the query receives a tagged stream that represent the *deltas* (i.e., incremental changes) in the answer.

Example 3 This example demonstrates the semantics and syntax of SyncSQL queries. The temperature monitoring query Q_3 that is used in Section 1, is expressed in SyncSQL as follows:

```
select STREAMED RoomID, Temperature
from  $\mathcal{R}$ (RoomTempStr) R
where R.Temperature > 80
```

where `RoomTempStr` is the input tagged stream that is defined in Example 1. $\mathcal{R}(\text{RoomTempStr})$ is the corresponding time-varying relation. The keyword `STREAMED` indicates that the output needs to be another stream of modify operations. The output stream of this query includes *insert* (or *update*) operations for rooms that qualify the predicate “`R.Temperature > 80`” and/or *delete* operations for previously qualified rooms that disqualify the predicate due to a temperature update.

Running example. Assume the following `RoomTempStr`’s tuples have arrived at Q_3 : $\langle a, 78 \rangle 1$, $\langle a, 105 \rangle 2$, $\langle b, 70 \rangle 3$, $\langle c, 95 \rangle 4$, $\langle a, 76 \rangle 5$, $\langle c, 103 \rangle 6$. Figure 2 shows the input and output streams in Q_3 are as follows. The input tuple $\langle a, 78 \rangle 1$ does not result in producing any output tuples, while $\langle a, 105 \rangle 2$, which arrives at time 2, results in *inserting* Room “a” in the answer via the output tuple $\langle a, 105 \rangle 2$. Similarly, when $\langle c, 95 \rangle 4$ arrives, Room “c” is *inserted* in the query answer via $\langle c, 95 \rangle 4$. Later, when $\langle a, 76 \rangle 5$ arrives, Room “a” is *deleted* from the output via $\langle a \rangle 5$. Notice that the “Attributes” part of the delete tuple $\langle a \rangle 5$ specifies

only the key value which is enough to perform deletion. Finally, when $\langle c, 103 \rangle 6$ arrives, a corresponding tuple $\langle c, 103 \rangle 6$ is produced in the query answer to report that Room “c” still qualifies the query predicate, but with a new temperature.

Views over streams. The unified interpretation (as tagged streams) of SyncSQL query inputs and outputs enables SyncSQL to define and exploit views over streams. Basically, a view over streams is a named SyncSQL query expression that is defined once and, then, can be used as input in any other query. For example, a view, say V_i , can be used as input in a query, say Q_i , if Q_i ’s expression (or part of it) is *equivalent* or is *contained* in V_i ’s expression. In Section 5, we give an algorithm to deduce the containment relationships among SyncSQL expressions.

Example 4 This example demonstrates answering queries using views. Consider the following query, Q_4 (from the same temperature-monitoring application as Q_3): “Continuously keep track of the rooms that have temperature greater than 100”. Similar to Q_3 , Q_4 can be expressed over `RoomTempStr` as follows:

```
select STREAMED RoomID, Temperature
from  $\mathcal{R}$ (RoomTempStr) R
where R.Temperature > 100
```

It is obvious that Q_4 is contained in Q_3 . As a result we can benefit from query composition by defining Q_3 as a view, say `HotRooms1`, as follows:

```
create STREAMED view HotRooms1 as
select RoomID, Temperature
from  $\mathcal{R}$ (RoomTempStr) R
where R.Temperature > 80
```

Then, the two queries Q_3 and Q_4 can be re-written in terms of `HotRooms1` as follows:

```
 $Q_3$ : select STREAMED RoomID, Temperature
      from  $\mathcal{R}$ (HotRooms1) R
 $Q_4$ : select STREAMED RoomID, Temperature
      from  $\mathcal{R}$ (HotRooms1) R
      where R.Temperature > 100
```

Running example. Figure 3 shows the execution of Q_4 over the output of `HotRooms1`. Notice that the output stream from `HotRooms1` is the same output stream from Q_3 that is shown in Figure 2. Basically, when the tuple $\langle a, 105 \rangle 2$ arrives at Q_4 at time 2, a corresponding tuple $\langle a, 105 \rangle 2$ is produced in the output. In contrast, $\langle c, 95 \rangle 4$ does not result in producing any output tuples since 95 does not qualify Q_4 ’s predicate. Later, $\langle c, 103 \rangle 6$ results in inserting Room “c” in Q_4 ’s answer via $\langle c, 103 \rangle 6$.

2.2 Window Queries

In addition to expressing queries over non append-only streams, SyncSQL still can express sliding-window queries




Figure 3. Query Composition.

over append-only streams. The sliding-window query model is the most widely used window model in the existing streaming literature. A sliding window is defined by two parameters: (1) *range* that specifies the size of the window, and (2) *slide* that specifies the step by which the window moves over the stream.

Windows may be assigned to streams (e.g., [2, 8]) or to operators (e.g., [7, 25]). However, the same relational operator (e.g., join) may have different semantics under the different window usages. For example, if we consider the window-per-operator usage, a window join with window size w , joins the input stream tuples that are within at most w time units from each other [7]. On the other hand, if we consider the window-per-stream usage, a binary window join has two different window sizes, w_1 and w_2 , one for each stream [2].

The difference in window semantics makes it difficult for a language that is defined by one window semantics to express queries from the other window semantics. To overcome this difficulty, SyncSQL does not assume specific window semantics. Instead, SyncSQL uses a general window model that can be used to express the various windows.

2.2.1 Expressing Window Queries in SyncSQL

In SyncSQL, raw input streams that represent append-only relations are mapped to tagged streams of *insert* operations (e.g., the TempStr stream in Example 1). SyncSQL does not use specific constructs to express sliding windows over the append-only streams. Instead, SyncSQL employs the predicate-window query model [15] in which the window *range* is expressed as a regular predicate in the *where* clause of the query. The window’s *slide* is expressed using the synchronization principle as will be explained in Section 3.

The predicate-window model is a generalization of the existing window models, since all types of windows (e.g., window-per-stream, window-per-operator) can be expressed as predicate windows. A time-based sliding window over an append-only stream, say S , is expressed as a predicate over $\mathcal{R}(S)$ ’s TS attribute. For example, a win-

dow join between two streams, S_1 and S_2 , where two tuples are joined only if they are at most 5 time units apart, can be expressed by the following predicate:

$\mathcal{R}(S_2).TS - 5 < \mathcal{R}(S_1).TS < \mathcal{R}(S_2).TS + 5$. The window predicate can be expressed over any attribute in the input stream tuple (ordered or non-ordered). For example, the temperature monitoring query, Q_3 , is a predicate-window query in which the predicate is defined over the unordered Temperature attribute. Moreover, sliding-window queries in which a separate window is attached to each input stream can be expressed using predicate windows as shown by the following example.

Example 5 Consider a road-monitoring application in which sensors are distributed to report car identifiers for cars passing through a specified intersection. The input stream S of car identifiers represents an append-only relation. A sliding window over S of size 5 time units is essentially a *view* that, at any time point T , contains the car identifiers that are reported between times $T - 5$ and T . Such window view is expressed in SyncSQL as follows:

```
create STREAMED view FiveUnitsWindow as
select *
from  $\mathcal{R}(S)$  R
where Now - 5 < R.TS ≤ Now
```

The view FiveUnitsWindow is refreshed when either $\mathcal{R}(S)$ is modified or Now is changed. $\mathcal{R}(S)$ is modified by the arrival of S tuples where new S tuples produce *insert* operations in the view’s output. On the other hand, Now is continuously changing to indicate the current time, and, as a result, *delete* operations are produced in the output to represent expired tuples that fall behind the window boundaries. Notice that even if S consists of only insert operations, FiveUnitsWindow’s output stream includes both insert and delete operations. In Section 3.3 we show that the value of Now can be represented as a view that is continuously updated to reflect the current time.

Example 6 This example demonstrates query composition by using of FiveUnitsWindow as input in another continuous query. Assume the following continuous query from the road monitoring application, Q_5 : “Group the input cars by type (e.g., trucks, cars, or buses). Then continuously report the number of cars passed in the last five time units in each group”. The query Q_5 is expressed over FiveUnitsWindow as follows:

```
select STREAMED COUNT(*)
from  $\mathcal{R}(\text{FiveUnitsWindow})$ 
groupby CarType
```

CarCount’output is a stream of *update* operations that represents the *incremental* query answer. An *update* operation is produced for a group, G , only whenever a car enters and/or expires from G . Notice that if the same query is expressed using COMPLETE output, then whenever the query

is refreshed, the query issuer sees the non-incremental answer that includes the count of cars in each group independent from whether the group has been changed or not. The non-incremental output of aggregate queries is the approach that is followed by most of the existing systems to evaluate aggregates over data streams (e.g., [7, 20]).

3 The Synchronization Principle

If we follow the traditional materialized view semantics, a SyncSQL query answer is refreshed whenever any of the input relations is modified. Unlike materialized views, in streaming applications, modifications may arrive at high rates. Usually, a continuous query issuer is interested in having coarser refresh periods for the answer. For example, as we discussed in Section 1, the issuer of the query Q_2 may be interested in getting an update of the answer every five minutes independent of the rate of changes in the parking lot state. The coarser refresh period is achieved via special constructs in other query languages, for example, the *slide* parameter in the sliding-window query model [3, 20] and the *for loop* in [8].

In this section, we introduce the synchronization principle as a generalization for sliding windows. The idea of the synchronization principle is to formally specify synchronization time points at which the input stream tuples are processed by the query pipeline. Input tuples that arrive between two consecutive synchronization points are not propagated immediately to produce query outputs. Instead, the tuples are accumulated and are propagated simultaneously at the following synchronization point. In the rest of the paper, we show that the synchronization principle distinguishes SyncSQL by being able to: (1) express queries with arbitrary refresh conditions, and (2) formally reason about the containment relationships among continuous queries with different refresh periods.

3.1 Synchronized Relations

We introduce the *synchronization* principle as a means for expressing coarser refresh periods in SyncSQL. The purpose of the synchronization principle is to define specific synchronization time points at which the query answer is refreshed in response to the input stream tuples. Input stream tuples that arrive between two consecutive synchronization points are not propagated immediately to produce query outputs. Instead, the tuples are accumulated and propagated simultaneously at the following synchronization point.

Similar to the *slide* parameter, the synchronization time points are specified independently for each input stream in the query. Each input stream, say S , is mapped to a corresponding *synchronized relation* $\mathcal{R}_{Sync}(S)$ that is modified by the input stream tuples *only* at the time points that are




Figure 4. Illustrating Synchronized Relations.

specified by the synchronization stream, Sync. For example, a *slide* parameter of two time units is specified by the synchronization stream Sync₂: 0, 2, 4, 6, ... In Section 3.2 we show how to define and construct synchronization streams.

Definition 2. Synchronized relation. A synchronized relation $\mathcal{R}_{Sync}(S)$ is a time-varying relation such that $\mathcal{R}_{Sync}(S) = R[S(T)] \forall T \in Sync$.

Example 7 This example illustrates the mapping from the input stream RoomTempStr, to the corresponding synchronized relation $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$. We use the same input stream tuples as in Example 2. Figure 4 shows the synchronized relation $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$, that is modified by the input stream tuples at time points: 2, 4, 6, ... For example, at time 1, $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$ is empty while the insertion of Room “a” is not reflected in $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$ until time 2.

Example 8 For the temperature monitoring query Q_3 , to achieve the coarser refresh (every two minutes) we use the synchronization stream Sync₂. Then, the view HotRooms₁ is expressed as follows:

```
create STREAMED view HotRooms1 As
select RoomID, Temperature
from  $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$  R
where R.Temperature > 80
```

Due to the use of Sync₂, $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$ is modified every two minutes. As a result, HotRooms₁ is refreshed every two minutes as is originally requested by Q_3 .

Example 9 Figure 5 shows the execution of HotRooms₁ and the subsequent Q_4 when using the synchronization principle. For simplicity, we assume that the basic time unit is “minute”. Hence, HotRooms₁’s answer is refreshed every two time units. Assume that the following input stream S_{in} has arrived at HotRooms₁: +<a,105>1, +<b,110>3, +<c,97>4, +<d,75>5, u<a,75>7. In Figure 5, Sync₂ represents HotRooms₁’s synchronization stream




Figure 5. The Synchronization Principle.

while S_1 shows HotRooms_1 's output. The input tuple $\langle a, 105 \rangle 1$ that arrived at time 1 results in producing the tuple $\langle a, 105 \rangle 2$ at time 2, which is the first synchronization time point after 1. Similarly, $\langle b, 110 \rangle 3$ results in producing $\langle b, 110 \rangle 4$, and $\langle a, 75 \rangle 7$ results in producing $\langle a, 75 \rangle 8$.

Query composition. S_1 is used as input in Q_4 , which uses the synchronization stream Sync_4 : 0, 4, 8, ... As a result, tuple $\langle a, 105 \rangle 2$ that arrives at Q_4 at time 2 results in producing the tuple $\langle a, 105 \rangle 4$ at time 4 in S_2 . Other tuples are processed similarly by Q_4 's pipeline.

Timestamps of the output stream tuples. Timestamps need to be attached to the output tuples from a STREAMED view so that the output stream can be used as input in another continuous query. When considering the synchronization principle, an input tuple possesses two timestamps as follows. (1) The *Arrival* timestamp that is equal to the timestamp attribute of the tuple, and (2) The *Release* timestamp that is equal to the time at which the input tuple is reflected in the query. The *arrival* and *release* timestamps may not be equal for tuples that arrive between two consecutive synchronization points. However, the timestamp of an output tuple is constructed as a function of the *release* timestamp(s) of the input tuple(s) that caused this output because the output necessarily follows the *release* time point. For example, in Example 9, the input tuple $\langle a, 105 \rangle 2$ in Q_4 , that has arrival timestamp of value 2, has a release timestamp of value 4. As a result, $\langle a, 105 \rangle 2$ results in producing the output tuple $\langle a, 105 \rangle 4$ which has a timestamp equals to 4.

3.2 Synchronization Streams

Before proceeding to the algebraic foundations of SyncSQL, this section discusses synchronization streams in

more detail. Basically, a synchronization stream specifies a sequence of time points. However, the representation of a synchronization stream follows the tagged stream semantics in Section 2.1, and is treated as any other stream. A synchronization stream is characterized by the following. (a) The underlying stream schema has only one attribute, termed *Timepoint*, and (b) tuples in the stream are *insert* operations of the form “ $\langle \text{Timepoint} \rangle \text{Timepoint}$ ”. Like any other stream, a synchronization stream Sync has a corresponding time-varying relation $\mathcal{R}(\text{Sync})$ where each “ $\langle \text{Timepoint} \rangle \text{Timepoint}$ ” adds a new time point of value *Timepoint* to $\mathcal{R}(\text{Sync})$. The default clock stream, clockStr : $\langle 0 \rangle 0$, $\langle 1 \rangle 1$, $\langle 2 \rangle 2$, $\langle 3 \rangle 3$, ... is the finest granularity synchronization stream where there is a time point for every clock tick. Coarser synchronization streams can be constructed using SyncSQL expressions over clockStr .

Example 10 The synchronization stream that has a tick every two time points is constructed from clockStr using the following view expression:

```
create STREAMED view Sync2 as
select C.Timepoint
from  $\mathcal{R}(\text{clockStr})$  C
where C.Timepoint mod 2 = 0
```

A tuple is produced in the output of Sync_2 whenever a tuple, c , is inserted in $\mathcal{R}(\text{clockStr})$ and $c.\text{Timepoint}$ qualifies the predicate “ $c.\text{Timepoint} \bmod 2 = 0$ ”. The output of Sync_2 is as follows: $\langle 0 \rangle 0$, $\langle 2 \rangle 2$, $\langle 4 \rangle 4$, $\langle 6 \rangle 6$, ... which indicates the time points: 0, 2, 4, 6, ..., which is the same as Sync_2 that is used in Example 8.

Composition of synchronization streams. The fact that synchronization streams are treated as regular streams allows us to compose synchronization streams to define a larger class of synchronization streams. For example, a synchronization stream can be defined as the *union* or *intersection* of two or more streams.

Example 11 The following view expression produces a synchronization stream that is the union of two input synchronization streams (Note that *duplicate elimination* is required so that every time point exists only once in the output stream):

```
create STREAMED view UnionSyncStr as
select DISTINCT(Timepoint)
from  $\mathcal{R}(\text{Sync}_2)$  S2  $\cup$   $\mathcal{R}(\text{Sync}_5)$  S5
```

The output from UnionSyncStr includes a time point T whenever T belongs to either Sync_2 or Sync_5 .

Event-based synchronization: The synchronization principle enables SyncSQL to express a wider class of continuous queries including queries that use event-based refresh conditions. Synchronization streams for event-based

conditions can be constructed using SyncSQL expressions as in the following example.

Example 12 Consider another temperature monitoring query, Q_6 , that is similar to Q_4 except that Q_6 needs to be refreshed only whenever a room reports a temperature greater than 120. We use the tagged stream `TempStr`, which is defined in Example 1, to generate a synchronization stream, say `HotSync`, such that `HotSync` includes time points that corresponds to reporting a temperature greater than 120. As explained in Section 2, `TempStr` consists of only *insert* operations and its corresponding relation $\mathcal{R}(\text{TempStr})$ has three attribute: `RoomID`, `Temperature`, and `TS`. A synchronization stream, `HotSync`, can then be constructed by the following query over $\mathcal{R}(\text{TempStr})$:

```
create STREAMED view HotSync as
select R.TS
from  $\mathcal{R}(\text{TempStr})$  R
where R.Temperature > 120
```

An input tuple from `TempStr`, of the form “+<RoomID, Temperature>Timestamp”, results in an output tuple, “+<Timestamp>Timestamp”, if “Temperature” is greater than 120. `HotSync` can be, then, used as a synchronization stream for Q_6 .

3.3 The Now View

In Example 5, `FiveUnitsWindow`’s contents depend on the value of `Now`. In order to be consistent with the SyncSQL semantics, the value of `Now` is defined as a view that is continuously modified by the clock stream `clockStr`: +<0>0, +<1>1, +<2>2, ... Notice that $\mathcal{R}(\text{clockStr})$ is an append-only relation in which the value of the last inserted tuple indicates the current time, `Now`.

Example 13 The following view, `NowView`, over $\mathcal{R}(\text{clockStr})$ always contains the value of `Now`:

```
create STREAMED view NowView as
select 1 as KEY, MAX(T.Timepoint) as currTime
from  $\mathcal{R}(\text{clockStr})$  T
```

The output of `NowView` is a time-varying relation that has a primary key, `KEY`. The view always contains one tuple with key value 1, and the tuple is continuously updated in response to insertions in $\mathcal{R}(\text{clockStr})$. As tuples are appended to $\mathcal{R}(\text{clockStr})$, the function `MAX(T.Timepoint)` selects the last appended tuple that has a value equals to the current time, `Now`. The output stream from `NowView` is as follows: +<1, 0>0, u<1, 1>1, u<1, 2>2, u<1, 3>3, ..., where the tuple u<1, 3>3, for example, means update the record with `KEY` value 1, to have a `currTime` value 3. The view `FiveUnitsWindow` over stream S from Example 5 is

rewritten in terms of `NowView` as follows:

```
create STREAMED view FiveUnitsWindow as
select R.*
from  $\mathcal{R}(S)$  R,  $\mathcal{R}(\text{NowView})$  N
where N.currTime - 5 < R.TS ≤ N.currTime
```

Example 14 This example shows how to use SyncSQL to define a sliding window that is defined by both the *range* and *slide* parameters. Assume we extend the definition of the sliding window in Example 5 such that the window is refreshed every 2 time units instead of every point in time (this corresponds to a sliding window with range 5 units and slide 2 units). In a way similar to using `clockStr` to define `NowView`, we use the synchronization stream `Sync2` to define a view, say `TwoUnitsSlide`, as follows:

```
create STREAMED view TwoUnitsSlide as
select 1 as KEY, MAX(T.Timepoint) as currTime
from  $\mathcal{R}(\text{Sync}_2)$  T
```

The `TwoUnitsSlide` view consists of only one tuple that is updated by `Sync2`’s tuples. The `TwoUnitsSlide` view can, then, be used to express a sliding window of range 5 and slide 2 over a stream S as follows:

```
create STREAMED view RangeFiveSlideTwo as
select R.*
from  $\mathcal{R}_{\text{Sync}_2}(S)$  R,  $\mathcal{R}(\text{TwoUnitsSlide})$  N
where N.currTime - 5 < R.TS ≤ N.currTime
```

Only at the time points that belongs to `Sync2`, `RangeFiveSlideTwo`’s output is refreshed to include S ’s tuples that arrived in the last 5 time units.

4 SyncSQL Algebra

In this section, we lay the algebraic foundation for SyncSQL as the basis for efficient execution and optimization of SyncSQL queries. One of our goals while developing SyncSQL is to minimize the extensions over the well-known relational algebra. By leveraging the relational algebra, SyncSQL execution and optimization can benefit from rich literature of traditional databases. We achieved our goal by mapping continuous queries to the traditional materialized views. However, the synchronization principle differentiates continuous queries from materialized views. In this section, we introduce the data types and transformation rules that are imposed by the synchronization principle.

4.1 Data Types

As discussed in Section 2, although the inputs in a SyncSQL expressions are tagged streams, SyncSQL queries are expressed over the input streams’ corresponding relations. The output from a SyncSQL expression is another relation that can be mapped into a tagged stream. Basically, a synchronized relation is the main data type over

which SyncSQL expressions are expressed. A synchronized relation, $\mathcal{R}_{Sync}(S)$, possesses two logical properties:

- **Data** (or state) that is represented by the tuples in the relation and is extracted from the input stream S .
- **Time** that is represented by the time points at which the relation is modified by the underlying stream S and is extracted from the synchronization stream $Sync$.

The time points at which $\mathcal{R}_{Sync_i}(S_i)$ reflects **all** S_i 's tuples up to time T_i (i.e., $\mathcal{R}_{Sync_i}(S_i) = R[S_i(T_i)]$) are called "**full synchronization points**" for the relation. Basically, the time points $T_i \in Sync_i$ represent the full synchronization points for $\mathcal{R}_{Sync_i}(S_i)$. On the other hand, the time points at which $\mathcal{R}_{Sync_i}(S_i)$ does not reflect all S_i tuples are called "**partial synchronization points**". Basically, the time points that lies between two consecutive $Sync_i$ represent the partial synchronization points for $\mathcal{R}_{Sync_i}(S_i)$.

4.2 Operators

Operators in SyncSQL are classified into three classes: Stream-to-Relation (S2R), Relation-to-Relation (R2R), and Relation-to-Stream (R2S). This operator classification is similar to the classification used by CQL [2], but with different instantiations of operators in each class. Basically, the S2R class includes one operator that is used to express the desired synchronization points. The R2R class includes the traditional relational operators. Finally, the R2S class includes one operator that is used in a query to express the desire of an incremental output.

4.2.1 S2R and R2S Operators

The stream-to-relation operator \mathcal{R} . \mathcal{R} takes a tagged stream of modify operations, say S , as input and a synchronization stream, say $Sync$, as a parameter and produces a synchronized relation, $\mathcal{R}_{Sync}(S)$, as output. Similar to $\mathcal{R}(S)$, $\mathcal{R}_{Sync}(S)$'s schema consists of S 's underlying schema plus the timestamp attribute TS as explained in Section 2.1. Basically, \mathcal{R} performs the following: (1) buffers S 's tuples, (2) modifies the output relation by the buffered tuples at every $Sync$'s point, T . The output relation at $Sync$'s point T is denoted by $R[S(T)]$.

The relation-to-stream operator ξ . ξ takes a synchronized relation, $\mathcal{R}_{Sync}(S)$, as input and produces a tagged stream as output. ξ produces output tuples only when the input relation is modified (i.e., at the time points that belongs to $Sync$). Basically, at every $Sync$'s time point, T , the input relation is $R[S(T)]$ and ξ performs the following: (1) generates delta tuples that represent $\mathcal{R}_{Sync}(S)$'s modifications (i.e., +, u, or -) since the previous synchronization point, (2) assigns T as the timestamp of every generated




Figure 6. The Relation-to-Stream Operator.

tuple and produces the delta tuples in the output. Notice that non append-only relations can be mapped to streams according to the SyncSQL stream semantics.

Example 15 The functionality of the S2R operator, \mathcal{R} , has been demonstrated in Figure 4. In this example we demonstrate the functionality of the R2S operator, ξ . Figure 6 gives the mapping from $\mathcal{R}_{Sync_2}(\text{RoomTempStr})$, that is given in Figure 4, to the corresponding stream, S_{out} (i.e., $S_{out} = \xi(\mathcal{R}_{Sync_2}(\text{RoomTempStr}))$). For example, at time 4, ξ produces $+<c, 80, 3>4$ and $u<a, 95, 4>4$ as the differences since the previous synchronization point, 2. Notice that, ξ assigns timestamps to the output stream tuples so that the output stream can be used as input in another continuous query. Notice that every S 's tuple has a corresponding tuple in S_{out} , although because of the synchronization, the corresponding tuples in S and S_{out} may not have the same timestamps. Notice also that S_{out} 's schema differs from S 's schema by having an additional attribute that corresponds to the $Timestamp$ field of S 's tuples (e.g., the tuple $+<a>1$ in S is mapped to $+<a, 1>2$ in S_{out}). This additional attribute is due to the composition of \mathcal{R} and ξ operators. Recall that an additional TS attribute is added by \mathcal{R} when S is mapped to $\mathcal{R}_{Sync_2}(S)$. As a result, TS is produced as an attribute in S_{out} 's schema when ξ maps $\mathcal{R}_{Sync_2}(S)$ to S_{out} . TS can be eliminated from $\mathcal{R}_{Sync_2}(S)$ by using an R2R project operator, π .

4.2.2 Extended R2R Operators.

The R2R class of operators includes extended versions of the traditional relational operators (e.g., σ , π , \bowtie , \cup , \cap , and $-$). The semantics of R2R operators in SyncSQL are the same as in the traditional relational algebra. The difference in SyncSQL is that the operators are continuous (not snapshot). A continuous operator means that, inputs to the operator are continuously changing and the operator is continuously running to produce a new output whenever any of the inputs changes.

As with materialized views, the output from an R2R operator is refreshed whenever any of the input relations is modified. For a unary operator (e.g., σ , π), the output relation is modified at the input relation's synchronization points. In other words, the synchronization points (full and

partial) for the output relation are the same as those for the input relation. However, a problem arises in non-unary operators if the input relations have different synchronization points. Notice that operating over relations with different synchronization points is similar to operating over windowed streams with different *slide* parameters (the latter has not been discussed in the existing literature).

For example, consider a binary operator, say \circ , that has two input synchronized relations, $\mathcal{R}_{Sync_1}(S_1)$ and $\mathcal{R}_{Sync_2}(S_2)$. The input relation $\mathcal{R}_{Sync_1}(S_1)$ is modified at every time point in $Sync_1$ while $\mathcal{R}_{Sync_2}(S_2)$ is modified at every point in $Sync_2$. As a result, the output of \circ is modified at every point $T \in (Sync_1 \cup Sync_2)$. The output of \circ is interpreted as follows:

- For every time point $T_1 \in (Sync_1 - (Sync_1 \cap Sync_2))$, T_1 is a *full* synchronization point for $\mathcal{R}_{Sync_1}(S_1)$ (i.e., at time T_1 , $\mathcal{R}_{Sync_1}(S_1)$ reflects **all** S_1 tuples up to T_1). However, the same point T_1 is a *partial* synchronization point for $\mathcal{R}_{Sync_2}(S_2)$ (i.e., at T_1 , $\mathcal{R}_{Sync_2}(S_2)$ does not reflect all S_2 tuples up to T_1). Hence, as a result, T_1 is a *partial* synchronization point for the output of \circ because at time T_1 , the output of \circ does not reflect **all** input tuples from **all** input streams.
- Similarly, every time point $T_2 \in (Sync_2 - (Sync_1 \cap Sync_2))$ is a *partial* synchronization point for the output of \circ because it does not reflect all input tuples from all input streams.
- Every time point $T \in (Sync_1 \cap Sync_2)$ is a *full* synchronization point for the output of \circ since it reflects all input tuples from all input streams.

Proposition 1. Unary operators. The output of a unary R2R operator, say Θ , over a synchronized relation, say $\mathcal{R}_{Sync}(S)$, is another synchronized relation, denoted by $\Theta(\mathcal{R}_{Sync}(S))$, such that:

$\forall T \in Sync, T$ is a *full sync point*, and
 $\Theta(\mathcal{R}_{Sync}(S)) = \Theta(R[S(T)])$, while
 $\forall T \notin Sync, T$ is a *partial sync point*, and
 $\Theta(\mathcal{R}_{Sync}(S)) = \Theta(R[S(\tilde{T})])$
 where $\tilde{T} = \max(t \in Sync \text{ and } t < T)$

Proposition 2. Binary operators. The output of a binary R2R operator, say Θ , over two synchronized relations, say $\mathcal{R}_{Sync_1}(S_1)$ and $\mathcal{R}_{Sync_2}(S_2)$, is a synchronized relation, denoted by $\mathcal{R}_{Sync_1}(S_1) \Theta \mathcal{R}_{Sync_2}(S_2)$, such that:

(1) $\forall T \in Sync_1 \cap Sync_2$,
 T is a *full sync point*, and,
 $\mathcal{R}_{Sync_1}(S_1) \Theta \mathcal{R}_{Sync_2}(S_2) = R[S_1(T)] \Theta R[S_2(T)]$,
 (2) $\forall T \in (Sync_1 - (Sync_1 \cap Sync_2))$,

T is a *partial sync point*, and,
 $\mathcal{R}_{Sync_1}(S_1) \Theta \mathcal{R}_{Sync_2}(S_2) = R[S_1(T)] \Theta R[S_2(\tilde{T})]$,
 where $\tilde{T} = \max(t \in Sync_2 \text{ and } t < T)$,
 (3) $\forall T \in (Sync_2 - (Sync_1 \cap Sync_2))$,
 T is a *partial sync point*, and,
 $\mathcal{R}_{Sync_1}(S_1) \Theta \mathcal{R}_{Sync_2}(S_2) = R[S_1(\tilde{T})] \Theta R[S_2(T)]$,
 where $\tilde{T} = \max(t \in Sync_1 \text{ and } t < T)$

According to Proposition 1, at any time point, say \tilde{T} , that does not belong to the output synchronization stream, the output synchronized relation from a unary operator reflects the input stream only up to a time point \tilde{T} where $\tilde{T} < T$. Similarly, according to Proposition 2, at any time point, say \tilde{T} , that does not belong to the output synchronization stream, the output from a binary R2R operator reflects one input streams up to time point \tilde{T} while reflects the other input stream only up to time \tilde{T} where $\tilde{T} < T$.

Derived synchronized relations. Based on the previous discussion, the output of an R2R expression over synchronized relations is a derived synchronized relation. The output derived relation has the following logical properties:

- **State** that is derived from the input relations' states.
- **Time** that represents the time points at which the derived relation is modified. The derived relation is modified at every time point, say T , if T belongs to the *union* of the input relations' synchronization streams. The time points at which the derived relation are modified are further classified as follows: (1) *Full synchronization points*: a time point, say T , is a *full* synchronization point for the derived relation only if T is a *full* synchronization point for all the input relations. (2) *Partial synchronization points*: a time point, say T , is a *partial* synchronization point for the derived relation if T is a *partial* synchronization point for at least one of the input relations.

Query pipeline. In order to express a query over tagged stream, the SyncSQL expression is constructed as follows.

(1) S2R: transform each input stream to the corresponding synchronized relation via an \mathcal{R} operator using the desired synchronization. (2) R2R: using R2R operators, and in a way similar to traditional SQL, express the query over the synchronized relations. The output of is another synchronized relation. (3) R2S: the output synchronized relation is transformed into an incremental output via an ξ operator.

Example 16 This example shows the execution pipeline for a join query between two synchronized relations, $\mathcal{R}_{Sync_2}(S_2)$ and $\mathcal{R}_{Sync_3}(S_3)$, where $Sync_2$ ticks every 2 time units while $Sync_3$ ticks every 3 time units. The SyncSQL expression is as follows:

```
select STREAMED *
```

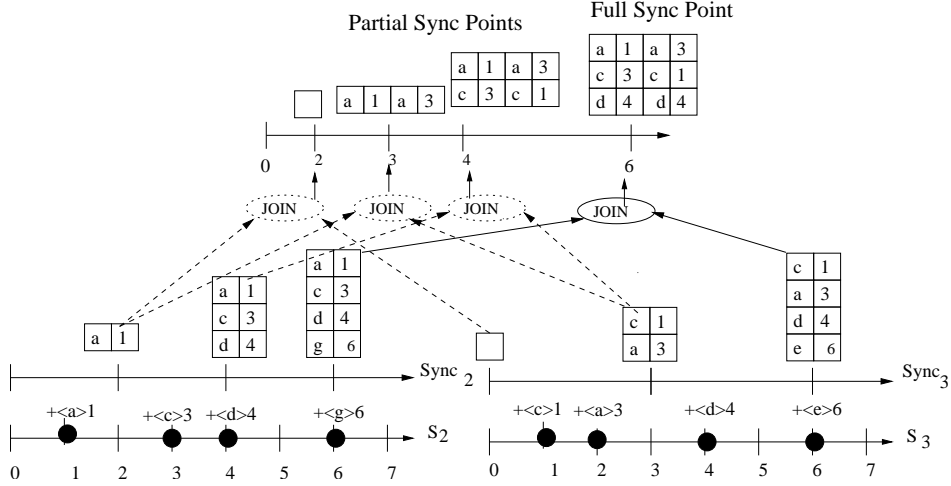


Figure 7. Joining Relations with Different Synchronization

from $\mathcal{R}_{\text{Sync}_2}(S_2) \bowtie \mathcal{R}_{\text{Sync}_3}(S_3)$
 where $R_2.ID = R_3.ID$

Figure 7 illustrates the pipeline and shows that the output of the join is refreshed at time points 2, 3, 4, and 6. The output at 2 is equal to $\mathcal{R}[S_2(2)] \bowtie \mathcal{R}[S_3(0)]$ and hence 2 is a *partial* synchronization point since it reflects S_3 only up to time 0. Similarly, 3 is a *partial* synchronization point since it reflects S_2 up to time 2. 4 also is a *partial* synchronization point since it reflects S_3 up to time 3. However, 6 is a *full* synchronization point for the output since it reflects *all* input tuples up to time 6.

4.3 Equivalences and Relationships

Achieving query composition is one of the main goals of SyncSQL. In order to achieve query composition, a query optimizer must be empowered by algorithms to reason about the equivalences and containment relationships among query expressions. In this section, we introduce preliminary relationships that are required by a query optimizer to enumerate the query plans and deduce query containment.

4.3.1 Containment Relationship among Synchronization Streams

A synchronization stream, say Sync_1 , is contained in another synchronization stream, say Sync_2 , if every time point in Sync_1 is also a time point in Sync_2 (i.e., $\mathcal{R}(\text{Sync}_1) \subseteq \mathcal{R}(\text{Sync}_2)$). Recall that, as explained in Section 3.2, a synchronization stream consists of only *insert* operations of the form $+\langle \text{Timepoint} \rangle \text{Timepoint}$. Containment relationships between synchronization streams can be deduced from the constructing SyncSQL expressions. For example,

a synchronization stream that is defined over `clockStr` by the predicate “Timepoint mod 4 = 0” (i.e., a stream that ticks every 4 time units) is contained in the synchronization stream that is defined by the predicate “Timepoint mod 2 = 0” (i.e., a stream that ticks every two time units).

Proposition 3. $\mathcal{R}(\text{Sync}_1) \subseteq \mathcal{R}(\text{Sync}_2)$ if $\forall I \in \text{Sync}_1 \Rightarrow I \in \text{Sync}_2$ where I is an insert operation of the form “ $+\langle T \rangle T$ ”.

4.3.2 Containment Relationships among Synchronized Relations

Reasoning about containment relationships between two synchronized relations must consider the two logical properties, state and time, of the relation. For example, consider two synchronized relations, $\mathcal{R}_{\text{Sync}_i}(S)$ and $\mathcal{R}_{\text{Sync}_j}(S)$, that are defined over the same stream S . Notice that, the *states* of $\mathcal{R}_{\text{Sync}_i}(S)$ and $\mathcal{R}_{\text{Sync}_j}(S)$ may not be equal at every time point if Sync_i and Sync_j are not the same. However, if Sync_i is contained in Sync_j , then $\mathcal{R}_{\text{Sync}_i}(S)$ is *contained* in $\mathcal{R}_{\text{Sync}_j}(S)$. The containment relationship means that every *full* synchronization time point of $\mathcal{R}_{\text{Sync}_i}(S)$ is also a *full* synchronization time point of $\mathcal{R}_{\text{Sync}_j}(S)$. The containment relationship is beneficial since $\mathcal{R}_{\text{Sync}_i}(S)$ can be computed from $\mathcal{R}_{\text{Sync}_j}(S)$ without accessing S . Notice that, the containment relationship is judged based only on the *full* synchronization time points of the relation because those are the time points of interest to the issuer of a query.

Theorem 1 For any stream S , a synchronized relation $\mathcal{R}_{\text{Sync}_i}(S)$ is contained in $\mathcal{R}_{\text{Sync}_j}(S)$ if $\mathcal{R}(\text{Sync}_i) \subseteq \mathcal{R}(\text{Sync}_j)$.

Proof:

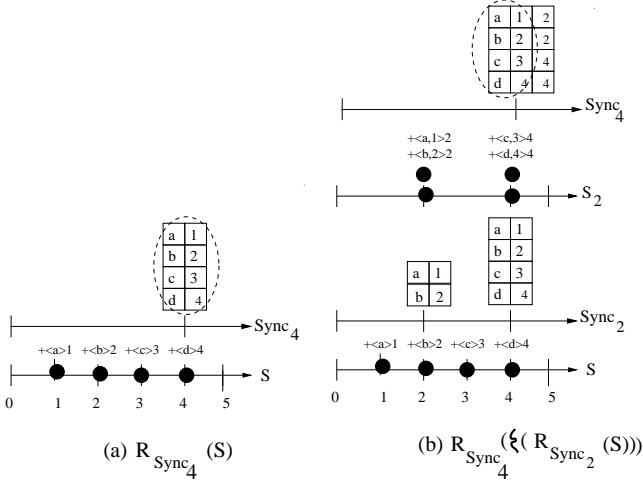


Figure 8. Relation Containment

1. Based on Definition 2:
 $\mathcal{R}_{Sync_j}(S) = \mathcal{R}[S(T)] \forall T \in \mathcal{R}(Sync_j)$;
2. Given that $\mathcal{R}(Sync_i) \subseteq \mathcal{R}(Sync_j)$, then, based on Proposition 3
 $\forall T \in \mathcal{R}(Sync_i) \Rightarrow T \in \mathcal{R}(Sync_j)$;
3. From 1 and 2 above,
 $\mathcal{R}_{Sync_j}(S) = \mathcal{R}[S(T)] \forall T \in \mathcal{R}(Sync_i)$;
4. Based on Definition 2,
 $\mathcal{R}_{Sync_i}(S) = \mathcal{R}[S(T)] \forall T \in \mathcal{R}(Sync_i)$;
5. From 3 and 4 above:
 $\mathcal{R}_{Sync_i}(S) = \mathcal{R}_{Sync_j}(S) = \mathcal{R}[S(T)]$
 $\forall T \in \mathcal{R}(Sync_i)$.

Corollary 1. If $\mathcal{R}(Sync_i) \subseteq \mathcal{R}(Sync_j)$, then $\mathcal{R}_{Sync_i}(S) \subseteq \mathcal{R}_{Sync_i}(\xi(\mathcal{R}_{Sync_j}(S)))$.

Corollary 1 means that $\mathcal{R}_{Sync_i}(S)$ can be constructed from $\mathcal{R}_{Sync_j}(S)$ without accessing S . This is done by applying $Sync_i$ over the output stream from $\xi(\mathcal{R}_{Sync_j}(S))$.

Example 17 This example illustrates Theorem 1 and Corollary 1. Consider two synchronization streams, $Sync_2$ and $Sync_4$, where $\mathcal{R}(Sync_4) \subset \mathcal{R}(Sync_2)$. Figure 8a gives the derivation of $\mathcal{R}_{Sync_4}(S)$ while Figure 8b gives the derivation of $\mathcal{R}_{Sync_4}(\xi(\mathcal{R}_{Sync_2}(S)))$. Notice that, all the *full* synchronization points for $\mathcal{R}_{Sync_4}(S)$ are also *full* synchronization points for $\mathcal{R}_{Sync_2}(S)$. Moreover, if only the *STREAMED* version of $\mathcal{R}_{Sync_2}(S)$ is available (i.e., $\xi(\mathcal{R}_{Sync_2}(S))$ or S_2 in Figure 8b), $\mathcal{R}_{Sync_4}(S)$ can be computed by applying $Sync_4$ over S_2 (i.e., $\mathcal{R}_{Sync_4}(S)$ at time 4 is contained in $\mathcal{R}_{Sync_4}(\xi(\mathcal{R}_{Sync_2}(S)))$ at time 4).

4.3.3 Commutability between Synchronization and R2R Operators

R2R operators in a $SyncSQL$ expression are executed over synchronized relations. In this section, we show that the order of applying the synchronization and R2R operators can be switched. The commutability between the synchronization and R2R operators allows executing the query pipeline over finest granularity relations and hence allows sharing the execution among queries that have similar R2R operators but with different synchronization points.

Corollary 2. For any *unary* R2R operator, say Θ , $\forall T$ such that T is a full synchronization point of $\Theta(\mathcal{R}_{Sync}(S))$, T is a full synchronization point of $\mathcal{R}_{Sync}(\xi(\Theta(\mathcal{R}(S))))$.

Corollary 3. For any *binary* R2R operator, say Θ , $\forall T$ such that T is a full synchronization point of $\mathcal{R}_{Sync_1}(S_1) \Theta \mathcal{R}_{Sync_2}(S_2)$, T is a full synchronization point of $\mathcal{R}_{Sync_1} \cap Sync_2(\xi(\mathcal{R}(S_1) \Theta \mathcal{R}(S_2)))$.

The main idea of Corollaries 2 and 3 is that we can pull the synchronization streams out of an R2R operator. Basically, an R2R operator can be executed over finest granularity relations and produce a finest granularity output. Then, the desired synchronization is applied over the fine granularity output. Notice that, Corollaries 2 and 3 can also be used in the opposite direction by a query optimizer to push the synchronization inside R2R operators and, hence, reducing the number of operator executions.

Based on Corollaries 2 and 3, a $SyncSQL$ expression can be executed as follows: (1) transform the input streams to the finest granularity synchronized relations, $\mathcal{R}(S)$, using the finest granularity synchronization stream (i.e., the clock stream), (2) execute the query pipeline over the finest granularity input producing a fine granularity output relation, (3) map the output relation to a stream using ξ , and finally (4) transform the output stream to the desired synchronized output using \mathcal{R} .

5 Shared Execution using Query Composition

In this section, we introduce a query matching algorithm for $SyncSQL$ expressions. The goal of the algorithm is that, given a $SyncSQL$ query, say Q_i , the algorithm determines whether Q_i (or a part of it) is contained in another view, say Q_j . If such Q_j exists, the algorithm re-writes Q_i in terms of Q_j in a way similar to answering queries using views in traditional databases.

5.1 Skinning $SyncSQL$ Expressions

To reason about containment of $SyncSQL$ expressions, we isolate the synchronization streams out of the expression. We term the resulting form of the expressions a

“skinned” form. The skinned form of a SyncSQL expression is an equivalent expression that consists of: (a) a global synchronization stream that specifies the *full* synchronization points of the expression, and (b) a SQL expression over finest granularity relations. Corollaries 2 and 3 are used to transform any SyncSQL expression into the corresponding skinned form.

Theorem 2 Any SyncSQL expression has an equivalent normal form.

Theorem 2 is proved using Corollaries 2 and 3.

Example 18 This example derives the normal form for the SyncSQL expression $Q = \sigma(\mathcal{R}_{Sync_1}(S_1) \bowtie \mathcal{R}_{Sync_2}(S_2))$. The derivation is performed in two steps as follows:

-Using Corollary 3, pull the synchronization streams out of the join operator.

$$Q = \sigma(\mathcal{R}_{Sync_1} \cap_{Sync_2}(\xi(\mathcal{R}(S_1) \bowtie \mathcal{R}(S_2))))$$

-Using Corollary 2, pull the synchronization stream out of the selection operator.

$$Q = \mathcal{R}_{Sync_1} \cap_{Sync_2}(\xi(\sigma(\mathcal{R}(S_1) \bowtie \mathcal{R}(S_2))))$$

The constructed normal form indicates that Q is equivalent to a synchronized relation with the following: (1) **Data:** $\sigma(\mathcal{R}(S_1) \bowtie \mathcal{R}(S_2))$, and (2) **Time:** $Sync_1 \cap Sync_2$.

5.2 Query Matching

SyncSQL query matching is similar to “view exploitation” in materialized views [16, 19]. However, SyncSQL queries differ from the traditional materialized views by the notion of synchronization. A matching algorithm for SyncSQL expressions matches the two parts of the skinned forms: the query expression and the global synchronization points.

After introducing the main tools, we now give the high-level steps of the query matching algorithm. The input to the algorithm is a SyncSQL query expression, say Q , and a set of skinned forms for the concurrent queries.

Algorithm SyncSQL-Expression-Matching:

1. Using Corollaries 2 and 3, transform Q to the corresponding normal form by constructing the two components: (1) Q ’s data, Q^d , and (2) Q ’s synchronization, Q^s ;
2. Match Q^d with data parts of the other input normal forms using a view matching algorithm from the materialized view literature (e.g., [16]). The result of the matching is a normal form (if any) for a matching expression, say \tilde{Q} , such that $Q^d \subset \tilde{Q}^d$;
3. If such \tilde{Q} exists, check whether $Q^s \subset \tilde{Q}^s$;

4. If $Q^s \subset \tilde{Q}^s$, then rewrite Q^d in terms of \tilde{Q}^d using the same algorithm used in Step 2 above. The output expression of the re-write operation is denoted as Q^D ;

5. The input query, Q , is then equivalent to the synchronized relation with: (1) **Data:** Q^D , and (2) **Time:** Q^s .

Notice that, the query matching algorithm is used to match an input query against a set of already existing views. On the other hand, if we know the whole set of queries in advance, the skinned forms are constructed using the greatest common divisor of all synchronization streams instead of the default clock stream.

Example 19 This example illustrates the steps performed to match the temperature monitoring query Q_4 with the view $HotRooms_1$ as explained in Example 4. Assume that the input expressions are as follows:

$$HotRooms_1 = \sigma_{Temp > 80}(\mathcal{R}_{Sync_2}(RoomTempStr))$$

$$Q_4 = \sigma_{Temp > 100}(\mathcal{R}_{Sync_4}(RoomTempStr))$$

The corresponding normal forms for the two expressions are as follows:

$$HotRooms_1 = \mathcal{R}_{Sync_2}(\xi(\sigma_{Temp > 80}(\mathcal{R}(RoomTempStr))))$$

$$Q_4 = \mathcal{R}_{Sync_4}(\xi(\sigma_{Temp > 100}(\mathcal{R}(RoomTempStr))))$$

By Comparing the two normal forms we can conclude that: (1) $\mathcal{R}(Sync_4) \subset \mathcal{R}(Sync_2)$, and (2) using a view matching algorithm (e.g., [16]) shows that the “Temp > 100” \subset “Temp > 80”. Then, the algorithm concludes that $Q_4 \subset HotRooms_1$. Then, Q_4 is re-written as follows: $Q_4 = \sigma_{Temp > 100}(\mathcal{R}_{Sync_4}(\xi(\mathcal{R}(HotRooms_1))))$.

6 Incremental Execution Model

Although the goal of this paper is to introduce the SyncSQL semantics for queries over data streams, in this section we briefly outline an execution model for SyncSQL queries. Detailed implementation and optimization techniques is beyond the scope of this paper.

As discussed in Section 2, a SyncSQL query over streams is semantically equivalent to a materialized view over the input streams’ relational views. Similar to materialized views, the straightforward way to keep the query answer (or view) consistent with the underlying relations is to re-evaluate the query expression whenever any of the inputs is modified. However, incremental approaches have been proposed to reduce the cost of maintaining the materialized views. In the incremental maintenance of materialized views, instead of re-evaluating the view expression, only the changes in the input relations are processed in order to produce a corresponding set of changes in the output. SyncSQL physical execution plans follows the incremental maintenance approach of materialized views. Basically,

at every synchronization time point, a differential operator processes only the modifications in the input relations and produce a corresponding set of modifications in the output.

As discussed in Section 4, inputs and outputs in any R2R operator are synchronized relations. According to SyncSQL algebra, a relational operation (e.g., σ or \bowtie) over an input stream S_{in} is executed as follows. At every synchronization time point, say T_1 , S_{in} is mapped to a corresponding relation, $R[S_{in}(T_1)]$. Then, the relational operation, say σ , is executed over $R[S_{in}(T_1)]$ and produce a corresponding output relation, say $R[S_{out}(T_1)]$. When the input relation is modified at a following synchronization time point, say T_2 , σ is re-executed over $R[S_{in}(T_2)]$ and produce the corresponding output relation $R[S_{out}(T_2)]$. If the output of σ is needed to be STREAMED, a ξ operator is executed at time T_2 to produce tuples in the output stream S_{out} that represent the deltas between $R[S_{out}(T_1)]$ and $R[S_{out}(T_2)]$. The delta tuples is a set of $+$, u or $-$ operations that need to be performed over $R[S_{out}(T_1)]$ in order to get $R[S_{out}(T_2)]$. In short, SyncSQL algebra assumes that an R2R operator is *re-executed* at every synchronization time point.

In contrast to the algebra, SyncSQL physical execution plans employs an incremental approach. At every synchronization time point, an incremental relational operator processes only the modifications in the input relations and produce a corresponding set of modifications in the output relation. For example, at a synchronization time point, T_2 , the incremental σ operator processes a set of delta tuples between $R[S_{in}(T_2)]$ and $R[S_{in}(T_1)]$ and produce another set of delta tuples between $R[S_{out}(T_2)]$ and $R[S_{out}(T_1)]$.

6.1 Derived Operators

The S2S counterparts of R2R operators. A SyncSQL execution plan consist of a set of S2S operators where each R2R operator (e.g., σ and \bowtie) has a corresponding incremental (or differential) S2S operator (e.g., σ^d and \bowtie^d). Basically, the functionality of an S2S operator is composed of three functions (S2R, R2R, then R2S) as follows: (1) S2R: takes an input modification tuple (i.e., $+$, u , or $-$) and apply the modification to the operator’s internal state. (2) R2R: perform the relational operator’s function over the operator’s internal state. (3) R2S: report the modifications in the internal state as an output tagged stream. Detailed implementation of S2S operators is addressed in [14].

The relationship between the input and output tagged streams from an S2S operator is defined algebraically by differential equations [17]. The functionality of a differential operator, say θ , is defined by two equations: one equation defines the modifications in θ ’s output in response to an *insert* in θ ’s input while the other equation defines the

changes in θ ’s output in response to a *delete* in θ ’s input. An *update* in θ ’s input is processed as a *deletion* of the old tuple followed by an *insertion* of the new tuple. For example, the functionality of the differential σ is defined by the following equations:

$$\begin{aligned}\sigma_p(R+r) &= \sigma_p(R) + \sigma_p(r) \\ \sigma_p(R-r) &= \sigma_p(R) - \sigma_p(r)\end{aligned}$$

where $+r$ ($-r$) represents the insertion (deletion) of a tuple r into (from) σ ’s input relation R , while $+\sigma_p(r)$ ($-\sigma_p(r)$) represents the corresponding insertion (deletion) into σ ’s output relation, $\sigma_p(R)$. Algebra for the various differential operators is introduced in [17].

The S2S counterpart of the S2R operator. In order to apply the synchronization principle with S2S operators, we introduce the regulator operator, ζ as the S2S counterpart of \mathfrak{R} . Similar to \mathfrak{R} , ζ takes a stream, S_{in} , as input, a synchronization stream, $Sync$, as a parameter and produces another stream, S_{out} , as output where

$$S_{out} = \zeta_{Sync}(S_{in}) = \xi(\mathfrak{R}_{Sync}(S_{in})).$$

Notice that, as discussed in Section 4.2 and Example 15, the schema of the resulting stream from $\xi(\mathfrak{R}_{Sync}(S_{in}))$ differs from S_{in} ’s underlying schema by having an additional timestamp attribute that corresponds to the *arrival* timestamp of S_{in} ’s tuples. The additional timestamp attribute is used to evaluate time-based predicates (if any) over S_{in} and is also included in the output stream, S_{out} , from ζ . Basically, ζ works as follows: buffers the input stream tuples and at every synchronization time point, say T , ζ performs the following for each buffered input tuple of the form “Type<Attributes>Timestamp”: (1) constructs a corresponding tuple of the form “Type<Attributes, Timestamp>”, by pushing the arrival timestamp, $Timestamp$, inside the tuple’s schema, and (2) assigns a timestamp to the tuple that is equal to the release time, or T . As a result, ζ ’s output tuples will have the form “Type<Attributes, Timestamp>syncTimestamp”.

Handling timestamps by the physical operators. An output tuple from ζ has two timestamps as follows: (1) $Timestamp$ that is equal to the tuple’s *arrival* timestamp and is used by the subsequent R2R differential operators to evaluate time-based predicates (if any), and (2) $syncTimestamp$ that is equal to the tuple’s *release* timestamp and is propagated by the subsequent R2R operators to the corresponding output tuples.

Example 20 This example shows the functionality of the regulator operator, ζ . Consider the same S and $Sync_2$ as those used in Example 7. Figure 9 shows S and the corresponding $\zeta_{Sync_2}(S)$. ζ transforms, for example, $+\langle a \rangle 1$ into $+\langle a, 1 \rangle 2$ by pushing the arrival timestamp of value 1 into the schema and attaching the release timestamp of value 2 as the timestamp of the output tuple. Figures 6 and 9 show that $\zeta_{Sync_2}(S) = \xi(\mathfrak{R}_{Sync_2}(S))$.




Figure 9. The Regulator, ζ , Operator.

7 Related Work

Continuous queries over data streams. Many research efforts have developed semantics and query languages for continuous queries over data streams, e.g., [2, 6, 7, 8, 11, 25]. The existing continuous query languages restrict the stream definition to the representation of an append-only relation. The restricted stream definition limits the set of queries that can produce streams as output. This is because, even if the input streams represent append-only relations, a continuous query may produce non-append only output. Different approaches have been followed by the existing languages to handle the non append-only *outputs* as follows:

-Restricted expressibility: To guarantee that the output of the query can be incrementally produced as a stream, a language restricts the set of operators that can be used to express queries over data streams. The restricted set of operators includes, for example, Select, Project, and Union. Sliding windows with the window-per-stream usage, for example, are not allowed since they produce non-append only output. Examples of systems that follow this approach include Aurora [7], Cougar [6], and GigaScope [11].

-Non-incremental output streams: Produce the output of the query in a *non-incremental* manner by representing the output as a relation then periodically stream out the relation. Notice that this non-incremental output stream does not follow the input stream definition and, hence, cannot be used as input in another query. Examples of systems that follow this approach include TelegraphCQ [8], and the RStream operator in CQL [2].

-Non-incremental output relations: Does not allow queries that produce non append-only output to produce streams. Instead, such queries produce concrete views as outputs. Moreover, only *snapshot* queries are allowed over the view. A snapshot query has to be re-issued in order to know the modifications in the view. This approach is followed by ESL [25].

-Divided output: CQL [2] divides the query into two separate queries that produce append-only streams such that one query produces a stream, *IStream*, to represent the inserted tuples and the other query produces a stream,

DStream, to represent the deleted tuples. It is unclear how to compose the two streams in order to produce a single output stream that can be used as input in another query.

SyncSQL semantics avoids these previous limitations by allowing the output of any continuous query to be produced incrementally in a single stream.

There are two SQL-based languages that are closest to SyncSQL: CQL [2] and ESL [25]. SyncSQL uses the same three classes of operators (i.e., S2R, R2R, and R2S) as that of CQL but use a different instantiation of operators in each class. CQL defines two types of sliding windows (time-based and tuple-based) and defines the window as an S2R operator. However, there are no algebraic or transformation rules to show how the window operator interacts with the other (R2R) operators in the pipeline. Moreover, semantics of non-unary operators on two streams with different *slide* parameters is not discussed. ESL [25] is another SQL-based continuous query language that is designed mainly for data mining and time-series queries. Only unary operators (e.g., selection and projection) can be used in queries to produce output streams. On the other hand, since a window function produces a non append-only output, window queries produce concrete views as output. Streams can be joined with the concrete views, but in this case, the modifications in the view do not affect the already produced stream tuples but they affect only the incoming stream tuples. ESL focuses on aggregate queries but does not thoroughly address set-based operators and queries.

Positive and negative tuples. Streams of positive and negative tuples (i.e., insert and delete tuples) are frequently used when addressing continuous query processing [1, 5, 13, 14]. However, query languages do not consider expressing queries over these modify streams. This conflict between the language and internal streams is the main obstacle in achieving continuous query composition. SyncSQL overcomes this obstacle by unifying the stream definition between the language and the execution model.

Continuous queries in traditional databases. Continuous queries are used in traditional databases before being used over data streams. Examples of systems that support continuous queries over database tables include Tapestry [24] and OpenCQ [21]. In these systems, both inputs and outputs of the continuous query are relations. Although the input relations in Tapestry are append-only, queries may produce non append-only output if the query includes either a reference to the current time (e.g., GetDate()), or a set-difference between two relations. In order to guarantee the append-only output, Tapestry uses a query transformation to transform a given query into the minimum bounding append-only query. The coarser refresh of the query is achieved via a “*FOREVER DO, SLEEP*” clause where the query is re-executed after every *SLEEP* period. On the other hand, in OpenCQ, input and output relations

can be modified by general modify operations. A continuous query is periodically re-executed and the output is produced as the delta between two consecutive query executions. Triggers are used to schedule the query re-execution. Our notion of synchronization time points is similar to OpenCQ's Triggers, but synchronization streams are distinguished by the fact that they can be generated using regular queries. Unlike Tapestry and OpenCQ, SyncSQL assumes that query inputs and outputs are streams and hence requires special handling of the timestamps. Moreover, we introduce an algebraic framework and address composition of SyncSQL expressions, which is not addressed by the previous systems.

Shared query execution. A typical streaming environment has a large number of concurrent continuous queries. Sharing the query execution is a primary task for query optimizers to address scalability. The current efforts for shared query execution focus on sharing the execution at the operator level. Shared aggregates are addressed in [4] where an aggregate operator is shared among multiple queries with different window *ranges*. Shared window join is addressed in [18]. NiagraCQ [10] proposes a framework for shared execution of non-windowed SPJ queries. Shared predicate indexing is used in [9, 10] to enhance the performance of a continuous query processor. Our approach for shared execution is distinguished from the existing approaches in that: (1) based on query composition; (2) matches window queries that differ in both the *range* and *slide* parameters, and (3) queries are examined for sharing based on a whole query expression not only at the operator level.

Materialized views: Our definitions of synchronized relations and predicate-windows enable us to benefit from the existing literature in materialized view. However, we extend the materialized view algorithms to work with synchronized relations. Our query matching algorithm extends the traditional view exploitation algorithms (e.g., [16]) by matching the synchronization time points in addition to matching the query expression. Moreover, the physical design of SyncSQL execution pipelines follows the incremental maintenance of materialized views [17].

8 Concluding Remarks

This paper provides the first language, SyncSQL, to express continuous queries over streams of modify operations. Modify streams are general since they can represent both raw input streams and streams that are generated as output from executing continuous queries. The unified definition of query inputs and outputs enables the composition of SyncSQL expressions. The paper provides the first shared execution algorithm for continuous queries that is based on query composition. Shared execution decisions are based on a query matching algorithm that is able

to reason about the equivalence and containment relationships among SyncSQL expressions. Efficient execution of SyncSQL queries is an important issue. We outlined an execution model to incrementally evaluate a SyncSQL query. Detailed implementation and optimization techniques will be reported in a separate paper.

References

- [1] D. Abadi. et al. The Design of the Borealis Stream Processing Engine. In *CIDR*, 2005.
- [2] A. Arasu, S. Babu, and J. Widom. The CQL Continuous Query Language: Semantic Foundations and Query Execution. *VLDBJ*, to appear.
- [3] A. Arasu and J. Widom. A Denotational Semantics for Continuous Queries over Streams and Relations. *SIGMOD Record*, 33(3):6–12, 2004.
- [4] A. Arasu and J. Widom. Resource Sharing in Continuous Sliding-Window Aggregates. In *VLDB*, 2004.
- [5] S. Babu, K. Munagala, J. Widom, and R. Motwani. Adaptive Caching for Continuous Queries. In *ICDE*, 2005.
- [6] P. Bonnet, J. E. Gehrke, and P. Seshadri. Towards Sensor Database Systems. In *MDM*, 2001.
- [7] D. Carney. et al. Monitoring Streams - A New Class of Data Management Applications. In *VLDB*, 2002.
- [8] S. Chandrasekaran. et al. TelegraphCQ: Continuous Dataflow Processing for an Uncertain World. In *CIDR*, 2003.
- [9] S. Chandrasekaran and M. J. Franklin. PSoup: A System for Streaming Queries over Streaming Data. *VLDBJ*, 12(2):140–156, 2003.
- [10] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang. NiagaraCQ: A Scalable Continuous Query System for Internet Databases. In *SIGMOD*, 2000.
- [11] C. D. Cranor. et al. Gigascope: A Stream Database for Network Applications. In *SIGMOD*, 2003.
- [12] A. Eisenberg. et al., SQL:2003 Has Been Published. *SIGMOD Record*, 33(1):119–126, 2004.
- [13] S. Ganguly. et al. Processing Set Expressions over Continuous Update Streams. In *SIGMOD*, 2003.
- [14] T. M. Ghanem. et al. Incremental Evaluation of Sliding-Window Queries over Data Streams. In *Purdue University Technical Report, CSD TR 04-040*.
- [15] T. M. Ghanem, W. G. Aref, and A. K. Elmagarmid. Exploiting Predicate-Window Semantics over Data Streams. *SIGMOD Record*, 35(1):3–8, 2006.
- [16] J. Goldstein and P.-Å. Larson. Optimizing Queries Using Materialized Views: A Practical, Scalable Solution. In *SIGMOD*, 2001.
- [17] T. Griffin and L. Libkin. Incremental Maintenance of Views with Duplicates. In *SIGMOD*, 1995.
- [18] M. A. Hammad, M. J. Franklin, W. G. Aref, and A. E. Elmagarmid. Scheduling for Shared Window Joins over Data Streams. In *VLDB*, 2003.
- [19] P.-Å. Larson and H. Z. Yang. Computing Queries from Derived Relations. In *VLDB*, 1985.
- [20] J. Li, D. Maier, K. Tuft, V. Papadimos, and P. Tucker. Semantics and Evaluation Techniques for Window Aggregates in Data Streams. In *SIGMOD*, 2005.
- [21] L. Liu, C. Pu, and W. Tang. Continual Queries for Internet Scale Event-Driven Information Delivery. *TKDE*, 11(4):610–628, 1999.
- [22] D. Maier, J. Li, P. Tucker, K. Tuft, and V. Papadimos. Semantics of Data Streams and Operators. In *ICDT*, 2005.

- [23] E. Ryvkina. et al. Revision Processing in a Stream Processing Engine: A High-Level Design. In *ICDE*, 2006.
- [24] D. B. Terry, D. Goldberg, D. Nichols, and B. M. Oki. Continuous Queries over Append-Only Databases. In *SIGMOD*, 1992.
- [25] WEB Information System Laboratory, UCLA, CS Department. An introduction to the Expressive Stream Language (ESL). <http://wis.cs.ucla.edu/stream-mill>.