



Evaluating Complex Queries on Streaming Graphs

Anil Pacaci, Angela Bonifati, M. Tamer Ozsu

► To cite this version:

Anil Pacaci, Angela Bonifati, M. Tamer Ozsu. Evaluating Complex Queries on Streaming Graphs. 2022 IEEE 38th International Conference on Data Engineering (ICDE), May 2022, Kuala Lumpur, Malaysia. pp.272-285, 10.1109/ICDE53745.2022.00025 . hal-03907496

HAL Id: hal-03907496

<https://hal.science/hal-03907496>

Submitted on 20 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evaluating Complex Queries on Streaming Graphs

Anil Pacaci
University of Waterloo
apacaci@uwaterloo.ca

Angela Bonifati
Lyon 1 University
angela.bonifati@univ-lyon1.fr

M. Tamer Özsu
University of Waterloo
tamer.ozsu@uwaterloo.ca

ABSTRACT

We study the problem of evaluating persistent queries over streaming graphs in a principled fashion. These queries need to be evaluated over unbounded and very high speed graph streams. We define a streaming graph data model and query model incorporating navigational queries, subgraph queries and paths as first-class citizens. To support this full-fledged query model we develop a streaming graph algebra that describes the precise semantics of persistent graph queries with their complex constructs. We present transformation rules and describe query formulation and plan generation for persistent graph queries over streaming graphs. Our implementation of a streaming graph query processor shows the feasibility of our approach and allows us to gauge the high performance gains obtained for query processing over streaming graphs.

1 INTRODUCTION

Many modern applications in various domains now operate on very high speed streaming graphs [62]. For example, Twitter’s recommendation system ingests 12K events/sec on average [31], Alibaba transaction graph processes 30K edges/sec at its peak [60]. Efficient querying of these streaming graphs is a crucial task for these applications that monitor complex graph patterns and relationships.

Existing graph DBMSs mostly follow the traditional database paradigm where data is persistent and queries are transient; consequently, they do not support persistent query semantics where queries are registered into the system and results are generated incrementally as the graph edges arrive. Persistent queries on streaming graphs enable users to continuously obtain new results on rapidly changing data, supporting online analysis and real-time query processing, the latter being an important functionality of future graph processing engines [64]. This is demonstrated in the following example.

EXAMPLE 1. *In many online social networking applications users post original content, sometimes link this to other users’ content and react to each other’s posts. We say that a user u_2 is a recentLiker for another user u_1 if u_2 has recently liked posts that are created by u_1 and u_2 , and u_1 are following each other. The goal of the recommendation service is to notify users, in real-time, of new content that are posted by others that are connected by a path of recentLiker relationship – these constraints are modeled as a complex graph pattern as the one shown in Figure 1. The service might provide the context for its recommendations by returning the full paths of people who are recent likers such as the path between users u_1 and u_k . This real-time notification task is an example of a persistent query over the streaming graph of user interactions that returns the recommended content in real-time.*

Real-world applications that feature complex graph patterns as the one shown in the above example require:

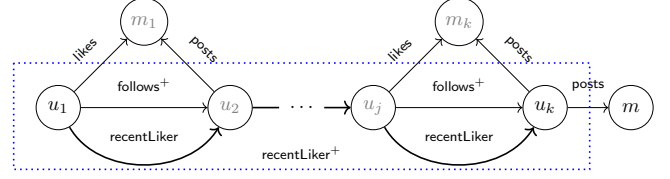


Figure 1: Complex graph pattern representing the query in Example 1

- **(R1)** subgraph queries that find matches of a given graph pattern (e.g. in Figure 1 the triangle pattern involving posts, likes and transitive closure of the follows relationship)
- **(R2)** path navigation queries that traverse paths based on user specified constraints (e.g. in Figure 1 arbitrary-length paths of the recentLiker relationship); and
- **(R3)** the ability to treat paths as first-class citizens of the data model, hence to manipulate and return paths (e.g. in Figure 1 the query returns the full paths of recentLiker).

Even in the context of one-time queries over static graphs, these requirements are poorly addressed by existing graph DBMSs and their languages. Existing query languages (e.g., PGQL, SPARQL v1.1, Cypher) address the first two issues by replacing edge labels of a subgraph pattern with regular expressions – this is known as *unions of conjunctive RPQs* (UCRPQ) [14, 72]. However, UCRPQ lacks algebraic closure and composability, limiting optimization opportunities. Furthermore, the output of a path navigation query is typically a set of pairs of vertices that are connected by a path under the constraints of a given regular expression. Hence, UCRPQ-based query languages limit path navigation queries to boolean reachability, without the ability to return and manipulate paths. G-CORE [6] addresses these limitations at the language specification level and has influenced the standardization efforts for a query language for graph DBMSs (see <https://www.gqlstandards.org/>). To the best of our knowledge, there is no work that uniformly addresses all three requirements.

The issues become more complex in the context of persistent query processing over streaming graphs, which is the focus of this paper. A number of specialized algorithms focus on evaluating subgraph pattern queries on streaming graphs [5, 18, 35, 44, 60], and our previous work focuses on path navigation queries and introduces the first streaming algorithms for RPQ evaluation [57], a limited subset of UCRPQ. However, a general-purpose model and framework that address the above discussed requirements of real-world applications (which feature complex graph patterns) in a uniform and principled manner is currently missing.

Querying streaming data in real-time imposes additional and novel requirements:

- **(R4)** graph streams are unbounded, making it infeasible to employ batch algorithms on the entire stream; and
- **(R5)** graph edges arrive at a very high rate and real-time answers are required as the graph emerges.

Existing graph DBMSs based on the *snapshot* model are not able to keep up with the high arrival rates [58]. The unsuitability of existing graph DBMSs for querying streaming data has motivated the design of specialized systems addressing singular application needs (e.g., [18, 31, 35, 44, 60]). The lack of systematic support for query processing over streaming graphs hinders the development of a general-purpose query processor for streaming graphs.

In this paper, we study the design of a general-purpose query processor for streaming graphs that addresses all of the above-discussed requirements in a uniform and principled manner. In analogy to traditional DBMSs, our framework provides the foundational tools to realize the well-known steps of query processing for streaming graph queries as follows:

- (1) a streaming graph query expressed in a declarative, high-level user language is translated into a query plan that consists of logical operators with precise semantics;
- (2) algebraic transformation rules are used to explore the plan space through query rewrites to find a “good” one, paving the way for query optimization;
- (3) the execution plan is built by selecting appropriate physical implementations of logical operators that are incremental and non-blocking;
- (4) the execution engine continuously executes the persistent query upon arrival of new edges to obtain new results.

Our focus in this paper is on the following key aspects of this framework:

- (1) We introduce the formal *Streaming Graph Query* model (SGQ – Section 4) based on a streaming graph model (Section 3). SGQ is based on the Regular Query (RQ) model¹, and provide precise semantics for persistent graph queries with subgraph patterns, path navigations and windowing constructs. Most of all, SGQ treats paths as first-class citizens, enabling queries to return and manipulate paths. Moreover, SGQ is closed under transitive closure and composable.
- (2) We present the *Streaming Graph Algebra* (SGA – Section 5.1) as a foundational basis for evaluating SGQ and provide an algorithm for translating SGQ into SGA expressions (Section 5.2). Utilizing SGA’s closedness (Section 5.3), we introduce transformation rules for novel SGA operators for systematic exploration of the plan space (Section 5.4).
- (3) We describe query execution plans (Section 6.1) and non-blocking physical operator implementations (Section 6.2).
- (4) We describe a prototype implementation of a streaming graph query processor based on *Timely Dataflow* [54], i.e., the dataflow computational model (Section 6).

Unlike existing work on streaming graphs that relies on ad hoc algorithms, our algebraic approach provides the foundational framework to precisely describe the semantics of complex persistent graph queries over streaming graphs and to optimize such queries

by query optimizers. SGA unifies path navigation and subgraph pattern queries in a structured manner (**R1** & **R2**), i.e., it attains composability by properly closing UCRPQ under recursion. In addition, our framework treats paths as first-class citizens of its data model, enabling the proposed algebra to express queries that return and manipulate paths (**R3**). Extensive experimental analysis over real and synthetic datasets of our prototype implementation that incorporates incremental, non-blocking algorithms as physical operators (**R4** & **R5**) demonstrates the feasibility and the performance gains of our approach. To the best of our knowledge, this is the first work to study the design of a streaming graph query processor, and to introduce an end-to-end solution to evaluating streaming graph queries with complex constructs.

2 RELATED WORK

2.1 Stream Processing Systems

Early research on stream processing primarily adapt the relational model and its query operators to the streaming setting (e.g., STREAM [8], Aurora [2], Borealis [1]). In contrast, modern Data Stream Processing Systems (DSPS) such as Storm [68], Heron [40], Flink [17] do not necessarily offer a full set of DBMS functionality. Existing literature (as surveyed by Hirzel et al. [33]) heavily focus on general-purpose systems and do not consider core graph querying functionality such as *subgraph pattern matching* and *path navigation*.

Existing work on streaming graph systems, by and large, target graph analytics workloads with little or no focus on graph querying functionalities on which we focus in this paper (e.g., STINGER [21], GraphOne [41, 42], GraphIn [66], GraphTau [34], GraPu [67], GraphBolt [49]). The primary focus of these systems is to build and maintain a sequence of snapshots for iterative graph analytics workloads under a stream of updates. SPARQL extensions have also been proposed for persistent query processing over RDF streams (e.g., C-SPARQL [12], CQELS [43], SPARQL_{stream} [16] and RSP-QL [20]). These are the most similar to our setting, but they are designed for SPARQLv1.0 and cannot formulate path expressions such as RPQs that cover more than 99% of all recursive queries found in massive Wikidata query logs [15]. Also, query processors of these systems do not employ incremental operators, except Sparkwave [37] that focuses on stream reasoning. Our proposed framework supports complex graph patterns arising in existing graph query languages, including SPARQL v1.1 property paths, and introduces non-blocking operators optimized for streaming workloads.

2.2 Incremental View Maintenance

A persistent query over sliding windows can be formulated as an IVM problem, where the view definition is the query itself and window movements are updates to the underlying database. Although conceptually similar, the algorithms (e.g., the classical *Counting* [32] algorithm) and systems (e.g., DBToaster [36], F-IVM [56], ViewDF [74]) in this category primarily target non-recursive queries. General-purpose IVM techniques for recursive queries (such as *DRed* [32] and *Absorption Provenance* [48]) ignore the structure of graph queries and the inherent temporal patterns of streaming graphs, resulting in significant computational overhead to compute and store a large number of derivations. Differential

¹A computationally well-behaved subset of Datalog with tractable query evaluation and decidable query containment, similar to relational algebra [61].

Dataflow (DD) [51] is a state-of-the-art distributed system for incremental maintenance of possibly cyclic dataflows. Graphsurge [63] uses DD as the underlying execution layer to share computation across multiple, possibly recursive views over static graphs. As we show in this paper (Section 6.2), DD can also be used for evaluating SGQ. Our framework, in contrast to these, exploits the structure of graph queries and the temporal patterns of sliding windows to minimize the cost of evaluating complex queries on streaming graphs (Section 7.2.2).

2.3 Graph Algorithms

TriAL [45] and Temporal Graph Algebra (TGA) [53] adopt an algebraic approach to graph query processing, similar to ours. TriAL is designed to be used for one-time navigational queries over static triplestores, and it cannot be used as a standalone graph query language. In contrast, our proposed SGA is designed as a standalone language for persistent graph queries over streaming graphs, and it can express complex graph patterns expressed in high-level user languages (Section 4.2). TGA extends temporal relational operators to PGM, and its Spark implementation introduces physical operator implementations for its algebraic primitives [4]. However, it is designed for exploratory graph analytics over the entire history of changes. In contrast, our framework is designed to continuously evaluate graph queries as the underlying (potentially unbounded) streaming graph changes.

The closest to ours are *specialized* algorithms on dynamic and streaming graphs [5, 18, 35, 44, 57, 60]. Some of these [18, 35] study the incremental evaluation of subgraph pattern queries. Their focus is developing efficient incremental algorithms to maintain matches of a given subgraph pattern as the underlying graph changes. Amar et al. [5] present distributed worst-case-optimal join algorithms for subgraph pattern queries. Li et al. [44] study subgraph isomorphism search over streaming graphs with timing-order constraints. GraphS [60] introduces efficient index structures that are optimized for cycle detection queries. In previous work [57], we study the design space of algorithms for path navigation over streaming graphs and provide algorithms for persistent evaluation of RPQ, a subset of the class of queries that we address in this paper. In contrast we target a *general* query model that captures the precise semantics of streaming graph queries.

3 DATA MODEL: STREAMING GRAPHS

3.1 Preliminaries

DEFINITION 1 (GRAPH). A *directed labeled graph* is a quintuple $G = (V, E, \Sigma, \psi, \phi)$ where V is a set of vertices, E is a set of edges, Σ is a set of labels, $\psi : E \rightarrow V \times V$ is an incidence function and $\phi : E \rightarrow \Sigma$ is an edge labelling function.

DEFINITION 2 (PATH AND PATH LABEL). Given $u, v \in V$, a *path* p from u to v in graph G is a sequence of edges $u \xrightarrow{p} v : \langle e_1, \dots, e_n \rangle$. The *label sequence* of a path p is defined as the concatenation of edge labels, i.e., $\phi^p(p) = \phi(e_1) \dots \phi(e_n) \in \Sigma^*$.

We use $\mathbb{T} = (\mathcal{T}, \leq)$ to define a discrete, total ordered time domain and use timestamps $t \in \mathcal{T}$ to denote time instants. Without loss of generality, the rest of the paper uses non-negative integers to represent timestamps.

DEFINITION 3 (STREAMING GRAPH EDGE). A *streaming graph edge* (sge) is a quadruple (src, trg, l, t) where $src, trg \in V$ are endpoints of an edge $e \in E$, $l \in \Sigma$ is the label of the edge e , and $t \in \mathcal{T}$ is the event (application) timestamp assigned by the source, i.e., $\psi(e) = (src, trg)$ and $\phi(e) = l$.²

DEFINITION 4 (INPUT GRAPH STREAM). An *input graph stream* is a continuously growing sequence of streaming graph edges $S^I = [sge_1, sge_2, \dots]$ where each sge represents an edge in graph G and sges are non-decreasingly ordered by their timestamps.³

Input graph streams represent external sources that provide the system with the graph-structured data. Our proposed framework uses a different format that generalizes Definition 4 to also represent intermediate results and outputs of persistent queries (Definition 7).

DEFINITION 5 (VALIDITY INTERVAL). A *validity interval* is a half-open time interval $[ts, exp)$ consisting of all distinct time instants $t \in \mathcal{T}$ for which $ts \leq t < exp$.

Timestamps are commonly used to represent the time instant in which the interaction represented by the sge occurred [44, 57, 60]. Alternatively, we use intervals to represent the period of validity of sges. In this paper, we argue that using *validity intervals* leads to a succinct representation and simplifies operator semantics by separating the specification of window constructs from operator implementation. As an example, each sge with timestamp t can be assigned a validity interval $[t, t + 1)$ that corresponds to a single time unit with smallest granularity that cannot be decomposed into smaller time units.⁴ Similarly, an sge $e = (u, v, l, [ts, exp))$ with a validity interval is equivalent to a set of sges $\{(u, v, l, t_1), \dots, (u, v, l, t_n)\}$ where $t_1 = ts$ and $t_n = exp - 1$. Windowing operator (to be precisely defined momentarily in Section 5.1) are used to assign validity intervals based on the windowing specifications of a given query.

3.2 Streaming Graphs

We now describe the logical representation of streaming graphs that is used throughout the paper. First, we extend the directed labeled graph model with materialized paths to represent paths as first-class citizens of the data model. As per Definition 2, a path between vertices u and v is a sequence of edges $u \xrightarrow{p} v : \langle e_1, \dots, e_n \rangle$ that connects vertices u and v , i.e., the path p defines a higher-order relationship between vertices u and v through a sequence of edges. By treating paths as first-class citizens like vertices and edges, the materialized path graph model enables queries to have paths as inputs and outputs. In addition, it enables edges and paths to be stitched together to form complex graph patterns as will be shown in Section 5.1.

DEFINITION 6 (MATERIALIZED PATH GRAPH). A *materialized path graph* is a 7-tuple $G = (V, E, P, \Sigma, \psi, \rho, \phi)$ where V is a set of vertices, E is a set of edges, P is a set of paths, Σ is a set of labels, $\psi : E \rightarrow V \times V$ is an incidence function, $\rho : P \rightarrow E \times \dots \times E$ is a total function that

²We assume that sges are generated by a single source and arrive in order, and leave out-of-order arrival as future work.

³We use $[]$ to denote ordered streams throughout the paper

⁴Commonly referred as NOW windows as described in Section 5.1.

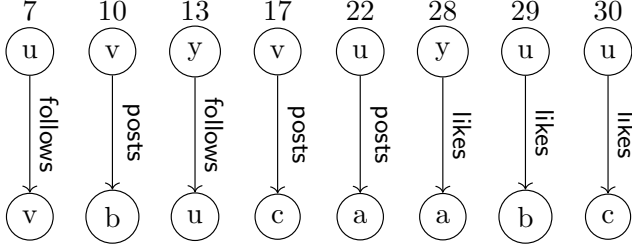


Figure 2: The input graph stream from an online social network of Example 1

assigns each path to a finite, ordered sequence of edges in E , and $\phi : (E \cup P) \rightarrow \Sigma$ is a labeling function, where images of E and P under ϕ are disjoint, i.e., $\phi(E) \cap \phi(P) = \emptyset$.

The function ρ assigns to each $p : u \xrightarrow{p} v \in P$ an actual path $\langle e_1, \dots, e_n \rangle$ in graph G satisfying: for every $i \in [1, \dots, n]$, $\psi(e_i) = (src_i, trg_i)$, $trg_i = src_{i+1}$, and $src_1 = u, trg_n = v$. Materialized path graph is a strict generalization of the directed labeled graph model (Definition 1), i.e., each directed labeled graph G is also a materialized path graph where $P = \emptyset$. We now generalize the notion of streaming graph edges (Definition 3) as follows:

DEFINITION 7 (STREAMING GRAPH TUPLE). A streaming graph tuple (sgt) is a quintuple $sgt = (src, trg, l, [ts, exp], \mathcal{D})$ where distinguished (explicit) attributes $src, trg \in V$ are endpoints of an edge $e \in E$ or a path $p \in P$ in graph G and $l \in \Sigma$ is its label, and non-distinguished (implicit) attributes $[ts, exp] \in \mathcal{T} \times \mathcal{T}$ is a half-open time-interval representing t 's validity and \mathcal{D} is a payload consists of edges in E that participated in the generation of the tuple t .

Streaming graph tuples generalize sges (Definition 3) to represent, in addition to input graph edges, derived edges (new edges as operator and query results that are not necessarily part of the input graph) and paths (sequence of edges as operator and query results). We use the notation $E^I \subset E$ to denote the set of input graph edges, and $\phi(E^I)$ to denote the fixed set of labels that are reserved for input graph edges. Additionally, non-distinguished (implicit) attribute \mathcal{D} of an sgt t captures the path p , i.e., sequence of edges, in case the sgt t represents a path. Otherwise, \mathcal{D} is the edge e that the sgt t represents.

DEFINITION 8 (STREAMING GRAPH). A streaming graph S is a continuously growing sequence of streaming graph tuples $S = [t_1, t_2, \dots]$ in which each tuple t_i arrives at a particular time t_i ($t_i < t_j$ for $i < j$).

Figure 2 depicts an excerpt of the input graph stream of the application in Example 1, where each tuple represent an interaction and each timestamp represent the time instant that the interaction occurs. Its corresponding streaming graph where each tuple is assigned a time interval is shown in Figure 3 (validity intervals are assigned by WSCAN operator – see Section 5.1).

Unless otherwise specified, we consider streaming graphs to be *append-only*, i.e., each sgt represents an insertion, and use the *direct approach* to process expirations due to window movements. Explicit deletions of previously arrived sgts can be supported by explicitly manipulating the validity interval of a previously arrived

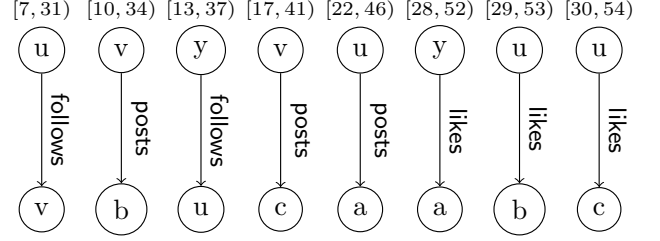


Figure 3: Streaming graph obtained from the input graph stream in Figure 2 where the validity interval of each element is set based on a 24h window

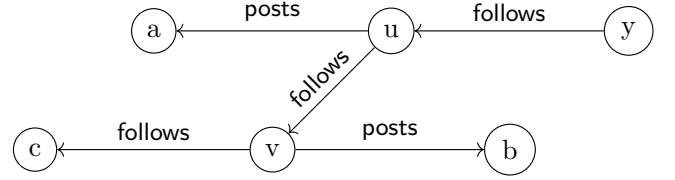


Figure 4: Snapshot graph of the streaming graph in Figure 3 at $t = 25$.

sgt [39]. This corresponds to the *negative tuple* approach [25, 29]. Section 6.2 describes the processing of insertions, deletions and expirations under alternative window semantics for physical operator implementations.

DEFINITION 9 (LOGICAL PARTITIONING). A logical partitioning of a streaming graph S is a label-based partitioning of its tuples and it produces a set of disjoint streaming graphs $\{S_{l_1}, \dots, S_{l_n}\}$ where each S_{l_i} consists of sgts of S with the label l_i , i.e., $S = \bigcup_{l \in \Sigma} (S_l)$.

This label-based partitioning of streaming graphs provides a coherent representation for inputs and outputs of operators in logical algebra (Section 5.1). At the logical level, it can be performed by the *filter* operator of the logical algebra (precisely defined in Definition 17), and logical operators of our algebra process logically partitioned streaming graphs as their inputs and outputs unless otherwise specified.

DEFINITION 10 (VALUE-EQUIVALENCE). Sgts $t_1 = (u_1, v_1, l_1, [ts_1, exp_1], \mathcal{D}_1)$ and $t_2 = (u_2, v_2, l_2, [ts_2, exp_2], \mathcal{D}_2)$ are value-equivalent iff their distinguished attributes are equal, i.e., they both represent the same edge or the same path possibly with different validity intervals. Formally, $t_1 = t_2 \Leftrightarrow u_1 = u_2, v_1 = v_2, l_1 = l_2$.

Value-equivalence is used for temporal coalescing of tuples with adjacent or overlapping validity intervals [47]. We extend the *coalesce* primitive from temporal database literature [19] to sgts with an aggregation function over the non-distinguished payload attribute, \mathcal{D} , as shown below:

DEFINITION 11 (COALESCE PRIMITIVE). The *coalesce primitive* merges a set of value-equivalent sgts $\{t_1, \dots, t_n\}$, $t_i = (src, trg, l, [ts_i, exp_i], \mathcal{D}_i)$ for $1 \leq i \leq n$ with overlapping or adjacent validity intervals using an operator-specific aggregation function f_{agg}

over the payload attribute \mathcal{D} :

$$\text{coalesce}_{f_{agg}}(\{t_1 \dots, t_n\}) = (\text{src}, \text{trg}, l, [\min_{1 \leq i \leq n} (ts_i), \max_{1 \leq i \leq n} (exp_i)], f_{agg}(\mathcal{D}_1, \dots, \mathcal{D}_n))$$

Distinguished attributes src , trg and the label l of sgts in a streaming graph S represent the topology of a materialized path graph. Hence, a finite subset of a streaming graph S corresponds to a materialized path graph over the set of edges and paths that are in the streaming graph and the set of vertices that are adjacent to these. We now use this to define snapshot graphs and the property of *snapshot reducibility*.

DEFINITION 12 (SNAPSHOT GRAPH). A snapshot graph G_t of a streaming graph S is defined by a mapping from each time instant in \mathcal{T} to a finite set of sgts in S . At any given time t , the content of a mapping $\tau_t(S)$ defines a snapshot graph $G_t = (V_t, E_t, P_t, \Sigma_t, \psi, \rho, \phi)$ where $E_t = \{e_i \mid e_i.ts \leq t < e_i.exp\}$ is the set of all edges that are valid at time t , $P_t = \{p_i \mid p_i.ts \leq t < p_i.exp\}$, and V_t is the set of all vertices that are endpoints of edges and paths in E_t and P_t , respectively.

Value-equivalence (Definition 10) and the coalesce primitive (Definition 11) ensure that snapshot graphs have the *set semantics*, i.e., at any point in time t , the snapshot graph G_t of a streaming graph S , a vertex, edge and path exists at most once.

4 STREAMING GRAPH QUERIES

This section presents our *streaming graph query* (SGQ) model. We first provide a formal definition of SGQ using Datalog, enabling the specification of precise SGQ semantics and to reason about its expressiveness. We then describe how SGQ captures a significant subset of existing graph query languages and provide concrete examples on how to formulate SGQ using a slight extension of G-CORE.

4.1 Query Model

We formally describe SGQ based on a streaming generalization of the Regular Query (RQ) model [61]. Informally, RQ corresponds to *binary, non-recursive* subset of Datalog with transitive closure and provides a principled way to combine subgraph patterns and path navigations. RQ provides a good basis for building a general-purpose framework for persistent query evaluation over streaming graphs, because (i) unlike UCRPQ, it is closed under transitive closure and therefore composable, (ii) it has more expressive power than the existing graph query languages such as SPARQL v1.1, Cypher, PGQL – RQ strictly subsumes UCRPQ on which these are based, and (iii) its query evaluation and containment complexity is reasonable [61]. Due to its well-defined semantics and computational behaviour, RQ has been gaining popularity as a logical foundation for graph queries, both in theory [13, 14] and in practice [6]. Indeed, RQ captures the core of the contemporary graph query language G-CORE that we use throughout this paper.

DEFINITION 13 (REGULAR QUERIES (RQ) – FOLLOWING [61]). The class of Regular Queries is the subset of non-recursive Datalog with a

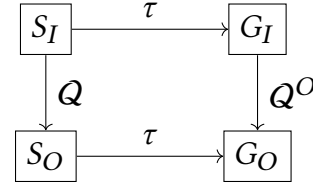


Figure 5: Snapshot reducibility (adapted from [39]).

finite set of rules where each rule is of the form:⁵

$$\text{head} \leftarrow \text{body}_1, \dots, \text{body}_n$$

Each body_i is either (i) a binary predicate $l(\text{src}, \text{trg})$ where $l \in \Sigma$ is a label, or (ii) $l^*(\text{src}, \text{trg})$ as d , which is a transitive closure over $l(\text{src}, \text{trg})$ for a label $l \in \Sigma$, $d \in \Sigma \setminus \phi(E)$, and each head predicate (head) is a binary predicate with $d(\text{src}, \text{trg})$ for a label $d \in \Sigma \setminus \phi(E)$ except the reserved predicate Answer $\notin \Sigma$ of an arbitrary arity.

In other words, an RQ is a binary, non-recursive Datalog program extended with the transitive closure of binary predicates where input graph edges with a label $l \in \phi(E^I)$ correspond to instances of the extensional schema (EDB) and derived edges and paths with a label $l \in \Sigma \setminus \phi(E^I)$ correspond to instances of the intensional schema (IDB). EDBs are predicates that appear only on the right-hand-side of the rules, which correspond to stored relations in Datalog [3]. Similarly, we define IDBs as predicates that appear in the rule heads, which correspond to output relations in Datalog.

EXAMPLE 2 (REGULAR QUERY). Consider the real-time notification query in Example 1 and its graph pattern in Figure 1. The one-time query for the same graph pattern corresponds to the following RQ:

$$RL(u_1, u_2) \leftarrow l(u_1, m_1), f^+(u_1, u_2) \text{ as } FP, p(u_2, m_1)$$

$$\text{Notify}(u, m) \leftarrow RL^+(u, v) \text{ as } RLP, p(v, m)$$

$$\text{Answer}(u, m) \leftarrow \text{Notify}(u, m)$$

where predicates l, f, FP, p, RL, RLP represent labels likes, follows, followsPath, post, recentLiker and recentLikerPath, respectively.

Next, we define the notion of snapshot reducibility that enables us to precisely define the semantics of streaming queries and operators using their non-streaming counterparts. Snapshot reducibility is used in temporal databases to generalize non-temporal queries and operators to temporal ones [19].

DEFINITION 14 (SNAPSHOT-REDUCIBILITY). Let S be a streaming graph, Q a streaming graph query and Q^O its non-streaming, one-time counterpart. Snapshot reducibility states that each snapshot of the result of applying Q to S is equal to the result of applying its non-streaming version Q^O over the corresponding snapshots of S , i.e., $\forall t \in \mathcal{T}, \tau_t(Q(S)) = Q^O(\tau_t(S))$.

Following existing research [27], we define the semantics of persistent evaluation of SGQ using the notion of snapshot reducibility (Definition 14). It is known that for many operations such as joins

⁵The *dependency graph* of a Datalog program is a directed graph whose vertices are its predicates and edges represent dependencies between predicates, i.e., there is an edge from p to q if q appears in the body of rule with head predicate p . A Datalog program is *non-recursive* iff its dependency graph is acyclic, i.e., no predicate depends recursively on itself.

and aggregation, exact results cannot be computed with a finite memory over unbounded streams [9]. In streaming systems, a common solution for bounding the space requirement is to evaluate queries on a window of data from the stream. The windowed evaluation model provides a tool to process unbounded streams with bounded memory, and restricts the scope of queries to recent data, a desired feature in many applications[9, 27]. Additionally, as opposed to streaming approximation techniques that trade off exact answers in favour of bounding the space requirements, window-based query evaluation enables exact query answers w.r.t. window specifications. Hence, we adopt the *time-based sliding window* for SGQ in the remainder.

DEFINITION 15 (STREAMING GRAPH QUERY – SGQ). An SGQ query Q is an RQ defined over a streaming graph S and a time-based sliding window $\mathcal{W}_{\mathcal{T}}$ whose semantics is defined using the corresponding one-time RQ Q^O and the notion of snapshot reducibility (Definition 14):

$$\forall t \in \mathcal{T}, \quad \tau_t(Q(S, \mathcal{W}_{\mathcal{T}})) = Q^O(\tau_t(\mathcal{W}_{\mathcal{T}}(S)))$$

Figure 5 illustrates the correspondence between streaming and one-time graph queries. A direct consequence of such a relationship is that SGQ can be evaluated by repeatedly executing the corresponding one-time query, known as *query re-evaluation* [25]. Specifically, the resulting streaming graph of an SGQ can be obtained from the sequence of snapshots that is the result of repeated evaluation of the corresponding one-time query at every time instant: an sgt $(u, v, l, [ts, expiry], D)$ is in the resulting streaming graph for an SGQ $Q \tau_t(Q(S, \mathcal{W}_{\mathcal{T}}))$ if there is an edge $e = (u, v, l)$ or a path $p : u \xrightarrow{p} v$ with $l = \phi^p(p)$ in the resulting snapshot graph of the corresponding one-time query $G_{t_i} = Q^O(\tau_t(\mathcal{W}_{\mathcal{T}}(S)))$ for $ts \leq t < exp$. However, such a strategy is wasteful as the input differences between two consecutive instants are likely to be small. Alternatively, *incremental evaluation* computes the changes in the output as new sgts arrive and old sgts expire due to window movements. In this paper, we focus on the *incremental evaluation* method and use the concept of *snapshot reducibility* to ensure correct evaluation semantics.

4.2 SGQ in Practice

The SGQ model formalizes an important class of graph queries in the streaming model. It captures the core features of current graph query languages such as subgraph pattern and reachability-based path queries. In this section, we demonstrate the SGQ’s expressive power by mapping core G-CORE constructs to SGQ. We choose G-CORE for demonstration due to the following reasons. G-CORE fulfills all three requirements of graph querying (**R1**, **R2** & **R3** in Section 1). Other existing languages (e.g., SPARQL v1.1, Cypher, PGQL) can only partially satisfy these requirements due to (i) the lack of algebraic closure and composability, and (ii) limited path navigation capability [14]. Moreover, G-CORE supports SGQ capabilities such as the treatment of paths as first-class citizens and returning graphs. Finally, G-CORE is one of the more prominent language specifications influencing the ongoing standardization process of a graph query language GQL (<https://www.gqlstandards.org>).

```

PATH RL = (u1)-/ <: follows^*> /->(u2),
           (u1)-[: likes]->(m1)<-[: posts]-(u2)
CONSTRUCT (u) -[: notify]-> (m)
MATCH (u)-/ p<~RL*> /->(v),
        (v)-[: posts]->(m),
ON social_stream WINDOW(24h) SLIDE(1h)

```

Figure 6: G-CORE representation of the SGQ in Example 1.

```

GRAPH VIEW rec_stream AS (
  CONSTRUCT (u1) -[: recommendation]-> (p)
  MATCH (u1)
    OPTIONAL (u1)-[: follows]->(u2)
    OPTIONAL (u1)-[: likes]->(m)<-[: posts]-(u2)
  ON social_stream WINDOW(24 hours)
  MATCH (c) -[: purchase]->(p)
  ON tx_stream WINDOW(30d) SLIDE(1d)
  WHERE (u2) = (c) )

```

Figure 7: G-CORE representation of the query in Example 4

Since we have not yet included properties in the data model, we focus on a subset of G-CORE where queries do not contain predicates or aggregation over property values. This particular subset already covers many important key features: **(a)** returning graphs, **(b)** ASCII-art syntax for pattern matching, **(c)** joins over multiple graphs, **(d)** views and optionals, **(e)** RPQ-based reachability queries, **(f)** and powerful path patterns as demonstrated in the two examples we discuss below.

G-CORE is originally targeted for one-time queries over static property graphs and it does not provide native windowing constructs. We slightly extend the **ON** clause with a **WINDOW** clause to incorporate window specifications. In particular, a time-based sliding window is defined by the newly introduced **WINDOW** clause that specifies the window length, and an optional **SLIDE** clause that specifies the slide interval, following a streaming graph reference in the **ON** clause.

EXAMPLE 3. The G-CORE query in Figure 6 represents the real-time notification example in Example 1 (its corresponding RQ is already given in Example 2). Its **PATH** and **MATCH** clauses use ASCII-art syntax **(b)** to define complex graph patterns **(f)** with RPQ-based reachability **(e)**, and its **CONSTRUCT** clause returns a streaming graph of notify edges **(a)**.

EXAMPLE 4. Consider the G-CORE query in Figure 7 that combines streaming information from a social network of user interactions and a transaction network of customer purchases to drive product recommendations. It defines a view over the resulting streaming graph of recommendation edges **(d)** by joining patterns from two streaming graphs **(c)**, and its **MATCH** clause features optional predicates to incorporate two alternative social interactions **(d)**. Its graph pattern

corresponds to the following RQ:

$$\begin{aligned} ACQ(u_1, u_2) &\leftarrow l(u_1, m_1), p(u_2, m_1) \\ ACQ(u_1, u_2) &\leftarrow f(u_1, u_2) \\ REC(u, p) &\leftarrow ACQ(u_1, u_2), pur(u_2, p) \\ Answer(u, p) &\leftarrow REC(u, p) \end{aligned}$$

where predicates l, f, p, pur, ACQ, REC represent labels likes, follows, post, purchase, acquaintance, and recommendation, respectively.

5 STREAMING GRAPH ALGEBRA

This section presents the logical foundation of our streaming graph query processing framework. We first introduce streaming graph algebra (SGA) and the semantics of its operators (Section 5.1). The main motivation behind the proposed SGA is similar to that of relational DBMSs: it enables us to formulate and represent query plans independent of specific physical implementations. We subsequently describe how to transform a given SGQ (Definition 15) into its canonical SGA expression and illustrate logical query plans (Section 5.2) and prove its closedness and composability (Section 5.3). Finally, we present some transformation rules for SGA to demonstrate alternative logical plan generation opportunities (Section 5.4).

5.1 SGA Operators

For ease of exposition, the rest of the paper assumes that inputs to each SGA operator are partitioned into one more streaming graphs S_a based on tuple labels where each S_a contains sgts with the same label $a \in \Sigma$. The output of each operator is also a streaming graph S_o where each sgt has the label $o \in \Sigma \setminus \phi(E^I)$.⁶ SGA contains the following operators: windowing (Definition 16), filter (Definition 17), union (Definition 18), subgraph pattern (Definition 19), and path navigation (Definition 20).

DEFINITION 16 (WSCAN). The **windowing operator** \mathcal{W} transforms a given input graph stream S^I to a streaming graph S by adjusting the validity interval of each sgt based on the window size \mathcal{T} and the optional slide interval β , i.e., $\mathcal{W}_{\mathcal{T}, \beta}(S^I) := S : [(u, v, l, [t, exp], \mathcal{D}) : e(u, v, l)) \mid (u, v, l, t) \in S^I \wedge exp = \lfloor t/\beta \rfloor \cdot \beta + \mathcal{T}]$.

The window size \mathcal{T} determines the length of the validity interval of sgts and the slide interval β controls the granularity at which the time-based sliding window progresses [8, 57]. If β is not provided, default is $\beta = 1$, i.e., single time instant with the smallest granularity, and it defines a sliding window that progresses at every time instant.

The WSCAN operator defines the semantics of time-based sliding windows. It acts as an interface between the external streaming graph sources and the query plans and it is responsible for providing data from input graph streams to a query plan, similar to the *scan* operator in relational systems. WSCAN manipulates the implicit temporal attribute and associates a time interval to each sgt representing its validity. Our model of representing streaming graphs (Definition 8) provides a concise representation of validity intervals and enables operators to treat time differently than the data stored in the graph. This enables us to distinguish operator

⁶ $\phi(E^I) \subset \Sigma$ is reserved for input graph edges and cannot be used by operators as labels for resulting sgts. In other words, $\phi(E^I) \subset \Sigma$ corresponds to EDBs in Datalog as described in Section 4.

semantics from window semantics and eliminates the redundancy caused by integrating sliding window constructs into each operator of the algebra. SGA operators access and manipulate validity intervals implicitly, generalizing their non-streaming counterparts with implicit handling of time.

EXAMPLE 5. Consider the real-time notification task of Example 1 with a 24-hour window of interest. WSCAN \mathcal{W}_{24} sets validity intervals of sgts of the input graph stream and produces a streaming graph where each sgt is valid for 24 hours, as shown in Figure 3.

DEFINITION 17 (FILTER). **Filter operator** $\sigma_\Phi(S)$ is defined over a streaming graph S and a boolean predicate Φ involving the distinguished attributes of sgts, and its output stream consists of sgts of S on which Φ evaluates to true. Formally:

$$\sigma_\Phi(S) = [(u, v, l, [ts, exp], \mathcal{D}) \mid (src, trg, l, [ts, exp], \mathcal{D}) \in S \wedge \Phi((src, trg, l, \mathcal{D}))].$$

DEFINITION 18 (UNION). Union $\cup^{[d]}$ with an optional output label $d \in \Sigma \setminus \phi(E^I)$ merges sgts of two streaming graphs S_1 and S_2 , and assigns the new label d if provided. Formally:

$$S_1 \cup^{[d]} S_2 = [t \mid t \in S_1 \vee t \in S_2]$$

DEFINITION 19 (PATTERN). The streaming **subgraph pattern operator** is defined as $\bowtie_\Phi^{src, trg, d}(S_{l_1}, \dots, S_{l_n})$ where each S_{l_i} is a streaming graph, Φ is a conjunction of a finite number of terms in the form $pos_i = pos_j$ for $pos_i, pos_j \in \{src_1, trg_1, \dots, src_n, trg_n\}$ where src_i, trg_i are endpoints of sgts in S_{l_i} , and $src, trg \in \{src_1, trg_1, \dots, src_n, trg_n\}$ are the endpoints of resulting sgts, and $d \in \Sigma \setminus \phi(E^I)$ represent the label of the resulting sgts. Formally:

$$\begin{aligned} \bowtie_\Phi^{src, trg, d}(S_{l_1}, \dots, S_{l_n}) &= [(u, v, d, [ts, exp], \mathcal{D}) : e(u, v, l)) \mid \\ &\exists t_i = (src_i, trg_i, l_i, [ts_i, exp_i], \mathcal{D}_i) \in S_{l_i}, 1 \leq i \leq n \\ &\wedge \Phi((src_1, trg_1, \dots, src_n, trg_n)) \wedge \\ &u = src \wedge v = trg \wedge \bigcap_{1 \leq i \leq n} [ts_i, exp_i] \neq \emptyset \wedge \\ &ts = \max_{1 \leq i \leq n} (ts_i) \wedge exp = \min_{1 \leq i \leq n} (exp_i)]. \end{aligned}$$

Given a subgraph pattern expressed as a conjunctive query, PATTERN finds a mapping from vertices in the stream to free variables where (i) all query predicates hold over the mapping, and (ii) there exists a time instant at which each edge in the mapping is valid.

EXAMPLE 6. Consider the real-time notification query given in Example 1; the recentLiker relationship defined in the form of a triangle pattern can be represented with PATTERN $\bowtie_\phi^{src1, src4, RL}$ where $\phi = (trg_1 = trg_2 \wedge src_1 = src_3 \wedge src_2 = trg_3)$. Its output over the streaming graph, given in Figure 3, consists of sgts $(y, RL, u, [28, 37])$, (y, RL, u) and $(u, RL, v, [29, 31])$, (u, RL, v) that correspond to derived edges with label recentLiker.

SGA operators may produce multiple value-equivalent sgts with adjacent or overlapping validity intervals. Unless otherwise specified, such sgts in resulting streaming graphs of SGA operators are coalesced to maintain the set semantics of streaming graphs and their snapshots (Definition 10). To illustrate, consider PATTERN in the above example: over the streaming graph given in Figure 3, the

PATTERN operator finds two distinct subgraphs with vertices (u, b, v) and (u, c, v) . Consequently, it produces two value-equivalent tuples $(u, RL, v, [29, 31])$, (u, RL, v) and $(u, RL, v, [30, 31])$, (u, RL, v) , which are coalesced into a single sgt by merging their validity intervals.

DEFINITION 20 (PATH). *The streaming path navigation operator is defined as $\mathcal{P}_R^d(S_{l_1}, \dots, S_{l_n})$ where R is a regular expression over the alphabet $\{l_1, \dots, l_n\} \subseteq \Sigma$, and $d \in \Sigma \setminus \phi(E^I)$ designates the label of the resulting sgts. The sgt $t = (u, v, l, [ts, exp], \mathcal{D} : p)$ is an answer for \mathcal{P}_R^d if there exists a path p between u and v in the snapshot of S at time t , i.e., $p : u \xrightarrow{p} v \in \tau_t(S) = G_t$, and the label sequence of the path p , $\phi^p(p)$ is a word in the regular language $L(R)$. Formally:*

$$\begin{aligned} \mathcal{P}_R^d(S_{l_1}, \dots, S_{l_n}) &= \{(u, v, d, [ts, exp], \mathcal{D}) \mid \exists p : u \xrightarrow{p} v \wedge \\ &\quad \forall e_i \in p, \exists t_i = (src_i, trg_i, l_i, [ts_i, exp_i], \mathcal{D}_i) \in S_{l_i} \wedge \\ &\quad \phi^p(p) \in L(R) \wedge \bigcap_{t \in p} [ts, t.exp] \neq \emptyset \wedge \\ &\quad ts = \max_{t \in p} (t.ts) \wedge exp = \min_{t \in p} (t.exp) \wedge \mathcal{D} = p\}. \end{aligned}$$

PATH finds pairs of vertices that are connected by a path where (i) each edge in the path is valid at the same time instant, and (ii) path label is a word in the regular language defined by the query. This closely follows the RPQ model where path constraints are expressed using a regular expression over the set of labels [72]. Path navigation queries in the RPQ model are evaluated under *arbitrary* and *simple* path semantics. The former allows a path to traverse the same vertex multiple times, whereas under the latter semantics a path cannot traverse the same vertex more than once [7, 10, 72]. In this paper, we adopt the arbitrary path semantics due to its widespread adoption in modern graph query languages [6, 7, 70], and the tractability of the corresponding evaluation problem [10].

EXAMPLE 7. *In the example of Figure 1 the path navigation over the derived recentLiker edges is represented by PATH \mathcal{P}_{RL}^{RLP} . Its output over the resulting streaming graph of PATTERN of Example 6 consists of sgts $(y, RLP, u, [28, 37])$, (y, RL, u) , $(u, RLP, v, [29, 31])$, (u, RL, v) , and $(y, RLP, v, [29, 31])$, $((y, RL, u), (u, RL, v))$ that correspond to materialized paths with label recentLikerPath of length one and two.*

Most existing work on the RPQ model focuses on the problem of determining reachability between pairs of vertices connected by a path conforming to given regular expression [38, 46, 52, 57]. By adapting the materialized path graph model (Definition 6), we pinpoint that PATH is equipped with the ability to return paths, i.e., each resulting sgt contains the actual sequence of edges that form the path with a label sequence conforming to given regular expression.

SGA builds on the Regular Property Graph Algebra (RPGA) [14], which is itself based on Regular Queries (RQ). Of course, both RPGA and RQ formulate graph queries over static graphs, while SGA operators are defined over streaming graphs (Definition 8), and they access and manipulate validity intervals implicitly. Thus they generalize their non-streaming counterparts with implicit handling of time. This follows from the fact that the semantics of SGA operators are defined through *snapshot reducibility* (Definition 14), that is, the snapshot of the result of a streaming operator on a streaming graph S at time t is equal to the result of the corresponding non-streaming operator on the snapshot of the streaming graph S at time t .

5.2 Formulating Query Plans in SGA

SGA can express all queries that can be specified by SGQ (Section 4). This section provides an algorithm for the conversion.

Given a SGQ $Q(S, \mathcal{W}_{\mathcal{T}})$ over a streaming graph and a time-based sliding window definition, Algorithm **SGQParser** produces the canonical SGA expression. The algorithm processes the predicates of a given SGQ and generates the corresponding SGA expression in a bottom-up manner: each EDB l corresponds to a WSCAN over an input streaming graph S_l^I , each application of transitive closure corresponds to a PATH, each IDB d corresponds to a UNION or PATTERN based on the body of the corresponding rule.

Algorithm SGQParser:

```

input : Streaming Graph Query  $Q(S, \mathcal{W}^T)$ 
output: SGA Expression  $e$ 
1  $G_Q \leftarrow \text{Graph}(Q)$  // dependency graph
2  $[r_1, \dots, r_n] \leftarrow \text{TopSort}(G_Q)$  // topological sort
3  $exp \leftarrow []$  // empty mapping
4 for  $1 \leq i \leq n$  do
5   switch  $r_i$  do
6     case  $l(src, trg), l \in \phi(E^I)$  do
7        $exp[l] \leftarrow \mathcal{W}^T(S_l)$ 
8     case  $l^*(x, y)$  do
9        $exp[d] \leftarrow \mathcal{P}_{l^*}^d(exp[l])$ 
10    otherwise do
11       $d(src, trg) \leftarrow r_i.head, [b_1, \dots, b_n] \leftarrow r_i.body$ 
12       $\Phi \leftarrow \text{GenPred}(r_i.body)$ 
13       $e \leftarrow \bigvee_{\Phi}^{src, trg, d} (exp[b_1], \dots, exp[b_n])$ 
14      if  $exp[d] \neq \emptyset$  then
15         $exp[d] \leftarrow exp[d] \cup e$ 
16      else
17         $exp[d] \leftarrow e$ 
18 return  $exp[\text{Answer}]$ 

```

THEOREM 1. *There exists a SGA expression $e \in \text{SGA}$ for any $Q \in \text{SGQ}$.*

PROOF. The dependency graph of an RQ is acyclic as RQ is non-recursive (Definition 13); hence, Line 2 is guaranteed to define a partial order over Q 's predicates. Algorithm **SGQParser** generates an SGA expression for each predicate in this order and caches it in exp array. In particular, Line 7 generates an SGA expression for each EDB predicate and Line 9 generates a PATH expression for each body predicate with a Kleene star. For each rule $d(src, trg) := l_1(src_1, trg_1), \dots, l_n(src_n, trg_n)$, Line 13 generates a PATTERN expression. Finally, Line 14 generates a UNION expression if there are multiple rules with the same head predicate d . As each predicate is processed based on the partial order defined by the dependency graph G_Q (Line 4), exp is guaranteed to have SGA expressions for each predicate r_j ($1 \leq j \leq i$) when processing predicate r_i . Once all predicates are processed, Line 18 returns the SGA expression of the *Answer* predicate. Hence, Algorithm **SGQParser** correctly constructs an SGA expression for a given SGQ. \square

The complexity of evaluating SGA expressions is the same as RQ given their relationship noted above: NP-complete in combined complexity and NLogspace-complete in data complexity [14, 61].

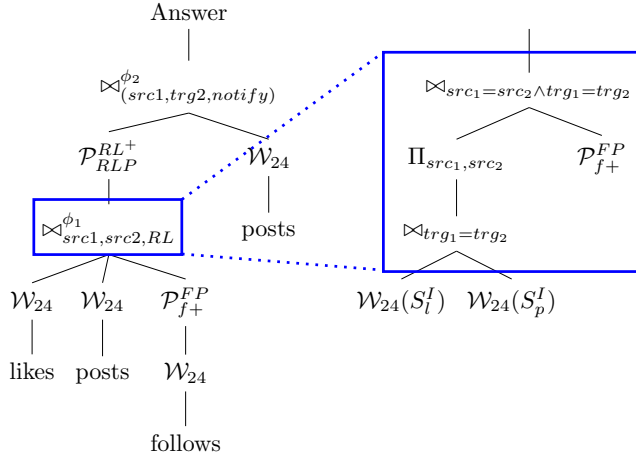


Figure 8: (left) Logical plan for the SGA expression in Example 8, and (right) binary join tree for its PATTERN.

EXAMPLE 8 (CANONICAL TRANSLATION). *For the real-time notification task in Example 1 and its corresponding RQ in Example 2, Algorithm SGQParser generates the following canonical SGA expression for its corresponding SGQ with a sliding window of 24 hours:*

$$\bowtie_{\phi_2}^{(src1, trg2, notify)} \left(\begin{aligned} &\mathcal{P}_{RLP}^{RL+} \left(\bowtie_{\phi_1}^{src1, src2, RL} (W_{24}(S_l), W_{24}(S_p), \mathcal{P}_{f+}^{FP}(W_{24}(S_f))) \right), \\ &W_{24}(S_p) \end{aligned} \right)$$

$$\phi_1 = (trg_1 = trg_2 \wedge src_1 = src_3 \wedge src_2 = trg_3),$$

$$\phi_2 = (trg_1 = src_2)$$

Figure 8 (left) illustrates the logical plan for the same SGQ that consists of SGA operators.

5.3 Closedness and Composability

Algebraic closure is a required property of any query algebra as it enables query rewriting (Section 5.4) and query optimization. Composability is a desired feature for a declarative query language as it facilitates query decomposition, view-based query evaluation, query rewriting etc. SGA operators are closed over streaming graphs as defined in Section 3; that is, the output of an SGA operator is a valid streaming graph if its inputs are valid streaming graphs. Thus SGA queries are composable, i.e., the output of one query can be used as input of another query.

SGQ language is also closed (Theorem 1) – each query takes one or more streaming graphs as input and produces a streaming graph as output. It is also composable as the output of a query can be the input of the subsequent query. As such, G-CORE variation that we use as our user-level query language example (Section 4.2) attains composability exactly as its original version is composable over property graphs [6]. This is in contrast to the other graph query languages that lack an algebraic basis like ours, e.g., SPARQL and Cypher are not composable and may not be closed. Cypher 9 requires graphs as input, but produces tables as output so the

language is neither closed nor composable – the results of a Cypher query cannot be used as input to a subsequent one without additional processing. SPARQL can produce graphs as output using the **CONSTRUCT** clause, and is therefore closed; however, it requires query results to be made persistent and therefore not easily composable [14].

5.4 Transformation Rules

As noted above, closedness of an algebra is important for query rewriting to explore the space of equivalent plans; this is a key component of query planning and optimization. Although optimization problems are beyond the scope of this paper, in this section we highlight possible transformation rules that enable the systematic exploration of the plan space in order to demonstrate the potential of our SGA.

Some of the traditional relational transformation strategies such as join ordering, predicate push down are applicable to UNION, FILTER and PATTERN due to snapshot-reducibility. UNION and FILTER operators are streaming generalizations of corresponding relational union and selection operators, and the PATTERN operator can be represented using a series of equijoins. Below, we describe transformation rules involving the other novel SGA operators:

Transformation Rules for WSCAN: WSCAN (\mathcal{W}^T) commutes with operators that do not alter the validity intervals of sgts, i.e., UNION and FILTER. Pushing FILTER down the WSCAN operator can potentially reduce the rate of sgts and consequently the amount of state the windowing operator needs to maintain. Formally:

- (1) $\mathcal{W}^T(\sigma^\phi(S)) = \sigma^\phi(\mathcal{W}^T(S))$
- (2) $\mathcal{W}^T(S_1 \cup S_2) = \mathcal{W}^T(S_1) \cup \mathcal{W}^T(S_2)$

Transformation Rules for PATH: We identify two transformation rules for PATH:

- (1) Alternation: $\mathcal{P}_{a|b}^d(S_a, S_b) = \bigcup^d(S_a, S_b)$
- (2) Concatenation: $\mathcal{P}_{a \cdot b}^d(S_a, S_b) = \bowtie_{trg_1=src_2}^{src1, trg2, d}(S_a, S_b)$

These transformation rules enable the exploration of a rich plan space for SGQ that are represented by SGA. In particular, PATH and its transformation rules enable the integration of existing approaches for RPQ evaluation with standard optimization techniques such as join ordering and pushing down selection in a principled manner. Traditionally, path query evaluation follows two approaches [10, 23, 38, 57, 65, 73]: graph traversals guided by finite automata, or relational algebra extended with transitive closure, i.e., *alpha*-RA. Yakovets et al. introduce a hybrid approach (Waveguide) and model the cost factors that impact the efficiency of RPQ evaluation on static graphs [73]. SGA enables the representation of these approaches in a uniform manner, and the above transformation rules enable the exploration of a rich plan space that subsumes these existing plans. An example application of these transformation rules and a micro-benchmark demonstrating the potential benefits of query optimizations through plan space exploration is given in Section 7.4.

6 IMPLEMENTATION

We implemented a prototype streaming graph query processor (<https://dsg-uwatool.github.io/s-graffito/>) based on the algebraic framework we propose in this paper. Our goal is to build a plausible

conceptual framework for expressing and evaluating SGQ – not to have a complete system. Hence, we focus on the construction of query execution plans and physical implementations of logical SGA operators.

6.1 Query Processor

Conceptually, SGQ can be evaluated by repeatedly evaluating from scratch the corresponding one-time query at each point in time (Section 4.1). Albeit semantically correct, such a re-execution strategy is, of course, infeasible. In streaming systems, the focus is on *incremental* evaluation of persistent queries where the goal is to avoid re-computing the entire result by only computing the changes to the output in real-time as new input arrives. Consequently, queries are executed in a data-driven (push-based) manner as opposed demand-driven (pull-based) query processing employed in traditional relational DBMSs [30]. As such, one can implement the framework proposed in this paper over existing streaming systems such as Apache Flink, Spark Streaming and Timely Dataflow. For our prototype, we use Timely Dataflow (TD) as the underlying execution engine.

Applications in TD are expressed as a directed graph of operations where vertices correspond to user-defined computations and edges correspond to the flow of data between them. In executing an SGQ, the query processor first creates a logical plan from the canonical SGA expression of the given query. The physical execution plan in the form of a dataflow graph is constructed by: (i) creating source vertices for leaves of the logical query plan that receives input graph streams, (ii) replacing logical SGA operators with physical operator implementations, and (iii) creating a sink vertex for the root of the logical query plan that pushes results back to the user. Section 6.2 describes physical operator implementations in detail.

TD associates each input data with a logical timestamp that enables fine-grained synchronization and progress tracking. We represent each input graph stream as an evolving collection where each item represents an sge (Definition 4) and event timestamps assigned by the source are used as logical timestamps. Upon the arrival of a new edge, TD propagates the corresponding sge through the physical execution plan and computes the new output at the given logical timestamp.

6.2 Physical Operator Algebra

6.2.1 Overview. Physical operator implementations for streaming systems have two requirements: they should be push-based and non-blocking, so they do not need the entire input to be available before producing the first result. The standard dataflow implementations of stateless FILTER and UNION operators can be directly used in SGA, and WSCAN can be implemented via the standard *map* operator that adjusts the validity intervals of sgts based on window specifications. We focus on the stateful operators PATTERN and PATH that need to maintain an internal operator state that is accessed during query processing. This state is updated as new sgts enter the window and old sgts expire. As discussed earlier, time-based sliding windows ensure that the portion of the input that may contribute to any future result is finite, making incremental, non-blocking computation possible.

TD’s Differential Dataflow (DD) layer [51] provides a set of built-in, high-level programming primitives (operators) that can be used to compose arbitrary dataflows for general-purpose computations, and it automatically incrementalizes these. Consequently, DD can be asked to evaluate SGQ by (i) creating a dataflow of DD operators for a given SGQ and (ii) maintaining the window content as an evolving collection. Indeed, we use such a strategy as a competitive baseline in our experimental evaluation (Section 7.2.2). However, DD’s generality comes at a performance cost for evaluating SGQ, as we show in Section 7.2.2. Below, we describe how to devise physical operator implementations specific to SGQ by utilizing the properties of the SGQ model. Of course, these are not the only physical operator implementations that are possible for SGA; other implementations can possibly be developed, and these are exemplars to demonstrate the implementability of the SGA operators. They are also what we use in the experiments.

6.2.2 Pattern Matching. Implementation of PATTERN models subgraph patterns as conjunctive queries that can be evaluated using a series of joins. There is a rich literature of streaming join implementations that can be used. Symmetric hash join [71] is commonly used to implement non-blocking joins in the streaming model: a hash table is built for each input stream and upon arrival (expiration) of a tuple, it is inserted into (removed from) its corresponding hash table and other tables are probed for matches [28, 69]. This produces an append-only stream of results for internal windows that do not invalidate previously reported results upon expiration of their participating tuples [26]. For external windows that require eviction of old results as the windows slide forward, expired results can be determined by maintaining expiration timestamps: a join result expires when one of its participating tuples expire. Our use of validity intervals enables the user or the application to adopt both window semantics without the need for explicit processing of expired input tuples.

Given a subgraph pattern, we take the standard approach of creating a binary join tree where leafs represent streaming graphs as input streams and internal nodes represent pipelined hash join operators. For instance, Figure 8 (right) shows the logical plan for the query in Example 1 and the join tree for its PATTERN. In our prototype, we use the ordering of predicates in PATTERN to construct the join tree and leave the problem of finding efficient join plans (e.g. using *worst-case optimal joins* [55]) for future investigation.

6.2.3 Path Navigation. DD’s *iterate* allows constructing cyclic dataflows that can model arbitrary nested iterations, and it can be used to evaluate PATH and its recursive path expressions. However, the use of recursion in SGQ is limited to transitive closure (Section 4), and the RPQ-based semantics of PATH is sufficient to evaluate this limited form of recursion (Theorem 1). In previous work [57], we study the design space of streaming algorithms for persistent RPQ under arbitrary path semantics. In brief, the streaming RPQ algorithm follows the automata-based RPQ evaluation method [14] and maintains a spanning forest-based data structure, called Δ -tree index, that enables the compact representation of partial path segments. In our prototype, we employ the streaming RPQ algorithm in [57] as the non-blocking, physical implementation of PATH operator, which eliminates the need for cycles in physical execution plans. Additionally, a spanning-tree-based representation of intermediate

results enables us to recover actual paths and allows queries to return and manipulate paths as first-class citizens.

6.2.4 Customized PATH Implementation. We now describe a novel algorithm *Streaming Path Navigation (S-PATH)* that can be used as an alternative physical operator for the PATH operator (Definition 20). In contrast to our previous work [57], **S-PATH** utilizes the validity intervals of path segments to simplify the state maintenance in the absence of explicit deletions. Algorithms in [57] are based on the *negative tuple* approach; expirations due to window movements are processed using the same machinery as explicit deletions. Upon expiration (deletion) of an edge, their algorithm first finds all results that are affected by the expiration (deletion), then it traverses the snapshot graph to ensure that there is no alternative path leading to the same result. This corresponds to re-derivation step of *DRed* [32], optimized for RPQ evaluation on streaming graphs. Instead, **S-PATH** utilizes the temporal pattern of sliding window movements and adopt the *direct* approach, i.e., it can *directly* determine expired tuples based on their validity intervals. This is possible due to the separation of the implementation of sliding windows from operator semantics via an explicit WSCAN operator.

Algorithm S-PATH:

input : Input streaming graph S , Regular expression R , output label o
output: Output streaming graph S_O

- 1 $A(S, \Sigma, \delta, s_0, F) \leftarrow \text{ConstructDFA}(R)$
- 2 Initialize $\Delta - \text{PATH}$
- 3 $S_O \leftarrow \emptyset$
- 4 $R \leftarrow \emptyset$
- 5 **foreach** $(u, v, l, [ts, exp], \mathcal{D}) \in S$ **do**
- 6 **foreach** $s, t \in S$ where $t = \delta(s, l)$ **do**
- 7 **if** $s = s_0 \wedge T_u \notin \Delta - \text{PATH}$ **then**
- 8 | add T_u with root node (u, s_0)
- 9 **if** $s = s_0$ **then**
- 10 | **if** $(v, t) \notin T_u$ **then**
- 11 | $R \leftarrow R + \text{Expand}(T_u, (u, s_0), (v, t), e(u, v))$
- 12 | **else if** $(v, t).exp < exp$ **then**
- 13 | $R \leftarrow R + \text{Propagate}(T_u, (u, s_0), (v, t), e = (u, v))$
- 14 $T \leftarrow \text{ExpandableTrees}(\Delta - \text{PATH}, (u, s), ts)$
- 15 **foreach** $T_x \in T$ **do**
- 16 | **if** $(v, t) \notin T_x$ **then**
- 17 | $R \leftarrow R + \text{Expand}(T_x, (u, s), (v, t), e(u, v))$
- 18 | **else if** $(v, t).exp < \min((u, s).exp, exp)$ **then**
- 19 | $R \leftarrow R + \text{Propagate}(T_x, (u, s), (v, t), e = (u, v))$
- 20 **foreach** $sgt\ t \in R$ **do**
- 21 | push t to S_O

Algorithm **S-PATH** incrementally performs a traversal of the underlying snapshot graph under the constraints of a given RPQ as sgts arrive. It first constructs a DFA from the regular expression of a PATH operator, and initializes a spanning forest-based data structure, called $\Delta - \text{PATH}$, that is used as the internal operator state during query processing. $\Delta - \text{PATH}$ is used to maintain a path segment, i.e., a partial result, between each pair of vertices in the form a spanning forest under the constraints of a given RPQ, consistent with Definition 20. Upon the arrival of an sgt, Algorithm

Algorithm Expand:

input : Spanning Tree T_x rooted at (x, s_0) , parent (u, s) , child (v, t) , edge $e(u, v)$
output: Set of results R

- 1 $R \leftarrow \emptyset$
- 2 Insert (v, t) as (u, s) 's child
- 3 $(v, t).ts = \max(e.ts, (u, s).ts)$
- 4 $(v, t).exp = \min(e.exp, (u, s).exp)$
- 5 **if** $t \in F$ **then**
- 6 | $p \leftarrow \text{PATH}(T_x, (v, t))$
- 7 | $R \leftarrow R + (x, v, O, [(v, t).ts, (v, t).exp], p)$
- 8 **foreach** edge $e(v, w) \in G_{ts}$ s.t. $\delta(t, \phi(e)) = q$ **do**
- 9 | **if** $(w, q) \notin T_x$ **then**
- 10 | $R \leftarrow R + \text{Expand}(T_x, (v, t), (w, q), e(v, w))$
- 11 | **else if** $(w, q).exp < \min((v, t).exp, e.exp)$ **then**
- 12 | $R \leftarrow R + \text{Propagate}(T_x, (v, t), (w, q), e(v, w))$
- 13 **return** R

Algorithm Propagate:

input : Spanning Tree T_x rooted at (x, s_0) , parent (u, s) , child (v, t) , edge $e(u, v)$
output: Set of results R

- 1 $R \leftarrow \emptyset$
- 2 $(v, t).pt = (u, s)$
- 3 $(v, t).ts = \min((v, t).ts, \max(e.ts, (u, s).ts))$
- 4 $(v, t).exp = \max((v, t).exp, \min(e.exp, (u, s).exp))$
- 5 **if** $t \in F$ **then**
- 6 | $p \leftarrow \text{PATH}(T_x, (v, t))$
- 7 | $R \leftarrow R + (x, v, O, [(v, t).ts, (v, t).exp], p)$
- 8 **foreach** edge $e = (v, w) \in G_{ts}$ s.t. $\delta(t, \phi(e)) = q$ **do**
- 9 | **if** $(w, q).exp < \min((v, t).exp, e.exp)$ **then**
- 10 | $R \leftarrow R + \text{Propagate}(T_x, (v, t), (w, q), e(v, w))$
- 11 **return** R

S-PATH probes $\Delta - \text{PATH}$ to retrieve partial path segments that can be extended with the edge (or a path segment) of the incoming sgt. Each partial path segment is extended with the incoming sgt, and Algorithm **S-PATH** traverses the snapshot graph G_t until no further expansion is possible.

DEFINITION 21 (SPANNING TREE T_x). Given an automaton A for the regular expression R of a PATH operator \mathcal{P}_d^R and a streaming graph S at time t , a spanning tree T_x forms a compact representation of valid path segments that are reachable from the vertex $x \in G_t$ under the constraints of a given RPQ, i.e., a vertex-state pair (u, s) is in T_x at time t if there exists a path $p \in G_t$ from x to u with label $\phi^p(p)$ such that $s = \delta^*(s_0, \phi^p(p))$.

A node $(u, s) \in T_x$ indicates that there is a path p in the snapshot graph with label $\phi^p(p)$ such that $s = \delta^*(s_0, \phi^p(p))$, and this path can simply be constructed by following parent pointers $((u, s).pt)$ in T_x . Under the arbitrary path semantics, there are potentially infinitely many path segments between a pair of vertices that conform to a given RPQ due to the presence of cycles in the snapshot graph and a Kleene star in the given RPQ. Among those, **S-PATH** materializes the path segment with the largest expiry timestamp, that is, the path segment that will expire furthest in the future.

Consequently, for each node $(u, s) \in T_x$, the sequence of vertices in the path from the root node to (u, s) corresponds to the path from x to u in the snapshot graph with the largest expiry timestamp. This is achieved by the coalesce primitive (Definition 11) with an aggregation function \max over the expiry timestamp of path segments.⁷ Upon expiration of a node (u, s) in T_x and its corresponding path segment in the snapshot graph, this guarantees that there cannot be an alternative path segment between x and u that have not yet expired. Hence, we can *directly* find expired tuples based on their expiry timestamps. This is based on the observation that expirations have a temporal order unlike explicit deletions, and S-PATH utilizes these temporal patterns to simplify window maintenance.

DEFINITION 22 (Δ - PATH INDEX). *Given an automaton A for the regular expression R of a PATH operator \mathcal{P}_d^R and a streaming graph S at time t , Δ - PATH is a collection of spanning trees (Definition 21) where each tree T_x is rooted at a vertex $x \in G_t$ for which there is an sgt $t \in S(t)$ with a label l such that $\delta(s_0, l) \neq \emptyset$ and $\text{src} = x$.*

Δ - PATH encodes a single entry for each pair of vertices under the constraints of a given query, consistent with the set semantics of snapshot graphs (Section 3). Due to spanning-tree construction (Definition 21), actual paths can easily be recovered by following the parent pointers; hence, Δ - PATH constitutes a compact representation of intermediate results for path navigation queries over materialized path graphs. Our implementation models Δ - PATH as a hash-based inverted index from vertex-state pairs to spanning trees, enabling quick look-up to locate all spanning trees that contain a particular vertex-state pair. Upon arrival of an sgt $t = (u, v, l, [ts, \text{exp}], \mathcal{D})$, Algorithm **S-PATH** probes this inverted index of Δ - PATH to retrieve all path segments that can be extended with the incoming sgt, that is, spanning trees that have the node (u, s) with an expiry timestamp smaller than ts for any state $s \in \{s \in S \mid \delta(s, l) \neq \emptyset\}$ (Line 14). If the target node (v, t) for $t = \delta(s, l)$ is not in the spanning tree T_x , Algorithm **Expand** is invoked to expand the existing path segment from $(x, 0)$ to (u, s) with the node (v, t) and to create a new leaf node as a child of (u, s) . In case there already exists a path segment between vertices $(x, 0)$ and (v, t) in Δ - PATH, i.e., the target node (v, t) is already in T_x , Algorithm **S-PATH** compares its expiry timestamp with the new candidate (Line 18). If the extension of the existing path segment from $(x, 0)$ to (u, s) with (v, t) results in a larger expiry timestamp than $(v, t).\text{exp}$, Algorithm **Propagate** is invoked to update the expiry timestamp of (v, t) and its children in T_x . Algorithms **Expand** and **Propagate** traverse the snapshot graph until no further update is possible. The following example illustrates the behaviour of Algorithm **S-PATH** on our running example.

EXAMPLE 9. *Consider the real-time notification query in Example 1 whose SGA expression is given in Example 8. Figure 9a shows an excerpt of the streaming graph input to the PATH operator \mathcal{P}_{RLP}^{RL+} . Figure 9b depicts a spanning tree $T_x \in \Delta$ -PATH at $t = 27$. Upon arrival of the sgt $(y, u, RL, [28, 37], \mathcal{D} = \{(y, RL, u)\})$ at $t = 28$, Algorithm **S-PATH** extends the path segment from $(x, 0)$ to $(y, 1)$ with $(u, 1)$, and compares its expiry timestamp with that of $(u, 1)$ that is already in T_x . As the new extension has larger expiry timestamp, the validity interval*

*and the parent pointer of $(u, 1) \in T_x$ is updated (Line 18 in Algorithm **S-PATH**). Then, incoming sgts at times $t = 28$ and $t = 29$ are processed by Algorithm **Expand** as corresponding target nodes $(v, 1)$ and $(s, 1)$ are not in T_x , adding $(v, 1)$ and $(s, 1)$ as children of $(u, 1)$. At $t = 30$, the incoming sgt $(w, v, RL, [30, 39], \mathcal{D} = \{(w, RL, v)\})$ might extend the path segment from $(x, 0)$ to $(v, 1)$ with expiry timestamp 33. However, Algorithm **S-PATH** does not make any modification to Δ - PATH, as $(v, 1)$ is already in T_x with a larger timestamp (Line 18). Figure 9c depicts the resulting spanning tree T_x at $t = 30$.*

As described earlier, Δ - PATH stores a *parent* pointer for each node pointing to its parent node in the corresponding spanning tree, and Algorithm **Propagate** updates these pointers during processing. By traversing these parent pointers for each resulting sgt, Algorithm **Expand** can construct the actual path (Line 7) and return it as a part of the resulting sgt, i.e., it populates the implicit payload attribute \mathcal{D} of the resulting sgt with the sequence of edges that forms the resulting path. The cost of this operation is $O(l)$ where l is the length of the resulting path.

Δ - PATH guarantees that the expiry timestamp of a node (u, s) in T_x is equal to largest expiry timestamp of all paths between x and u in the snapshot graph with a label l such that $s = \delta^*(s_0, l)$. Consequently, for a node $(u, s) \in T_x$ with expiry timestamp smaller than t , there cannot be another path from x to u with an equivalent label that is valid at time t . Consider the spanning tree given in Figure 9c. Algorithm **S-PATH** can directly determine, without additional processing, that nodes $(z, 1)$ and $(t, 1)$ are expired as their expiry timestamp is 31. Thus, at any given time t , Algorithm **S-PATH** can simply ignore a node $(u, s) \in T_x$ with expiry timestamps smaller than t (Line 18) and such nodes can be removed from Δ - PATH. To prevent Δ - PATH from growing unboundedly due to expired tuples, a background process periodically purges expired tuples from Δ - PATH.

The following example illustrates how the *negative tuple* and *direct* approaches differ, respectively for [57] and our proposed algorithm.

EXAMPLE 10. *Consider the same real-time notification query as in Example 9. Both approaches behave similarly until $t = 28$ as all vertex-state pairs in T_x have a single derivation at $t = 27$ (Figure 9b). Upon arrival of the sgt $(y, u, HI, [28, 37], \mathcal{D} = \{(y, HI, u)\})$ at $t = 28$, the negative tuple approach as in [57] does not update T_x as $(u, 1)$ is already in T_x , whereas the direct approach as used in this paper updates the validity interval and the parent pointer of $(u, 1) \in T_x$ (Line 18 in Algorithm **S-PATH**). Then, incoming sgts at times $t = 28$ and $t = 29$ are processed similarly, adding $(v, 1)$ and $(s, 1)$ as children of $(u, 1)$. Figures 9c and 9d depict the corresponding spanning trees at $t = 30$ for the direct and the negative tuple approaches, respectively. Note that in Figure 9c, the validity intervals of nodes in the subtree rooted at node $(u, 1)$ reflects the newly discovered path from x to u through y in G_{30} . The negative tuple and the direct approach differs at $t = 31$ as multiple nodes expire. The negative tuple used as in [57] marks the entire subtree of $(z, 1)$ as potentially expired (Figure 9d), and performs a traversal of the snapshot graph G_{31} to find alternative, valid paths for expired nodes. These traversals undo the effect of expired sgts via explicit deletions. Upon discovering alternative paths for nodes $(u, 1)$, $(v, 1)$ and $(s, 1)$ that are valid at time $t = 31$, they are re-inserted into T_x . Instead, our proposed algorithm can directly*

⁷Arbitrary path semantics provides the flexibility for the aggregation function f_{agg} of the coalesce primitive.

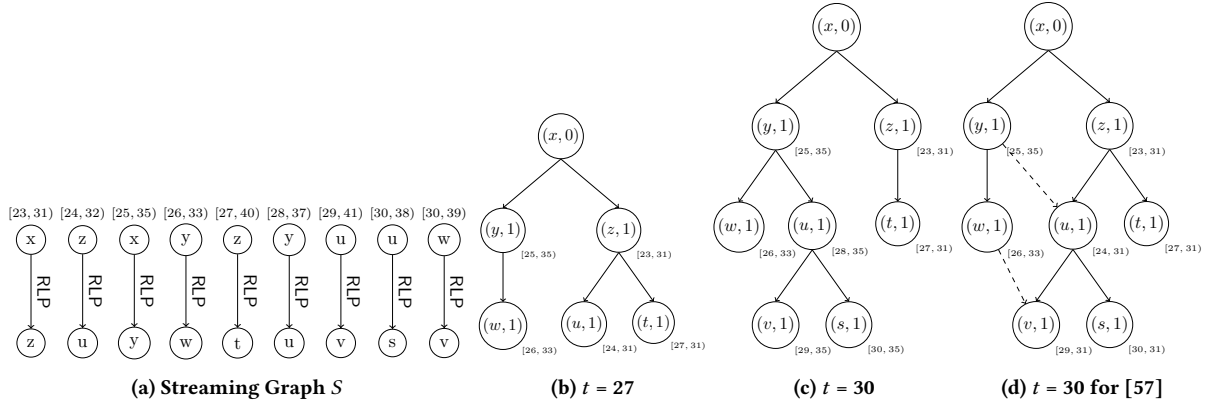


Figure 9: (a) A streaming graph S_{RLP} as the input for PATH operator, (b) spanning tree T_x at $t = 28$, (c) spanning tree T_x at $t = 30$ of the proposed algorithm following the *direct* approach, and (d) spanning tree T_x at $t = 30$ of [57] following the *negative tuple* approach.

determine the expired nodes based on the validity intervals (nodes $(z, 1)$ and $(t, 1)$ as shown in Figure 9c) without additional processing.

6.2.5 Explicit Deletions. Physical implementations of PATTERN and PATH rely on the *direct* approach that utilizes the temporal pattern of sliding windows for state maintenance; the expiration timestamps are used to *directly* locate expired tuples. On append-only streaming graphs, existing sgts only expire due to window movements. Albeit rare, certain applications might require explicit deletions of previously inserted sgts, which necessitates the use of the *negative tuple* approach to handle such explicit deletions.⁸ Deleted sgt, with an additional flag to denote deletion – i.e., a negative tuple, is used to undo the effect of the original sgt on the operator state and to invalidate previously reported results, if necessary. For pipelined hash join, which is used in the implementation of PATTERN, processing of negative tuples is the same as original input tuples: a negative tuple is removed from its corresponding hash table and other tables are probed to find corresponding deleted results.

The use of negative tuples for explicit deletions in the context of RPQ evaluation is first proposed by Pacaci et al [57]. Here, we describe how to adopt the *negative tuple* approach for explicit deletions of sgts in Algorithm **S-PATH**. In brief, upon explicit deletion of an sgt, we first identify tree-edges that disconnect spanning trees in Δ – PATH. For each such tree-edge, we first mark the nodes in the subtree that are disconnected due to the explicit deletion. Then, for each node in this subtree, we use a Dijkstra-based traversal over the snapshot graph to find an alternative path with the largest expiry timestamp. Dijkstra’s algorithm guarantees that we can efficiently find the path with the largest expiry timestamp for each marked node, consistent with Definition 22. A marked node is removed from the spanning tree only if there is no alternative valid path. Deletion of a non-tree edge does not require any modification as it leaves spanning trees unchanged.

⁸The *negative tuple* approach can also be used to signal expirations through an explicit deletion of the corresponding sgt to undo its effect [25, 29, 57]. Indeed, the negative-tuple approach can be used for incremental evaluation of an arbitrary computation, whereas direct approach is applicable to negation-free queries over append-only streaming graphs [29].

7 EXPERIMENTAL ANALYSIS

Our objective is to demonstrate the feasibility of implementing a performant system that incorporates the algebraic framework we propose in this paper. Using the prototype implementation described in §6, we first provide an end-to-end performance analysis of our algebraic approach for persistent evaluation of streaming graph queries (§7.2). Then, we assess the scalability by varying the window size \mathcal{T} and the slide interval β (§7.3). Finally, we highlight the benefits of the proposed SGA in exploring the rich plan space through transformation rules and demonstrate the potential performance improvements of exploring this plan space (§7.4).

7.1 Experimental Setup

7.1.1 Setup. Experiments are run on a Linux server with 32 physical cores and 256GB memory. For each query and configuration, we report the tail latency of each window slide, i.e., the total time to process all arriving and expired sgts upon window movement and to produce new results, and the average throughput after ten minutes of processing on warm caches.

7.1.2 Datasets. We use **Stackoverflow** (SO) and **LDBC SNB** (SNB) graphs for our experimental analysis; these are publicly available, large-scale graphs with labelled and timestamped edges on which persistent queries with complex graph patterns can be formulated. SO is a temporal graph of user interactions on the stackoverflow website containing 63M interactions (edges) of 2.2M users (vertices), spanning 8 years [59], and SNB is a synthetic social network graph that simulates the interactions of an online social network [22]. We extract the update stream of the LDBC workload that contains 8 different types of interactions, and we use replyOf, hasCreator and likes edges between users and posts, and knows edges between users. We use a scale factor of 10 with 7.2M users and posts (vertices) and 40M user interactions (edges). SO contains only a single type of vertex and 3 different edge labels, and its cyclic nature causes a high number of intermediate results and resulting paths; so it is the most challenging one for the proposed algorithms. We set the

Table 1: $Q_1 - Q_4$ correspond to common RPQ observed in real-world query logs [15], and $Q_5 - Q_7$ are Datalog encodings of RQ-based complex graph patterns that we use to define streaming graph queries. Q_5 and Q_6 correspond to complex graph patterns of LDBC SNB queries *IS7* and *IC7* [22], respectively, and Q_7 corresponds to the complex graph pattern given in Example 1 that is defined as a recursive path query over the graph pattern of Q_6 . a, b and c correspond to edge predicates that are instantiated based on the dataset characteristics.

Name	Query
Q_1	$?x, ?y \leftarrow ?x a^* ?y$
Q_2	$?x, ?y \leftarrow ?x a \circ b^* ?y$
Q_3	$?x, ?y \leftarrow ?x a \circ b^* \circ c^* ?y$
Q_4	$?x, ?y \leftarrow ?x (a \circ b \circ c)^+ ?y$
Q_5	$RR(m1, m2) \leftarrow a(x, y), b(m1, x), b(m2, y), c(m2, m1)$
Q_6	$RL(x, y) \leftarrow a^+(x, y), b(x, m), c(m, y)$
Q_7	$RL(x, y) \leftarrow a^+(x, y), b(x, m), c(m, y)$ $Ans(x, m) \leftarrow RL^+(x, y), c(m, y)$

window size \mathcal{T} to 1 month and the slide interval β to 1 day unless specified otherwise.

7.1.3 Workloads. To the best of our knowledge, no current benchmark exists featuring RQ for graph DBMSs. The existing benchmarks are limited to UCRPQ, thus not capturing the full expressivity of RQ even for static graphs. Streaming RDF benchmarks such as LSBench (<https://code.google.com/archive/p/lbench/>) and Stream WatDiv [24] only focus on SPARQL v1.0 (thus not even including simple RPQs), and their workloads do not contain any recursive queries. Hence, we formulate a set SGQ from existing UCRPQ-based workloads as follows: we collect a set of graph patterns in the form of UCRPQ from existing benchmarks and studies [11, 15, 22, 57, 73], and we compose a set of complex graph patterns from those by applying a Kleene star over each graph pattern. Table 1 lists the set of graph patterns of increasing expressivity (from RPQ to complex RQ with complex graph patterns) that we use to define streaming graph queries. $Q_1 - Q_4$ are commonly used RPQs in existing studies [15, 57, 73], and we use those to test our PATH operator. Q_5 & Q_6 are CRPQ-based complex graph patterns based on SNB queries *IS7* and *IC7* [22]. For instance, $Q_6 - IC7$ of SNB – with edge labels knows, likes and hasCreator asks for recent likers of a person’s messages that are also connected by a path of friends. $Q_7 - Ex. 1$ – is the most expressive RQ-based complex graph patterns we use to demonstrate the abilities of the proposed SGA to unify subgraph pattern and path navigation queries in a structured manner and to treat paths as first-class citizens. It defines a path query over the complex graph pattern of Q_6 ; it finds arbitrary length paths where users are connected by the recentLiker pattern. Note that this query cannot be expressed in existing graph query languages such as Cypher and SPARQL. Finally, for each dataset, we instantiate the query workload from these graph patterns by choosing appropriate predicates, i.e., edge labels, for each query edge and by setting the duration of time-based sliding windows $\mathcal{W}_{\mathcal{T}}$ as described above.

7.2 Query Processing Performance

7.2.1 Throughput & Tail Latency. Table 2 (SGA) shows the aggregated throughput and tail latency of our streaming graph query processor for all queries in Table 1. We discard each streaming graph edge whose label is not in a given SGQ. Tail latencies reflect the 99th percentile latency of processing a window slide and produce the corresponding resulting sgts. Across queries, the performance is lower for SO graph because it is dense and cyclic. The throughput ranges from hundreds of edges-per-second for the SO to hundreds of thousands of edges-per-second for SNB.

7.2.2 Comparative Analysis. Existing work on query processing over streaming data such as data stream management systems and streaming RDF systems cannot process queries in Table 1 as they focus on relational queries and SPARQL v1.0, respectively (§1). To the best of our knowledge, TD with its DD layer is the only general-purpose system that can be used to incrementally evaluate recursive computations that are modelled as cyclic dataflows. Table 2 (DD) reports the throughput and tail latency of DD dataflows for all queries in Table 1. Overall, our SGA-based query processor outperforms the DD baseline on SO and provides a competitive performance on the SNB dataset. On SNB, Q_6 & Q_7 do not have the Kleene-plus over a as it causes DD to timeout. Due to highly cyclic structure of SO, there are many alternative paths between each pair of vertices, and our streaming RPQ algorithm for PATH implementation maintains a compact representation of valid path segments and utilizes the temporal patterns of sliding window movements to simplify expirations (§6.2). DD-based query processor provides better performance on linear path queries $Q_1 - Q_4$ on SNB, but not others. This is due to the tree-shaped structure of replyOf edges in SNB, where there is only one path between a pair of vertices, so PATH specific optimizations do not apply. Performance variations on SNB suggest optimization opportunities for recursive graph queries when selecting physical operator implementations, as in the case for streaming relational joins [29]. These results demonstrate the feasibility of our algebraic approach for evaluating SGQ and our physical operator implementations.

7.3 Sensitivity Analysis

In this section, we analyze the impact of the window size \mathcal{T} and the slide interval β on end-to-end query performance of the proposed streaming graph query processor. We use SO graph for this experiment as its dense, cyclic structure stresses our operator implementations. Fig. 10a reports the aggregate throughput and the tail latency for each query across various window sizes. As expected, the throughput of all tested queries decreases with increasing \mathcal{T} , as a larger window size increases the # of sgts in each window. Similarly, the tail latency of each window slide increases with the increasing window size.

We also assess the impact of the slide interval β on performance. As previously mentioned, the slide interval β controls the time-granularity at which the sliding window progresses, and our prototype implementation uses β to control the input batch size. Figure 10b shows that the aggregate throughput and the tail latency for each query remain stable across varying slide intervals. This is due to tuple-oriented implementation of physical operators of SGA; SGA operators are designed to process each incoming tuple eagerly in favour of minimizing tuple-processing latency, and they do not

Table 2: (Tput) The throughput (edges/s) and (TL) the tail latency (s) of SGA and DD for queries in Table 1 on SO and SNB graphs with $|W| = 30$ days and $\beta = 1$ day.

Graph	System	Q ₁		Q ₂		Q ₃		Q ₄		Q ₅		Q ₆		Q ₇	
		Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL
SO	SGA	2762	3.3	8513	4.3	413	120	379	102.4	231064	0.3	374	52.7	376	56.3
	DD	1209	6.3	4512	5.8	368	121.7	374	82.8	63330	1	283	72.6	275	74
SNB	SGA	97187	1.5	237313	1.9	245766	1.9	277475	0.4	13345	79.1	428592	0.8	131250	10.2
	DD	121133	0.8	299245	1.2	316267	1.1	303068	0.2	12053	109.5	402048	0.9	21284	141

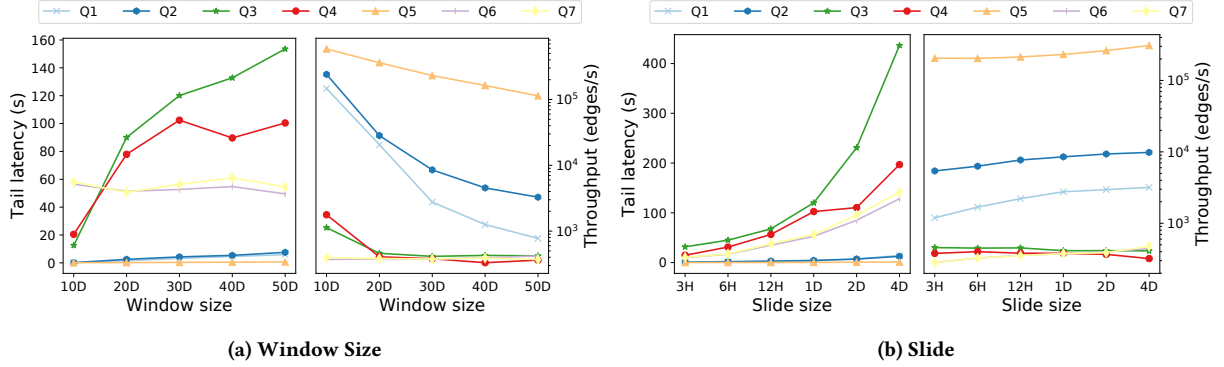


Figure 10: The tail latency of each window slide and the aggregate throughput of the proposed streaming graph query processor with increasing (a) window size \mathcal{T} and (b) slide interval β on SO graph.

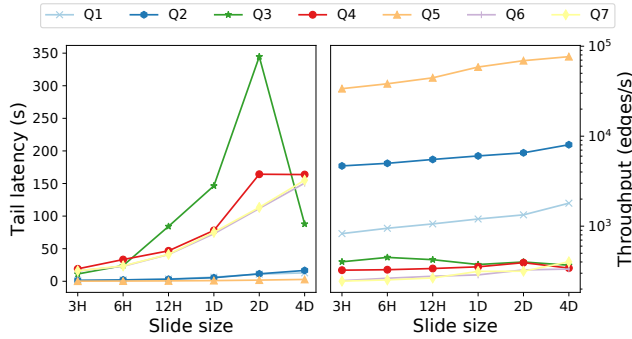


Figure 11: The tail latency of each window slide and the aggregate throughput of SGQ evaluation on DD with increasing slide interval β on SO graph.

utilize batching to improve throughput with larger batch sizes. Consequently, the tail latency of window movements increases with increasing slide interval. This is in contrast to DD whose throughput increases with increasing β as shown in Figure 11. DD and its underlying indexing mechanism, i.e., shared arrangements [50], are designed to utilize batching and improve throughput with increasing batching size: all sgts that arrive within one interval are batched together with a single logical timestamp (epoch) and DD operators can explore the latency vs throughput trade-off by changing the granularity of each epoch. The investigation of batching within SGA operators and the identification of other optimization opportunities is a topic of future work.

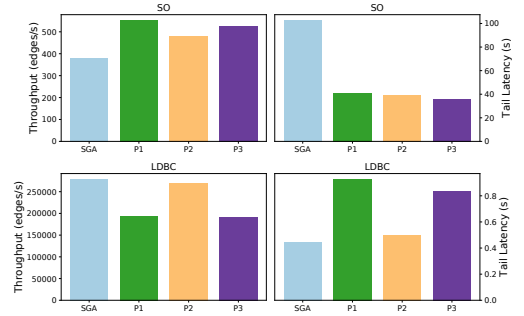


Figure 12: The throughput and tail latency of (a) Q_4 on (top) SO and (bottom) SNB for the default SGA plan and an alternative equivalent physical plan generated via SGA transformation rules (Section 5.4).

7.4 Exploring the Plan Space

SGA proposed in this paper enables a rich foundation for logical SGQ optimization through query rewriting as previously discussed (§5.4). In this paper, we do not address the full scope of optimization issues (full optimizer development is the topic of ongoing research), but we design a micro-benchmark to highlight the possibilities provided by SGA. In particular, we choose Q_4 (Table 1) as its linear pattern combined with a Kleene plus demonstrates the potential benefits of SGA transformation rules involving SGA's novel PATH operator. Fig. 12 demonstrates the throughput and the tail latency of different plans obtained from following equivalent SGA expressions for Q_4 :

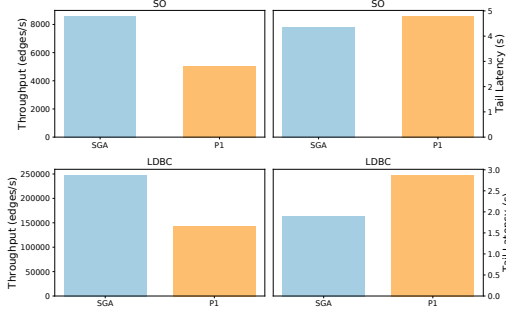


Figure 13: The throughput and tail latency of (a) Q_2 on (top) SO and (bottom) SNB for the default SGA plan and an alternative equivalent physical plan generated via SGA transformation rules (Section 5.4).

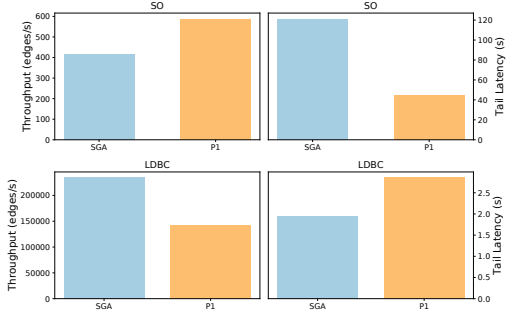


Figure 14: The throughput and tail latency of (a) Q_3 on (top) SO and (bottom) SNB for the default SGA plan and an alternative equivalent physical plan generated via SGA transformation rules (Section 5.4).

- SGA: $\mathcal{P}_{d^+}^l(\bowtie_{trg_1=src_2 \wedge trg_2=src_3}^{src_1, trg_3, d}(S_a, S_b, S_c))$
- P1: $\mathcal{P}_{(a \cdot b \cdot c)^+}^l(S_a, S_b, S_c)$
- P2: $\mathcal{P}_{(a \cdot d)^+}^l(S_a, \bowtie_{trg_1=src_2}^{src_1, trg_2, d}(S_b, S_c))$
- P3: $\mathcal{P}_{(d \cdot c)^+}^l(S_c, \bowtie_{trg_1=src_2}^{src_1, trg_2, d}(S_a, S_b))$

The first expression, **SGA**, is the canonical SGA expression for Q_4 that is generated by the Algorithm **SGQParser**. Such plans are called *loop-caching* in literature as they enable re-use of the intermediate results for the base pattern $(a \cdot b \cdot c)$ [73]⁹. **P1**, **P2** and **P3** are obtained from the canonical SGA expression using the transformation rules given in §5.4, and represent novel plans that are possible due to novel PATH operator. Fig. 12 clearly illustrates the potential benefits of exploring the rich plan space offered by SGA: some of the newly computed plans provide up to 60% increase in throughput and 60% reduction in the latency. We observe a similar behaviour on other path queries Q_2 and Q_3 (up to 50% difference in throughput). These results suggest further optimization opportunities for logical query optimization as query rewrites that are generated by SGA

⁹It is also the only plan available to DD baseline that performs a fixed-point iteration over the base pattern $(a \cdot b \cdot c)$.

transformation rules can provide significant performance benefits for evaluating SGQ over streaming graphs.

7.5 Alternative Physical Operators

Our prototype implementation employs the automata-based streaming RPQ algorithm of [57] as the physical implementation of the PATH operator. Section 6.2.4 describes an alternative physical implementation (*Streaming Path Navigation – S-PATH*) for the novel PATH operator of our proposed SGA based on the *direct approach*. Table 3 compares its impact on the performance of our proposed streaming graph query processor. The performance impact of adopting the Algorithm **S-PATH** for PATH varies depending on the workload characteristics: it provides improvements in throughput for most queries on the SO graph, whereas the performance differences on the SNB graph is small. This can be attributed to the cyclic nature of the SO graph. A higher number of intermediate path segments increases the size of the internal operator state that the physical implementations of the PATH operator need to maintain. Overall, these results illustrate the potential benefits of exploring alternative operator implementations. In particular, cost-based optimization techniques for determining suitable operator implementations based on data and query characteristics might provide significant performance benefits for evaluating SGQ over streaming graphs.

8 CONCLUSION AND FUTURE WORK

This paper introduces a general-purpose query processing framework for streaming graphs that consists of (i) streaming graph query model and algebra with well-founded semantics, and (ii) a prototype streaming graph query processor as an embodiment of the proposed framework. The SGQ model addresses the requirements we set forth for querying streaming graphs by capturing both common features of existing graph query languages and novel features that are beyond the power of existing languages (such as the treatment of paths as first-class citizens, powerful path patterns and composability). SGA and its transformation rules provide the foundational framework to express and evaluate streaming graph queries. Our prototype streaming graph query processor exemplifies physical operator implementations specific to streaming graph queries due to this framework. Experimental analyses on real-world and synthetic streaming graphs demonstrate the feasibility and the potential performance gains of our approach.

SGQ and SGA establish the basis of systematic study of query processing issues over streaming graphs. We are now engaged in research along two dimensions: (i) designing an SGA-based query optimizer for the systematic exploration of the rich plan space using SGA’s transformation rules, and (ii) incorporating attribute-based predicates to fully support the property graph model.

Additional work that would enrich the implementation include the development of alternative physical operators and additional transformation rules for plan space enumeration. Given that the SGQ-SGA framework is complete with a clear semantics (as captured in RQ model), and given that SGQ incorporates the capabilities of existing and emerging query languages, future work can focus on efficient implementation.

Table 3: The impact of S-PATH on the performance of the throughput (edges/s) and (TL) the tail latency (s) of the queries in Table 1 on SO and SNB graphs with $|W| = 30$ days and $\beta = 1$ day.

Graph	System	Q_1		Q_2		Q_3		Q_4		Q_5		Q_6		Q_7	
		Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL	Tput	TL
SO	S-PATH	2884	4	9074	4.9	391	177	348	94.9	234058	0.4	625	51.4	353	52.6
	Improvement	4.4%	-21.2%	6.5%	-13.9%	-5.3%	-47.5 %	-8.1%	7.3%	1.2%	-33.3 %	67.1%	2.4%	-6.1%	6.5%
SNB	S-PATH	95903	1.4	244653	1.8	224342	1.9	278647	0.4	14000	79.5	450957	0.8	130651	10.8
	Improvement	-1.3%	6.6%	3.1 %	5.2%	-8.7%	-	0.4%	-	4.9%	-0.5%	5.2%	-	-0.4%	-5.8%

REFERENCES

- [1] Daniel J. Abadi, Yanif Ahmad, Magdalena Balazinska, Ugur Çetintemel, Mitch Cherniack, Jeong-Hyon Hwang, Wolfgang Lindner, Anurag Maskey, Alex Rasin, Esther Ryzkina, Nesime Tatbul, Ying Xing, and Stanley B. Zdonik. 2005. The Design of the Borealis Stream Processing Engine. In *Proc. 2nd Biennial Conf. on Innovative Data Systems Research*. 277–289.
- [2] Daniel J. Abadi, Don Carney, Ugur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. 2003. Aurora: a new model and architecture for data stream management. *Vldb J.* 12, 2 (2003), 120–139.
- [3] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of databases*. Vol. 8. Addison Wesley.
- [4] Amir Aghasadeghi, Vera Zaychik Moffitt, Sebastian Schelter, and Julia Stoyanovich. 2020. Zooming Out on an Evolving Graph.. In *Proc. 23rd Int. Conf. on Extending Database Technology*. 25–36.
- [5] Khaled Ammar, Frank McSherry, Semih Salihoglu, and Manas Joglekar. 2018. Distributed Evaluation of Subgraph Queries Using Worst-case Optimal and Low-Memory Dataflows. *Proc. VLDB Endowment* 11, 6 (2018), 691–704. <https://doi.org/10.14778/3184470.3184473>
- [6] Renzo Angles, Marcelo Arenas, Pablo Barcelo, Peter Boncz, George Fletcher, Claudio Gutierrez, Tobias Lindacker, Marcus Paradies, Stefan Plantikow, Juan Sequeda, et al. 2018. G-CORE: A core for future graph query languages. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 1421–1432.
- [7] Renzo Angles, Marcelo Arenas, Pablo Barcelo, Aidan Hogan, Juan Reutter, and Domagoj Vrgoč. 2017. Foundations of modern query languages for graph databases. *ACM Comput. Surv.* 50, 5 (2017), 68.
- [8] A. Arasu, S. Babu, and J. Widom. 2006. The CQL Continuous Query Language: Semantic Foundations and Query Execution. *Vldb J.* 15, 2 (2006), 121–142.
- [9] B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. 2002. Models and Issues in Data Stream Systems. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*. 1–16.
- [10] Pablo Barcelo Baeza. 2013. Querying graph databases. In *Proc. 32nd ACM SIGACT-SIGMOD-SIGART Symp. on Principles of Database Systems*. 175–188.
- [11] Guillaume Bagan, Angela Bonifati, Radu Ciucanu, George HL Fletcher, Aurélien Lemay, and Nicky Advokaat. 2016. gMark: schema-driven generation of graphs and queries. *IEEE Trans. Knowl. and Data Eng.* 29, 4 (2016), 856–869.
- [12] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, and Michael Grossniklaus. 2009. C-SPARQL: SPARQL for continuous querying. In *Proc. 18th Int. World Wide Web Conf.* 1061–1062.
- [13] Angela Bonifati and Stefania Dumbra. 2019. Graph queries: From theory to practice. *ACM SIGMOD Rec.* 47, 4 (2019), 5–16.
- [14] Angela Bonifati, George Fletcher, Hannes Voigt, and Nikolay Yakovets. 2018. Querying Graphs. *Synthesis Lectures on Data Management* 10, 3 (2018), 1–184.
- [15] Angela Bonifati, Wim Martens, and Thomas Timm. 2019. Navigating the Maze of Wikidata Query Logs. In *Proc. 28th Int. World Wide Web Conf.* 127–138.
- [16] Jean-Paul Calbimonte, Oscar Corcho, and Alasdair JG Gray. 2010. Enabling ontology-based access to streaming data sources. In *Proc. 9th Int. Semantic Web Conf.* 96–111.
- [17] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink™: Stream and Batch Processing in a Single Engine. *Q. Bull. IEEE TC on Data Eng.* 38, 4 (2015), 28–38. <http://sites.computer.org/debull/A15dec/p28.pdf>
- [18] Sutanay Choudhury, Lawrence B. Holder, George Chin Jr, Khushbu Agarwal, and John Feo. 2015. A Selectivity based approach to Continuous Pattern Detection in Streaming Graphs. In *Proc. 18th Int. Conf. on Extending Database Technology*. 157–168. <https://doi.org/10.5441/002/edbt.2015.15>
- [19] Christian S Jensen James Clifford, Ramez Elmasri, Curtis Dyreson, Fabio Grandi Wolfgang K&fer Nick Kline, Nikos Lorentzos, Yamsis Mitsopoulos, Angelo Montanari, Daniel Nonen Elisa Peressi Barbara Pernici, John F Roddick Nandlal L Sarda, and Maria Rita Scalas Arie Segev. 1994. A consensus glossary of temporal database concepts. *ACM SIGMOD Rec.* 23, 1 (1994).
- [20] Daniele Dell’Aglio, Jean-Paul Calbimonte, Emanuele Della Valle, and Oscar Corcho. 2015. Towards a unified language for RDF stream query processing. In *Proc. 12th Extended Semantic Web Conf.* 353–363.
- [21] David Ediger, Rob McColl, Jason Riedy, and David A Bader. 2012. Stinger: High performance data structure for streaming graphs. In *Proc. 2012 IEEE Conf. on High Performance Extreme Comp.* IEEE, 1–5.
- [22] Orri Erling, Alex Averbuch, Josep Larriba-Pey, Hassan Chafi, Andrey Gubichev, Arnau Prat, Minh-Duc Pham, and Peter Boncz. 2015. The LDBC Social Network Benchmark: Interactive Workload. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 619–630. <https://doi.org/10.1145/2723372.2742786>
- [23] Wenfei Fan, Chunming Hu, and Chao Tian. 2017. Incremental graph computations: Doable and undoable. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 155–169.
- [24] Libo Gao, Lukasz Golab, M. Tamer Özsu, and Gunes Aluc. 2018. Stream WatDiv – A Streaming RDF Benchmark. In *Proc. ACM SIGMOD Workshop on Semantic Big Data*. 3:1–3:6.
- [25] Thanaa M Ghanem, Moustafa A Hammad, Mohamed F Mokbel, Walid G Aref, and Ahmed K Elmagarmid. 2006. Incremental evaluation of sliding-window queries over data streams. *IEEE Trans. Knowl. and Data Eng.* 19, 1 (2006), 57–72.
- [26] L. Golab. 2006. *Sliding Window Query Processing over Data Streams*. Ph.D. Dissertation. University of Waterloo.
- [27] Lukasz Golab and M. Tamer Özsu. 2003. Issues in data stream management. *ACM SIGMOD Rec.* 32, 2 (2003), 5–14.
- [28] Lukasz Golab and M. Tamer Özsu. 2003. Processing Sliding Window Multi-Joins in Continuous Queries over Data Streams. In *Proc. 29th Int. Conf. on Very Large Data Bases*. 500–511.
- [29] Lukasz Golab and M. Tamer Özsu. 2005. Update-Pattern-Aware Modeling and Processing of Continuous Queries. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 658–669.
- [30] G. Graefe. 1993. Query Evaluation Techniques for Large Databases. *ACM Comput. Surv.* 25, 2 (1993), 73–170.
- [31] Ajeet Grewal, Jerry Jiang, Gary Lam, Tristan Jung, Lohith Vuddemmarri, Quannan Li, Aaditya Landge, and Jimmy Lin. 2018. RecService: Multi-Tenant Distributed Real-Time Graph Processing at Twitter. In *Proc. 10th USENIX Workshop on Hot Topics in Cloud Computing*.
- [32] A. Gupta, I. S. Mumick, and V. S. Subrahmanian. 1993. Maintaining Views Incrementally. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 157–166.
- [33] Martin Hürzel, Guillaume Baudart, Angela Bonifati, Emanuele Della Valle, Sherif Sakr, and Akrivi Vlachou. 2018. Stream Processing Languages in the Big Data Era. *ACM SIGMOD Rec.* 47, 2 (2018), 29–40.
- [34] Anand Padmanabha Iyer, Li Erran Li, Athagata Das, and Ion Stoica. 2016. Time-evolving graph processing at scale. In *Proc. 4th Int. Workshop on Graph Data Management Experiences and Systems*. 1–6.
- [35] Kyoungmin Kim, In Seo, Wook-Shin Han, Jeong-Hoon Lee, Sungpack Hong, Hassan Chafi, Hyungyu Shin, and Geonhwa Jeong. 2018. TurboFlux: A Fast Continuous Subgraph Matching System for Streaming Graph Data. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 411–426.
- [36] Christoph Koch, Yanif Ahmad, Oliver Kennedy, Milos Nikolic, Andres Nötzli, Daniel Lupei, and Amir Shaikhha. 2014. DBToaster: Higher-Order Delta Processing for Dynamic, Frequently Fresh Views. *Vldb J.* 23, 2 (2014), 253–278.
- [37] Srdjan Komazec, Davide Cerri, and Dieter Fensel. 2012. Sparkwave: continuous schema-enhanced pattern matching over RDF data streams. In *Proc. 6th Int. Conf. Distributed Event-Based Systems*. 58–68.
- [38] André Koschmieder and Ulf Leser. 2012. Regular path queries on large graphs. In *SSDBM12*. 177–194.
- [39] Jürgen Krämer and Bernhard Seeger. 2009. Semantics and implementation of continuous sliding window queries over data streams. *ACM Trans. Database Syst.* 34, 1 (2009), 1–49.
- [40] Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg, Suresh Mittal, Jignesh M. Patel, Karthik Ramasamy, and Siddarth Taneja. 2015. Twitter Heron: Stream Processing at Scale. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 239–250. <https://doi.org/10.1145/2723372.2742788>
- [41] Pradeep Kumar and H Howie Huang. 2019. GraphOne: A data store for real-time analytics on evolving graphs. In *Proc. 17th USENIX Conf. on File and Storage Technologies*. 249–263.

- [42] Pradeep Kumar and H Howie Huang. 2020. GraphOne: A Data Store for Real-time Analytics on Evolving Graphs. *ACM Trans. Storage* 15, 4 (2020), 1–40.
- [43] Danh Le-Phuoc, Minh Dao-Tran, Josiane Xavier Parreira, and Manfred Hauswirth. 2011. A native and adaptive approach for unified processing of linked streams and linked data. In *Proc. 10th Int. Semantic Web Conf.* 370–388.
- [44] Youhuan Li, Lei Zou, M Tamer Özsu, and Dongyan Zhao. 2019. Time constrained continuous subgraph search over streaming graphs. In *Proc. 35th Int. Conf. on Data Engineering*. IEEE Press, 1082–1093.
- [45] Leonid Libkin, Juan L Reutter, Adrian Soto, and Domagoj Vrgoč. 2018. TriAL: A navigational algebra for RDF triplestores. *ACM Trans. Database Syst.* 43, 1 (2018), 1–46.
- [46] Leonid Libkin and Domagoj Vrgoč. 2012. Regular path queries on graphs with data. In *Proc. 15th Int. Conf. on Database Theory*. 74–85.
- [47] Ling Liu and M. Tamer Özsu (Eds.). 2009. *Encyclopedia of Database Systems*. Springer.
- [48] Mengmeng Liu, N.E. Taylor, Wenchao Zhou, Z.G. Ives, and Boon Thau Loo. 2009. Recursive Computation of Regions and Connectivity in Networks. In *Proc. 25th Int. Conf. on Data Engineering*. 1108–1119. <https://doi.org/10.1109/ICDE.2009.36>
- [49] Mugilan Mariappan and Keval Vora. 2019. GraphBolt: Dependency-driven synchronous processing of streaming graphs. In *Proc. 14th ACM SIGOPS/EuroSys European Conf. on Comp. Syst.* 1–16.
- [50] Frank McSherry, Andrea Lattuada, Malte Schwarzkopf, and Timothy Roscoe. 2020. Shared Arrangements: Practical Inter-Query Sharing for Streaming Dataflows. *Proc. VLDB Endowment* 13, 10 (2020), 1793–1806.
- [51] Frank McSherry, Derek Gordon Murray, Rebecca Isaacs, and Michael Isard. 2013. Differential Dataflow. In *Proc. 6th Biennial Conf. on Innovative Data Systems Research*.
- [52] Alberto O Mendelzon and Peter T Wood. 1995. Finding regular simple paths in graph databases. *SIAM J. on Comput.* 24, 6 (1995), 1235–1258.
- [53] Vera Zaychik Moffitt and Julia Stoyanovich. 2017. Temporal graph algebra. In *Proc. 16th Int. Symposium on Database Programming Languages*. 1–12.
- [54] Derek G Murray, Frank McSherry, Rebecca Isaacs, Michael Isard, Paul Barham, and Martin Abadi. 2013. Naiad: a timely dataflow system. In *Proc. 24th ACM Symp. on Operating System Principles*. 439–455.
- [55] Hung Q Ngo, Christopher Ré, and Atri Rudra. 2014. Skew strikes back. *ACM SIGMOD Rec.* 42, 4 (Feb 2014), 5–16. <https://doi.org/10.1145/2590989.2590991>
- [56] Milos Nikolic and Dan Olteanu. 2018. Incremental view maintenance with triple lock factorization benefits. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 365–380.
- [57] Anil Pacaci, Angela Bonifati, and M. Tamer Özsu. 2020. Regular Path Query Evaluation on Streaming Graphs. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. ACM, 1415–1430. <https://doi.org/10.1145/3318464.3389733>
- [58] Anil Pacaci, Alice Zhou, Jimmy Lin, and M. Tamer Özsu. 2017. Do We Need Specialized Graph Databases?: Benchmarking Real-Time Social Networking Applications. In *Proc. 5th Int. Workshop on Graph Data Management Experiences and Systems* (Chicago, IL, USA). Article 12, 7 pages. <https://doi.org/10.1145/3078447.3078459>
- [59] Ashwin Paranjape, Austin R Benson, and Jure Leskovec. 2017. Motifs in temporal networks. In *Proc. 10th ACM Int. Conf. Web Search and Data Mining*. 601–610.
- [60] Xiafei Qiu, Wubin Cen, Zhengping Qian, You Peng, Ying Zhang, Xuemin Lin, and Jingren Zhou. 2018. Real-time constrained cycle detection in large dynamic graphs. *Proc. VLDB Endowment* 11, 12 (2018), 1876–1888.
- [61] Juan L Reutter, Miguel Romero, and Moshe Y Vardi. 2017. Regular queries on graph databases. *Theory of Comput. Syst.* 61, 1 (2017), 31–83.
- [62] Siddhartha Sahu, Amine Mhedhbi, Semih Salihoglu, Jimmy Lin, and M. Tamer Özsu. 2020. The Ubiquity of Large Graphs and Surprising Challenges of Graph Processing. *VLDB J.* 29 (2020), 595–618. <https://doi.org/10.1007/s00778-019-00548-x>
- [63] Siddhartha Sahu and Semih Salihoglu. 2020. Graphsurge: Graph Analytics on View Collections Using Differential Computation. *arXiv preprint arXiv:2004.05297* (2020).
- [64] Sherif Sakr, Angela Bonifati, Hannes Voigt, Alexandru Iosup, Khaled Ammar, Renzo Angles, Walid G. Aref, Marcelo Arenas, Maciej Besta, Peter A. Boncz, Khuzaima Daudjee, Emanuele Della Valle, Stefania Dumbra, Olaf Hartig, Bernhard Haslhofer, Tim Hegeman, Jan Hidders, Katja Hose, Adriana Iamnitchi, Vasiliki Kalavri, Hugo Kapp, Wim Martens, M. Tamer Özsu, Eric Peukert, Stefan Plantikow, Mohamed Ragab, Matei Ripeanu, Semih Salihoglu, Christian Schulz, Petra Selmer, Juan F. Sequeda, Joshua Shinavier, Gábor Szárnyas, Riccardo Tommasini, Antonino Tumeo, Alexandru Uta, Ana Lucia Varbanescu, Hsiang-Yun Wu, Nikolay Yakovets, Da Yan, and Eiko Yoneki. 2020. The Future is Big Graphs! A Community View on Graph Processing Systems. *CoRR abs/2012.06171* (2020).
- [65] Semih Salihoglu and Nikolay Yakovets. 2019. Graph query processing. *Encyclopedia of Big Data Technologies* (2019), 890–898.
- [66] Dipanjan Sengupta, Narayanan Sundaram, Xia Zhu, Theodore L Willke, Jeffrey Young, Matthew Wolf, and Karsten Schwan. 2016. Graphin: An online high performance incremental graph processing framework. In *European Conference on Parallel Processing*. Springer, 319–333.
- [67] Feng Sheng, Qiang Cao, Haoran Cai, Jie Yao, and Changsheng Xie. 2018. GraPU: Accelerate streaming graph analysis through preprocessing buffered updates. In *Proc. 9th ACM Symp. on Cloud Computing*. 301–312.
- [68] Ankit Toshniwal, Siddarth Taneja, Amit Shukla, Karthik Ramasamy, Jignesh M Patel, Sanjeev Kulkarni, Jason Jackson, Krishna Gade, Maosong Fu, Jake Donham, et al. 2014. Storm@ twitter. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 147–156.
- [69] Tolga Urhan and M.J. Franklin. 2000. XJoin: A Reactively-Scheduled Pipelined Join Operator. *Q. Bull. IEEE TC on Data Eng.* 23 (2000), 27.
- [70] Oskar van Rest, Sungpack Hong, Jinha Kim, Xuming Meng, and Hassan Chafi. 2016. PGQL: a property graph query language. In *Proc. 4th Int. Workshop on Graph Data Management Experiences and Systems*. 7.
- [71] A.N. Wilschut and P.M.G. Apers. 1991. Dataflow query execution in a parallel main-memory environment. In *Proc. 1st Int. Conf. on Parallel and Distributed Information Systems*. 68–77.
- [72] Peter T Wood. 2012. Query languages for graph databases. *ACM SIGMOD Rec.* 41, 1 (2012), 50–60.
- [73] Nikolay Yakovets, Parke Godfrey, and Jarek Gryz. 2016. Query planning for evaluating SPARQL property paths. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 1875–1889.
- [74] Yuke Yang, Lukasz Golab, and M. Tamer Özsu. 2017. ViewDF: Declarative Incremental View Maintenance for Streaming Data. *Inf. Syst.* 71 (2017), 55–67. <https://doi.org/10.1016/j.is.2017.07.002>