


An Overview of Continuous Querying in (Modern) Data System

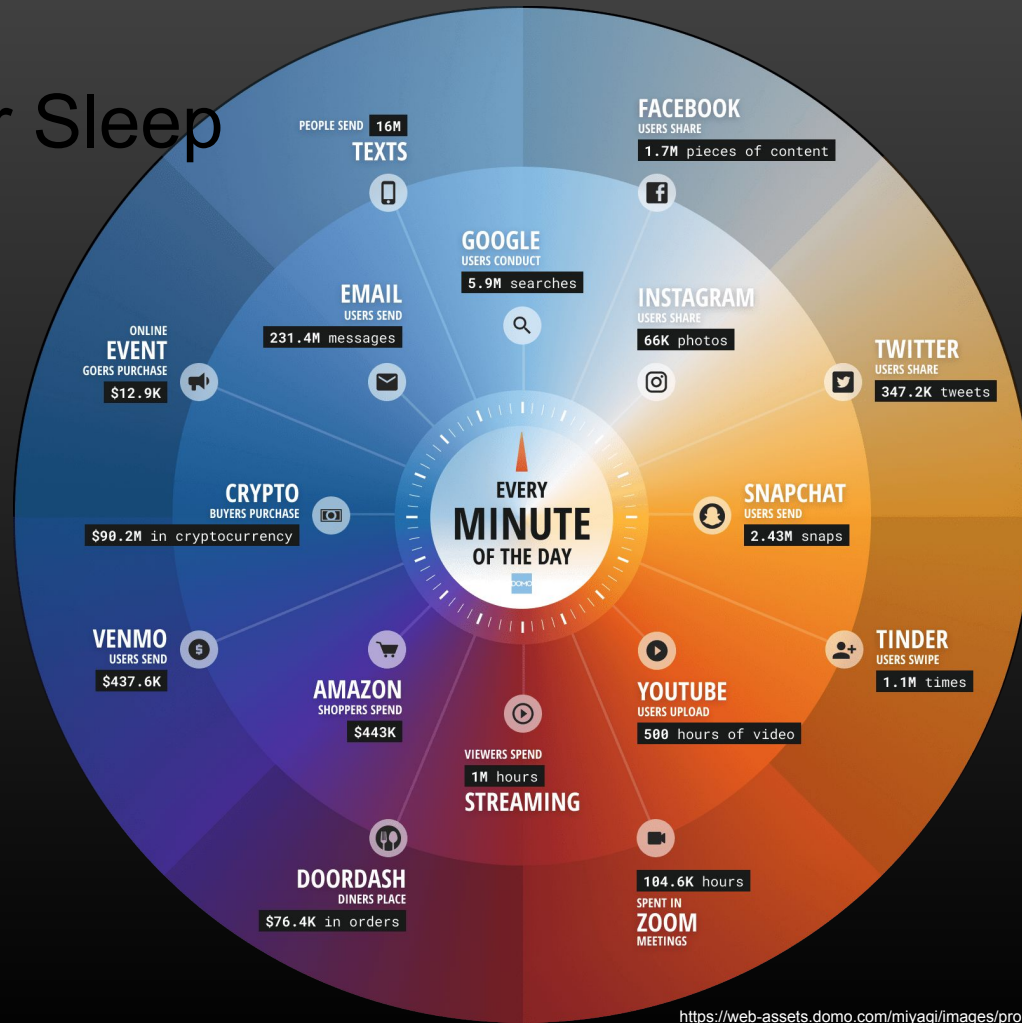
Riccardo Tommasini, INSA Lyon, CNRS Liris (France)
Angela Bonifati, Lyon 1 University, CNRS Liris, IUF (France)






Slides at



Data Never Sleep




Data



The Burden of Unboundedness


...requires a paradigm shift



Unboundedness

The Burden of

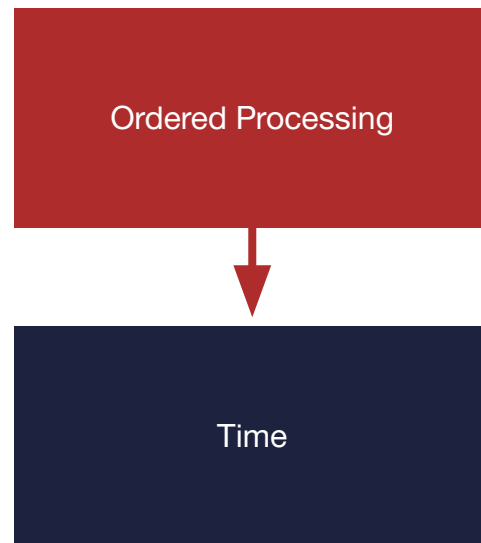
- A program expressed over an **infinite** input may **not terminate**
- Unless we can **reformulate** the notion of “**termination**”
- From an **infinite input** we may observe an **infinite output**
- We still need a way to **determine** what **part** of the **input** maps to the **output**



Unboundedness

The Burden of

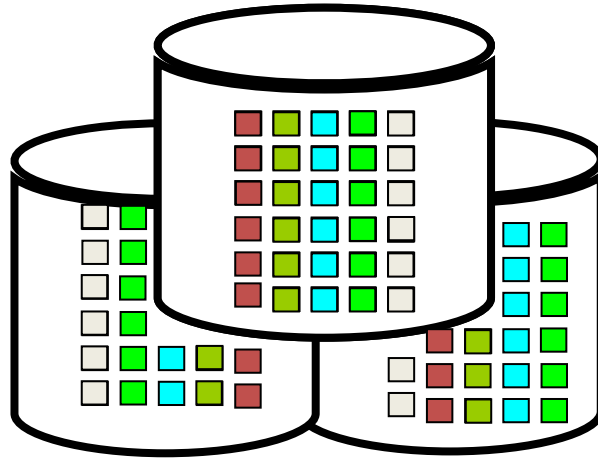


- An **infinite dataset** poses the problem of where to start **computing**
- **Recency*** is a form of **temporality** that enables also **reactivity**
- **Temporality**** may assume other forms
 - About Time (Temporal Data)
 - Through Time (Versioned Data)
 - In-Time (Streaming Data)



*Akidau, Tyler, et al. "The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing." PVLDB (2015).


** Polleres, Axel, et al. "How does knowledge evolve in open knowledge graphs?." *Transactions on Graph Data and Knowledge* 1.1 (2023): 11-1.

Data



What is a Stream?


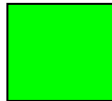
An unbounded partially ordered sequence of data points



What is an Event?

- Event: time-based notification of a known fact defined by
- p a key-value payload
- τ , a type
- t , a timestamp
- d , an optional duration

- payload: 520 - 565 mm
- type: green
- timestamp: t
- duration: 0




time

Continuous Queries on append only databases

how many boxes **green color observations** there are ?

(, ???)


AODB



time


Continuous Queries

on append only databases



Continuous Queries (Monotonic)

on append only databases



Monotonicity is not enough!

Can we achieve more?




Continuous Aggregation

How many **red colored boxes** are in the last minute?

 7


1 minute wide window




time

Top-K

What are the **top-2** most popular colors?




ALL Stream




time

Skyline Continuous Queries

Given the most popular colors in the last minute, what is dominant shade?



1 minute wide window




time

Complex Event Recognition

Is there a **primary cool** colour followed by a **secondary warm** one in the last minute?

yes,  followed by 

1 minute wide window




time


Graph Stream Processing

Are there any **triangles** with warm colours are in the last minute?

Yes,

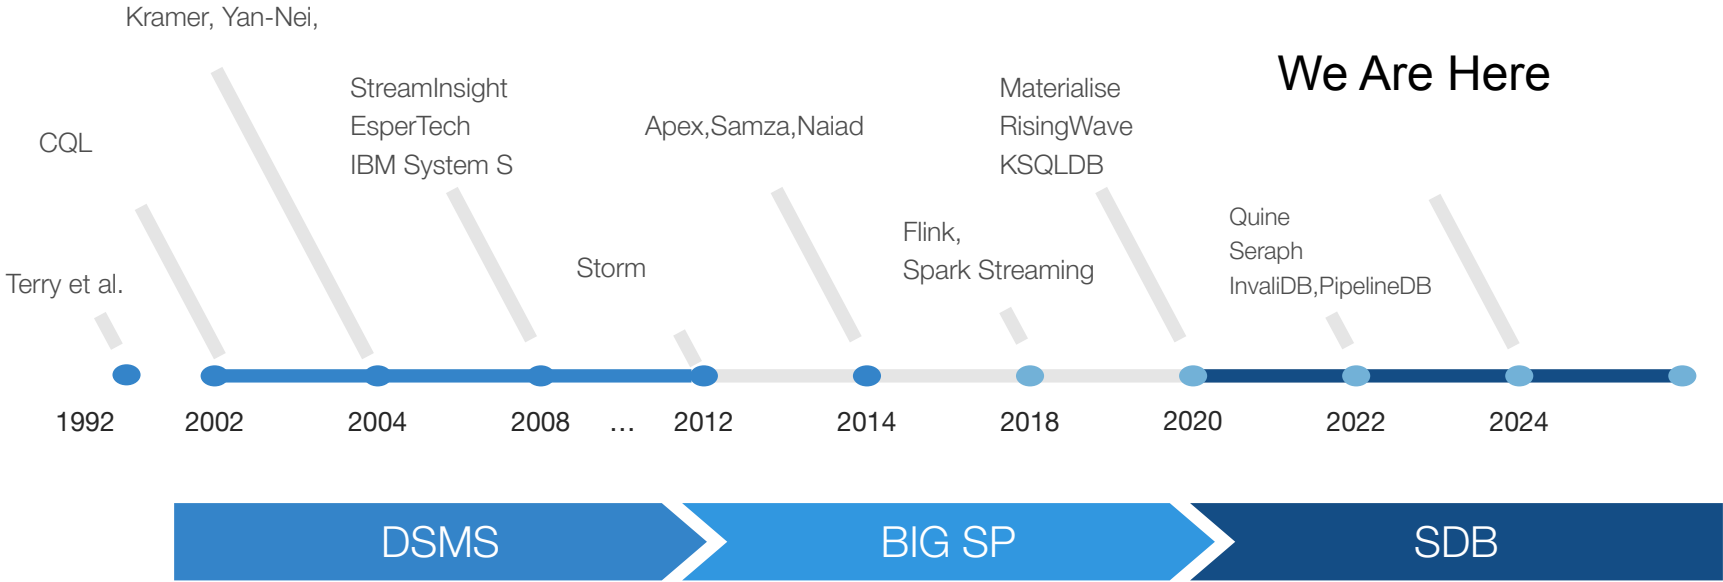


1 minute wide window

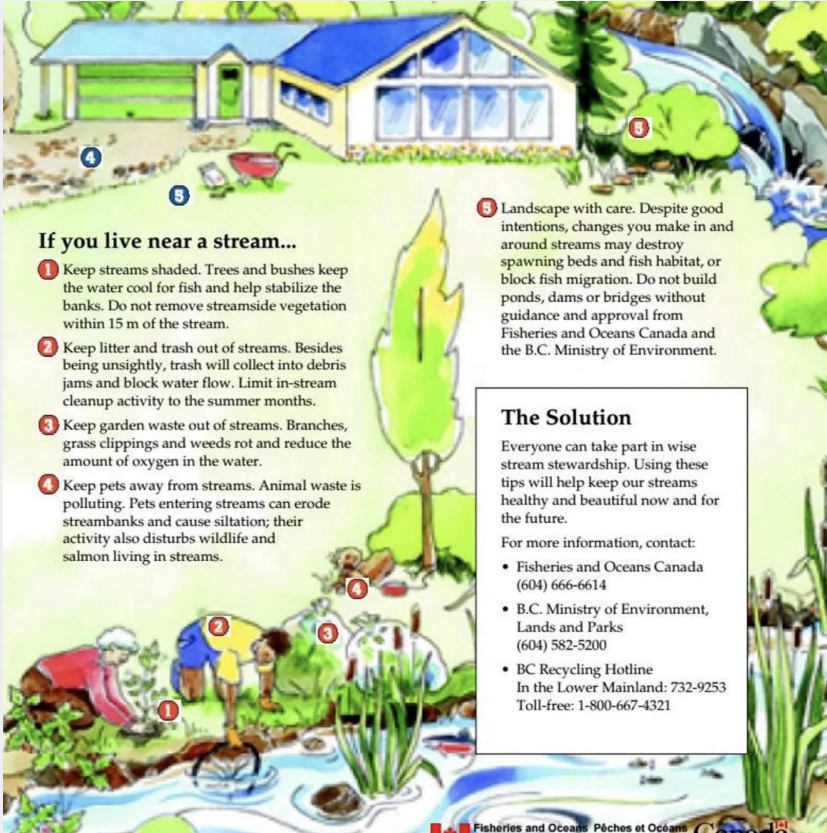


time

Historical Notes



Data Stream Management Systems



If you live near a stream...

- 1 Keep streams shaded. Trees and bushes keep the water cool for fish and help stabilize the banks. Do not remove streamside vegetation within 15 m of the stream.
- 2 Keep litter and trash out of streams. Besides being unsightly, trash will collect into debris jams and block water flow. Limit in-stream cleanup activity to the summer months.
- 3 Keep garden waste out of streams. Branches, grass clippings and weeds rot and reduce the amount of oxygen in the water.
- 4 Keep pets away from streams. Animal waste is polluting. Pets entering streams can erode streambanks and cause siltation; their activity also disturbs wildlife and salmon living in streams.
- 5 Landscape with care. Despite good intentions, changes you make in and around streams may destroy spawning beds and fish habitat, or block fish migration. Do not build ponds, dams or bridges without guidance and approval from Fisheries and Oceans Canada and the B.C. Ministry of Environment.

The Solution


Everyone can take part in wise stream stewardship. Using these tips will help keep our streams healthy and beautiful now and for the future.

For more information, contact:


- Fisheries and Oceans Canada (604) 666-6614
- B.C. Ministry of Environment, Lands and Parks (604) 582-5200
- BC Recycling Hotline
In the Lower Mainland: 732-9253
Toll-free: 1-800-667-4321

Fisheries and Oceans Pêches et Océans

CQ@DBMS




CQ@DBMS

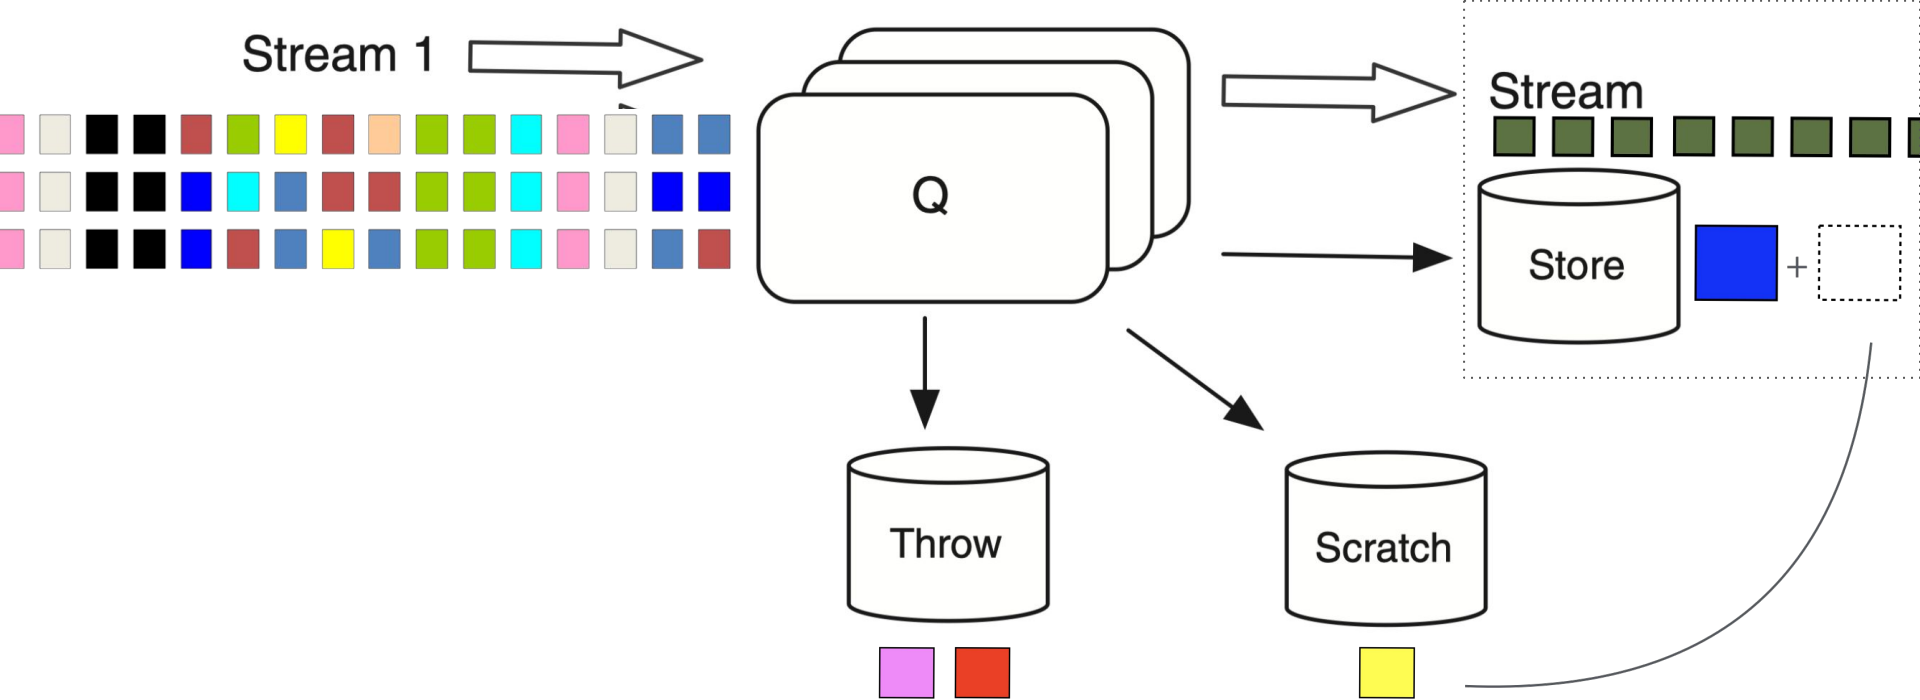


Continuous Queries

on append only databases




CQ@DBMS

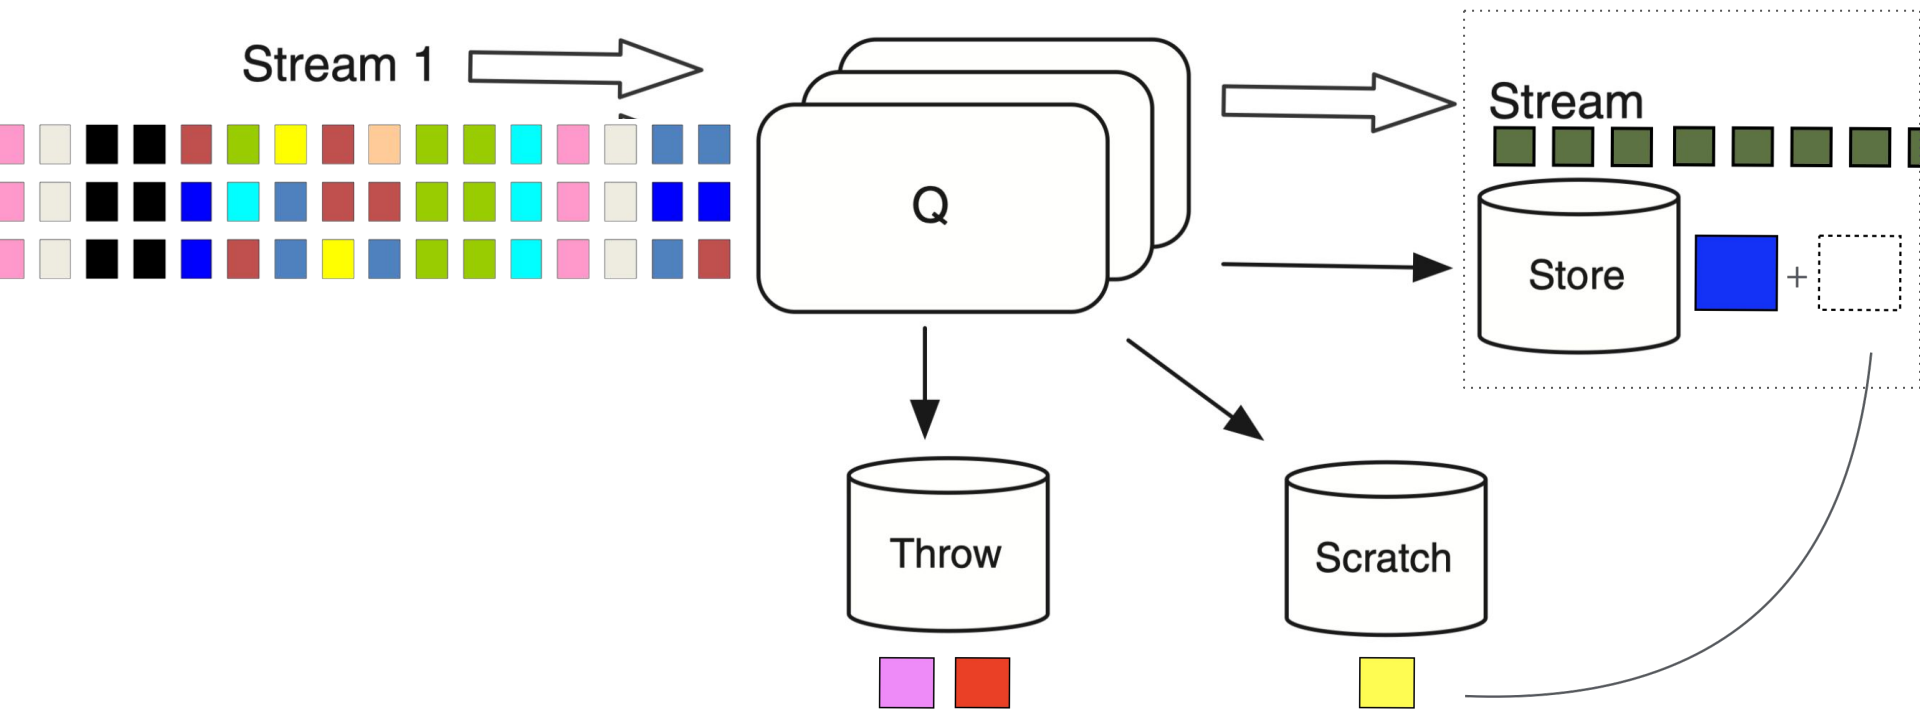


Continuous Queries (Monotonic)

on append only databases



CQ@DBMS



Monotonicity Explained

- **Monotonic** queries produce an **append-only** output stream and therefore do not incur deletions from their answer set.
 - A query is monotonic if for two instances of the database $S1$ and $S2$ such that $S1 \subseteq S2$ then $Q(S1) \subseteq Q(S2)$, where $Q(Si)$ denotes the set of tuples that satisfy Q when applied to the instance Si .
- Only stateless **operators** over infinite streams (projection, selection, time-wise union, and distributive aggregates) can give rise to **monotonic** queries.
- Hence, non **monotonicity** is caused by so called **blocking** operators, i.e., operators that need to see the whole stream to report


Non-monotonicity

Causes of

The concept has taken many names

- stateless/stateful **functions**
- blocking/non blocking **operators**
- event-level/stream-level **semantics**

All notions share the intuition of “memory”, how far do I have to know to answer?



Interval Strategy

- If there are no deletions in the database, the answer to a continuous, non-monotonic query Q can be approximated as $Q = P - N$
- $P(k\tau) - P((k-1)\tau)$ at every interval $[(k-1)\tau, k\tau)$
- In the worst case, this approximation gives a superset of set of data items in the right hand side of the equation below

$$Q(k\tau) - Q((k-1)\tau) = \bigcup_{t_0 \leq s \leq k\tau} Q(s) - \bigcup_{t_0 \leq s \leq (k-1)\tau} Q(s).$$

Fixed-Structure Rewriting

- A rewriting $Q = P - N$ is fixed structure if the following conditions are true:
- P is monotonic
- Interval strategy holds while
 - P holding within k_1 and Inf , N holding within k_2, k_3 , with $k_1 < k_2$

$$P(k\tau) - P((k-1)\tau) = \bigcup_{t_0 \leq s \leq k\tau} Q(s) - \bigcup_{t_0 \leq s \leq (k-1)\tau} Q(s).$$

Admitting Deletions (append-only dbs)

- A query Q is deletion sensitive iff for $S1$ and $S2$ such that $S2 \subseteq S1$, then there exists an item $D1 \in S1$, $D1 \in S2$ such that one of the following is true:
 - $D1 \in Q(S1)(t)$ AND $D1 \notin Q(S2)(t)$
 - $D1 \notin Q(S1)(t)$ AND $D1 \in Q(S2)(t)$

The semantics evolves as

$$P(k\tau) - P((k-1)\tau) = \bigcup_{t_0 \leq s \leq k\tau} Q(s) - \bigcup_{t_0 \leq s \leq (k-1)\tau} Q(s) - D_\tau^k.$$

Fix the
Size of
the
Answer



Fixing the Size of the Answer

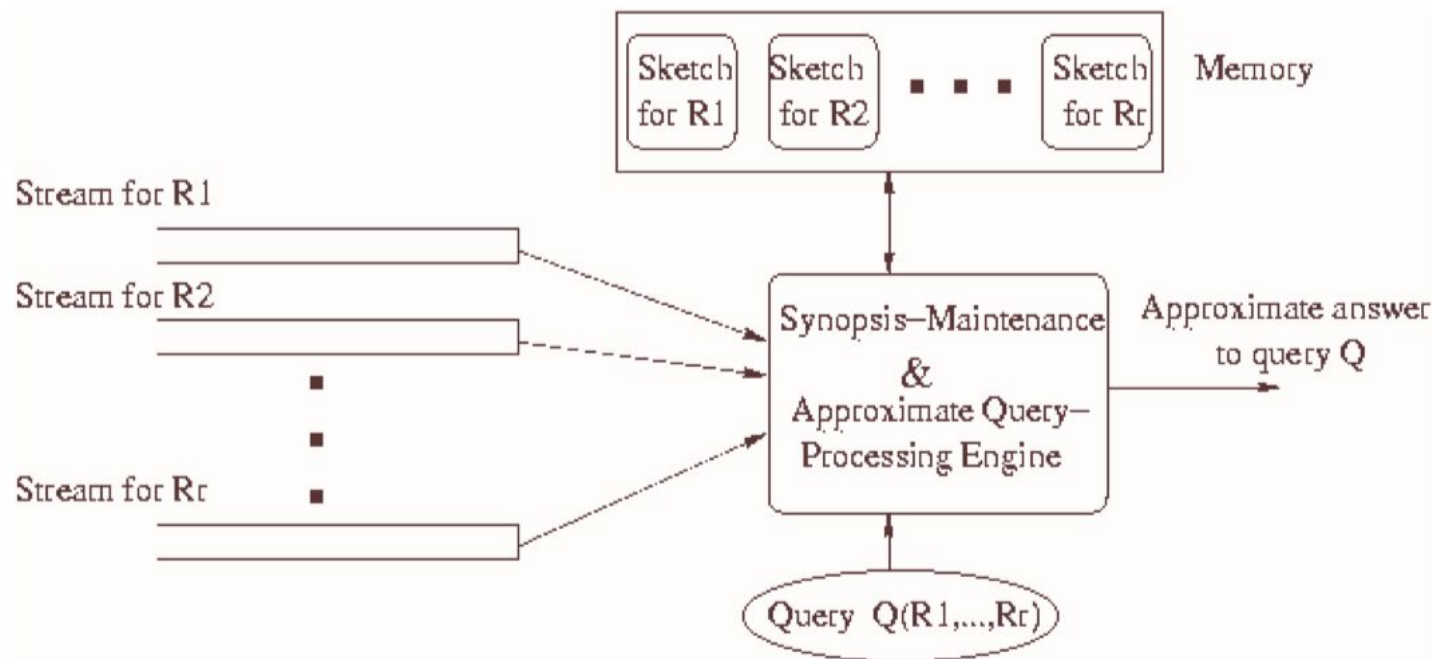
Idea: summarize the characteristics of a stream reducing the memory footprint

A **histograms** summarize a dataset by grouping the data values into buckets and compute for each bucket a set of summary statistics

Wavelet transform the data to represent the most significant features in a frequency domain

Sketches, data structures or algorithms that provide approximate answers to given queries.

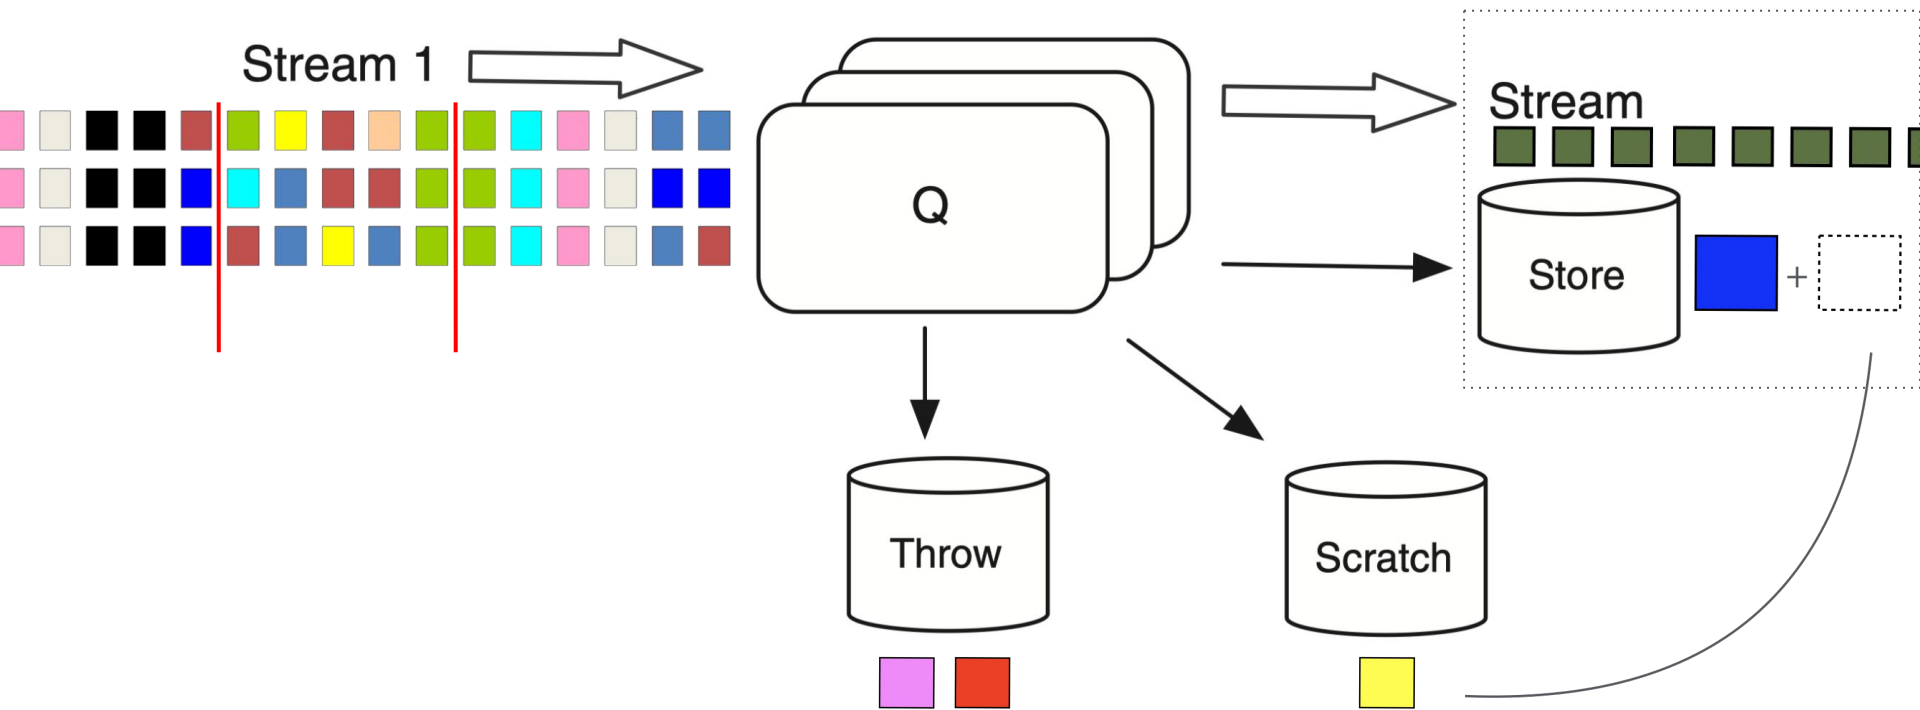
Synopsis-Based DSMS



A man with a beard, wearing a blue long-sleeved shirt and khaki pants, is bent over on a sandy bank of a river. He is holding a clear plastic water bottle and pouring water into the stream. The water is splashing as it hits the surface. The background is a lush, green forest with large trees and a river flowing through it. The scene is lit with warm, golden light, suggesting late afternoon or early morning. The overall mood is peaceful and natural.

Modify the
Input Stream

CQ+Pubctuations@DBMS



Punctuations


- A punctuation is a **predicate** on stream elements that **must** evaluate to **false** for **every element** following the punctuation (Boolean functions).
- In the **original** papers, they are presented as a simple **grammar**
 - [*****,**+**,value,[], range]
 - A tuple matches the punctuation if each of its attributes matches the corresponding pattern
- A punctuated stream is a data stream that contains additional information describing a (possibly empty) subset of data over the domain of the stream

Punctuations Correctness

- A **punctuated stream** S is **grammatical** if for all i , for all $j > i$, if the punctuation $p \in S[i]$ and the tuple $t \in S[i \rightarrow j]$, t does not match p .
- **Safety**: That is, we never emit output unless we can be sure it will not conflict with any later input.
- **Completeness**: we always emit an output if it will necessarily be generated by the relational operator under any additional input, including no input.

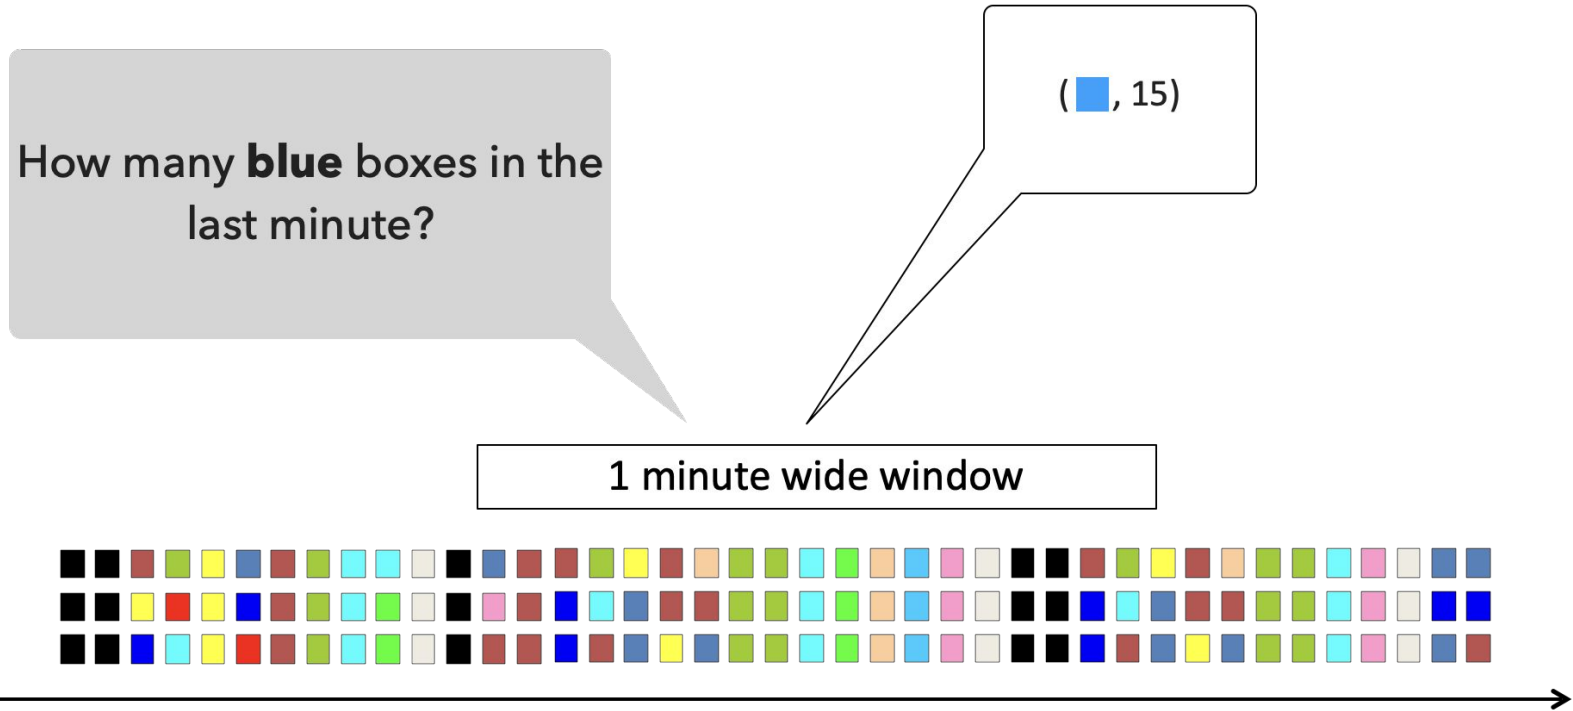
Advantages/Caveats of Punctuations

- Simple to implement
- All **result** data items for a query will eventually be output.
- **Cleanses**. Every data item that resides in the state for any operator in the query **will** eventually be removed.
- **Who is going to emit the punctuation?**
 - Sources?
 - Other operators?
- **Different data models have different punctuation semantics**
- **We still do not know what queries benefits from a given punctuation**

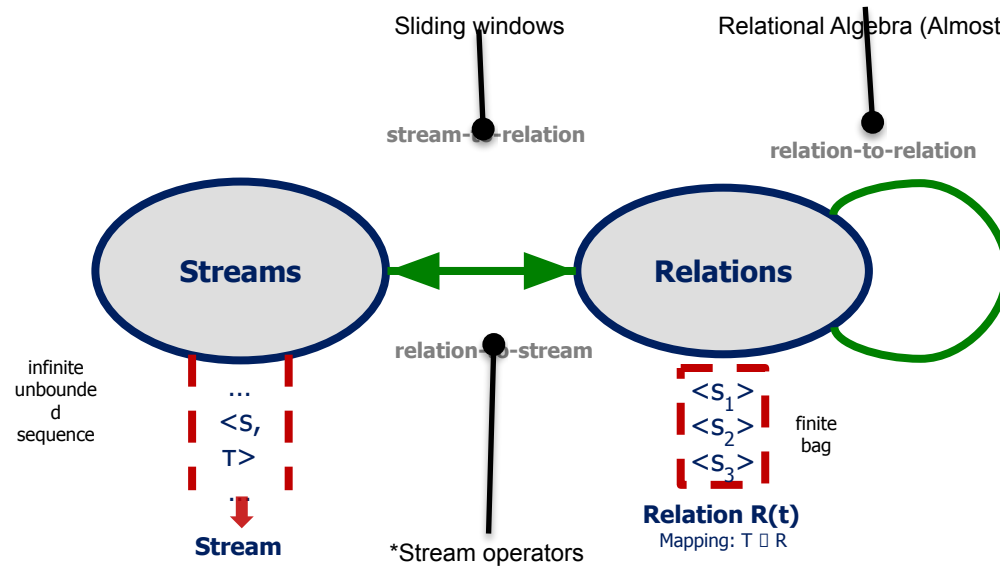
A man with a beard, wearing a dark long-sleeved shirt and pants, is bent over a stream in a lush forest. He is pouring water from a metal bucket into another metal bucket. A wicker basket and another metal bucket are on the ground nearby. The scene is misty and sunlit, with large trees and a rocky stream bed.

Modify the
Output Stream
a.k.a. Modify the
Query

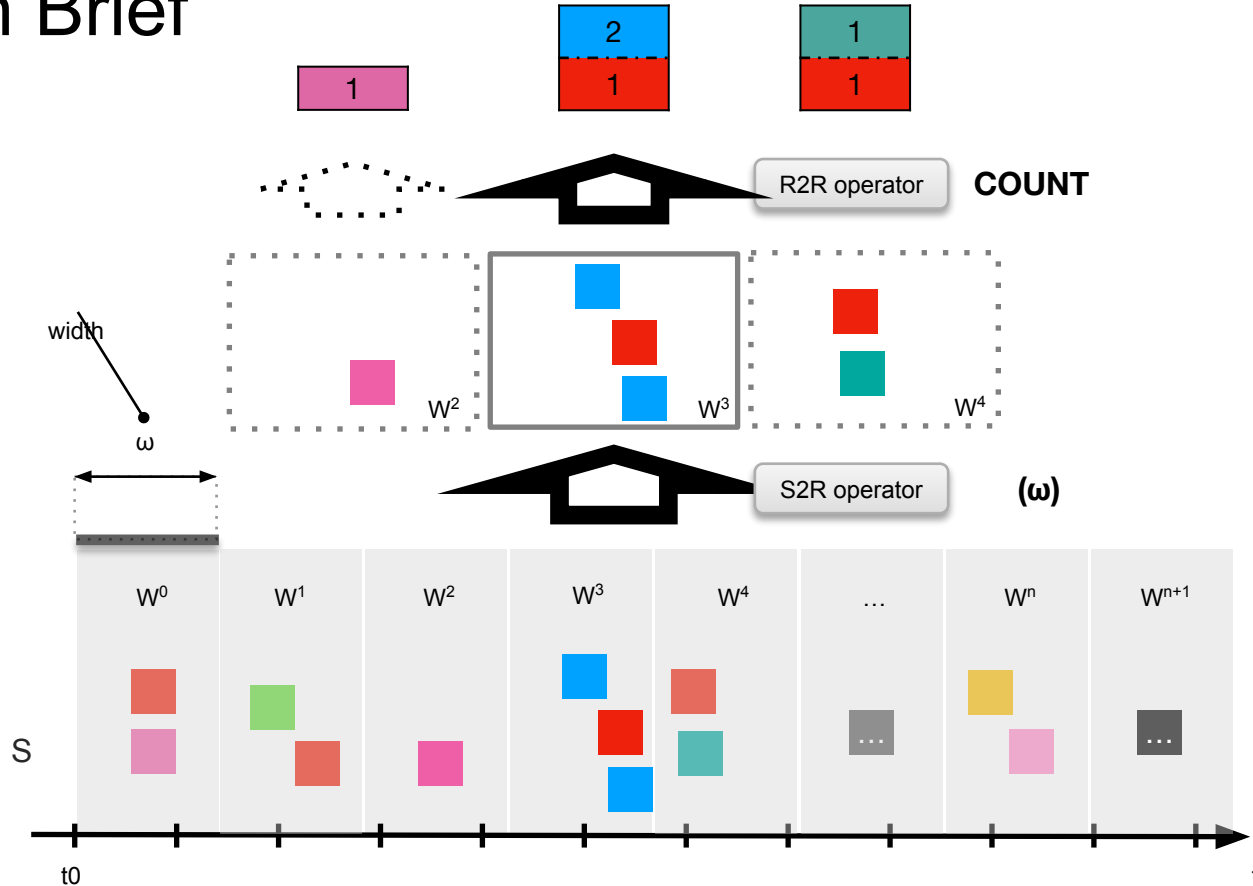
Window-Based Continuous Querying



Continuous Query Language




CQL In Brief



Types of Windows

Window Types	Parameters
Sliding	width
Hopping	Width, slide
Tumbling	Width == slides
Session	Inactivity
...	...



Windows and Monotonicity


- **Weakest non-monotonic** queries do not store state and do not reorder incoming tuples during processing; tuples are either dropped or appended to the output stream immediately.
 - Projection and selection over a single sliding window are weakest non-monotonic
- **Weak non-monotonic** have the property that the expiration time of each result tuple can be determined without generating negative tuples on the output stream.
 - join, duplicate elimination, and groupby.
- **Strict non-monotonic** queries have the property that at least some of their results expire at unpredictable times.
 - Negation over two windows is one example.

Advantages/Caveats of Windows

- SP over sliding Windows are very easy to optimise
- Advantages for parallelised computations
- Enable efficient aggregation and possibly synopsis [paper]
- Users need to know the data semantics
 - Which may be hard given their non-monotonic behaviour
- Out-of-order processing requires sophisticated strategies
- Currently are system-dependent, and harm query portability

Correctness

In Search of



A Temporal Foundation



Data Model

Physical Stream (PS) is infinite sequence of tuples $(e, [ts, te])$ with the same schema. Two elements i, j , are value-equivalent iff $e_i = e_j$


Logical Stream (LS) is a possibly infinite multiset of triples (e, t, n) composed of a record $e \in \Omega$, a point in time $t \in T$, and a multiplicity $n \in \mathbb{N}$.

Physical to Logical (**planning**): For each tuple $(e, [ts, te]) \in LS$, we split the associated time interval into points of time at finest time granularity.

Logical to Physical (**execution**)

- Map each logical stream element (e, t, n) into a physical element $(e, [t, t+1])$
- Coalesce value-equivalent elements that are close to each other (maximal validity)

Snapshot Reducibility



Snapshot Reducibility

For a given logical stream LS and a specified point in time t , the timeslice operation returns a non-temporal multiset of all records in LS that are valid at time instant t

(Snapshot-Reducibility) A **logical** stream operator op_T is **snapshot-reducible** to its non-temporal counterpart op over multisets, if for any point in time $t \in T$ and for all logical input streams LS_1, \dots, LS_n $n \in \mathbb{N}$,

Operators over logical streams

Kramer et al introduce the following operations on logical stream: filter(σ), map (μ), Cartesian product (\times), duplicate elimination (δ), difference ($-$), group (γ), aggregation (α), union (\cup) and window (ω).

Unfortunately, windows, unions, and are again non-window reducible...

Are we condemned to observe this kinds of results...?

The course of querying streams

Several people tries to overcome the issue of “memory”...

Yan-Nei et al showed us that if we use **sequences as our basic data model**, non-monotonic queries become so dominant that only basic project/select operations can be expressed as continuous queries.


The non-blocking subset of relational algebra (NB-RA) and SQL (NB-SQL) are not NB-complete i.e., it cannot express every monotonic set function

A New Hope

User defined Aggregates

A query language that supports non-blocking UDAs and set union can express all monotonic set functions on data streams.

While **UDAs** makes the query language NB-complete they also make it **turing complete** on classical tables



Time and Approximation

However, when moving to timestamped data stream (ordered set), we are losing again the monotonicity of binary operators (.e.g, join).

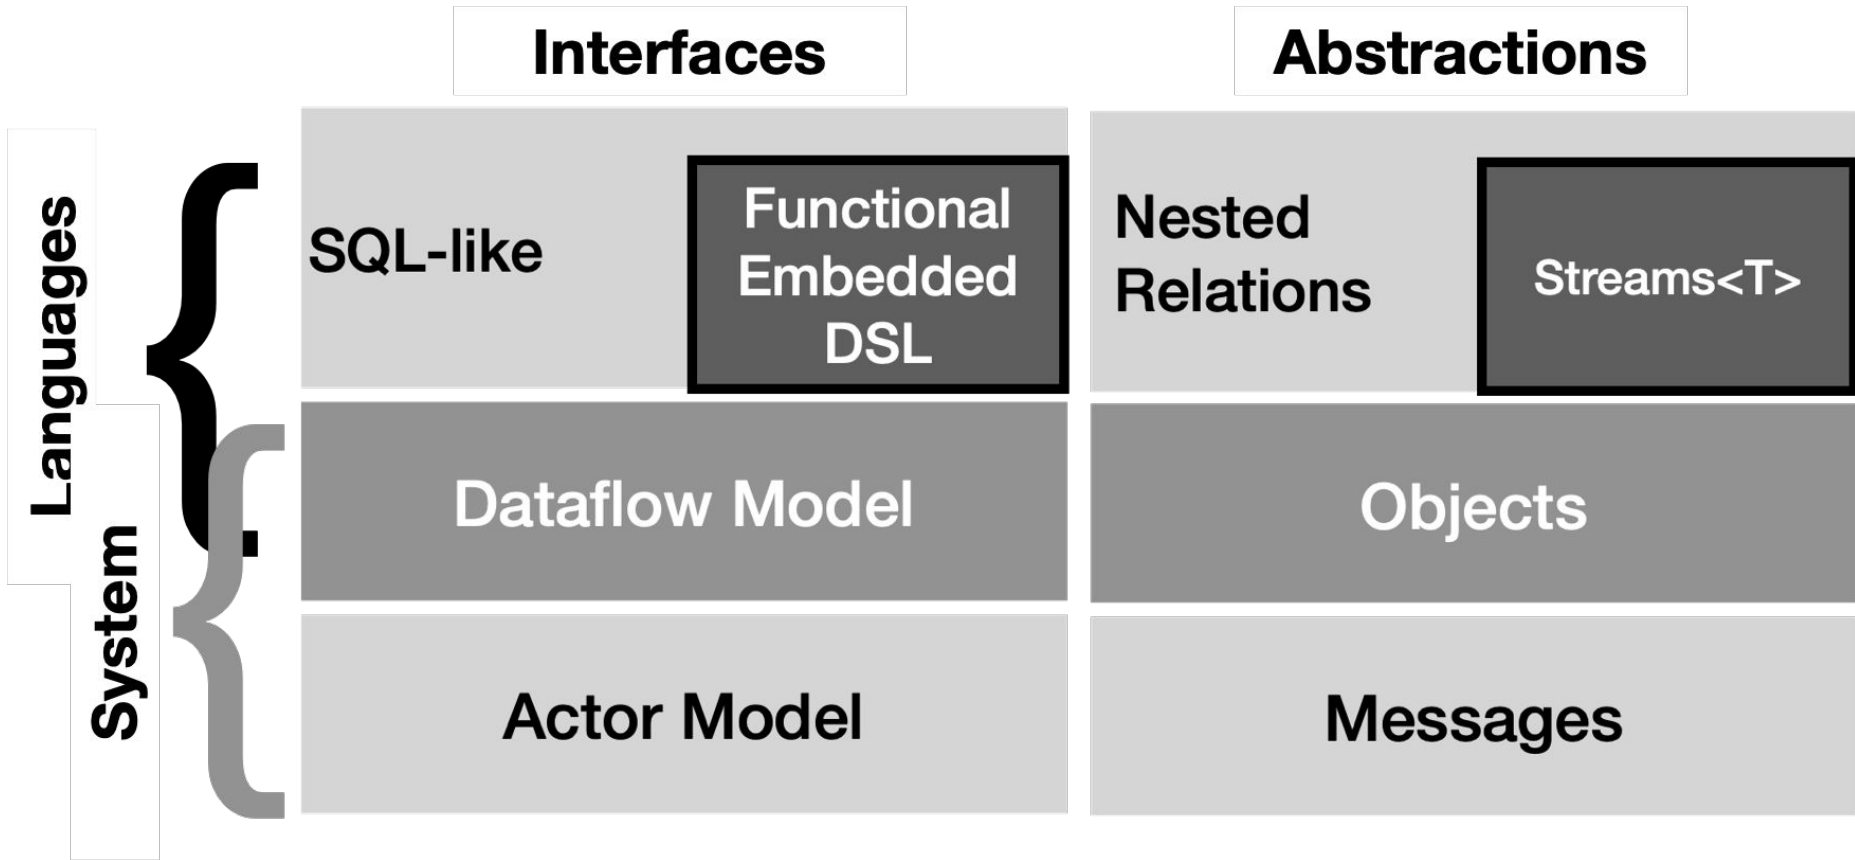
- **Can at least have a monotonic approximation?**
 - Yes, if we consider the recent sub-portion of the data stream: t-operations (union, product, difference) only look at the data up to t
- Focusing on **streams that have no delay** (not true in practice):
 - a query language that has UDA and a version of union that is **t-approximated** is NB-complete.

Streaming Systems

For Big Data

STREAMING & MESSAGING





Continuous Queries



SQL-Like Languages

- New trend is hiding the complexity of the processing behind SQL
- Alternative design debates on how to extend the languages
 - WINDOW Clauses
 - CEP Operations
 - Extended GroupBy
 - Report Controlling

One SQL to Rule Them All: An Efficient and Syntactically Idiomatic Approach to Management of Streams and Tables

An Industrial Paper

Edmon Begoli
Oak Ridge National Laboratory /
Apache Calcite
Oak Ridge, Tennessee, USA
begoli@apache.org

Tyler Akidau
Google Inc. / Apache Beam
Seattle, WA, USA
takidau@apache.org

Fabian Hueske
Veriverica / Apache Flink
Berlin, Germany
fhueske@apache.org

Julian Hyde
Looker Inc. / Apache Calcite
San Francisco, California, USA
jhyde@apache.org

Kathryn Knight
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA
knightke@ornl.gov

Kenneth Knowles
Google Inc. / Apache Beam
Seattle, WA, USA
kenn@apache.org

ABSTRACT

Real-time data analysis and management are increasingly critical for today's businesses. SQL is the de facto *lingua franca* for these endeavors, yet support for robust streaming analysis and management with SQL remains limited. Many approaches restrict semantics to a reduced subset of features and/or require a suite of non-standard constructs. Additionally, use of event timestamps to provide native support for analyzing events according to when they actually occurred is not pervasive, and often comes with important limitations.

We present a three-part proposal for integrating robust streaming into the SQL standard, namely: (1) time-varying relations as a foundation for classical tables as well as streaming data, (2) event time semantics, (3) a limited set of optional keyword extensions to control the materialization of time-varying query results. Motivated and illustrated using examples and lessons learned from implementations in Apache Calcite, Apache Flink, and Apache Beam, we show how with these minimal additions it is possible to utilize the complete suite of standard SQL semantics to perform robust stream processing.

Publication rights licensed to ACM. ACM acknowledges that this contribution is licensed under CC BY 4.0.

CCS CONCEPTS

• Information systems → Stream management; Query languages;

KEYWORDS

stream processing, data management, query processing

ACM Reference Format:


Edmon Begoli, Tyler Akidau, Fabian Hueske, Julian Hyde, Kathryn Knight, and Kenneth Knowles. 2019. One SQL to Rule Them All: An Efficient and Syntactically Idiomatic Approach to Management of Streams and Tables. In *An Industrial Paper. In 2019 International Conference on Management of Data (SIGMOD '19)*, June 30–July 5, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3299869.3314040>

1 INTRODUCTION

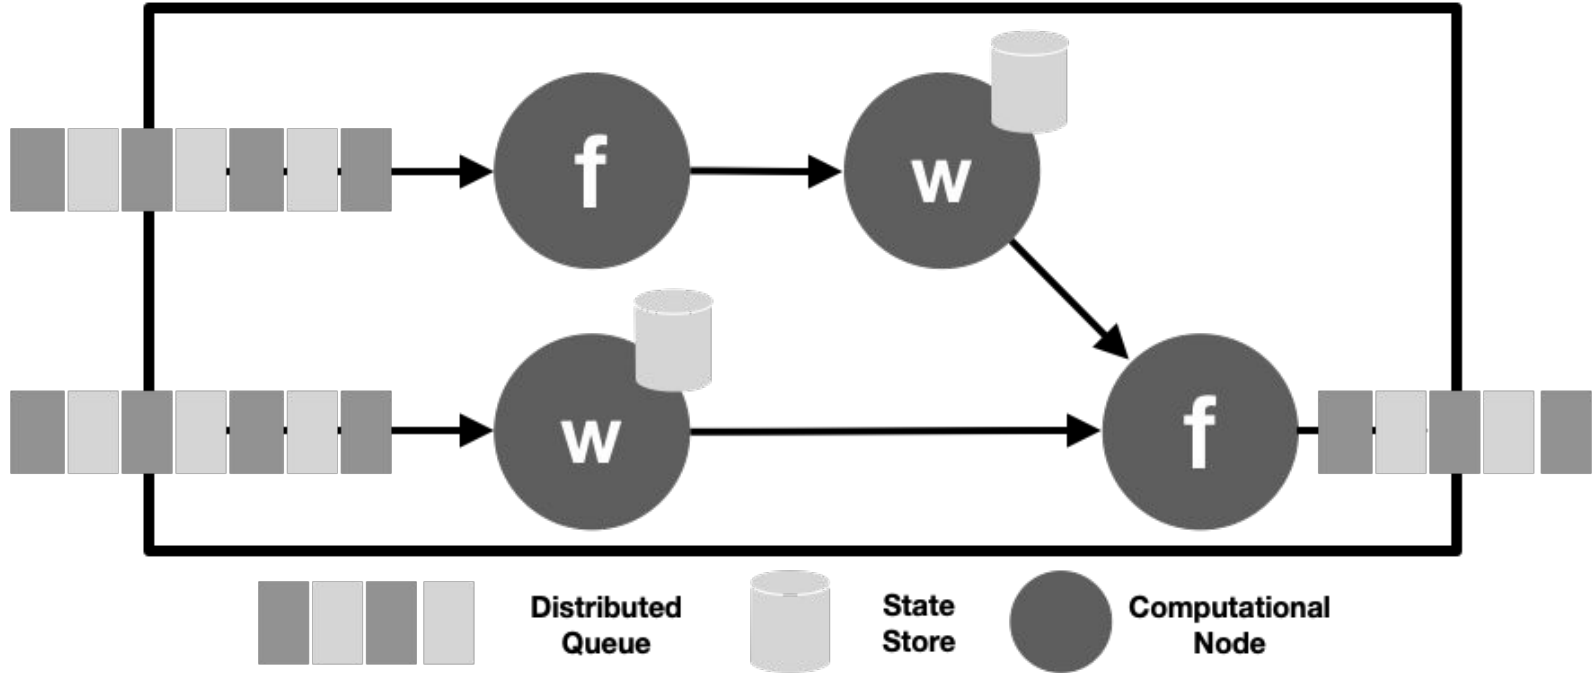
The thesis of this paper, supported by experience developing large open-source frameworks supporting real-world streaming use cases, is that the SQL language and relational model, as-is and with minor non-intrusive extensions, can be very effective for manipulation of streaming data.

Our motivation is two-fold. First, we want to share our ob-

Functional DSL



Operator Topology





Continuous Queries in the Modern Data Landscape

THE 2023 MAD (MACHINE LEARNING, ARTIFICIAL INTELLIGENCE & DATA) LANDSCAPE

This section contains a dense grid of logos for various companies. The logos are organized into several main categories:

- INFRASTRUCTURE:** Includes logos for storage (AWS, Microsoft, Google), MPP DBs (Teradata, Vertica, Exasol), data lakes/warehouses (Dremio, Databricks, Snowflake), data warehouses (Amazon Redshift, Google BigQuery, Microsoft Azure Synapse), streaming/in-memory (Kafka, Flink, Apache Pulsar), BI platforms (Looker, Tableau, Microsoft Power BI), data science notebooks (Databricks, Amazon SageMaker), machine learning/artificial intelligence (IBM, AWS, Google), enterprise ML platforms (Databricks, AWS), sales (Salesforce), marketing (HubSpot, Marketo), customer experience (Salesforce, Adobe), human capital (Gigamonks), automation & operations (UiPath, Automation Anywhere), and decision & optimization (Alteryx, SAS).
- ANALYTICS:** Includes logos for visualization (Tableau, Microsoft Power BI), data science notebooks (Databricks, Amazon SageMaker), machine learning/artificial intelligence (IBM, AWS, Google), enterprise ML platforms (Databricks, AWS), sales (Salesforce), marketing (HubSpot, Marketo), customer experience (Salesforce, Adobe), human capital (Gigamonks), automation & operations (UiPath, Automation Anywhere), and decision & optimization (Alteryx, SAS).
- MACHINE LEARNING & ARTIFICIAL INTELLIGENCE:** Includes logos for machine learning/artificial intelligence (IBM, AWS, Google), enterprise ML platforms (Databricks, AWS), sales (Salesforce), marketing (HubSpot, Marketo), customer experience (Salesforce, Adobe), human capital (Gigamonks), automation & operations (UiPath, Automation Anywhere), and decision & optimization (Alteryx, SAS).
- APPLICATIONS - ENTERPRISE:** Includes logos for customer experience (Salesforce, Adobe), human capital (Gigamonks), automation & operations (UiPath, Automation Anywhere), and decision & optimization (Alteryx, SAS).
- APPLICATIONS - HORIZONTAL:** Includes logos for text (Grammarly, ProtonMail), audio & voice (Zoom, Microsoft Teams), image (Adobe Photoshop, Adobe Lightroom), video editing (Adobe Premiere Pro, Adobe After Effects), animation & 3D (Autodesk Maya, Autodesk Houdini), and search (Elasticsearch, Algolia).
- APPLICATIONS - INDUSTRY:** Includes logos for finance & insurance (Kensho, Lupton), healthcare (Tempus, Flatiron), life sciences (Moderna, AstraZeneca), transportation (Uber, Lyft), agriculture (John Deere, Case IH), industrial & logistics (DHL, FedEx), and govt & intelligence (Palantir, BlackSky).

This section contains a row of logos for various open source infrastructure companies, including:

- Frameworks: Spring, Substratum, etc.
- Format: Avro, Parquet, etc.
- Query/Data Flow: Spark, Flink, etc.
- Data Access: Dremio, etc.
- Databases: PostgreSQL, MySQL, etc.
- OLAP: Databricks, etc.
- Orchestration: Airflow, etc.
- Infra-Structure: Kubernetes, etc.
- Data Ops: Databricks, etc.
- Streaming & Messaging: Kafka, Flink, etc.
- Stats Tools & Languages: Prometheus, etc.
- ML/ops & AI Infra: MLflow, etc.
- AI Frameworks & Libraries: TensorFlow, PyTorch, etc.
- AI Models & Architectures: OpenAI, etc.
- Search: Elasticsearch, etc.
- Logging & Monitoring: Prometheus, etc.
- Visualization: Grafana, etc.
- Collaboration: Slack, etc.

This section contains a row of logos for various data sources, marketplaces, and consulting firms, including:

- DATA SOURCES & APIS:** Includes logos for data marketplaces (Bloomberg, Dow Jones), financial & market data (iStock, Shutterstock), and other data sources (DataCamp, etc.).
- DATA MARKETPLACES & DISCOVERY:** Includes logos for data marketplaces (Bloomberg, Dow Jones), financial & market data (iStock, Shutterstock), and other data sources (DataCamp, etc.).
- DATA & AI CONSULTING:** Includes logos for consulting firms (Deloitte, PwC, EY, etc.).

This section contains a comprehensive grid of 20+ categories of data and AI companies. The categories include:

- INFRASTRUCTURE:** STORAGE (AWS, Microsoft, etc.), MPP DBs (beradata, Vertica, etc.), DATA LAKES / WAREHOUSES (dremio, databricks, etc.), DATA WAREHOUSES (Snowflake, Oracle, etc.), STREAMING / IN-MEMORY (Apache, etc.), GPU DATABASES (kinetica, etc.), VECTOR DATABASES (Pinecone, etc.), REAL TIME DATABASES (Cockroach, etc.), GRAPH DBs (neo4j, etc.), NOSQL DATABASES (MongoDB, etc.), NEWSQL DATABASES (Cockroach, etc.), EL / ETL DATA TRANSFORMATION (dbt, etc.), REVERSE ETL (Census, etc.), DATA INTEGRATION (boomi, etc.), DATA GOVERNANCE & COMPLIANCE (Alteryx, etc.), PRODUCT ANALYTICS (Mixpanel, etc.), LOG ANALYTICS (Splunk, etc.), CRYPTO / WEB3 ANALYTICS (Chainalysis, etc.), AI HARDWARE (Google, etc.), GPU CLOUD (Paperforce, etc.), CLOSED SOURCE MODELS (OpenAI, etc.), FRAMEWORKS (PyTorch, etc.), FORMAT (Parquet, etc.), QUERY / DATA FLOW (Spark, etc.), DATA ACCESS (Snowflake, etc.), DATABASES (Cockroach, etc.), OLAP (DuckDB, etc.), ORCHESTRATION (Airflow, etc.), INFRA-STRUCTURE (Kubernetes, etc.), DATA OPS (Databricks, etc.), STAT TOOL LANGUAGES (R, etc.), DATA MARKETPLACES & DISCOVERY (Bloomberg, etc.), FINANCIAL & MARKET DATA (Bloomberg, etc.), AIR / SPACE / SEA (Orbital, etc.), PEOPLE / ENTITIES (Zoom, etc.), LOCATION INTELLIGENCE (Mapbox, etc.), DATA SOURCES & APIS (Estuary, etc.), APPLICATIONS - ENTERPRISE (Sales, Marketing, etc.), APPLICATIONS - HORIZONTAL (Code & Documentation, Text, etc.), APPLICATIONS - INDUSTRY (Finance & Insurance, Healthcare, etc.), ORCHESTRATION (Data Quality & Observability, Fully Managed, etc.), MGMT / MONITORING (AWS, etc.), PRIVACY & SECURITY (OneTrust, etc.), COMPUTE (AWS, etc.), QUERY ENGINE (Presto, etc.), ENTERPRISE SEARCH (Alibaba, etc.), AI HARDWARE (Google, etc.), GPU CLOUD (Paperforce, etc.), CLOSED SOURCE MODELS (OpenAI, etc.), FRAMEWORKS (PyTorch, etc.), FORMAT (Parquet, etc.), QUERY / DATA FLOW (Spark, etc.), DATA ACCESS (Snowflake, etc.), DATABASES (Cockroach, etc.), OLAP (DuckDB, etc.), ORCHESTRATION (Airflow, etc.), INFRA-STRUCTURE (Kubernetes, etc.), DATA OPS (Databricks, etc.), STAT TOOL LANGUAGES (R, etc.), DATA MARKETPLACES & DISCOVERY (Bloomberg, etc.), FINANCIAL & MARKET DATA (Bloomberg, etc.), AIR / SPACE / SEA (Orbital, etc.), PEOPLE / ENTITIES (Zoom, etc.), LOCATION INTELLIGENCE (Mapbox, etc.), DATA SOURCES & APIS (Estuary, etc.), APPLICATIONS - ENTERPRISE (Sales, Marketing, etc.), APPLICATIONS - HORIZONTAL (Code & Documentation, Text, etc.), APPLICATIONS - INDUSTRY (Finance & Insurance, Healthcare, etc.),

STREAMING & MESSAGING

Companies listed in this category include: Spark Streaming, kafka, beam, PULSAR, Flink, STORM, RabbitMQ, Apache RocketMQ, samza, nifi, and debezium.

Modern Data Landscape

STREAMING / IN-MEMORY

aws Amazon Kinesis | Google Cloud Dataflow | CONFLUENT

databricks | ORACLE Coherence | SAP HANA Cloud

stream | GIGASPACE | GridGain | decodable

HAZELCAST | meroxa | VOLTRON DATA

DELTASTREAM | ESTUARY | Redpanda

tinybird | StreamNative | MotherDuck

ververica | conductor | aiven | EuiX

RisingWave Labs | bytewax | up

arcion | ably

REAL TIME DATABASES

SingleStore

redis

imply

ROCKSET

ClickHouse

Materialize

star-tree

Alinity

Macrometa

ReadySet

STREAMING & MESSAGING

Spark Streaming | kafka | beam | PULSAR

Flink | STORM | RabbitMQ

Apache RocketMQ | samza | nifi | debezium

ESTUARY Flow | Gazette | memphis | bytewax

Modern Landscape

- Real-Time Databases
 - focus on OLAP
 - long time series
- Streaming Databases
 - in-database windowing
 - continuous computation



MATERIALIZED



Incremental View

Maintenance

a long standing problem

- Beyond conjunctive queries, there are studies on IVM for intersection joins, Datalog, Differential Datalog, and DBSP (best paper VLDB 2023)
- The maintenance of complex analytics over evolving databases, which includes linear algebra computation, collection programming, and in-database machine learning.



MATERIALIZED




A network graph visualization with blue nodes and edges on a dark blue background. The nodes are of varying sizes, and the edges are thin lines connecting them. The text "Streaming Graphs" is overlaid on the left side of the image.

Streaming Graphs

The Future is Big Graphs

In 2021 already

- A Community Vision of the role of graph in the future years
- Streaming Graph processing is core
 - Complex Query Execution
 - Incremental Graph Analytics
 - Temporal Dynamic Graphs

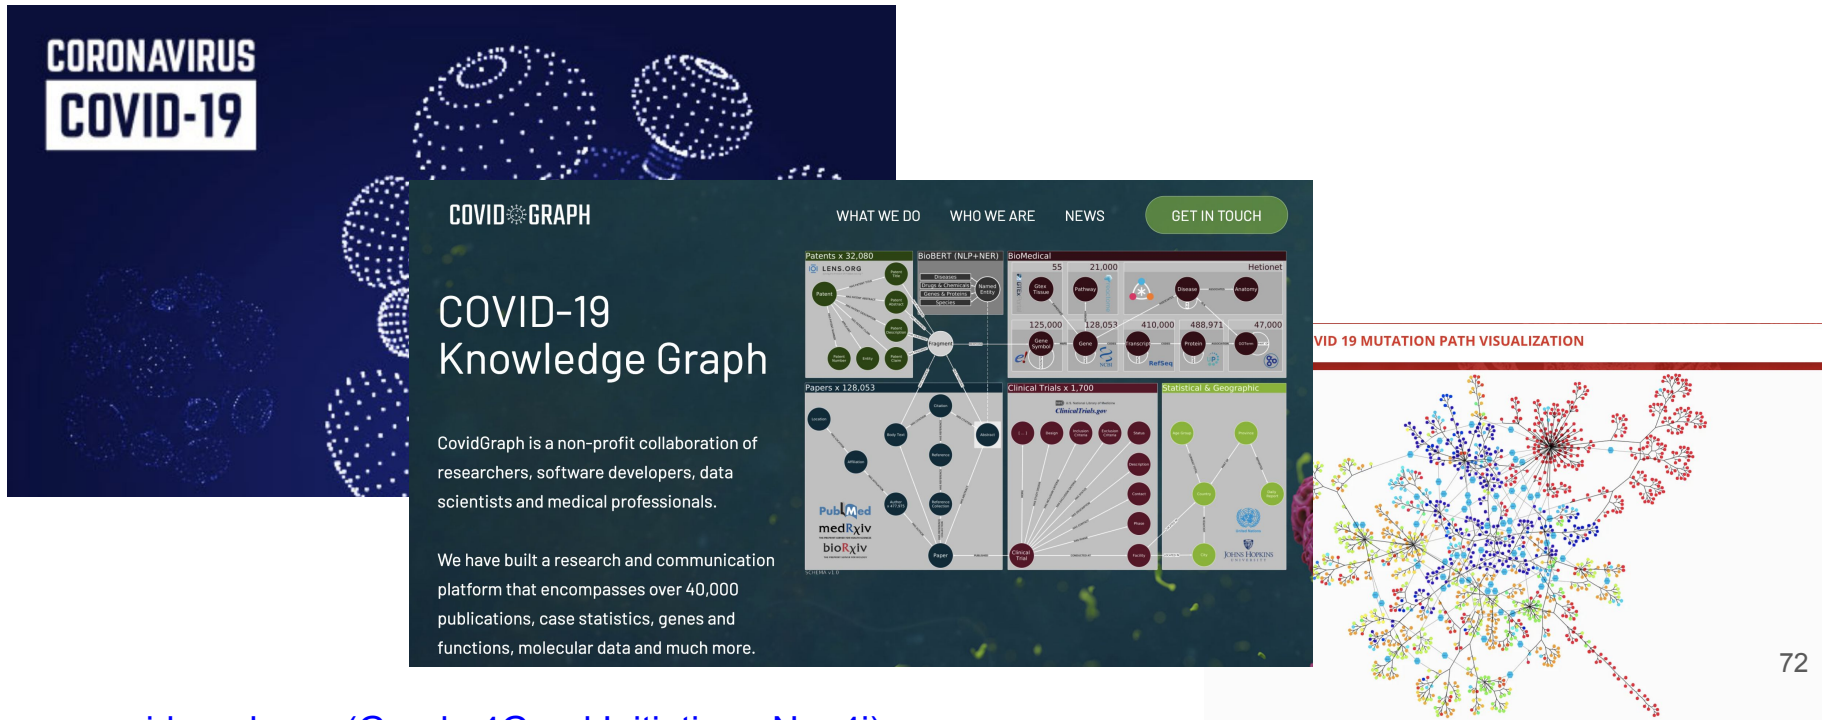


Challenges for Next-generation Graph Processing Systems

- Ch1. A lattice of graph data models and graph algebras
- Ch2. Complex data management ecosystems
- Ch3. Performance and benchmarking

Graphs are ubiquitous across diverse applications

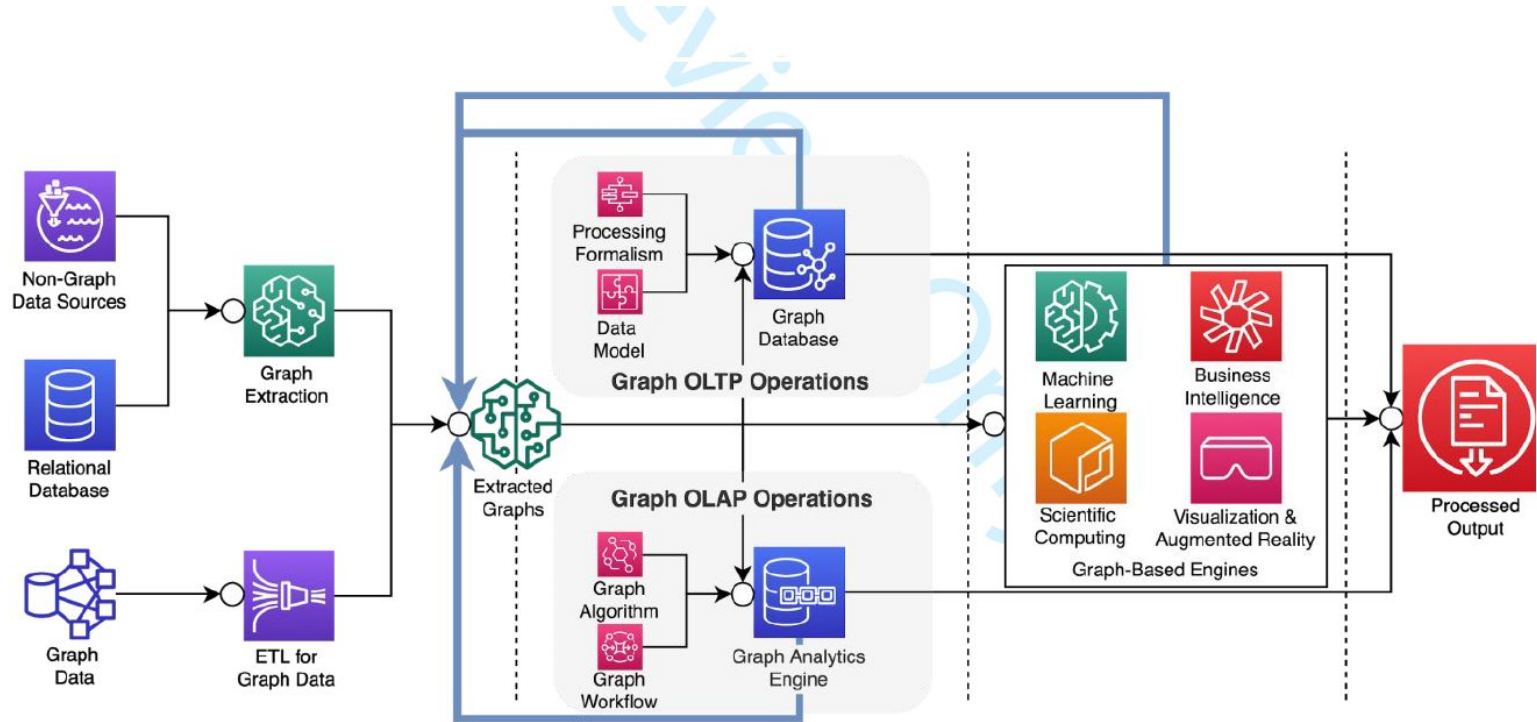
- Several killer applications exist, e.g. financial, logistic, scientific, fraud detection, cybersecurity, supply chain management etc.



Ch1. Expressivity of the graphs/queries

- Dependence on the chosen data model
- How do humans conceptualize graphs?
- The interoperability issues (due to multiple heterogeneous data sources) are to be taken into account
- A data model lattice to navigate across data models, balancing understandability and expressive power
- A new algebra for the variety of graph workloads

Ch2. A complex data management ecosystem



Ch3. Performance and benchmarking

- The need for new, reproducible experimental methodologies to facilitate quick yet meaningful performance-testing?
- How to define more faithful metrics for executing a graph algorithm, query, program, or workflow?
- How to generate workloads with combined operations, covering temporal, spatial, and streaming aspects?
- How to benchmark pipelines including machine learning and simulation?

Dell'Aglio et al

RSP-QL

- Input data: RDF Triples (s,p,o)
- Semantics is based on denotation
 - extends SPARQL with window functions (outside algebraic structure)
 - derived from CQL
- output: time-annotated binding
 - or graphs

```
REGISTER RSTREAM :outStream AS
SELECT ?green
FROM NAMED WINDOW :window ON :colorStream [RANGE PT15S STEP PT5S]
WHERE {
  WINDOW :window {
    ?green a color:Green.
  }
}
```


Streaming Graphs

Towards Stream Graph Processing Systems


- Dynamic graphs are graphs that can accommodate updates (insertions, deletions, changes) and allow querying on the new/old state
- Streaming graphs are graphs that are unbound as new data arrives at high-speed.
- Current systems and libraries (Gelly/Apache Flink) focus on aggregates/projections
- However, more complex query processing operators taking into account recursion, path-oriented semantics etc. need to be investigated
- Graph processing systems are also inherently dynamic and need to respond to all these challenges

Streaming Graphs

Building the underlying graph one edge at a time




- Combines two difficult problems: streaming+graphs
- **Unbounded** \Rightarrow don't see entire graph
- **Streaming rates can be very high**




Streaming graph models

- Window-based semantics (use window to batch edges)
- Continuous semantics (edges are batched as they come)
- Complex vs. Simple operations




Streaming Regular Path Queries (RPQs)

Reflecting the different semantics of graph queries




Simple paths



Arbitrary paths

Towards a Streaming Graph Query Processor

Based on LDBC SNB Interactive Query 7:



G-CORE representation

```
PATH RL = (x)-[:follows]->(y),  
          (x)-[:likes]->(m1)<-[:posts]-(y)  
CONSTRUCT (p1) -[:notify]-> (m)  
MATCH (p1)-/ <~RL+> /->(p2),  
       (p2)-[:posts]->(m)  
ON ldbc_stream WINDOW(24 hours)
```

Datalog program:

```
RL( $u_1, u_2$ )  $\leftarrow$  l( $u_1, m_1$ ), f( $u_1, u_2$ ), p( $u_2, m_1$ )  
Answer( $u, m$ )  $\leftarrow$  RL+( $u, u_2$ ), p( $u_2, m$ )
```

Streaming Graph Algebra

A common foundation for streaming graph query engines

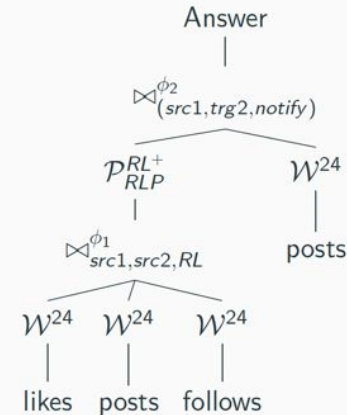
G-CORE Query:

```
PATH RL = (x)-[:follows]->(y),
          (x)-[:likes]->(m1)<-[:posts]-(y)
CONSTRUCT (p1) -[:notify]-> (m)
MATCH (p1)-/ <~RL+> /->(p2),
      (p2)-[:posts]->(m)
ON ldbc_stream WINDOW(24 hours)
```

SGA Expression

$$S_l = \sigma_{I=likes}(\mathcal{W}^{24}(S))$$
$$S_f = \sigma_{I=follows}(\mathcal{W}^{24}(S))$$
$$S_p = \sigma_{I=posts}(\mathcal{W}^{24}(S))$$
$$S_{RecentLiker} = \bowtie_{\phi}^{src1,src3,RecentLiker} (S_{likes}, S_{follows}, S_{posts})$$
$$S_{Related} = \mathcal{P}_{RecentLiker+}^{Notify} (S_{RecentLiker})$$
$$Answer = \bowtie_{\phi}^{src1,trg2,Notify} (S_{Related}, S_p)$$

Logical query plan



Streaming graphs in graph query languages

- Input stream of edges (tuples)
- Snapshot graph is the query focus
- Query models is based on non-recursive Datalog + Kleene Star
- Semantics derived from snapshot reducibility (over graphs)
- Output: a graph (path) detected (as in standard query languages, such as GQL and SQL/PGQ)

G-CORE Query:

```
PATH RL = (x)-[:follows]->(y),  
          (x)-[:likes]->(m1)<-[:posts]-(y)  
CONSTRUCT (p1) -[:notify]-> (m)  
MATCH (p1)-/ <~RL+> /->(p2),  
       (p2)-[:posts]->(m)  
ON ldbc_stream WINDOW(24 hours)
```

How can we continuously process large graph streams as soon as they are discovered?


Can we design a declarative language that enables continuous graph querying?

Research Question

Research Goal

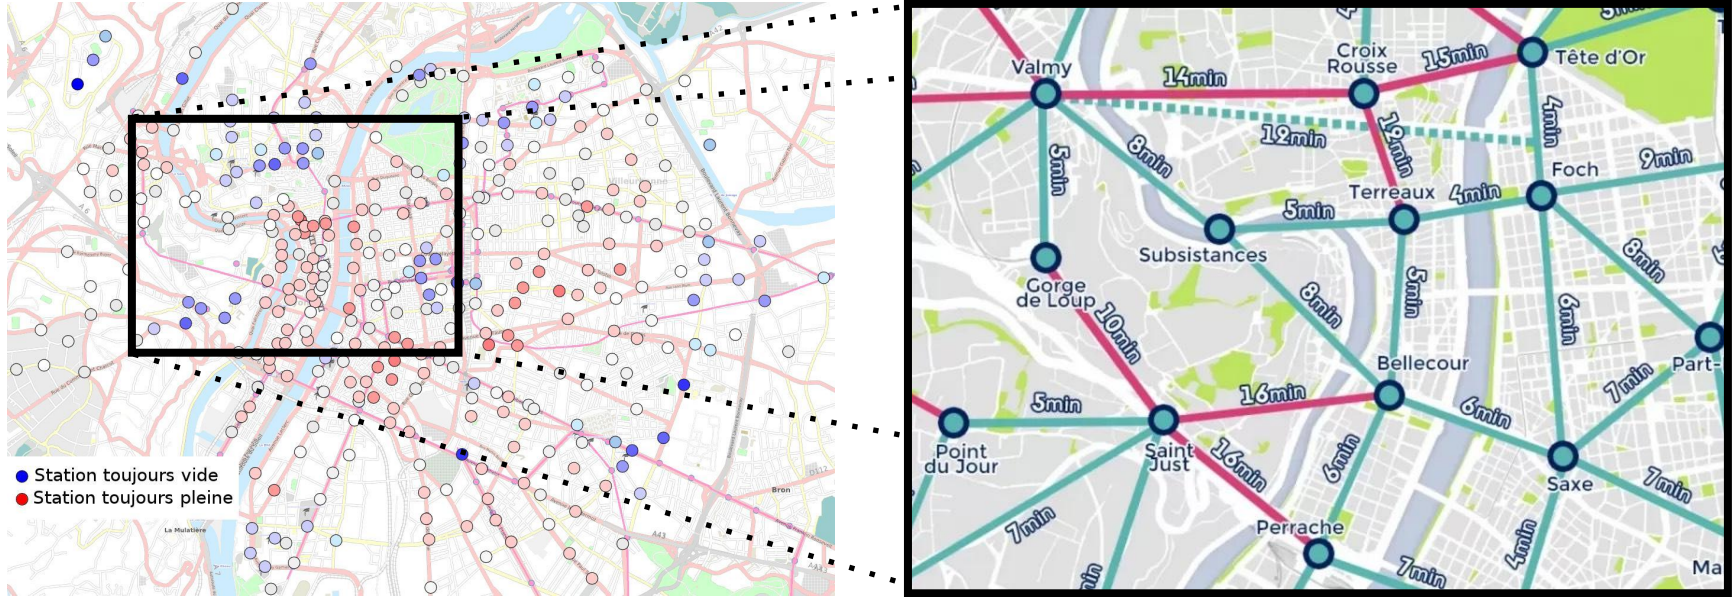
Running Riding Example

Detecting Free Riders in Smart Biking



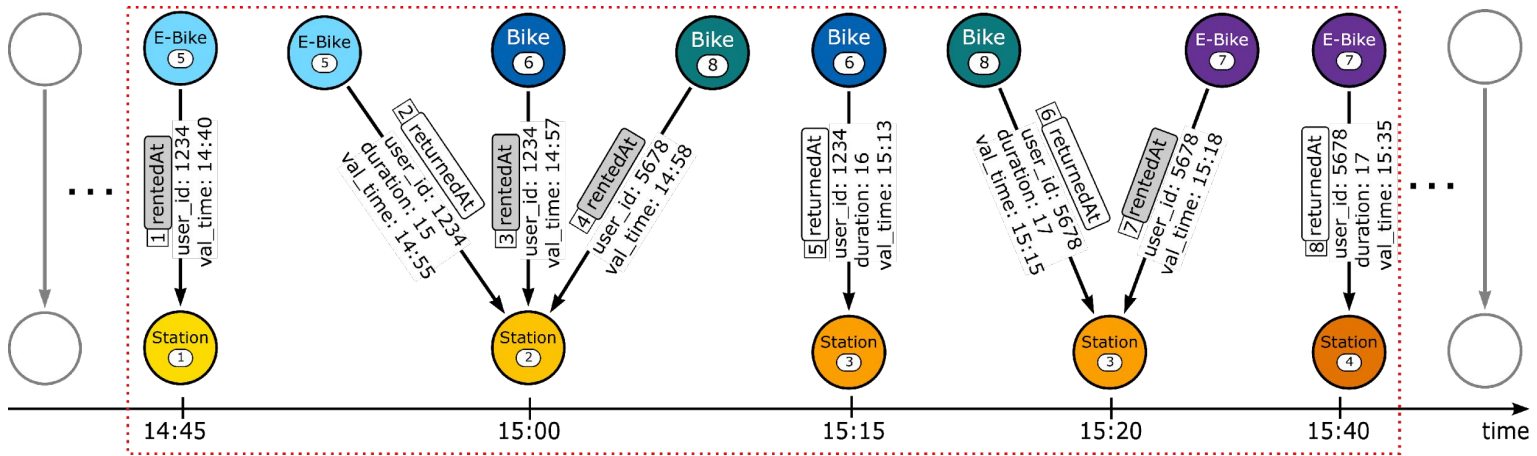
Running Riding Example

Detecting Free Riders in Smart Biking



Running Riding Example

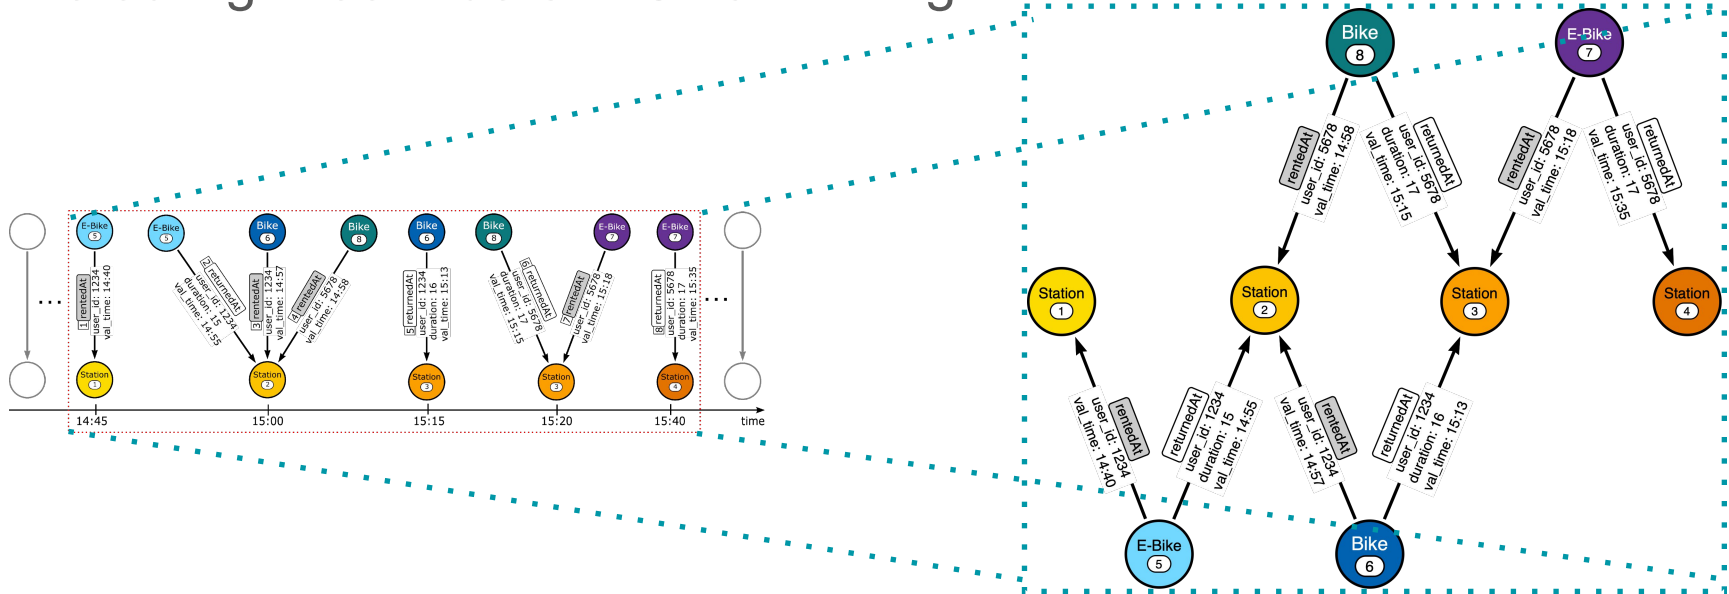
Detecting Free Riders in Smart Biking



What users have used the free period for subsequent rentals in the last hour?

Running Riding Example

Detecting Free Riders in Smart Biking




What users have used the free period for subsequent rentals in the last hour?

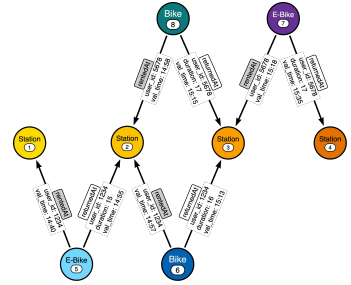
Running Riding Example

Detecting Free Riders in Smart Biking

What users have used the free period for subsequent rentals in the last hour?

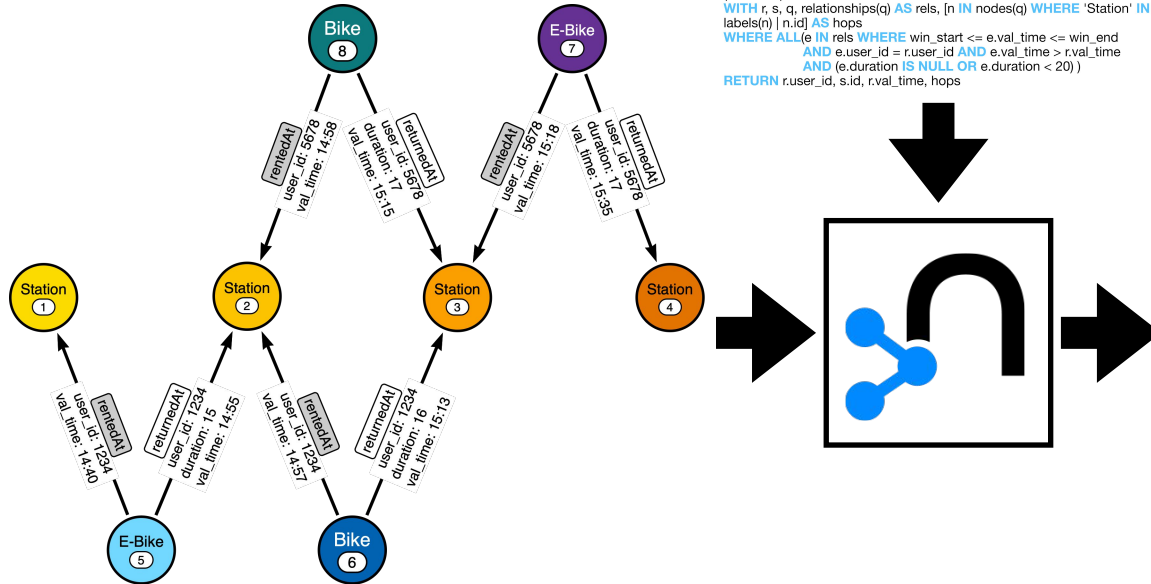


```
WITH datetime() - duration('PT60M') AS win_start, datetime() AS win_end
MATCH (:Bike)-[r:rentedAt]->(s:Station),
        q = (b)-[:returnedAt|rentedAt*3..]->(o:Station)
WITH r, s, q, relationships(q) AS rels,
        [n IN nodes(q) WHERE 'Station' IN labels(n) | n.id] AS hops
WHERE ALL(e IN rels WHERE win_start <= e.val_time <= win_end
            AND e.user_id = r.user_id AND e.val_time > r.val_time
            AND (e.duration IS NULL OR e.duration < 20) )
RETURN r.user_id, s.id, r.val_time, hops
```



Running Riding Example

Detecting Free Riders in Smart Biking



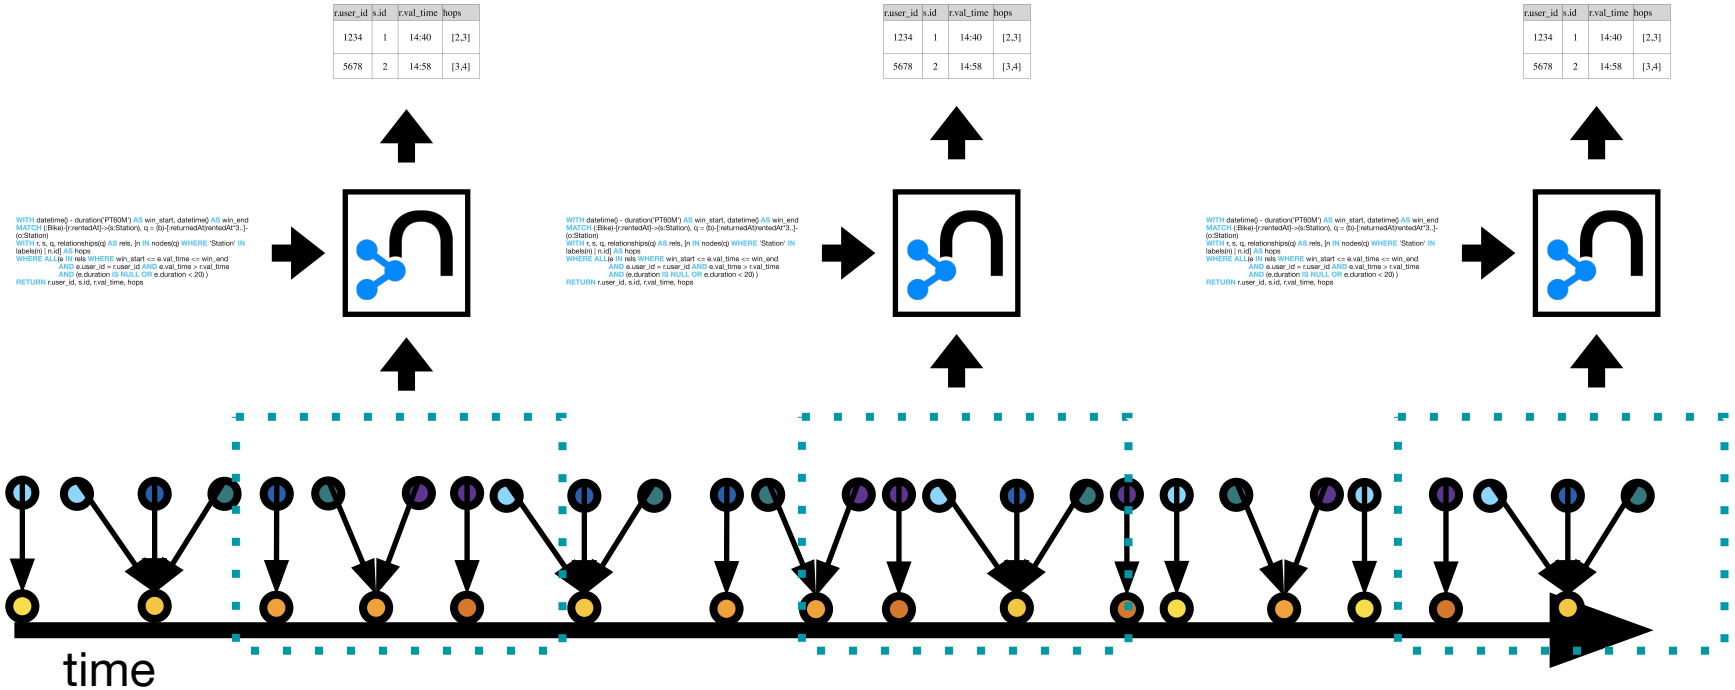
```

WITH datetime() - duration('PT60M') AS win_start, datetime() AS win_end
MATCH (:Bike)-[rentedAt]->(s:Station), q = (b)-[returnedAt][rentedAt*3..]-
(c:Station)
WITH r, s, q, relationships(q) AS rels, [n IN nodes(q) WHERE 'Station' IN
labels(n) | n.id] AS hops
WHERE ALL(e IN rels WHERE win_start <= e.val_time <= win_end
AND e.user_id = r.user_id AND e.val_time > r.val_time
AND (e.duration IS NULL OR e.duration < 20))
RETURN r.user_id, s.id, r.val_time, hops
    
```

r.user_id	s.id	r.val_time	hops
1234	1	14:40	[2,3]
5678	2	14:58	[3,4]

Running Riding Example

Detecting Free Riders in Smart Biking




Using Cypher for Streaming

Pros

- Declarative
- Interactive
- Intuitive
- Standard-ISH*

Cons

- Temporality is reduced to a selection
 - Verbose
 - Unoptimised
- “Now” = User Time
 - Not Reactive
- Results Reporting:
 - Now + Latency



You don't
really know
someone
your data
until you
fight **stream**
them

- **Declarative Semantics.** Seraph allows systems portability and optimisations, as well as adoption.
- **Continuous evaluation.** Seraph's operators allow the repeated evaluation over time, i.e., choosing a time interval and a sequence to evaluate the query.
- **Result emitting.** Seraph's operators allow controlling the report of results, i.e., what is part of the result and when it will be ready to be emitted.
- **Preserving expressiveness.** Seraph preserves openCypher's expressiveness



Seraph's Syntax

Before

```
WITH datetime() - duration('PT60M') AS
win_start, datetime() AS win_end
MATCH (:Bike)-[r:rentedAt]->(s:Station), q =
(b)-[:returnedAt|rentedAt*3..]- (o:Station)
WITH r, s, q, relationships(q) AS rels, [n IN
nodes(q) WHERE 'Station' IN labels(n) |
n.id] AS hops
WHERE ALL(e IN rels WHERE win_start
<= e.val_time <= win_end
AND e.user_id = r.user_id
AND e.val_time > r.val_time
AND (e.duration IS NULL
OR e.duration < 20) )
RETURN r.user_id, s.id, r.val_time, hops
```

After

```
REGISTER QUERY student_trick
STARTING AT 2022-10-14T14:45 {

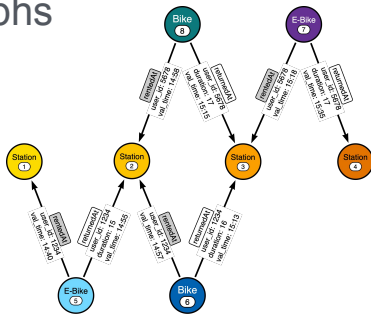
MATCH (:Bike)-[r:rentedAt]->(s:Station),
q = (b)-[:returnedAt|rentedAt*3..]- (o:Station)
WITHIN PT1H
WITH r, s, q, relationships(q) AS rels,
[n IN nodes(q) WHERE 'Station' IN labels(n) |
n.id] AS hops
WHERE ALL(e IN rels WHERE e.user_id =
r.user_id AND e.
val_time > r.val_time AND e.duration < 20)
EMIT r.user_id, s.id, r.val_time, hops
ON ENTERING EVERY PT5M
}
```

Seraph's Data Model

After

Before

- Property Graphs



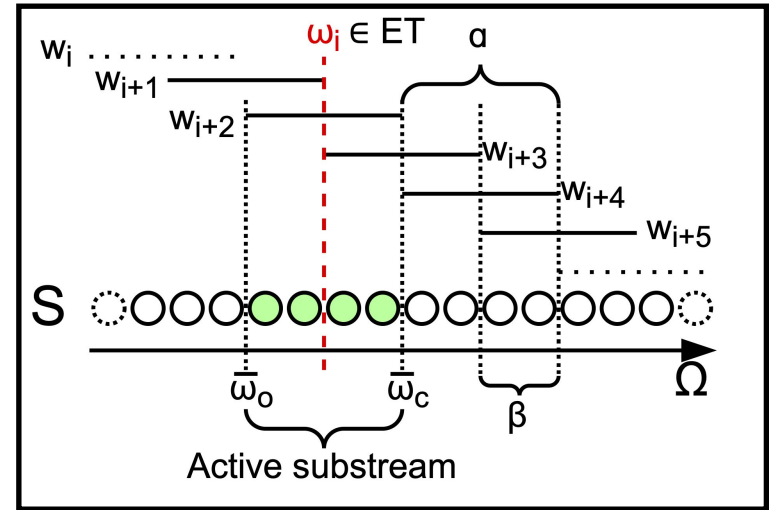
- Tables

r.user_id	s.id	r.val_time	hops
1234	1	14:40	[2,3]
5678	2	14:58	[3,4]

- Property Graph Stream
 - unbounded ordered sequence of pairs (G, ω) where G is a PG and ω a timestamp
- Snapshot Graph
 - Union of all the PGs within a finite sub-portion of a PGStream
- Time-annotated Tables
 - Extend Tables with temporal Bound
- Time-varying Table ($\Psi: \Omega \rightarrow T$)
 - a functional extension of the relational model to incorporate the time semantics


Seraph's Data Model

- Time-varying Table ($\Psi:\Omega\rightarrow T$)
 - Consistency, i.e., Ψ always identifies a time-annotated table
 - Chronologicality, i.e., Ψ always identifies the time-annotated table with the earliest (minimal) opening timestamp.
 - Monotonicity, i.e., Ψ always identifies subsequent time-annotated tables for subsequent time instants.



Seraph's Semantics

$$\begin{aligned}
 \llbracket \text{RETURN } * \rrbracket_{\tilde{G}}(\Psi, \omega) &= \Psi(\omega) \text{ where } \omega \in [\omega_o, \omega_c), \text{ and } \omega_o, \omega_c \text{ are the time annotations of } \Psi(\omega) \\
 \llbracket \text{EMIT } * \rrbracket_{\tilde{G}}(\Psi, \omega) &= \forall \omega_e \in ET \llbracket \text{RETURN } * \rrbracket_{\tilde{G}}(\Psi, \omega_e) \text{ a proposal} \\
 \llbracket \text{EMIT } * \text{ ON ENTERING} \rrbracket_{\tilde{G}}(\Psi, \omega) &= \llbracket \text{EMIT } * \rrbracket_{\tilde{G}}(\Psi, \omega), \text{ where } \Psi = \{\mu \mid \mu \in \Psi(\omega) \setminus \Psi(\omega - 1)\} \\
 \llbracket \text{EMIT } * \text{ ON EXIT} \rrbracket_{\tilde{G}}(\Psi, \omega) &= \llbracket \text{EMIT } * \rrbracket_{\tilde{G}}(\Psi, \omega), \text{ where } \Psi = \{\mu \mid \mu \in \Psi(\omega - 1) \setminus \Psi(\omega)\} \\
 \llbracket \text{EMIT } * \text{ SNAPSHOT} \rrbracket_{\tilde{G}}(\Psi, \omega) &= \llbracket \text{EMIT } * \rrbracket_{\tilde{G}}(\Psi, \omega), \text{ where } \Psi = \{\mu \mid \mu \in \Psi(\omega)\} \\
 \llbracket \text{WITH } * \rrbracket_{\tilde{G}}(\Psi, \omega) &= \Psi(\omega) \text{ if } \Psi(\omega) \text{ has at least one field} \\
 \llbracket \text{WITH ret WHERE expr} \rrbracket_{\tilde{G}}(\Psi, \omega) &= \Psi(\omega) \text{ if } \Psi(\omega) \text{ has at least one field} \\
 \llbracket \text{STARTING AT } \omega_0 \text{ MATCH } \pi \text{ WITHIN } \alpha \text{ EVERY } \beta \rrbracket_S &= \llbracket \text{MATCH } \pi \rrbracket_S^{W(\omega_0, \alpha, \beta)}(\Psi, \omega) \\
 &\hat{=} \llbracket \text{MATCH } \pi \rrbracket_{W(\omega_0, \alpha, \beta)(S)}(\Psi, \omega) \\
 &\hat{=} \llbracket \text{MATCH } \pi \rrbracket_{\tilde{S}_{\omega_o}^{\omega_c}(\omega)}(\Psi, \omega) \\
 &\hat{=} \llbracket \text{MATCH } \pi \rrbracket_{\tilde{G}_{\bar{w}}}(\Psi, \omega) \\
 &= \biguplus_{\mu \in \Psi(\omega)} \{\mu \cdot \mu' \mid \mu' \in \overline{\text{match}(\pi, \tilde{G}, \mu)}\}
 \end{aligned}$$



Graph Streaming

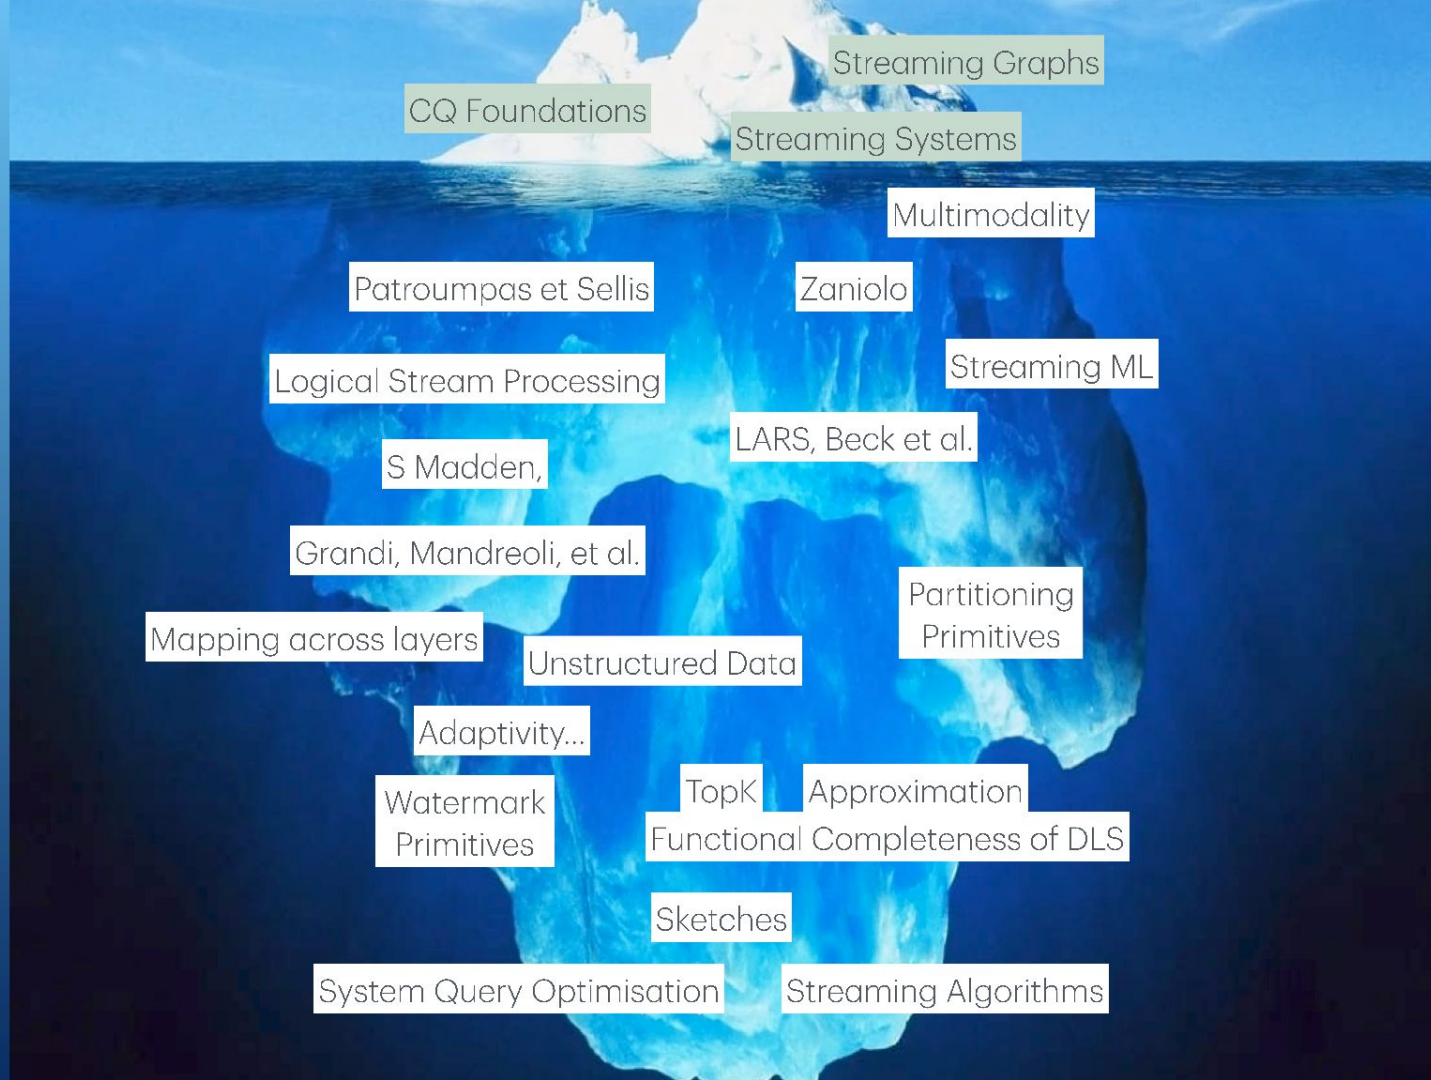
is in its infancy

- Need of considering standard graph query languages
- Need of adapting graph query semantics (trail, shortest path etc.)
- Need to make it efficient
- Related problems:
quality-aware streaming,
fairness-aware streaming



You don't
really know
someone
your graph
until you
fight **stream**
them





CQ Foundations

Streaming Graphs

Streaming Systems

Multimodality

Patrourmpas et Sellis

Zaniolo

Logical Stream Processing

Streaming ML

S Madden,

LARS, Beck et al.

Grandi, Mandreoli, et al.

Partitioning
Primitives

Mapping across layers

Unstructured Data

Adaptivity...

Watermark
Primitives

TopK

Approximation

Functional Completeness of DLS


Sketches

System Query Optimisation

Streaming Algorithms

Linguistic Maturity

- what are the the fundamental abstractions to enable continuous queries ?
- In the first “Stream Processing Era”, several foundational languages have been presented
 - Terry et al, CQL, Kramer et al. but the list continues.
- Are the language design principles shared?




Query Portability

what are the the fundamental abstractions to enable continuous queries ?


Surprisingly(?), it is not that easy to port continuous queries from a system to another

Intermediate representation like River, DBSP, Brooklet, Arc, go in such direction



Data Complexity

- Users are demanding more and more sophisticated view over data
- In the while data may seems simple, but information needs push for challenging this assumption
- How far can we push such challenge? to what data models SP generalise?
- Can we extend into unstructured multimodality?



Observations

Linguistic Maturity

- “windows” aside, what are the the fundamental abstractions to enable continuous queries ?
- There are evidence of industrial adoption and standardization seems possible
- Despite the variety of language design proposal, such languages are still “domain-specific”

Query Portability

- Streaming System internals remain largely custom, hindering query portability
- Can we dissect the modern and established approaches to uniform them?
- Intermediate representation like River, DBSP, Brooklet, Arc, go in such direction

Data Complexity

- Users are demanding more and more complex view over data
- How far can we push such challenge? to what data models SP generalise?
- How far can we push such challenge? to what data models SP generalise?

Q&A

An Overview of Continuous Querying in (Modern) Data System

Riccardo Tommasini, INSA Lyon, CNRS Liris (France)
Angela Bonifati, Lyon 1 University, CNRS Liris, IUF (France)



SIGMOD
PODS
2024

