



# Linked Stream Data Processing Part I: Basic Concepts & Modeling

Danh Le-Phuoc, Josiane X. Parreira, and Manfred Hauswirth  
DERI - National University of Ireland, Galway

Reasoning Web Summer School 2012

© Copyright 2011 Digital Enterprise Research Institute. All rights reserved.



Enabling networked knowledge


## ■ Part I: Basic Concepts & Modeling (Josi)

- Linked Stream Data
- Data models
- Query Languages and Operators
- Choices/Challenges when designing a Linked Stream Data processor

## ■ Part II: Building a Linked Stream Processing Engine (Danh)

- Analysis of available Linked Stream Processing Engines
  - Design choices, implementation
  - Performance comparison
  - Open Challenges

# Streams everywhere



## Biomonitoring

- A. Chemistry chip pendant (saliva sampling)
- B. Skin electrodes (cardiac & respiratory)
- C. Strain gauge & heart-rate monitor
- D. Accelerometer
- E. Phone / camera / GPS
- F. Chemistry ( $pO_2$ ,  $pCO_2$ , sugar)
- G. Local area network
- H. Pulse pressure



# Application Domains



Enterprise Environments



Smart Cities




Telehealth

# Sorry, I can't understand you...



- Heterogeneous data representations
- Lack of semantics
- A priori knowledge of data sources needed
- Disconnected

**Integration Problem!**




## ■ Semantic Web

- Collaborative movement to promote common data formats on the World Wide Web.
- Inclusion of semantic content in web pages
- From unstructured and semi-structured documents to a “Web of data”

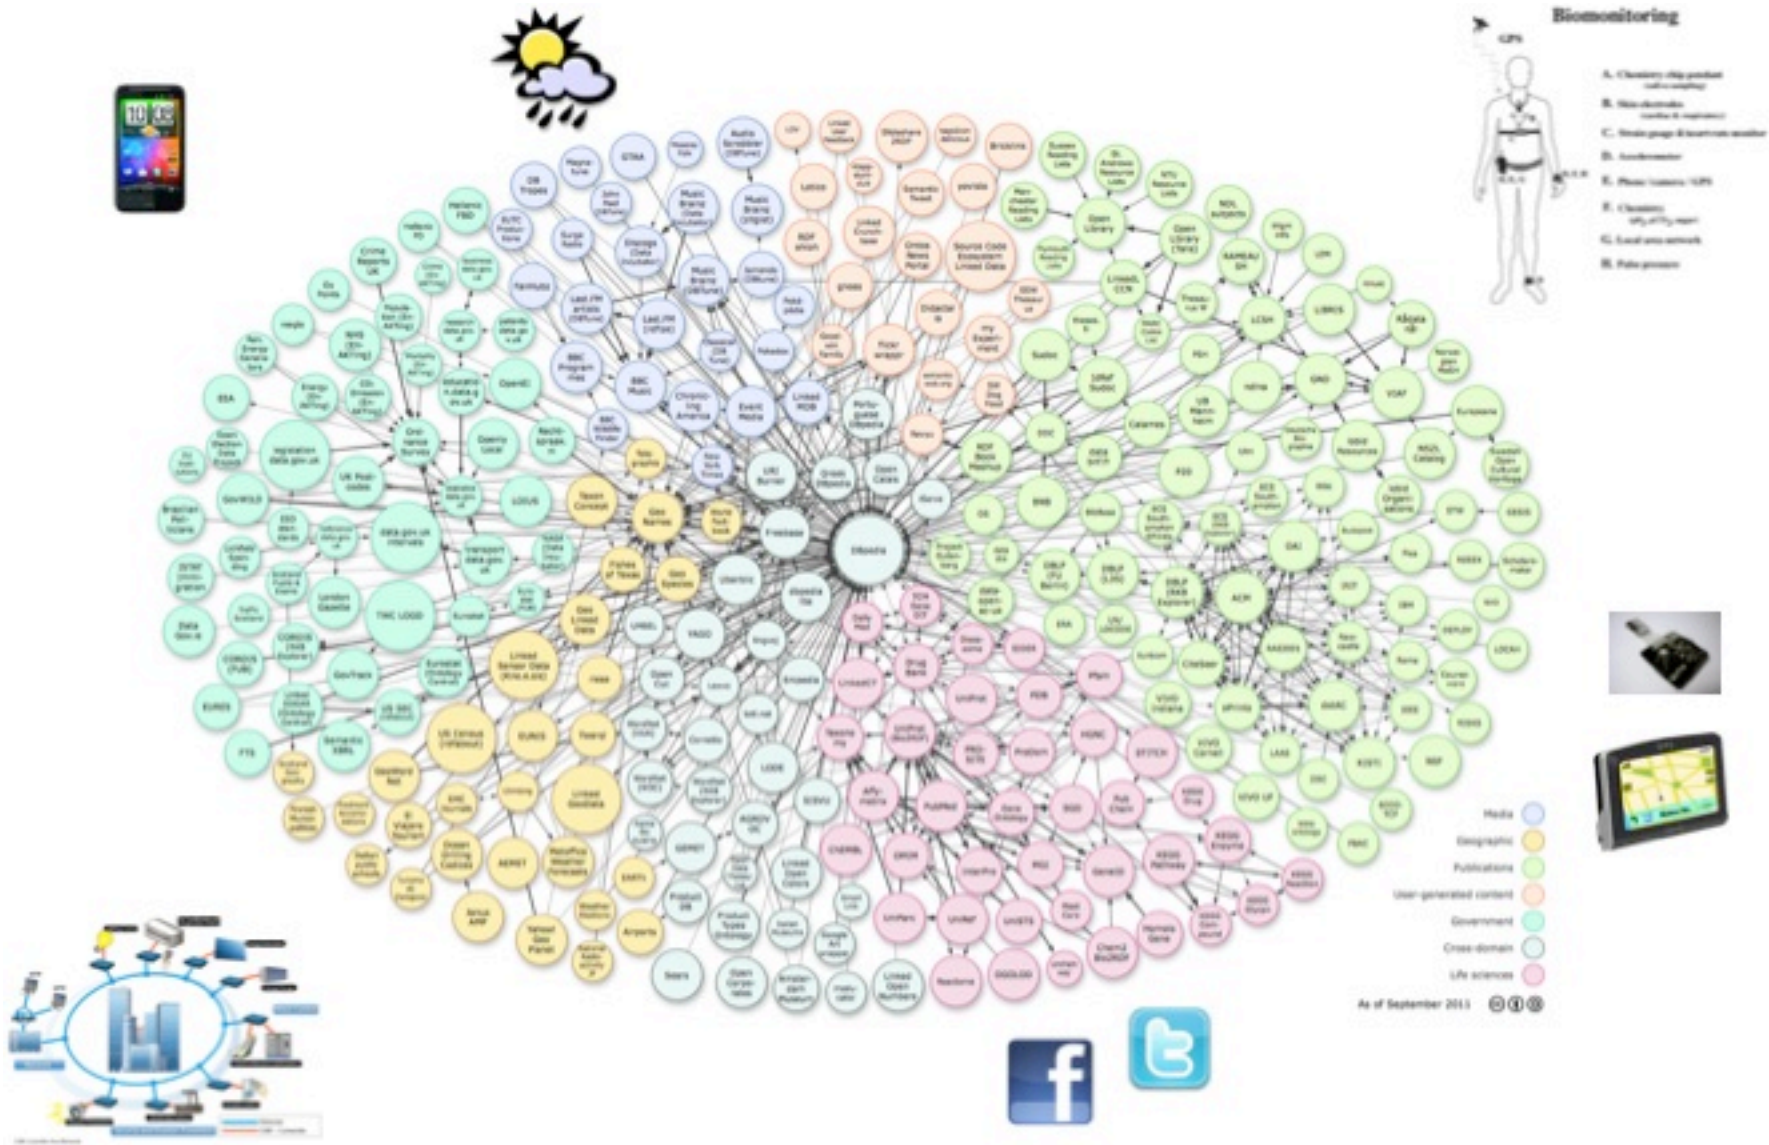
## ■ Linked Data

- Best practices to represent, publish, link data on the Semantic Web
- Linked Data Cloud: collection of datasets that have been published in Linked Data format

# LINKED STREAM DATA




# Linked Stream Data





# Linked Stream Data



- **Semantically enriched stream data**
- **Linked Stream Data examples**
  - W3C Semantic Sensor Network Incubator Group
  - RDF wrappers for Twitter, Facebook, etc
- **Data integration, connects dynamic and static data**
- **Linked Data + DSMS**
  - Stream Data representation/processing different from standard RDF/SPARQL
    - Temporal aspect, continuous query processing
  - DSMS use relational storage model
    - Efficient RDF processing requires heavy replication

# Running example



**Danh Le Phuoc**

List of publications from the DBLP Bibliographic Service - FIG

Pub views: ACM DL, CiteSpace, CSD - MetaPress - Google - Bing - Yahoo

Year	Journal/Conference	Title
2011	VLDB	Wuqiang Ai, Matthias Mann, Stefan Edelkamp, Danh Le Phuoc, Felix Jan Martin: A Peer-Based Approach to Privacy Preserving Content Management. <i>VLDB</i> 2011: 18-29
2011	VLDB	Danh Le Phuoc, Minh Quan Tran, Antonio Garcia-Soler, Manfred Spanthaler: A Native and Adaptive Approach for Conflict Processing of Linked Sources and Linked Data. <i>International Semantic Web Conference</i> 2011: 570-589
2011	VLDB	Danh Le Phuoc, Antonio Garcia-Soler, Michael Schweizer, Tommaso Di Noia, Manfred Spanthaler: Live-Infused open source workflow. <i>ISDEA/STOCK</i> 2011
2011	VLDB	Danh Le Phuoc, Antonio Garcia-Soler, Yannis Kostas, Manfred Spanthaler: RDF On the Go: RDF Storage and Query Processor for Mobile Devices. <i>ISWC/PersonalSemantics</i> 2011
2011	VLDB	Danh Le Phuoc, Axel Polleres, Manfred Spanthaler, Giovanni Tummarello, Christian Mathias: Rapid prototyping of semantic mash-up-based services. <i>Web pages</i> WWW 2011: 741-746
2011	VLDB	Christian Mathias, Danh Le Phuoc, Axel Polleres, Manfred Spanthaler, Giovanni Tummarello: Processing Semantic Web Types. <i>ISWC</i> 2011: 601-616
2011	VLDB	Liu, Wu, Zhang, Ding, Manfred Spanthaler, Danh Le Phuoc, Hong Ya, Liu, Zhang: Transferring semantic data to wireless mobile devices using semantic mash-up. <i>ISWC</i> 2011: 24-32

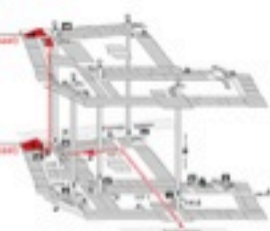
**Lageplan: HS B (Heinz Parkus HS), TU Wien,**  
in 2340 Wien, Karlupf 13, Hauptgebäude, Erdgeschoss und 1. Obergeschoss, Stock VI3

**Übersichtplan: TU Wien**




**Detailplan: Hauptgebäude**

**1. Obergeschoss**

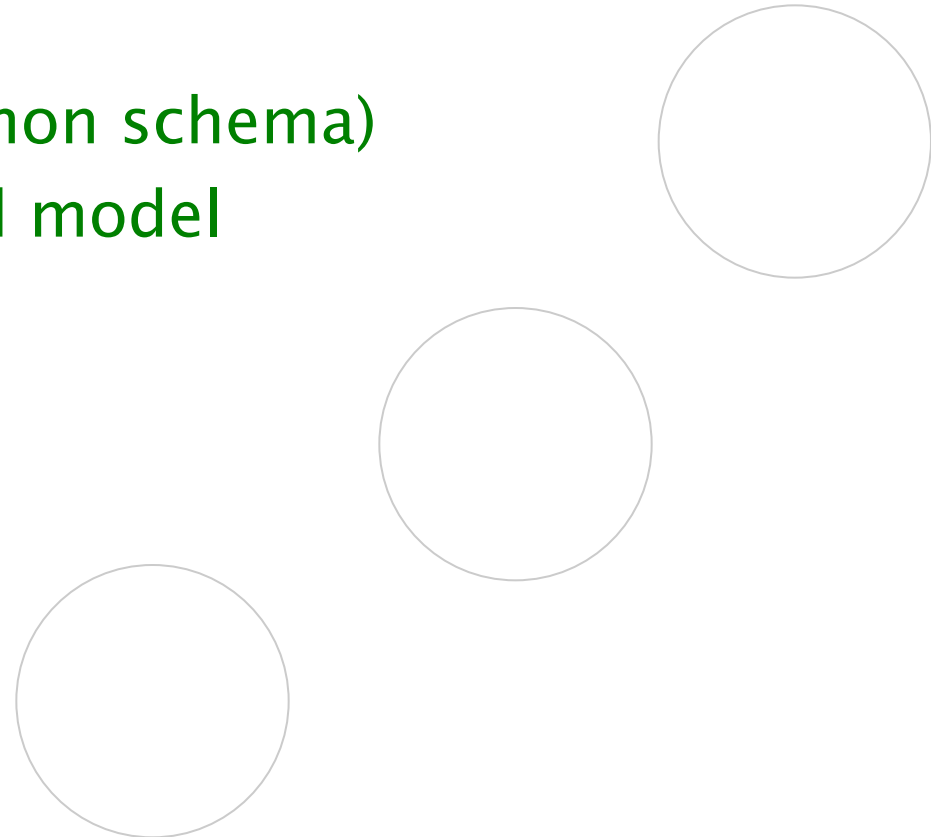


**Erdgeschoss**



Copyrights belong either to www.deri.ie or www.deri.ie/deri.org/deri.org

- Tracking system (e.g. RFID tags) : Stream data
- Attendees information (e.g. DBLP records, FOAF)
- Building information (e.g. layout, connections, room names)
- Different sources (no common schema)
- Linked data used as unified model



(Q1) Inform a participant about the name and description of the location he currently is




PREFIX lv: <http://deri.org/floorplan/>  
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

```
SELECT ?locName ?locDesc
FROM NAMED <http://deri.org/floorplan/>
WHERE {
  STREAM <http://deri.org/streams/rfid> [NOW] {?person lv:detectedat ?loc}
  GRAPH <http://deri.org/floorplan/>
    {?loc lv:name ?locName. ?loc lv:desc ?locDesc}
  ?person foaf:name "$Name$".
}
```

- Linked Data principles applied to stream data
- Extensions to deal with the temporal aspects
  - Data modeling
  - Query languages
  - Query operators
  - System design and architectures



# DATA MODELS, QUERY LANGUAGES AND OPERATORS



## ■ Extends the definition of RDF nodes and RDF triples

- RDF node:  $I$ ,  $B$ , and  $L$ , which are pair-wise disjoint infinite sets of Information Resource Identifiers (IRIs), blank nodes and literals
- RDF triple:  $(s, p, o) \in IB \times I \times IBL$ , where  $IL = I \cup L$ ,  $IB = I \cup B$  and  $IBL = I \cup B \cup L$

## ■ Stream element: RDF triple with temporal annotations

- Interval-based (e.g.  $\langle :John :at :office, [7,9] \rangle$ ) – Streaming SPARQL
- Point-based (e.g.  $\langle :John :at :office, 7 \rangle$ ,  $\langle :John :at :office, 8 \rangle$ ,  $\langle :John :at :office, 9 \rangle$ ) – EP-SPARQL, C-SPARQL, SPARQL<sub>Stream</sub>, CQELS
- Point-based (maybe) redundant, but instantaneous (more practical)



- **RDF Stream: bag of elements  $\langle (s,p,o) : [t] \rangle$** 
  - $(s,p,o) : \text{RDF triple}$
  - $t : \text{timestamp}$
  - stream elements from stream  $S$  with timestamp  $\leq t$   
$$S^{\leq t} = \{ \langle (s,p,o) : [t'] \rangle \in S \mid t' \leq t \}$$
- **Non-stream data (RDF datasets) also need to follow the Linked Stream Data model to allow integration →**  
Instantaneous RDF dataset:  $G(t)$
- **$G(t) : \text{set of RDF triples valid at time } t, \text{ called instantaneous RDF dataset.}$**
- **RDF dataset : sequence  $G = [G(t)], t \in \mathbb{N}, \text{ ordered by } t.$** 
  - Static RDF dataset ( $G^s$ ):  $G(t) = G(t+1)$  for all  $t \geq 0$

- Pattern matching as basic operator (extended from SPARQL)

- Mappings which are defined as partial functions

$$\mu : V \mapsto IBL$$

where  $V$  is an infinite set of variables disjoint from  $IBL$ , and  $\text{dom}(\mu)$  is the subset of  $V$  where  $\mu$  is defined.

- Compatible mappings

$$\mu_1 \doteq \mu_2 \iff \forall x \in \text{dom}(\mu_1) \cap \text{dom}(\mu_2) \Rightarrow \mu_1(x) = \mu_2(x)$$

- Join, union, different and left outer-join follow mappings ( $\Omega_1$  and  $\Omega_2$  are mapping sets)

$$\Omega_1 \bowtie \Omega_2 = \{\mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1 \wedge \mu_2 \in \Omega_2 \wedge \mu_1 \dot{\neq} \mu_2\}$$

$$\Omega_1 \cup \Omega_2 = \{\mu \mid \mu_1 \in \Omega_1 \vee \mu_2 \in \Omega_2\}$$

$$\Omega_1 \setminus \Omega_2 = \{\mu \in \Omega_1 \mid \neg \exists \mu' \in \Omega_2, \mu' \dot{\neq} \mu\}$$

$$\Omega_1 \bowtie \Omega_2 = (\Omega_1 \bowtie \Omega_2) \cup (\Omega_1 \setminus \Omega_2)$$

## ■ Triple matching operator

$$\llbracket P, t \rrbracket_G = \{ \mu \mid \text{dom}(\mu) = \text{var}(P) \wedge \mu(P) \in G(t) \}$$

- Triple pattern  $P \in (I \cup V) \times (I \cup V) \times (I \cup V)$
- $\mu(P)$ : triple obtained by replacing variables within  $P$  according to  $\mu$

## ■ Window matching operator

$$\llbracket P, t \rrbracket_S^\omega = \{ \mu \mid \text{dom}(\mu) = \text{var}(P) \wedge \langle \mu(P) : [t'] \rangle \in S \wedge t' \in \omega(t) \}$$

- $\omega(t)$ :  $\mathbb{N} \rightarrow 2^{\mathbb{N}}$  : function mapping a timestamp to a (possibly infinite) set of timestamps ( $\mathbb{N}$  : set of natural numbers)
- $\omega(t)$  will depend on the type of the window (e.g. time-based, tuple-based)

## ■ Sequential Operator

$$\begin{aligned} \llbracket P_1 \Rightarrow^t P_2 \rrbracket_S^\omega = & \{ \mu_1 \cup \mu_2 \mid \mu_1 \in \llbracket P_1, t \rrbracket_S^\omega \wedge \mu_2 \in \llbracket P_2, t \rrbracket_S^\omega \wedge \mu_1 \dot{=} \mu_2 \\ & \wedge \langle \mu_1(P_1) : [t'_1] \rangle \in S \wedge \langle \mu_2(P_2) : [t'_2] \rangle \in S \wedge t'_1 \leq t'_2 \} \end{aligned}$$

- AND, UNION, OPT, FILTER, AGG can be derived from operators introduced so far

- Extensions of SPARQL grammar for continuous queries
  - Few different languages have been proposed
  - Clauses to handle streams and to add window operators
- StreamingSPARQL: DatastreamClause, Window

```
SelectQuery ::= 'SELECT' ('DISTINCT' | 'REDUCED')?(Var | '*')(DatasetClause* |  
                DatastreamClause*)WhereClause SolutionModifier  
DatastreamClause ::= 'FROM' (DefaultStreamClause | NamedStreamClause)  
DefaultStreamClause ::= 'STREAM' SourceSelector Window  
NamedStreamClause ::= 'NAMED' 'STREAM' SourceSelector Window  
GroupGraphPattern ::= { TriplesBlock? ((GraphPatternNotTriples | Filter)'.')?  
                        TriplesBlock? }*(Window)?}  
Window ::= (SlidingDeltaWindow | SlidingTupleWindow | FixedWindow)  
skipSlidingDeltaWindow ::= 'WINDOW' 'RANGE' ValSpec 'SLIDE' ValSpec?  
FixedWindow ::= 'WINDOW' 'RANGE' ValSpec 'FIXED'  
SlidingTupleWindow ::= 'WINDOW'ELEMSINTEGER  
                        ValSpec ::= INTEGER | Timeunit?  
Timeunit ::= ('MS' | 'S' | 'MINUTE' | 'HOUR' | 'DAY' | 'WEEK')
```

## ■ C-SPARQL: FromStrClause, Window

*FromStrClause* → 'FROM' ['NAMED'] 'STREAM' *StreamIRI* '[RANGE' *Window*']

*Window* → *LogicalWindow* | *PhysicalWindow*

*LogicalWindow* → *Number* *TimeUnit* *WindowOverlap*

*TimeUnit* → 'd' | 'h' | 'm' | 's' | 'ms'

*WindowOverlap* → 'STEP' *Number* *TimeUnit* | 'TUMBLING'

## ■ CQELS: StreamGraphPattern (IRIs for streams)

*GraphPatternNotTriples* ::= *GroupOrUnionGraphPattern* | *OptionalGraphPattern*  
| *MinusGraphPattern* | *GraphGraphPattern*  
| **StreamGraphPattern** | *ServiceGraphPattern* | *Filter* | *Bind*

**StreamGraphPattern** ::= 'STREAM' '[' *Window* ']' *VarOrIRIref* '{' *TriplesTemplate* '}'

*Window* ::= *Range* | *Triple* | 'NOW' | 'ALL'

*Range* ::= 'RANGE' *Duration* ('SLIDE' *Duration* | 'TUMBLING')?

*Triple* ::= 'TRIPLES' INTEGER

*Duration* ::= (INTEGER 'd' | 'h' | 'm' | 's' | 'ms' | 'ns')<sup>+</sup>

# Query Example: 1 stream



(Q1) Inform a participant about the name and description of the location he just entered

■ **C-SPARQL**

```
SELECT ?locName ?locDesc
FROM STREAM <http://deri.org/streams/rfid> [NOW]
FROM NAMED <http://deri.org/floorplan/>
WHERE {
  ?person lv:detectedat ?loc. ?loc lv:name ?locName.
  ?loc lv:desc ?locDesc
  ?person foaf:name "$Name".
}
```

■ **CQELS**

```
SELECT ?locName ?locDesc
FROM NAMED <http://deri.org/floorplan/>
WHERE {
  STREAM <http://deri.org/streams/rfid> [NOW] {?person lv:detectedat ?loc}
  GRAPH <http://deri.org/floorplan/>
  {?loc lv:name ?locName. ?loc lv:desc ?locDesc} ?person foaf:name "$Name
  $ ".
}
```



(Q2) Notify two people when they can reach each other from two different and directly connected (from now on called nearby) locations.

- Streaming SPARQL and C-SPARQL don't allow multiple windows in one stream in their grammar
  - C-SPARQL solution: create two virtual streams
- CQELS

```
CONSTRUCT {?person1 lv:reachable ?person2}
FROM NAMED <http://deri.org/floorplan/>
WHERE {
  STREAM <http://deri.org/streams/rfid> [NOW] {?person1 lv:detectedat ?loc1}
  STREAM <http://deri.org/streams//rfid> [RANGE 3s] {?person2 lv:detectedat ?loc2}
  GRAPH <http://deri.org/floorplan/> {?loc1 lv:connected ?loc2}
}
```

- Different streams can provide the same pattern

Q3: Name of location of people nearby the DERI building

- CQELS (queries all streams that provide “nearby” info)

```
SELECT ?name ?locName
FROM NAMED <http://deri.org/floorplan/>
WHERE {
  STREAM ?streamURI [NOW] {?person lv:detectedat ?loc}
  GRAPH <http://deri.org/floorplan/>
  {
    ?streamURI lv:nearby :DERI_Building. ?loc lv:name ?locName.
    ?person foaf:name ?name.
  }
}
```

# Query Example: Timestamps



- EP-SPARQL and C-SPARQL allow functions to deal with timestamps
  - Timestamp can be retrieved and bound to a variable
  - Timestamp of a bound stream element can be retrieved
- Q4: Return pairs of people that were detected in a location in consecutive times (in the last 30min)
- EP-SPARQL

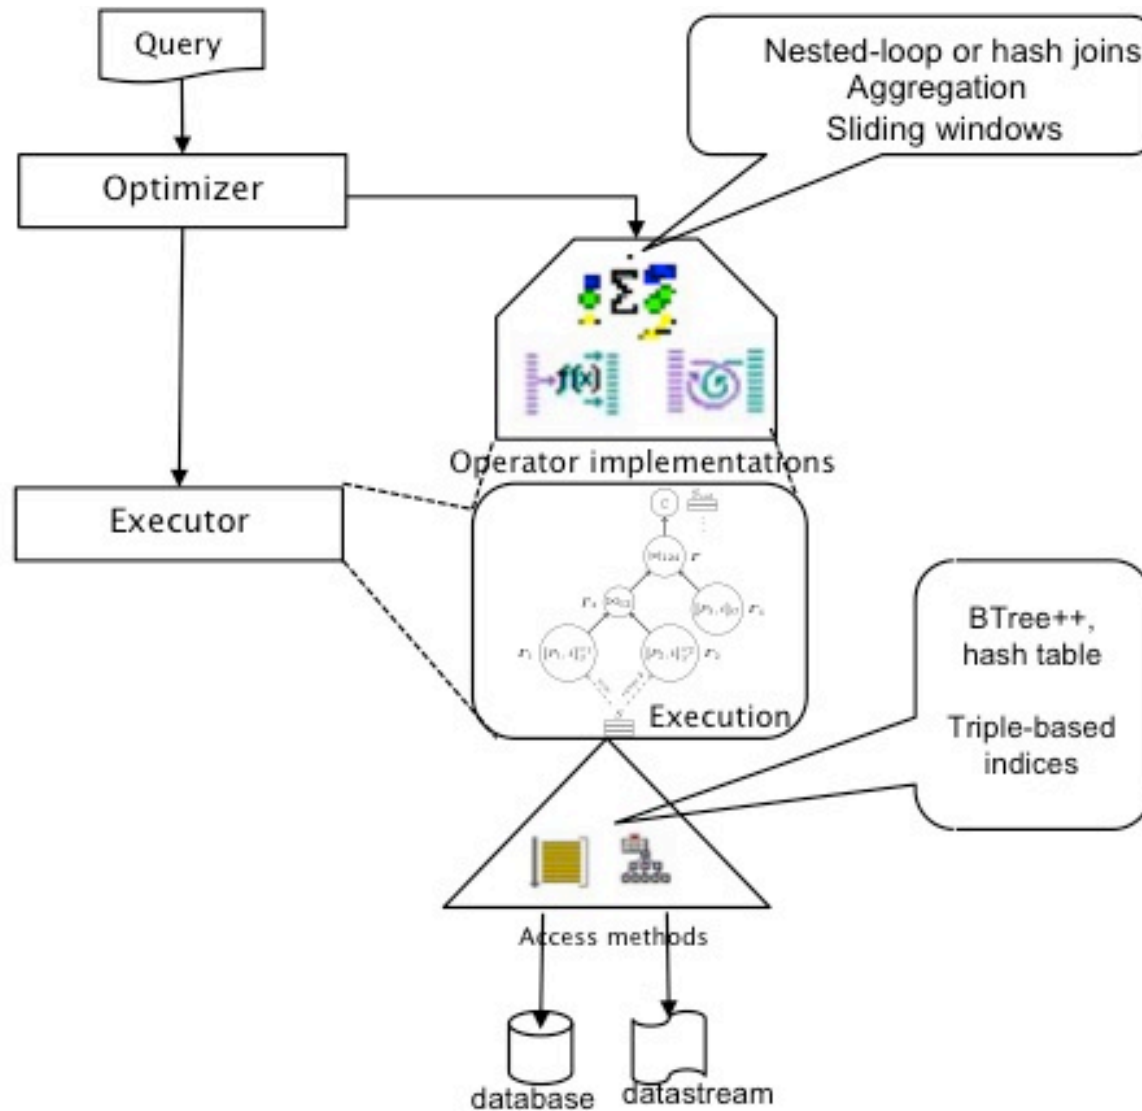
```
CONSTRUCT {?person2 lv:comesAfter ?person1} {  
  SELECT ?person1 ?person2  
  WHERE {  
    {?person1 lv:detectedat ?loc}  
    SEQ {?person2 lv:detectedat ?loc}  
  }  
  FILTER (getDURATION()<"P30m"^^xsd:duration)
```

# DESIGN CHOICES & CHALLENGES

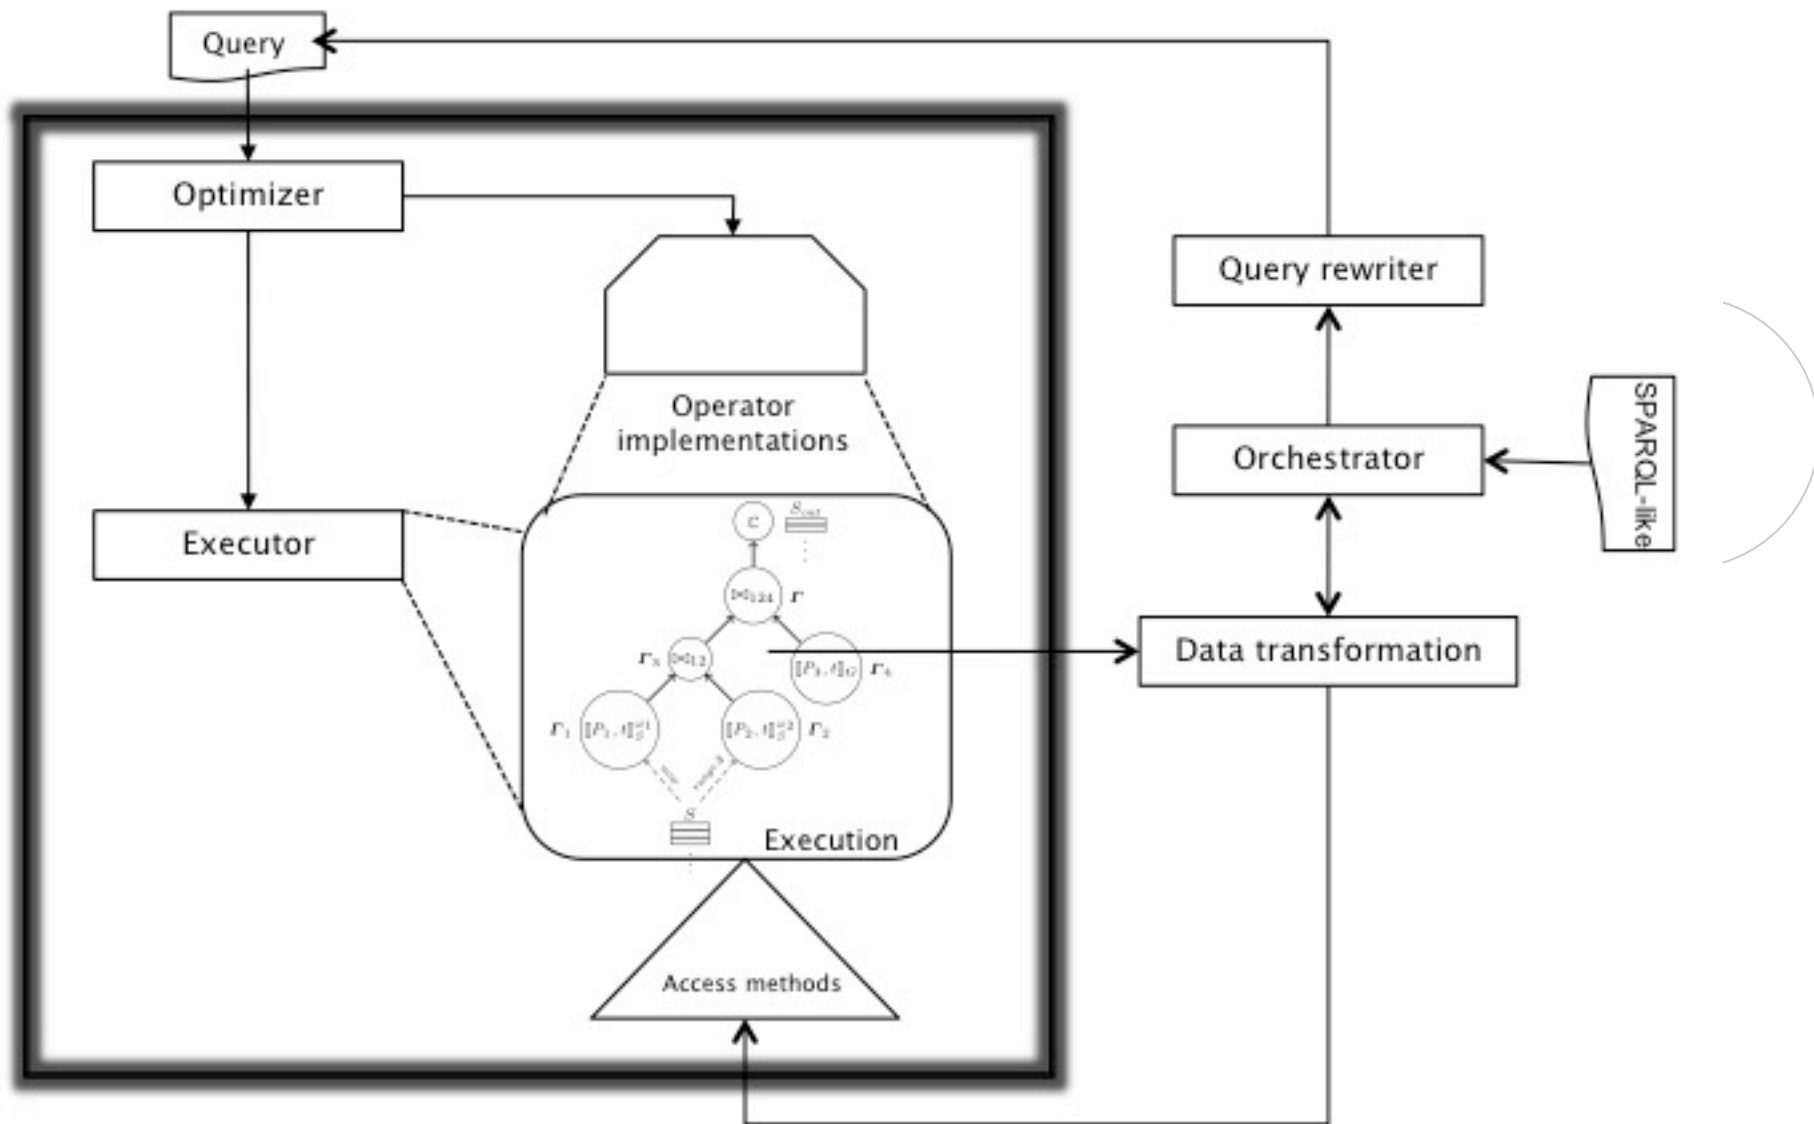


- **Current available systems can be classified into two categories based on their architecture design**
- **White box architecture**
  - Implements all required components
    - physical operators (e.g. windows, join, triple pattern matching)
    - data structures (e.g. B+-Trees, hashtables)
    - query generator/optimizer/executor
- **Black box architecture**
  - Uses existing RDF and data stream processing systems as sub-components
  - Query rewriter, data translator and orchestrator among sub-components is needed
- **Black box easier to implement, but no full-control and data transformation overhead**

# White box



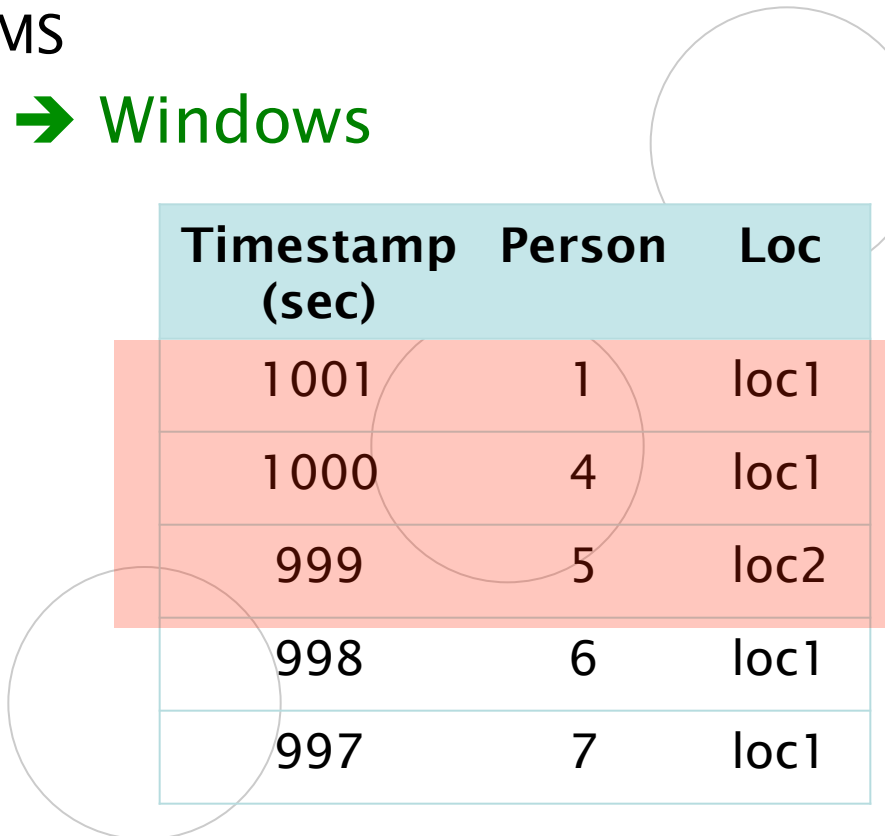
# Black box



- Current Linked Stream Data approaches follow/reuse operators from relational DSMS
- Continuous query
  - $Q(t)$ : query results up to time  $t$
  - $R(t)$  : unordered bag of tuples (relations) at time instant  $t$
  - Relation  $R$ : sequence  $R = [R(t)]$ ,  $t \in \mathbb{N}$ , ordered by  $t$ .
- Query algebra
  - Stream-to-stream (Streaming SPARQL): stream-to-stream operator
  - Mixed (C-SPARQL,  $\text{SPARQL}_{\text{Stream}}$ , CQELS): stream-to-relation, relation-to-relation and relation-to-stream operators

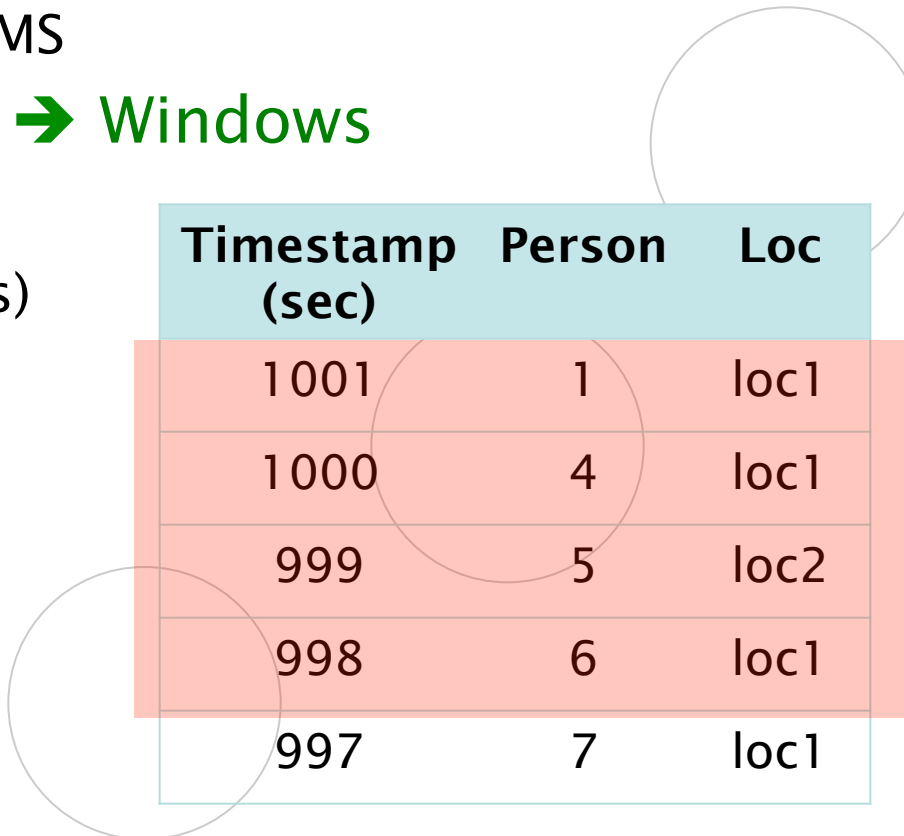


- **Stream-to-stream operator**
  - One-time queries in SQL that are continuously executed
- **Relation-to-relation operator**
  - As in traditional relational DBMS
- **Stream-to-relation operator → Windows**
  - Time-based (e.g. last 3 secs)

A diagram illustrating a sliding window over a stream of data. A table with five rows is shown. The top three rows (timestamps 1001, 1000, and 999) are highlighted with a light red background, representing the current window. The bottom two rows (timestamps 998 and 997) are not highlighted. A large white circle is drawn around the entire table, and a smaller white circle is drawn around the first row (timestamp 1001).

Timestamp (sec)	Person	Loc
1001	1	loc1
1000	4	loc1
999	5	loc2
998	6	loc1
997	7	loc1

- **Stream-to-stream operator**
  - One-time queries in SQL that are continuously executed
- **Relation-to-relation operator**
  - As in traditional relational DBMS
- **Stream-to-relation operator → Windows**
  - Time-based (e.g. last 3 secs)
  - Tuple-based (e.g. last 4 tuples)

A diagram illustrating a window in a stream. A table with five rows is shown. The first four rows are highlighted in light red, representing the current window. The fifth row is highlighted in light blue, representing the next tuple to enter the window. A large light blue circle highlights the entire table area. A smaller light blue circle highlights the first row of the table. Another light blue circle highlights the 'Person' column of the first row.

Timestamp (sec)	Person	Loc
1001	1	loc1
1000	4	loc1
999	5	loc2
998	6	loc1
997	7	loc1

- **Stream-to-stream operator**
  - One-time queries in SQL that are continuously executed
- **Relation-to-relation operator**
  - As in traditional relational DBMS
- **Stream-to-relation operator → Windows**
  - Time-based (e.g. last 3 secs)
  - Tuple-based (e.g. last 4 tuples)
  - Partitioned (e.g. Loc last tuple)

Timestamp (sec)	Person	Loc
1001	1	loc1
1000	4	loc1
999	5	loc2
998	6	loc1
997	7	loc1

## ■ Relation-to-stream operator: produces a stream from relation R

- Istream (insert stream): add element  $\langle s, t \rangle$  whenever  $s$  is in  $R(t) - R(t-1)$
- Dstream (delete stream): add element  $\langle s, t \rangle$  whenever  $s$  is in  $R(t-1) - R(t)$
- Rstream (relation stream): add element  $\langle s, t \rangle$  whenever  $s$  is in  $R$  at time  $t$ .

### Istream

```
SELECT Istream(*)  
FROM RFIDstream [RANGE Unbounded]  
WHERE signalstrength >= 85
```

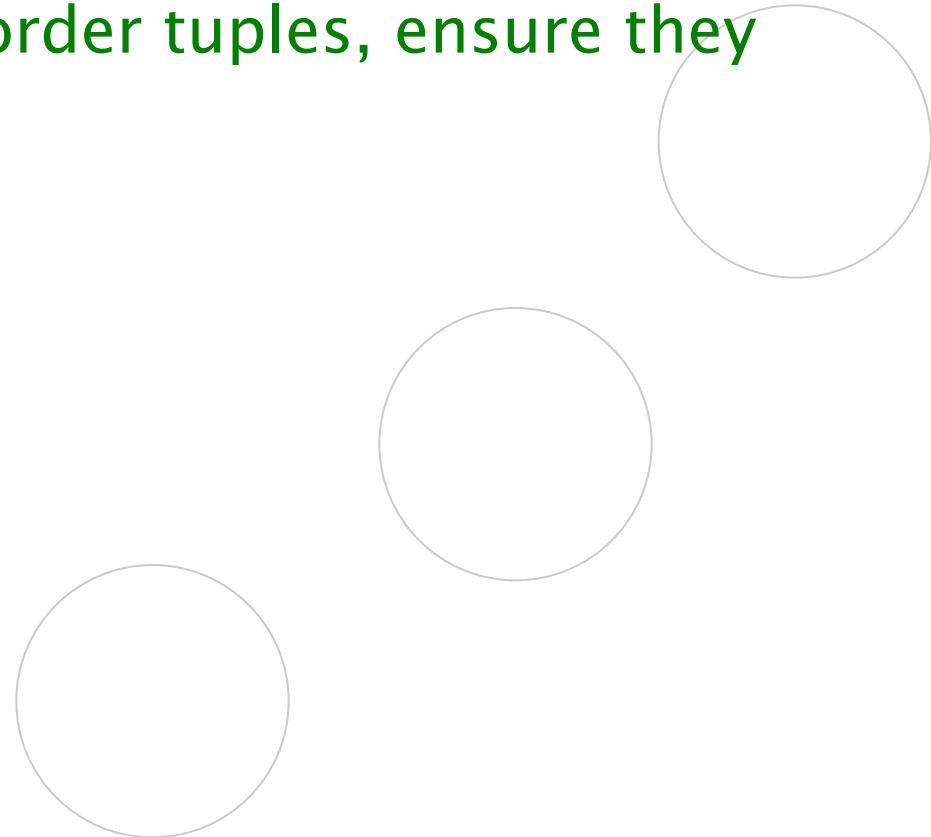
### Rstream

```
SELECT Rstream(*)  
FROM RFIDstream [NOW]  
WHERE signalstrength >= 85
```


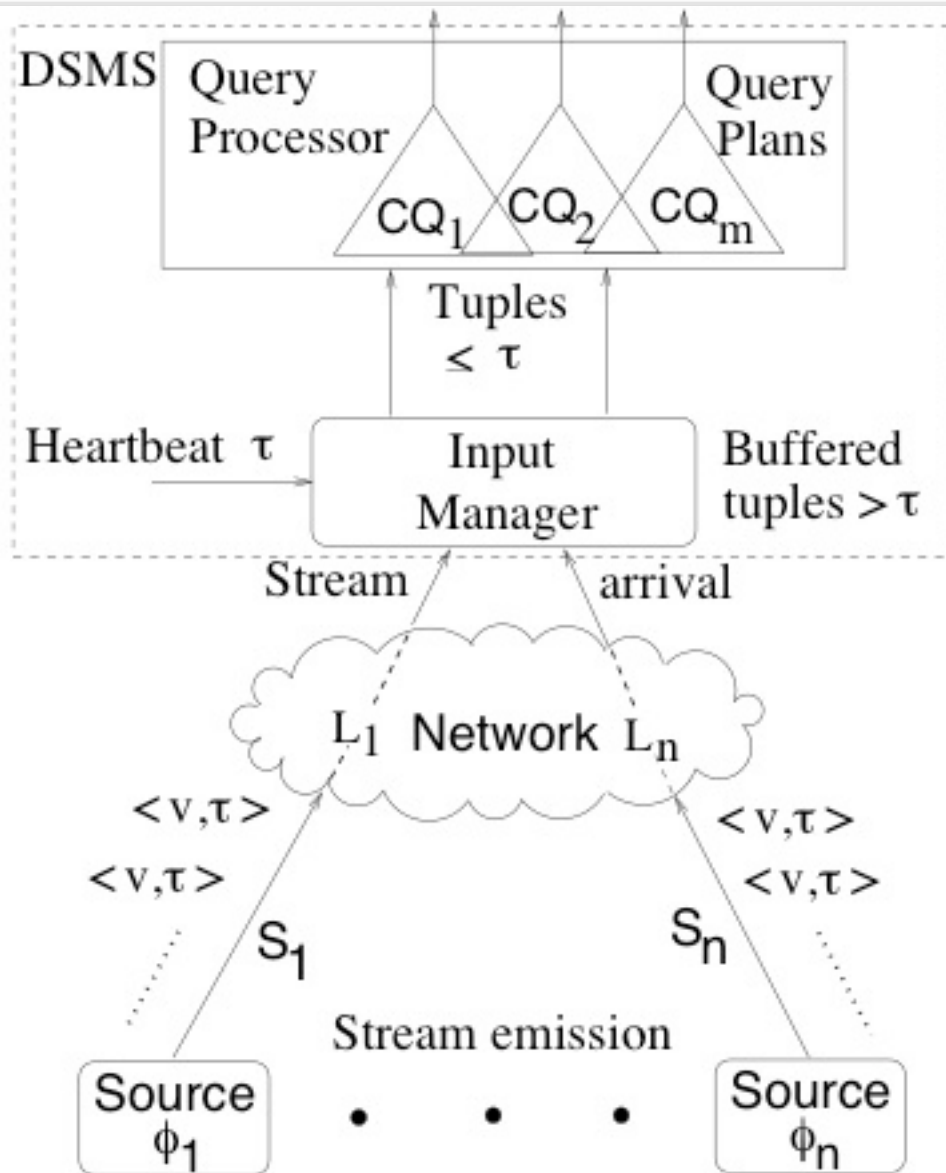
### Dstream

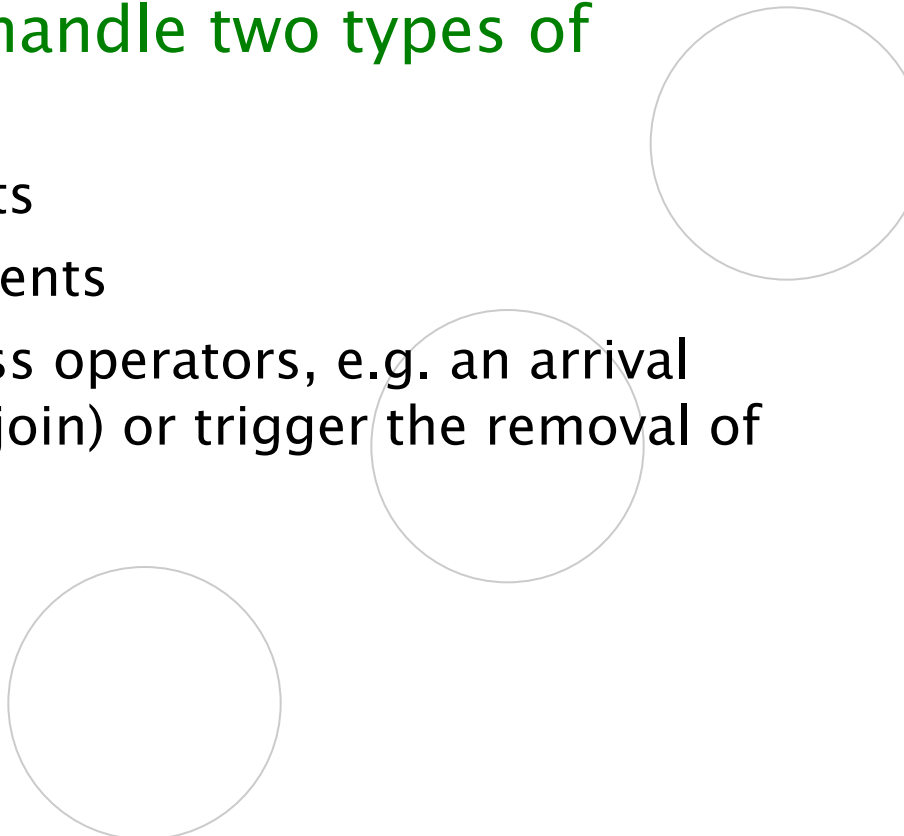
```
SELECT Dstream(tagid) FROM RFIDstream [60 seconds]
```

- Timestamps are necessary to order stream elements
- Application timestamp (source) vs. system timestamp (DSMS)
- Input manager: buffers to order tuples, ensure they are processed in order
  - Heartbeat (timestamp)
  - Punctuation (pattern)

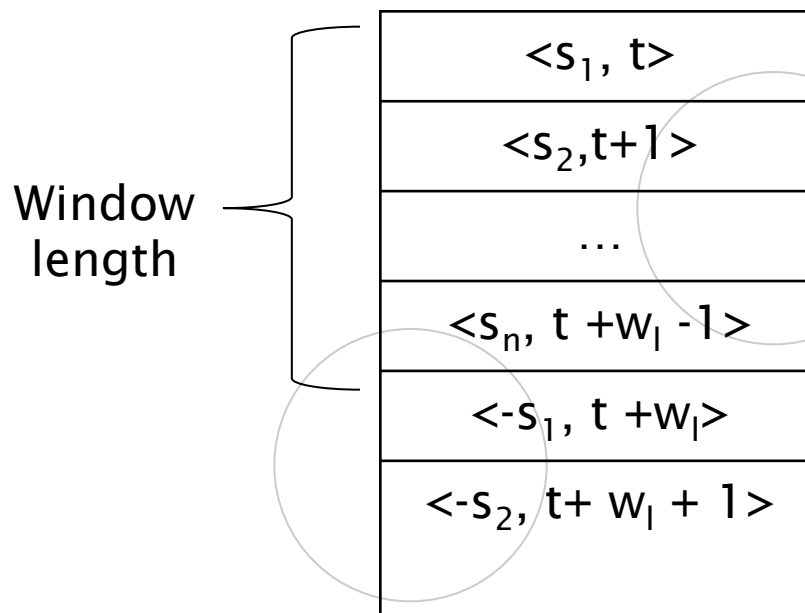


# Time Management




- Eager re-evaluation vs. period re-evaluation
    - Eager: too expensive if update rate is high
    - Periodic: might cause stale results
  - Query evaluation needs to handle two types of events
    - Arrival of new stream elements
    - Expiration of old stream elements
    - Action upon events vary across operators, e.g. an arrival might generate a new result (join) or trigger the removal of an existing result (negation)
- 
- Three empty circles are scattered on the right side of the slide. One is at the top right, one is in the middle right, and one is at the bottom center.

- Arrivals are triggered by stream source
- Expiration needs to be handle by the query processor
  - Timestamp
  - Negative tuple: for a window of length  $w_l$ , a tuple inserted at time  $t$  will generate a negative tuple at time  $t+w_l$





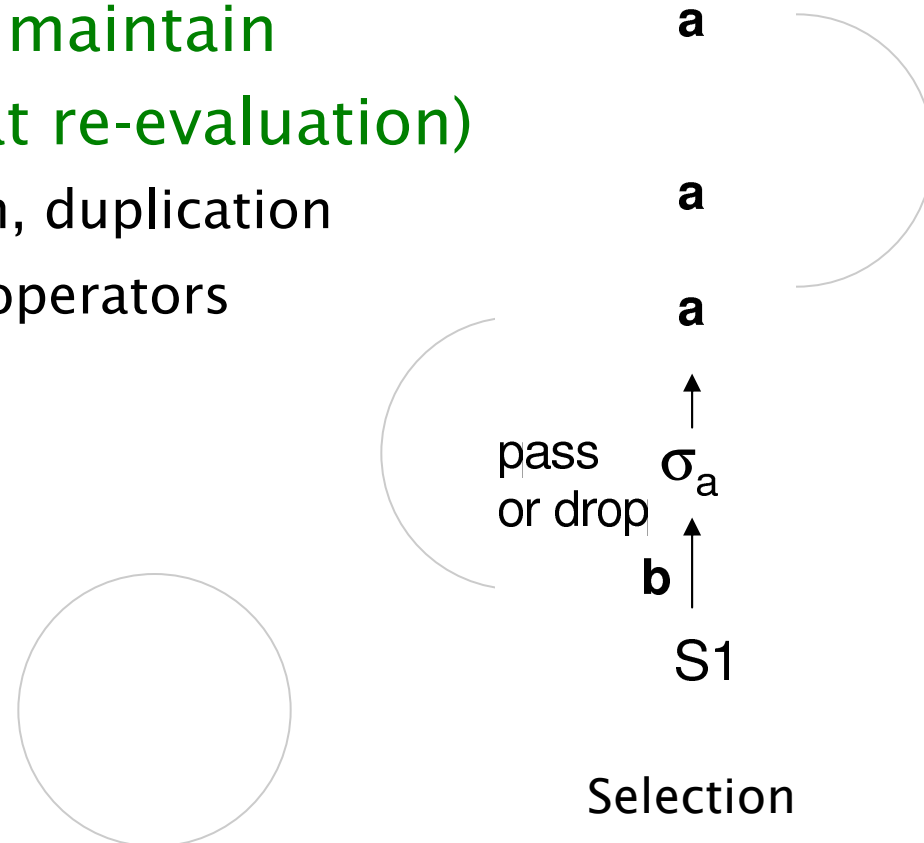
# Adding and evicting stream elements



# Query Evaluation

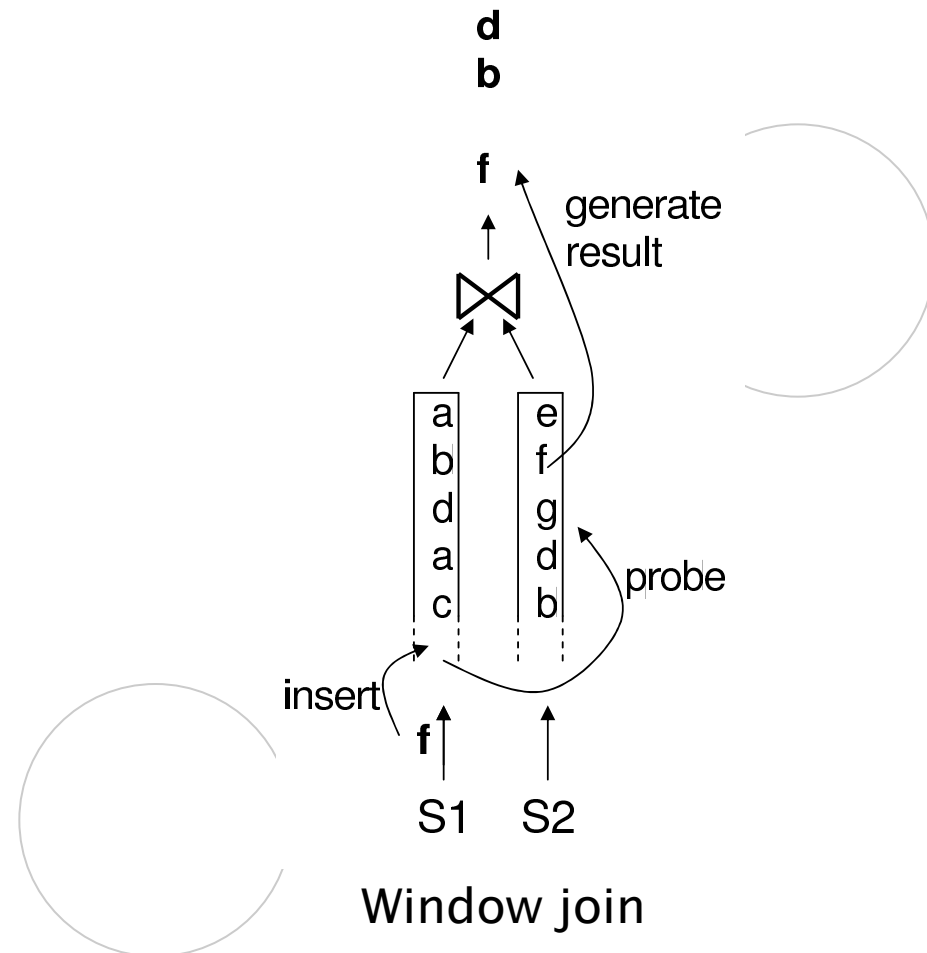


- **Stateless operators: processed “on the fly” (directly on stream)**
  - E.g. Selection, union.
- **Stateful operators: need to maintain processing states (probed at re-evaluation)**
  - E.g. window join, aggregation, duplication elimination, non-monotonic operators



Selection

- Window join : new arrival in one input triggers probing on the other input



## ■ Aggregation

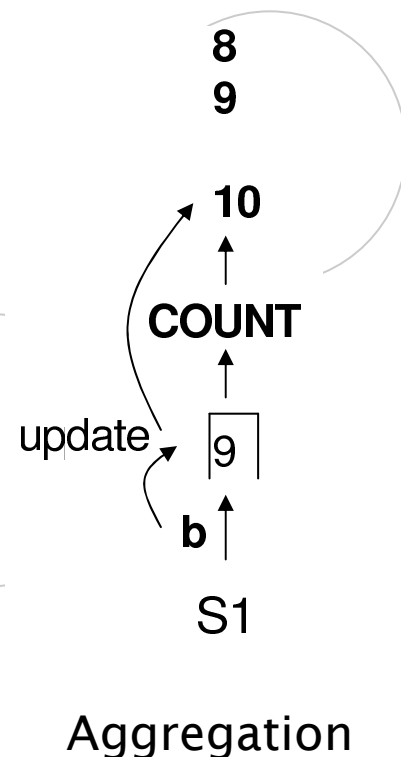
- Expirations must be dealt with immediately
- Time and space requirements depends on the aggregation function

## ■ Distributive aggregates

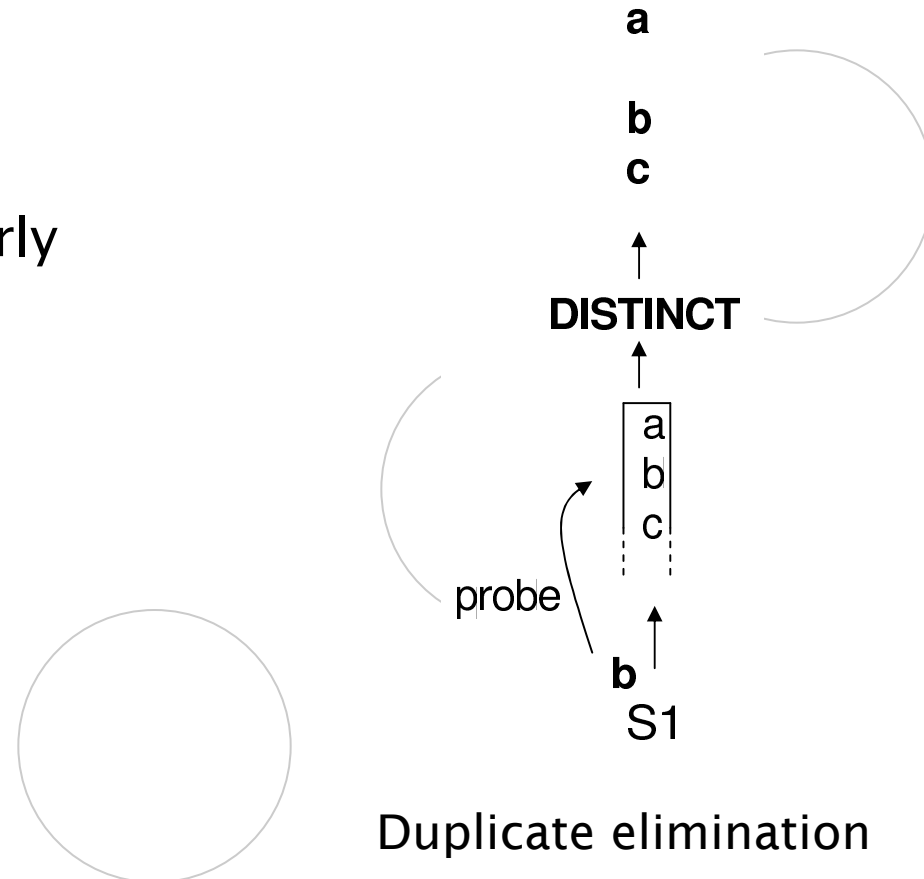
- Computed incrementally, constant time/space requirements
- E.g. COUNT, SUM, MAX, MIN

## ■ Algebraic aggregates

- Computed using values from distributive aggregates. Constant time/space requirements
- E.g. AVG (SUM/COUNT)

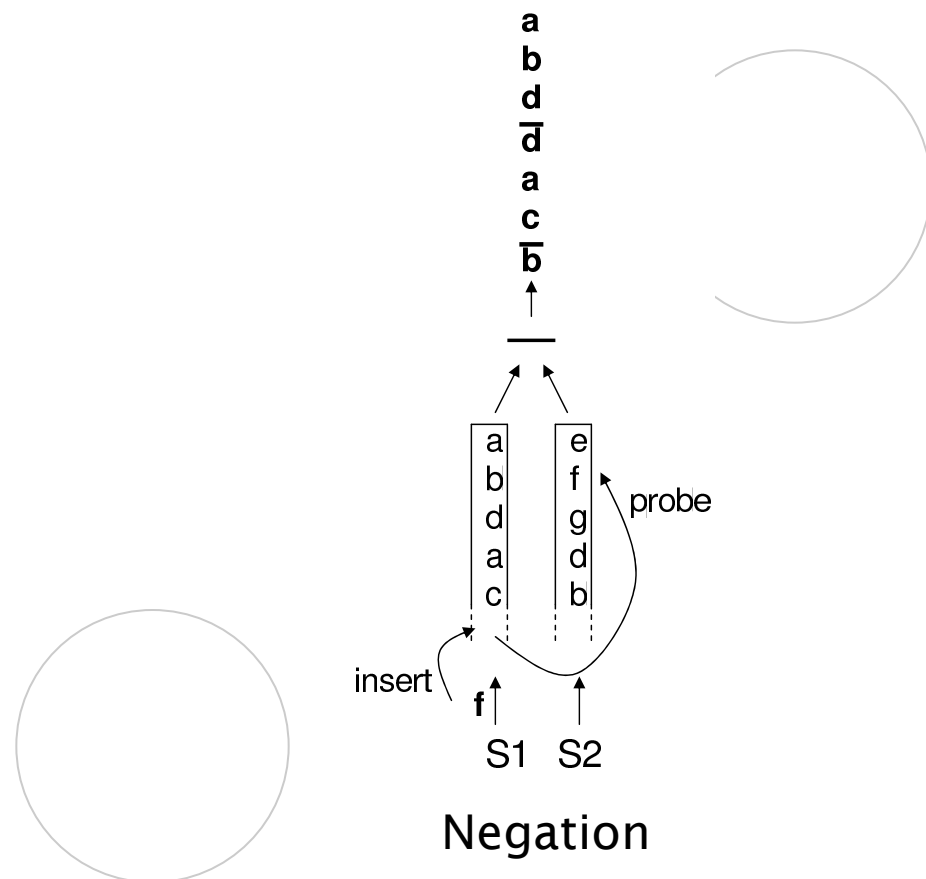


- **Holistic aggregates: space consumption linear to input sizes**
  - E.g. TOP-k, COUNT DISTINCT
- **Duplicate elimination**
  - Distinct values are kept
  - Expirations are handled eagerly






## ■ Non-monotonic operators

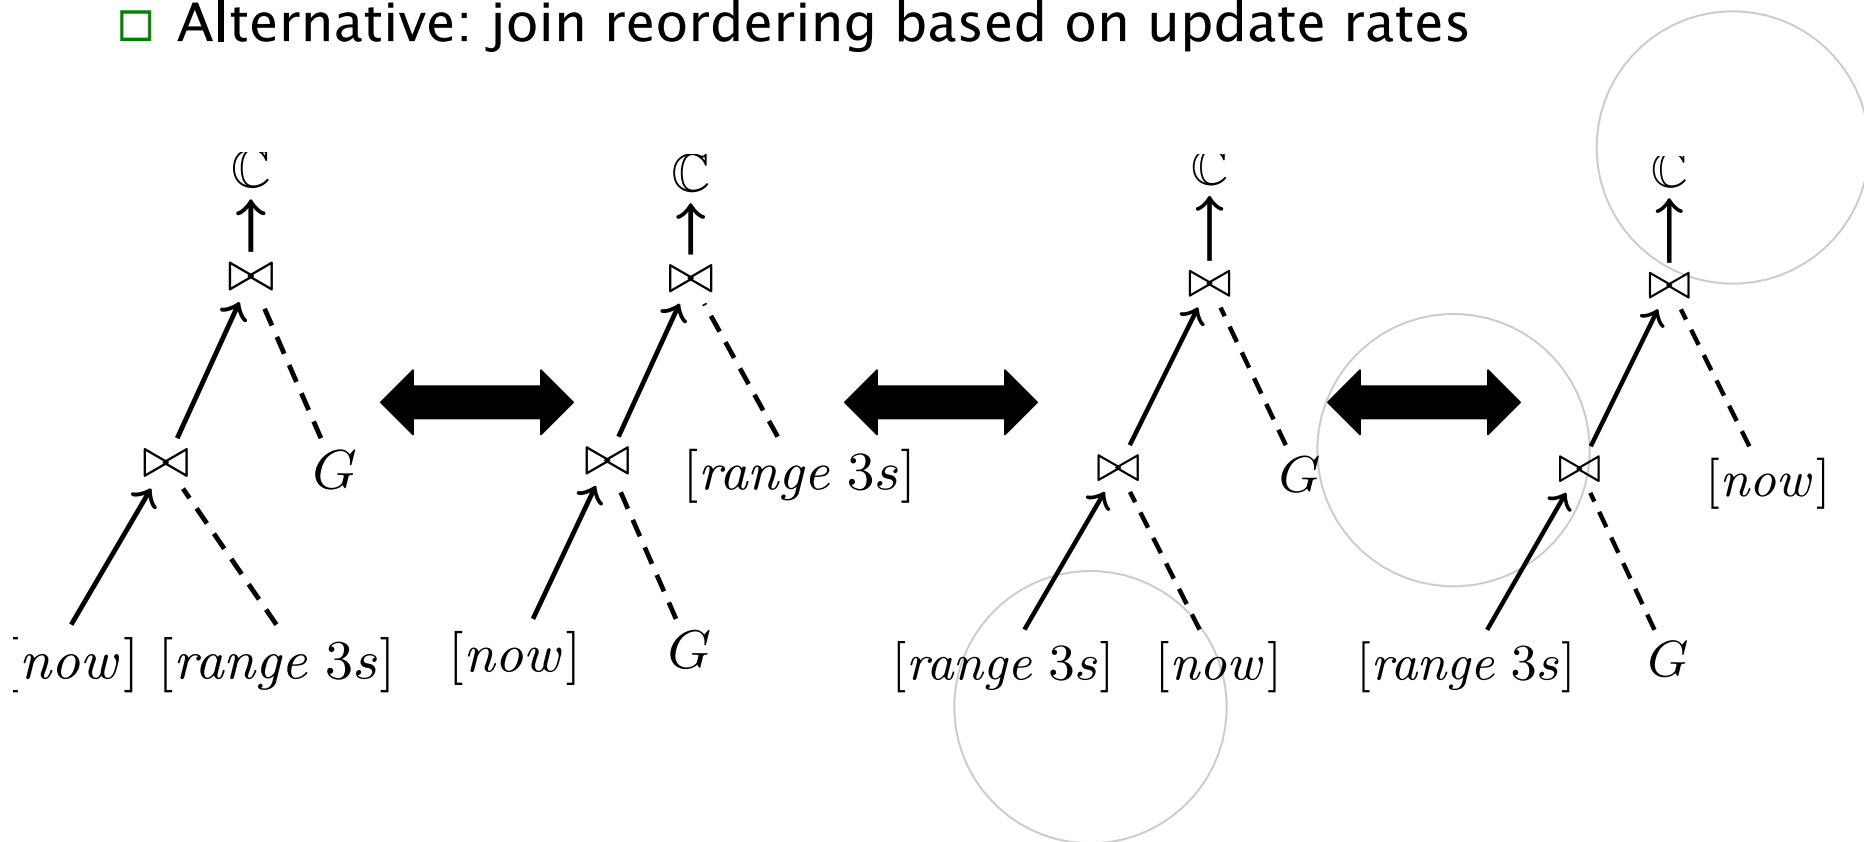
- Previous results removed when they no longer satisfy query
- E.g. negation
- Negative tuples can be used



- Some join operators already handles memory overflow by sending input partitions to disk.
- Use of secondary storage requires indexes
  - Expensive under high update rates
- **Alternative: Partition the data to make updates “local”**
  - Sort tuples chronologically
  - Inserts in newer partition only
  - Deletes in older partition only
  - Problem: search is not efficient. Assumes insertion/expiration order is the same
    - Sub-indexes
    - Doubly partitioned indexes



- Re-arrange query operators for more efficient execution
  - Traditional selectivity estimates can't be applied
  - Alternative: join reordering based on update rates





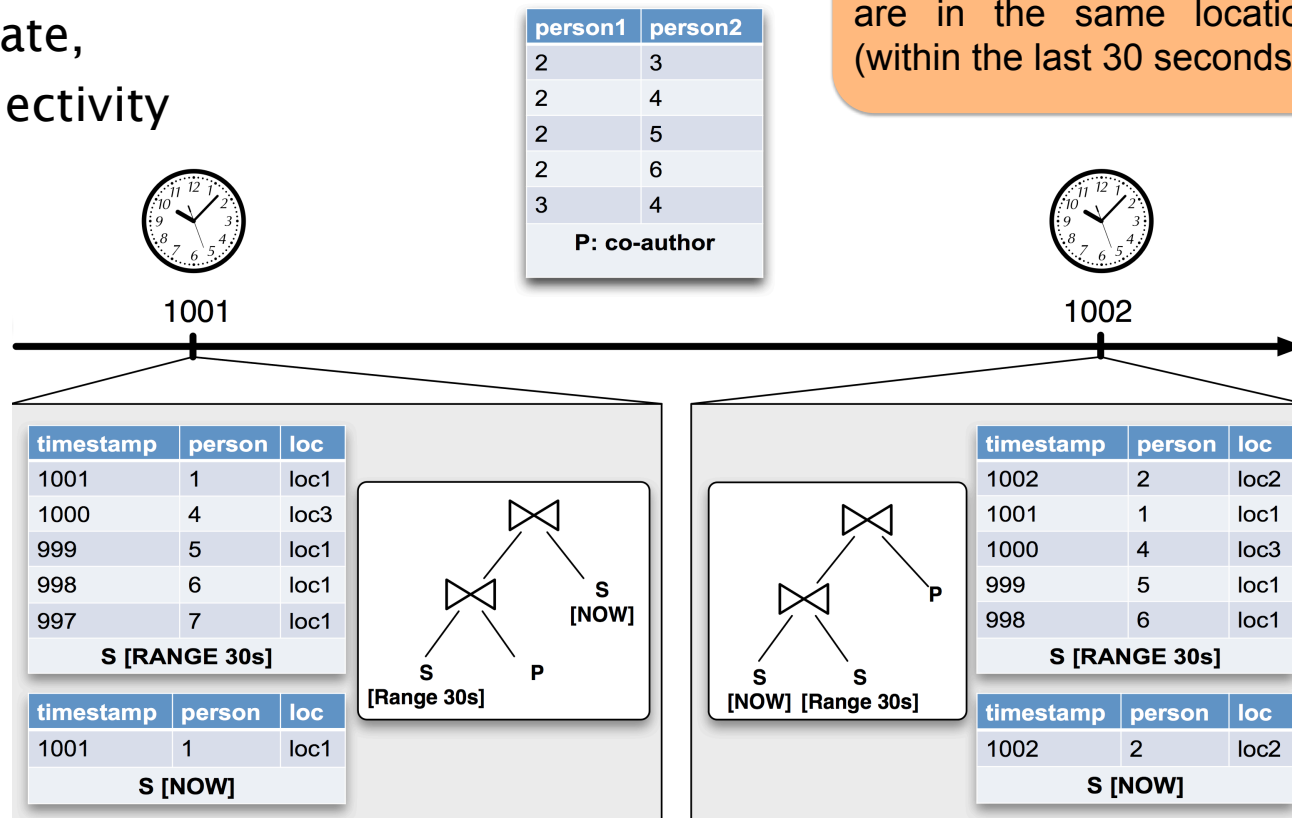
# Query Optimization



## ■ Adaptivity is key!

- Processor must be able to reorder query operators on the fly
- Changes in:
  - operator costs (processing time),
  - update rate,
  - input selectivity


“Notify two people who are co-authors of a paper if they are in the same location (within the last 30 seconds)”



# Query Optimization



- **Operators routing (instead of fixed query plan tree)**
  - Eddies: estimate which operators are faster/more selective
  - Overhead: migration of internal state of query plan
- **Continuous query: multi-query optimization possible**
  - Better memory usage
  - Trade-offs exists (e.g. join -> selection vs. selection -> join)



- Data first push into queues, then consumed by operators
- Scheduler determines which data in which queue to process next
  - Different scheduling strategies (e.g. round robin, arrival time, time slice)
  - Choice depends on factors such as stream arrival patterns, max/avg output latency.

