

A Query Language for Semantic Complex Event Processing: Syntax, Semantics and Implementation

Editor(s): Name Surname, University, Country
Solicited review(s): Name Surname, University, Country
Open review(s): Name Surname, University, Country

Syed Gillani ^a, Antoine Zimmermann ^b, Gauthier Picard ^b and Frédérique Laforest ^c

^a *Univ Lyon, INSA Lyon, France*

E-mail: syed.gillani@insa-lyon.fr

^b *Univ Lyon, MINES Saint-Étienne, CNRS, Laboratoire Hubert Curien UMR 5516, France*

E-mail: antoine.zimmermann@emse.fr, gauthier.picard@emse.fr

^c *Univ Lyon, UJM Saint-Étienne, CNRS, Laboratoire Hubert Curien UMR 5516, France*

E-mail: frederique.laforest@telecom-st-etienne.fr

Abstract

The field of Complex Event Processing (CEP) deals with the techniques and tools developed to efficiently process pattern-based queries over data streams. The Semantic Web, through its standards and technologies, is in constant pursuit to provide solutions for such paradigm while employing the RDF data model. The integration of Semantic Web technologies in this context can handle the heterogeneity, integration and interpretation of data streams at the semantic level. In this paper, we propose and implement a new query language, called SPASEQ, that extends SPARQL with new Semantic Complex Event Processing (SCEP) operators that can be evaluated over RDF graph-based events. The novelties of SPASEQ include (i) the separation of general graph pattern matching constructs and temporal operators; (ii) the support for RDF graph-based events and multiple RDF graph streams; (iii) the expressibility of temporal operators such as *Kleene+*, conjunction, disjunction and event selection strategies; and (iv) the operators to integrate background information and streaming RDF graph streams. Hence, SPASEQ enjoys good expressiveness compared with the existing solutions. Furthermore, we provide an efficient implementation of SPASEQ using a non-deterministic automata (NFA) model for an efficient evaluation of the SPASEQ queries. We provide the syntax and semantics of SPASEQ and based on this, we show how it can be implemented in an efficient manner. Moreover, we also present an experimental evaluation of its performance, showing that it improves over state-of-the-art approaches.

Keywords: Complex Event Processing, Query Language, Semantic Web, SPARQL, RDF Streams, RDF Stream Processing, Automata Model, Query Optimisation

1. Introduction

Stream processing has become an important paradigm for processing data at high speed and large scale, where query operators such as selection, aggregation, filtering of data are performed in a streaming fashion [1,2]. Complex Event Processing (CEP) systems, however, provide a different view and additional operators for

streaming applications: each data item within a data stream is considered as an event and predefined temporal patterns are used to generate actions to the systems, people and devices. CEP systems have demonstrated utility in a variety of applications including financial trading, security monitoring, healthcare, social and sensor network analysis [3,4,5]. In general,

CEP denotes algorithmic methods for making sense of events by deriving higher-level knowledge, or complex events from lower-level events in a timely fashion [3,6]. CEP applications commonly involve three requirements [3,6,7,8]:

- (i) complex predicates (filtering, correlation);
- (ii) temporal, order and sequential patterns; and
- (iii) transforming event(s) into more complex structures.

The past several years have seen a large number of CEP systems and query languages being developed by both academic and industrial worlds [5,8,9,10,11,12,13,14]. Following the trend of using RDF as a unified data model for integrating diverse data sources across heterogeneous domains, there is a great interest in using RDF data model for the CEP. This results in a new field of the Semantic CEP (SCEP) [13,15]. The use of RDF as an input data model provides the following attributes [16,17]. First, it offers a unified data model (directed labelled graphs). Second, it allows building upon countless domain knowledge formalisations, in the form of vocabularies, thesauri and ontologies. Third, RDF-based data integration systems benefit from the reasoning capabilities offered by the Semantic Web technologies to reason upon the context of detected events. Fourth, it makes it possible to leverage the huge knowledge base represented by the Web of Data, thereby opening up opportunities to discover related data sets, enrich data, and create added value by mashing up all this information.

The design of an efficient SCEP system requires carefully marrying the temporal operators with the RDF data model and the additional characteristic of event enrichment. Even though SCEP can be evolved from the common practice of stitching heterogeneous techniques and systems, a well organised query language is a vital part of SCEP: it not only allows users to specify known queries or patterns of events in an intuitive way, but also to showcase the expected answers of a query while hiding the implementation details.

While there does not exist a standard language for expressing continuous queries over RDF streams, a few options have been proposed. In particular, a first strand of research focuses on extending the scope of SPARQL to enable the stateless continuous evaluation of RDF triple streams called RDF Stream Processing (RSP) languages. These query languages include CQELS [18], C-SPARQL [19], SPARQL_{Stream} [20]: they are used to match query-defined graph patterns, aggregates and fil-

tering operators against RDF triple streams, i.e. a sequence of RDF triples $\langle \text{subject}, \text{predicate}, \text{object} \rangle$, each associated with a timestamp τ . Most of these languages do not provide any explicit temporal pattern matching operator and thus cannot be classified under the SCEP languages.

The second strand of research focused on SCEP languages to extend SPARQL with stateful operators, where few options have been proposed [21,22,23]. EP-SPARQL [21] – with its expressive language and framework – is the primary player in this field with other works focusing on a subset of operators and functionalities. EP-SPARQL extends SPARQL with sequence constructs that allow temporal ordering over triple streams and its semantics are derived from the ETALIS language [13]. Although EP-SPARQL is a pioneering work in the field of SCEP, it suffers from drawbacks such as mixing of sequence and graph pattern matching operators, single triple stream model, lack of explicit *Kleene+*, and event selection operators.

Considering these shortcomings, our contribution in this paper is twofold. First, we present a novel query language and system, called SPASEQ, to enable complex event processing over multiple streams using the RDF graph model. Second, we present an efficient framework to process SPASEQ queries. SPASEQ covers the aforementioned shortcomings of existing languages and systems, and provides a unified language for SCEP over RDF graph streams, while introducing expressive explicit temporal operators. The use of explicit operators lets the users specify complex queries at high level and enables the appropriate implementation details and optimisations at the domain specific level. The main features of the SPASEQ query language are summarised as follows:

- One of the most important features of SPASEQ is that it clearly separates the query components for describing temporal patterns over RDF graph events, from specifying the graph pattern matching over each RDF graph event. This results in language reusability and expressiveness. This separation distinguishes SPASEQ from other SCEP query languages.
- It provides an RDF graph stream model to support streaming of graphs instead of triples. This allows to structure more complex events in a stream. Furthermore, it provides event processing over multiple RDF graph streams, while following the SPARQL specification of named datasets [24].

- It is equipped with expressive temporal operators such as *Kleene+*, disjunction, conjunction and event selection strategies over events from multiple streams.
- It provides explicit operators to join background knowledge base and RDF graph streams.

These features stem from the general specifications described by the W3C RSP community group [25] and the use cases for the CEP [26,27,10,8,28,29].

As our second contribution, we provide an executional framework for SPASEQ using a non-deterministic finite automata (NFA) model called NFA_{scep} . Although a large number of techniques exist to process temporal operators [5,13,11], the NFA model offers higher expressiveness required for SPASEQ. Hence, leveraging the automata-based techniques, SPASEQ queries are compiled over equivalent NFAs: the compiled NFA is processed to incrementally produce the partial matches before a full match of the query is produced. Moreover, we employ various optimisation techniques on top of our system. It includes indexing and partitioning of incoming streams; pushing the stateful predicates and query windows; and lazy evaluation of the disjunction and conjunction operators.

Our main contributions are summarised as follows:

- We present SPASEQ, a novel SCEP language, including a full description of its syntax and semantics.
- We provide the detailed semantics of SPASEQ and its main operators.
- We provide the NFA-based framework to efficiently compile and evaluate SPASEQ queries.
- We provide system and operator-level optimisation strategies for SPASEQ queries.
- Using real-world and synthetic datasets, we demonstrate the effectiveness of our optimisation strategies and our experimental evaluations show that they outperform existing systems for the same use cases and datasets.

The rest of the paper is structured as follows. Section 2 presents the motivational examples to showcase the requirements for a SCEP language. Section 3 reviews the related work. Section 4 presents the data model and syntax of SPASEQ. Section 5 presents the semantics of the SPASEQ language. Section 6 presents the details about the compilation of SPASEQ queries onto the NFA_{scep} model, the evaluation of NFA_{scep} automata and the design of the SPASEQ query engine. Section 7 presents the optimisations techniques used

for the SPASEQ query engine. Section 8 provides experimental evaluations of the SPASEQ query engine. Section 9 concludes the paper.

2. Motivational Examples and Requirements

In the following, we use various use cases to show the kind of expressiveness and flexibility needed for a SCEP query language. These use cases are illustrative to showcase how a SCEP can be applicable over provided RDF graph streams.

UC 1 (Smart Grid Monitoring) *A smart grid monitoring application processes events from multiple distributed streams with the aim to notify the users of an online service to take a decision to improve the power usage when a defined pattern is detected [30,31]. An example of it is to notify the user to switch to a renewable power source instead of fuel-based power source, if the system observes the following pattern: (A) the price of electricity generated by a fuel-based power source is greater than a certain threshold; (B) weather conditions are favourable for renewable energy production (one or more events); and (C) the price of storage source attached to the renewable power source is less than the fuel-based power source. Moreover, a background knowledge base (KB) containing information about the homes' addresses, owners and measurement units, etc., can be used to further enrich the matched events.*

UC 2 (Trajectory Classification) *Trajectory classification uses the sequence of objects movement (trajectories) to determine their types [26,32]. For instance, finding the fishing boat by discovering the trajectory of a boat over some time interval. An example pattern to determine the trajectory of fishing boats represents the following pattern: A: vessel leaves the harbour, B: vessel travels by keeping steady speed and direction (one or more events), C: vessel arrives at the fishing area and stops. The information about the owner of the vessels, shipping company a vessel is operating under, address of the fishing area, etc., can be used from a background KB to enrich the matched events.*

UC 3 (Fraud Management) *Fraud management has become a central aspect in today's world, and it covers a large variety of domains. The goal in credit card fraud management is to detect fraud within 25 milliseconds, in order to prevent financial loss [27]. An important pattern in this context is to detect the "Big af-*

ter Small” fraudulent transaction [27,33]. That is, an attacker withdraws one or a series of small amounts from a credit card after withdrawing a large sum of money. The SCEP query in this context detects the following pattern: **(A)** the withdrawn amount gradually increases; **(B)** until it is considerably less than the earlier withdrawn (one or more events). Furthermore, a background KB containing the cards’ owner names, issuing banks, etc., can further enrich the context of the matched events.

UC 4 (Stock Market Analytics) A stock market analytics platform processes financial transactions to detect the patterns that signify the emergence of profit opportunities [10,8,28]. Examples of such pattern are head-and-shoulders [34], V-shaped [8,35], and W-shaped [29] patterns. From the set of these patterns, the head and shoulder pattern is defined as follows: **A** the left shoulder (increase in the stock price) that is formed by an uptrend; **B** detect the head being a peak higher than the left shoulder (one or more events); **C** the right shoulder that shows an increase but fails to take out the peak value of the head. Moreover, background KB containing the company related information, such as address, CEO, type, etc., can be joined with the matched events to extract the contextual information.

UC 5 (Traffic Management) With the growing popularity of Internet of Things (IoT) technologies, more and more cities are leaning towards the initiative of smart cities to provide robust traffic management [36,37]. A key use case in this context is to be proactive against the traffic jams due to an accident or a social event. Hence taking actions to keep the traffic jams as local as possible and stop its effect propagation to the main roads of the city [37]. A query for such use case would be: **A** If a road segment is congested followed by **B** (at least) one of its connected road segments is also congested then traffic should be directed to other directions. The information about the road segments’ congestion is gathered from a set of streams, each from a sensor located at a road segment. Further contextual information such as the exact location of the sensors, the type of the road, etc., can be extracted from an static background KB.

Using the aforementioned use cases, we identify the set of requirements for a SCEP.

2.1. RDF Graph-based Events

All the use cases discussed above provide events that contain multiple observations and hence require an RDF graph-based event model instead of RDF triples. This allows to structure more complex events in a stream, as opposed to plain triples. For instance, a weather event requires multiple triples to describe the current weather conditions such as temperature, humidity, wind speed, etc. Another reason to provide RDF graph-based events is the flexibility in timestamping, i.e. using source or system timestamp. This can only be possible if timestamps can be attached to an event structure. That cannot be achieved with plain triples [25,23].

2.2. Multiple Streams

The integration of multiple streams is essential in UC 1, 2 and 5. For instance, UC 1 detects patterns over events coming from a fuel-based power stream, a weather stream and a power-storage stream, while in UC 5 we employ multiple streams each for a road segment. Although one could merge all the streams together before processing, this would introduce the practical problems as described for RDF reification [24], where a pure triple data model is not adequate enough to represent meta-information about the RDF data [24,38]. In SPARQL, the RDF reification problem is addressed with the introduction of named graphs [24]. Furthermore, named graphs enforce the blank node scoping rules [39] with the global assumption that blank nodes cannot be shared between named graphs. Considering this, the SCEP should allow the definition of multiple named streams to provide the scope on the graph-based events. A similar approach is used in the existing RSP query languages, such as CQELS and C-SPARQL, and this can be achieved by extending the FROM NAMED clause of SPARQL.

2.3. Expressive Temporal Operators

The main aim of an SCEP language is to provide temporal operators on top of standard SPARQL graph pattern matching constructs. Thus, the list of temporal operators introduced for the CEP over relational models [4,9,40] should be supported in a SCEP language. The two main operators, apart from the traditional sequence operator, that are required in all the aforementioned use cases are *Kleene+* and event selection strategies. They are described as follows.

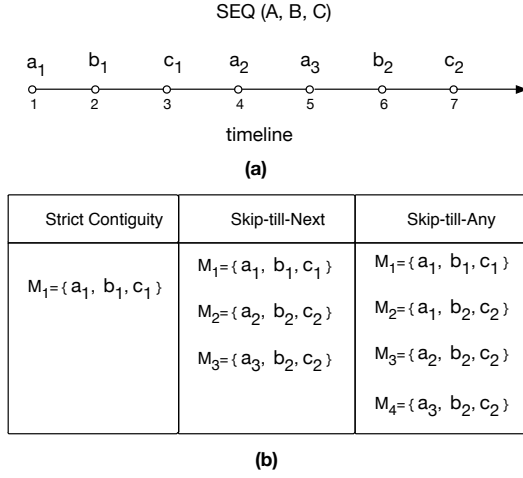


Figure 1. (a) Pattern to be matched and the input stream, (b) the table showcase all the set of matches that can be extracted from the input stream using different event selection strategies

Kleene+ Operator. This is a widely used operator in the CEP languages [5,41,8]. It is used to extract finite yet unbounded number of events with a particular property from the input streams. Most, if not all, of the aforementioned use cases cannot be expressed to its core without this operator. For example, UC 1 requires *Kleene+* to select one or more events that represent favourable weather conditions; UC 2 requires *Kleene+* on activity events of the vessels; UC 3 requires one or more lower withdrawn amounts using *Kleene+*.

Event Selection Strategies. In general, the *sequence temporal patterns* between events detect the occurrence of an event *followed-by* another. However, a deeper look reveals that it can represent various different circumstances using different event selection strategies [41,4]. Event selection strategies determine how the selected events for a pattern follow each other, hence mixing the relevant and irrelevant events from the input streams [41]. The three main event selection strategies described in the CEP literature [4,8,28] are *strict contiguity*, *skip-till-next* and *skip-till-any*. The first selection strategy states that the events for a match should be contiguous in the input stream, the second relaxes this by skipping the irrelevant events until the next relevant event arrives, while the third one selects all the possible patterns by skipping both relevant and irrelevant events. Figure 1 (a) shows the input stream and the pattern to be matched, while the table in Figure 1 (b) shows all the matches for the pattern for each event selection strategy. Note how having the same sequence

pattern but different event selection strategies results in diverse outputted matches. The importance of these operators can also be inferred from the aforementioned use cases. For instance, in UC 1, weather-related events generally repeat the observed values and can be skipped to find the most relevant event to be matched; in UC 4, we need to select the relevant events by ignoring local price fluctuations to preserve opportunities of detecting longer and thus more reliable patterns. Providing event selection strategies at the query level would enable the users to determine the semantic differences between them and tailor their usage according to the targeted use cases [23]. Furthermore, these operators are also described at the query level for most of the CEP languages [41,8].

2.4. Separation of Query Constructs

A SCEP query language has to express graph patterns to match events' content and temporal operators to determine the temporal relations between events. Hence, a separation of these two orthogonal constructs would enable language reusability and ease in implementation. This way, constructs for graph patterns can be borrowed directly from SPARQL, and temporal operators from the CEP works. For example, in UC 1, the price and weather conditions for events can be matched using the constructs from SPARQL 1.1, while the sequencing and *Kleene+* operators can be used to determine the defined sequences.

2.5. Background Knowledge Base

One of the main properties of the SCEP is to provide an extra knowledge about the situation of interest. Since the streams traditionally do not include the static information, such information can be used as a background KB. Hence a SCEP query language should provide operators to directly enrich events through background KB. For example, in UC1, user profiles and detailed information about the power sources; in UC 2, detailed information about the shipping vessels; in UC 4, detailed information about the company whose stock events are in question, etc. The operators for the background KB at the query-level also provide the control over which background information is required for a query, i.e. joining only the relevant information. For the same reason, existing RSP solutions [42,18] provide such operators at the query-level.

In this section, we pointed out various general requirements of a SCEP query language. In the next section, using them as a yardstick, we outline the limitations of existing languages.

3. Related Work

Existing languages for RDF stream processing systems differ from each other in a wide range of aspects, which include the executional semantics, data models and targeted use cases. In this section, we adopt the same classification criteria as used in [15], and divide the systems into two classes: RDF Stream Processing (RSP) systems, and Semantic Complex Event Processing (SCEP) systems. Their details are discussed as follows.

3.1. RDF Stream Processing (RSP) Systems

The standardisation of the RSP is still an ongoing debate and the W3C RSP community group [25] is an important initiative in this context. Most of the RSP systems [18,19,20,43] inherit the processing model of Data Stream Management Systems (DSMSs), but consider a semantically annotated data model, namely RDF triple streams. Query languages for these systems are inspired from CQL [44], where a continuous query is composed from three classes of operators, namely stream-to-relation (S2R), relation-to-relation (R2R), and relation-to-stream (R2S) operators. C-SPARQL [19] and CQELS [18] are among the first contributions, and often cited as a reference in this field. They support timestamped RDF triples and queries are continuously updated with the arrival of new triples. The query languages for both systems extend SPARQL with operators such as `FROM STREAM` and `WINDOW` to select the operational streams, and the most recent triples within sliding windows. They also support the integration of background KB to further enrich the incoming RDF triples. Unlike the aforementioned systems, recently, we proposed a system called SPECTRA [45] to process RDF graph streams. It provides various system and operator level optimisations and continuously processes the standard SPARQL queries over RDF graph streams.

All the aforementioned systems and various others [20,43] are mainly developed as real-time monitoring systems: the states of the events are not stored to implement temporal pattern matching among a set of events. For the same reason, their query languages do

not provide any operators for temporal pattern matching. Although, the sequence-based query is partially supported in C-SPARQL through a timestamp function in `Filter`, the query construction is not intuitive and the validity of results cannot be determined due to the missing semantics for such sequence pattern matching.

3.2. Semantic CEP Systems

Semantic CEP (SCEP) systems are inspired from the classical rule-based CEP systems, i.e. by integrating high-level knowledge representation and background KB. To the best of our knowledge, EP-SPARQL is the only system that provides a unified language and executional framework for processing semantically enriched events with the temporal ordering. Hence, it is the most relevant work w.r.t ours. EP-SPARQL extends SPARQL 1.0 with a set of four binary temporal operators: `SEQ`, `EQUALS`, `OPTIONAL-SEQ`, and `EQUALS-OPTIONAL`. Using these operators, complex events in EP-SPARQL are defined as basic graph patterns. Since this technique is similar to the `UNION` and `OPTIONAL` operators in SPARQL, events are not first class citizens.

Although EP-SPARQL is a pioneering work in the field of SCEP, it lacks various important features. These limitations are discussed as follows:

- *Triple-based Events*: EP-SPARQL only supports triples as events annotated with time-intervals. This restricts to structure more complex events as discussed before.
- *Single Stream*: EP-SPARQL data model is based on a single stream model and all the triples within a defined window are merged into a default graph and then queried for matches. Hence the meta-information of triples, such as the source and occurring time, cannot be captured. Furthermore, it raises several questions about the treatment of blank nodes: being part of the RDF specification, blank nodes are now a core aspect of Semantic Web technology and they are featured in several W3C standards, a wide range of tools, and hundreds of datasets across the Web [39].
- *Temporal Operators*: EP-SPARQL only supports a small subset of temporal operators. Operators such as *Kleene+* or event selection strategies are not supported. These operators are, however, important for many use cases as discussed before. Moreover, the conjunction and disjunction operators in EP-SPARQL are inspired from SPARQL

Table 1
Properties of the Existing SCEP Systems

SCEP Systems	Input Model	Operators	Available Implementation
EP-SPARQL [21]	Triple Streams	Sequence, Conjunction, Disjunction, Optional	✓
STARQL [22]	Triple Streams	Sequence	✗
RSEP-QL [23]	Graph Streams	Sequence, Event Selection Strategies	✗
CQELS-CEP [46]	Graph Streams	Sequence, Optional, Negation	✗

(UNION and AND). These operators do not provide the nesting over a set of events as described for CEP systems [5,40]. This leads to a design where the semantics of temporal operators and SPARQL graph patterns are mixed, hence it can be devious to extend it and construct advanced event processing patterns [47].

- *Enriching Events with Background Knowledge Base*: The static background KB is used to extract further implicit information from events. As a query language, EP-SPARQL does not provide any explicit operator to join graph patterns defined on an external knowledge and incoming RDF events. It, however, employs Prolog rules or RDFS rules within an ETALIS engine. This results in expensive reasoning process for each incoming event and the user is not able to select only the specific required information [15].

Apart from EP-SPARQL, recently, some other works also provide the intuition of SCEP. Some of these works are presented mainly for the purpose of theoretical analysis instead of practical implementation, while others take an approach for transforming queries over ontologies into relational ones, via ontology-based data access (OBDA) with temporal reasoning. Table 1 shows the properties of these systems. STARQL [48] uses the OBDA technique to determine ABox sequencing with a sorted first order logic on top of them. It provides a simple formalism/mapping to SQL for sequence operators and all the other operators (such as *Kleene+*, conjunction, disjunction, event selection strategies) are not part of its framework. Furthermore, it is not a freely available system and it does not provide operators for explicitly referencing different points in time [22]. CQELS recently proposed in [46] the integration of sequence and path negation operators inspired from EP-SPARQL. However, its sequence clause is evaluated over a single stream and its syntax and semantics does

not include event selection strategies. RSEP-QL [23] is a reference model to capture the behaviour of existing RSP solutions and to capture the semantics of EP-SPARQL's sequence operator. It is based on the RDF graph model; however, its main focus is to capture the event selection strategies and other complex operators are currently not supported. It defines the semantics of the three main event selection strategies supported by the ETALIS engine. It includes *recent*, *chronological* and *unrestricted*. The recent selection strategy can be mapped to the strict contiguity, while chronological and unrestricted are different variations of the skip-till-next strategy. However, due to the time-interval semantics these selection strategies cannot be directly mapped to ours.

4. The SPASEQ Query Language

Considering the shortcomings of existing languages, we propose a new language called SPASEQ. The design of SPASEQ is based on the following main principles: (1) support of an RDF graph stream model; (2) clear separation between the temporal and RDF graph operators; (3) adequate expressive power, i.e. not only based on core SPARQL constructs but also including general purpose temporal operators (inspired from the common CEP operators); (4) genericity, i.e. independent of the underlying evaluation techniques; (5) compositionality, i.e. the output of a query can be used as an input for another one. Hence SPASEQ provides all the required features as discussed in Section 2.

The most important feature of SPASEQ is that it clearly separates the query components for describing temporal patterns over RDF graph events, from specifying the graph pattern matching over each RDF graph event. This enables SPASEQ to employ expressive temporal operators, such as *Kleene+*, disjunction, conjunction and event selection strategies over RDF graph-

based events from multiple streams. In the following, we start with the data model over which SPASEQ queries are processed and then provide the details regarding its syntax and semantics.

4.1. Data Model

In this section, we introduce the structural data model of SPASEQ that captures the concept of RDF graph events: this serves as the basis of our query language. We use the RDF data model [49] to model an event. That is, we assume three pairwise disjoint and infinite sets of IRIs (\mathcal{I}), blank nodes (\mathcal{B}), and literals (\mathcal{L}). An RDF triple is a tuple $\langle s, p, o \rangle \in (\mathcal{I} \cup \mathcal{B}) \times \mathcal{I} \times (\mathcal{I} \cup \mathcal{B} \cup \mathcal{L})$ and an RDF graph is a set of RDF triples. Based on this, the concepts of RDF graph event and stream are defined as follows.

Definition 1 An **RDF graph event** (G_e) is a pair (τ, G) where G is an RDF graph, and τ is an associated timestamp that belongs to a one-dimensional, totally ordered metric space.

We do not make explicit what timestamps are because one may rely on, e.g., UNIX epoch, which is a discrete representation of time, while others could use `xsd:dateTime` which is arbitrarily precise.

In our setting, *streams* are sets of RDF graph events defined as follows:

Definition 2 An **RDF graph event stream** \mathcal{S} is a possibly infinite set of RDF graph events such that, for any given timestamps τ and τ' , there is a finite amount of events occurring between them, and there is at most one single graph associated with any given timestamp.

An RDF graph event stream can be seen as a sequence of chronologically ordered RDF graphs marked with timestamps. The constraints ensure that it is always possible to determine what unique event immediately precedes or succeeds a given timestamp. Without the first restriction, it would be possible to define a stream $\{(\frac{1}{n}, G) \mid n \neq 0 \text{ an integer}\}$ where there is no event immediately succeeding 0. In order to handle multiple streams, we identify each using an IRI and group them in a data model we call *RDF streamset*.

Definition 3 A **named stream** is a pair (u, \mathcal{S}) where u is an IRI, called the stream name, and \mathcal{S} is an RDF graph event stream. An **RDF graph streamset** Σ is a set of named streams such that stream names appear only once.

In the rest of the paper, we simply use the terms *graph* for RDF graph, *event* for RDF graph event, *stream* for RDF graph stream, and *streamset* for RDF graph streamset.

Example 1 Recall *UC 1*, here we extend it with our data model using three named streams. The first named stream (u_1, \mathcal{S}_1) provides the events about the power-related sources from a house, the second named stream (u_2, \mathcal{S}_2) provides the weather-related events for house, and the third named stream (u_3, \mathcal{S}_3) provides the power storage-related events. Figure 2 illustrates the general structure of the events from each source. An exemplary content of a named stream (u_1, \mathcal{S}_1) is describe as follow:

time	graph
10	:H1 :loc :L1 :H1 :pow :Pw1 :Pw1 :source 'solar' :Pw1 :fare 5 :Pw1 :watt 20
15	:H2 :loc :L2 :H2 :pow :Pw2 :Pw2 :source 'wind' :Pw2 :fare 6 :Pw2 :watt 20
...	...

Note that the above table describe the named stream in an informative way, in practice RDF 1.1 Trig or an Nquid format is used to represent such streams.

4.2. Syntax of SPASEQ

This section defines the abstract syntax of SPASEQ, where SPASEQ queries are meant to be evaluated over a streamset, and each query is built from the two main components: *graph pattern matching expression* (GPM) for specifying the SPARQL graph patterns over events; and *sequence expression* for selecting the sequence of a set of GPM expressions. For this discussion, we assume that the reader is familiar with the definition and the algebraic formalisation of SPARQL introduced in [50]. In particular, we rely on the notion of SPARQL graph pattern by considering operators AND, OPT, UNION, FILTER, and GRAPH.

Definition 4 The syntax of a SPASEQ *SELECT* query is a tuple $Q = (\mathcal{V}, \omega, \text{SeqExp})$, where \mathcal{V} is a set of vari-

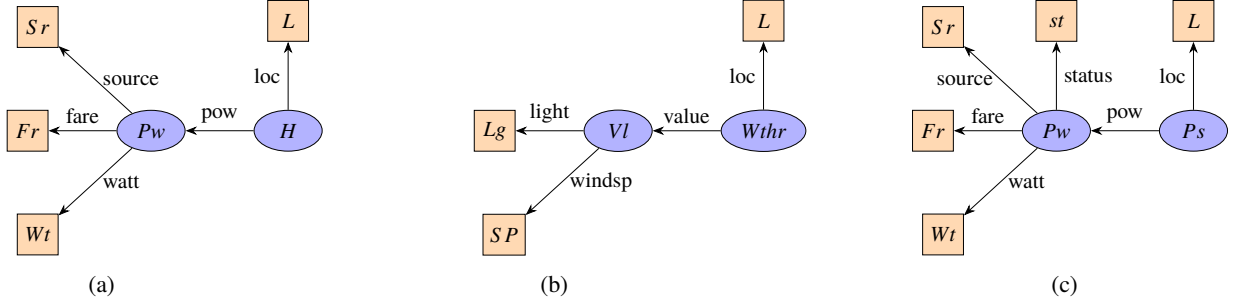


Figure 2. Structure of the Events from three Named Streams, (2a) (u_1, \mathcal{S}_1) Power Event, (2b) (u_2, \mathcal{S}_2) Weather Event, (2c) (u_3, \mathcal{S}_3) Power Storage Event

ables, ω is a duration, and SeqExp is a sequence expression defined according to the following grammar:

$$\begin{aligned} \text{SeqExp} &::= \text{Atom} \mid \text{SeqExp} \text{' ; ' } \text{Atom} \mid \text{SeqExp} \text{' ; ' } \text{Atom} \\ &\quad \mid \text{SeqExp} \text{' : ' } \text{Atom} \\ \text{Atom} &::= \text{GPM} \mid \text{GPM} \text{' + ' } \mid \text{BOP} \\ \text{BOP} &::= \text{GPM} \text{' \& ' } \mid \text{GPM} \text{' | ' } \\ \text{GPM} &::= (u, P) \mid (u, P) \text{ Graph } (u, P_D) \end{aligned}$$

where u is an IRI, P is a SPARQL graph pattern, (u, P) is called a graph pattern matching expression (GPM), and P_D is a SPARQL graph pattern defined for a static RDF graph, i.e. an external knowledge-base.

The concrete syntax of SPASEQ is illustrated in Query 1 which includes syntactic sugar that is close to SPARQL. It contains three GPM expressions each identified with a name (A, B, and C), which allows one to concisely refer to GPMs and to the named streams. These names are employed by the sequence expression to apply various temporal operators. The sequence expression in Query 1 is presented at line 9; the streams are described at lines 3-5; the GPM expressions on these streams start at lines 11, 19 and 27.

One of the main properties of the SPASEQ language is depicted in Query 1, i.e. the separation of sequence and GPM expressions. Herein, we first study how the sequence expression interacts with the graph pattern to enable temporal ordering between matched events. We start with the brief description of unary operator ($\{ '+' \}$), the event selection strategies ($\{ \text{' ; ' }, \text{' ; ' }, \text{' : ' } \}$) and binary operators ($\{ \text{' \& ' }, \text{' | ' } \}$). The details of these operators are covered during the description of their semantics in Section 5. Furthermore, we also present the Graph operator for the SPASEQ query language.

4.2.1. SPASEQ Unary Operators

The sequence expression SeqExp in SPASEQ is used to determine the sequence between the events matched

to the graph pattern P . The symbol $\{ '+' \}$ corresponds to the Kleene+ operator. It determines the occurrence of one or more events of the same kind. This means a series of events can be matched using this operator.

```

1 SELECT ?house ?fr1 ?fr2 ?house2 ?city ?oname
2 WITHIN 30 MINUTES
3 FROM STREAM S1 <http://smartgrid.org/mainSource>
4 FROM STREAM S2 <http://smartgrid.org/weather>
5 FROM STREAM S3 <http://smartgrid.org/storageSource>
6
7 WHERE {
8
9   SEQ (A; B+, C)
10
11  DEFINE GPM A ON S1 {
12    ?house :loc ?l.
13    ?house :pow :Pw.
14    :Pw :source "fuel".
15    :Pw :fare ?fr1.
16    FILTER(?fr1 > 20).
17
18    GRAPH <http://smartgrid.org/db> {
19      ?house :locatedIn ?city.
20      ?house :nearBy ?house2.
21      ?house :ownerName ?oname.
22    }
23  }
24
25  DEFINE GPM B ON S2 {
26    ?wthr :loc ?l.
27    ?wthr :value :Vl.
28    :Vl :light ?lt.
29    :Vl :windsp ?sp.
30    FILTER (?sp > 3 && ?lt > 40).
31  }
32
33  DEFINE GPM C ON S3 {
34    ?storage :loc ?l.
35    ?storage :pow :Pw.
36    :Pw :source "solar".
37    :Pw :fare ?fr2.
38    FILTER (?fr2 < ?fr1).
39  }
40 }

```

Query 1: A Sample SPASEQ Query for the UC 1

4.2.2. SPASEQ Event Selection Strategies

The event selection strategies overload the sequence operator with the constraints to define how to select the relevant events from an input stream, while mixing relevant and irrelevant events (as described in Section 2.3). The symbols $\{ \text{' ; ' }, \text{' ; ' }, \text{' : ' } \}$ are binary opera-

tors which describe the interpretations of the sequence between events. That is, ‘;’ represents strict contiguity, ‘;’ represents skip-till-next, ‘:’ represents skip-till-any.

4.2.3. SPASEQ Binary Operators

Conjunction and disjunction defined over the event streams constitute the binary operators. In SPASEQ, these operators are introduced within the sequence expression through symbols (‘&’) and (‘|’) respectively. They provide the intuitive way of determining if a set of events happen at the same time (conjunction) or at least one event among the set of events happens (disjunction).

Example 2 Consider the SPASEQ Query 1, which illustrates the UC 1. The sequence expression SEQ(A; B+, C) illustrates that the query returns a match: if an event of type A defined on a stream S1 matches the GPM expression A followed-by one or more events (using skip-till-next operator (‘;’) and (‘+’)) from stream S2 that match the GPM expression B, and finally immediately followed-by (using strict contiguity operator (‘;’)) an event from stream S3 that matches the GPM expression C. Notice that a GPM expression mainly utilises SPARQL graph patterns for the evaluation of each event.

4.2.4. SPASEQ Graph, Window and Stream Operator

The combination of streaming information in the form of RDF graph streams and other information from the static knowledge base can lead to novel semantics and information rich SCEP. The Graph operator in SPASEQ is designed to take advantage of static information available in the form of an RDF graph. Thus, a graph pattern P_D defined over the static RDF graph G_D is first evaluated and then the results are matched with the incoming stream S . This leads to a SCEP system, where detailed information regarding a context can be revealed with the help of already available static datasets. For instance, consider the SPASEQ Query 1, where we use static background KB about the house owners, their neighbours and city to have better understanding of the matched events.

In SPASEQ, the sequence expression is defined over a streamset. Thus, we use the FROM STREAM clause to define a set of streams. For instance, in Query 1 we use three streams identified as S1, S2 and S3 (lines 3-5). These stream names are used within the defined GPM expressions. Furthermore, since the sequence over a set of events is constrained by the temporal window, we use the WITHIN clause to define the temporal windows (line 2 in Query 1). In SPASEQ, windows can be defined

in seconds, minutes and hours. For instance, in Query 1 we use a 60 MINUTES window.

In this section, we presented the syntax of SPASEQ query language with a query for the UC 1. The queries for the remaining use cases are described in Appendix A. Using the defined syntax, we present the semantics of SPASEQ in the proceeding section.

5. Formal Semantics of SPASEQ

To formally define the semantics of SPASEQ queries, we reuse concepts from the semantics of SPARQL as defined in [50]. A mapping is a partial function from a set of variables to RDF terms ($\mathcal{B} \cup \mathcal{I} \cup \mathcal{L}$). The domain of a mapping μ , denoted $\text{dom}(\mu)$, is the set of variables that have an image via μ . We say that two mappings μ and μ' are compatible if they agree on all shared variables, i.e. if $\mu(x) = \mu'(x)$ for all $x \in \text{dom}(\mu) \cap \text{dom}(\mu')$. For a graph pattern P , we denote by $\text{vars}(P)$ the set of variables appearing in P and $\mu(P)$ is the graph pattern obtained by replacing each variable $v \in \text{vars}(P)$ by $\mu(v)$ whenever defined.

We repeat the definitions of join (\bowtie), union (\cup), minus (\setminus), left outer-join (\ltimes) and evaluation of graph patterns as in [50].

Definition 5 Let Ω_1 and Ω_2 be sets of mappings:

$$\begin{aligned} \Omega_1 \bowtie \Omega_2 &= \{ \mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1, \mu_2 \in \Omega_2, \mu_1 \text{ and } \mu_2 \text{ compatible} \} \\ \Omega_1 \cup \Omega_2 &= \{ \mu \mid \mu \in \Omega_1 \text{ or } \mu \in \Omega_2 \} \\ \Omega_1 \setminus \Omega_2 &= \{ \mu_1 \in \Omega_1 \mid \text{for all } \mu_2 \in \Omega_2, \mu_1 \text{ and } \mu_2 \text{ not compatible} \} \\ \Omega_1 \ltimes \Omega_2 &= (\Omega_1 \bowtie \Omega_2) \cup (\Omega_1 \setminus \Omega_2) \end{aligned}$$

Definition 6 Let t be a triple pattern, P, P_1, P_2 graph patterns and G an RDF graph, then the evaluation $\llbracket \cdot \rrbracket_G$ is recursively defined as follows:

$$\begin{aligned} \llbracket t \rrbracket_G &= \{ \mu \mid \text{dom}(\mu) = \text{vars}(t) \text{ and } \mu(t) \in G \} \\ \llbracket P_1 \text{ AND } P_2 \rrbracket_G &= \llbracket P_1 \rrbracket_G \bowtie \llbracket P_2 \rrbracket_G \\ \llbracket P_1 \text{ UNION } P_2 \rrbracket_G &= \llbracket P_1 \rrbracket_G \cup \llbracket P_2 \rrbracket_G \\ \llbracket P_1 \text{ OPTIONAL } P_2 \rrbracket_G &= \llbracket P_1 \rrbracket_G \ltimes \llbracket P_2 \rrbracket_G \\ \llbracket P \text{ FILTER } R \rrbracket_G &= \{ \mu \in \llbracket P \rrbracket_G \mid \mu(R) \text{ is true} \} \end{aligned}$$

In this section, we define the semantics of SPASEQ in a bottom-up manner, where we start with the semantics of GPM expressions by integrating the temporal aspects of events and streams. Note that, for the sake of brevity, we show the evaluation of GPM expressions over a streamset and the aspects of evaluating Graph operator (over RDF dataset) within GPM expressions are discussed later. This will aid us in highlighting the decisions we took to define the semantics of SPASEQ operators.

5.1. Evaluation of Graph Pattern Matching Expressions

In SPASEQ, Graph Pattern Matching expressions are evaluated against a streamset over a finite time interval that “tumbles” as time passes. Consequently, we constrain the evaluation function to a temporal boundary (i.e. a window), with a start time (τ_b) and an end time (τ_e). In addition, we use the notation $\Sigma(u)$ to select a stream of name u from a streamset, such that

$$\Sigma(u) = \begin{cases} \mathcal{S} & \text{if } (u, \mathcal{S}) \in \Sigma \\ \emptyset & \text{otherwise} \end{cases}$$

In order to define the evaluation of a GPM expression, we introduce a pair $(\tau, \llbracket P \rrbracket_G)$ which represents a set of mappings annotated with a timestamp τ of the matched event. The evaluation of a GPM expression is defined as follows.

Definition 7 The *evaluation of a GPM expression* (u, P) over the streamset Σ for the time boundaries $[\tau_b, \tau_e]$ is:

$$\llbracket (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \{(\tau, \llbracket P \rrbracket_G) \mid (\tau, G) \in \Sigma(u) \wedge \tau_b \leq \tau \leq \tau_e \wedge \llbracket P \rrbracket_G \neq \emptyset\}$$

The evaluation of the GPM expression (u, P) over an event within a streamset results in a pair $(\tau, \llbracket P \rrbracket_G)$. In the absence of a match, no results are returned for the considered timestamp.

Example 3 Consider a GPM expression $(u_1, P_1) := (u_1, \{(?h, pow, ?p), (?h, loc, ?l)\})$ and a power-related named stream $(u_1, \mathcal{S}_1) \in \Sigma$ with events as follows:

time	graph
10	:H1 :pow :Pw1 :H1 :loc :L1
15	:H2 :pow :Pw2 :H2 :loc :L2
25	:H3 :pow :Pw3 :H3 :loc :L3
...	...

The evaluation of (u_1, P_1) over Σ for the time boundaries $[5, 15]$, i.e. $\llbracket (u_1, P_1) \rrbracket_{\Sigma}^{[5, 15]}$, is described as follows:

time	?h	?p	?l
10	:H1	:Pw1	:L1
15	:H2	:Pw2	:L2

Notice that since the end time of the window is restricted at $\tau = 15$, only the events at $\tau = 10$ and $\tau = 15$ are included in the result. The event at $\tau = 25$ is outside the window and thus is not included in the results.

In order to define the semantics of sequence expressions, we use the notion of BOp expressions from Definition 4 for conjunction and disjunction operators. A BOp expression does not contain Kleene+ or event selection operators.

5.2. Evaluation of Binary Operators

Herein, we define the semantics of binary operators provided for SPASEQ, i.e. conjunction and disjunction of events.

Definition 8 Given two BOp expressions Ψ_1, Ψ_2 the *evaluation of the conjunction operator* over the streamset Σ and for the time boundaries $[\tau_b, \tau_e]$ is defined as follows:

$$\llbracket \Psi_1 \& \Psi_2 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \left\{ \begin{array}{l} (\tau, X \bowtie Y) \mid (\tau, X) \in \llbracket \Psi_1 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \\ (\tau, Y) \in \llbracket \Psi_2 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge X \bowtie Y \neq \emptyset \end{array} \right\}$$

The conjunction operator detects the presence of two or more events that match the defined GPM expressions and occur at the same time, i.e. containing the same timestamps.

Example 4 Consider the following, a GPM expression $(u_1, P_1) := (u_1, \{(?h, pow, ?p), (?h, loc, ?l)\})$ and a power-related named stream $(u_1, \mathcal{S}_1) \in \Sigma$ as follows:

time	graph
10	:H1 :pow :Pw1 :H1 :loc :L1
25	:H2 :pow :Pw2 :H2 :loc :L2
...	...

Now consider a GPM expression $(u_2, P_2) := (u_2, \{(?w, value, ?v), (?w, loc, ?l)\})$ and a weather-related named stream $(u_2, \mathcal{S}_2) \in \Sigma$ as follows:

time	graph
10	:W1 :value :V11 :W1 :loc :L1
20	:W2 :value :V12 :W2 :loc :L1
...	...

The evaluation of the conjunction operator over the aforementioned GPM expressions and named streams ($\llbracket (u_1, P_1) \& (u_2, P_2) \rrbracket_{\Sigma}^{[10,25]}$) for the time boundaries $[10,25]$ will result in the following sets of mappings.

time	?h	?p	?l	?w	?v
10	:H1	:Pw1	:L1	:W1	:V11

We now define the disjunction operator. It detects the occurrence of events that match to a GPM expression within the set of defined ones.

Definition 9 Given two BOp expressions Ψ_1 and Ψ_2 , the **evaluation of the disjunction operator** over the streamset Σ and for the time boundaries $[\tau_b, \tau_e]$ is defined as follows:

$$\llbracket \Psi_1 \mid \Psi_2 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \llbracket \Psi_1 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \cup \llbracket \Psi_2 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]}$$

Example 5 Consider the two GPM expressions (u_1, P_1) and (u_2, P_2) , and the two named streams (u_1, S_1) and $(u_2, S_2) \in \Sigma$ from Example 4. Then the evaluation of the disjunction sequence operator for the sequence expression $\llbracket ((u_1, P_1) \mid (u_2, P_2)) \rrbracket_{\Sigma}^{[10,25]}$, and for the time boundaries $[10,25]$ is as follows:

time	?h	?p	?l	?w	?v
10	:H1	:Pw1	:L1		
10			:L1	:W1	:V11
15			:L1	:W1	:V11
20			:L1	:W2	:V12
25	:H2	:Pw2	:L2		

Notice that the disjunction operator may generate several sets of mappings for the same timestamp, as we can see at time 10 in the example.

5.3. Evaluation of Event Selection Operators

Herein we present the semantics of the three event selection strategies namely strict contiguity, skip-till-next and skip-till-any.

Let σ be a sequence with a set of GPM expressions and binary/unary operators. The evaluation of the skip-till-any operator is defined as follows:

Definition 10 Given a sequence σ and a BOp expression Ψ , the **evaluation of the skip-till-any** (‘:’) sequence operator over a streamset Σ for the time boundaries $[\tau_b, \tau_e]$ is defined as follows:

$$\llbracket \sigma ; \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \left\{ \begin{array}{l} (\tau, X \bowtie Y) \mid \exists \tau', (\tau, Y) \in \llbracket \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \\ (\tau', X) \in \llbracket \sigma \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \tau' < \tau \wedge X \bowtie Y \neq \emptyset \end{array} \right\}$$

From the above definition, the evaluation of the skip-till-any operator is simply the join between the mapping sets from σ and the GPM expression. Its evaluation is explained in the following example.

Example 6 Consider the following, a GPM expression $(u_1, P_1) := (u_1, \{(?h, pow, ?p), (?h, loc, ?l)\})$ and a power-related named stream $(u_1, S_1) \in \Sigma$ as follows:

time	graph
10	:H1 :pow :Pw1 :H1 :loc :L1
15	:H2 :pow :Pw2 :H2 :loc :L1
...	...

A GPM expression as follows:

$$(u_2, P_2) := (u_2, \{(?w, value, ?v), (?w, loc, ?l)\})$$

and a weather-related named stream $(u_2, S_2) \in \Sigma$ as follows:

time	graph
15	:W1 :value :V11 :W1 :loc :L2
20	:W1 :value :V11 :W1 :loc :L1
25	:W2 :value :V12 :W2 :loc :L1
...	...

Then for the evaluation of the skip-till-any operator on these GPM expressions for the time boundaries $[10,25]$, i.e. $\llbracket (u_1, P_1); (u_2, P_2) \rrbracket_{\Sigma}^{[10,25]}$, we have the mappings as given below.

time	?h	?p	?l	?w	?v
20	:H1	:Pw1	L1	:W1	:V11
25	:H1	:Pw1	L1	:W2	:V12
20	:H2	:Pw2	L1	:W1	:V11
25	:H2	:Pw2	L1	:W1	:V11

Notice that there are four matches sequences: both the first and second events (at $\tau = 10$ and $\tau = 15$) from the power-related named stream matched with both events ($\tau = 20$ and $\tau = 25$) in the weather-related named stream. This is due to the skip-till-any operator and all the possible combinations of matches are produced.

The following definition shows the evaluation of the skip-till-next operator.

Definition 11 Given a sequence σ and an BOp expression Ψ , the **evaluation of the skip-till-next** ($;$) sequence operator over a streamset Σ for the time boundaries $[\tau_b, \tau_e]$ is defined as follows:

$$\llbracket \sigma; \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \left\{ \begin{array}{l} (\tau, X \bowtie Y) \mid \exists \tau', (\tau', X) \in \llbracket \sigma \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \tau' < \tau \wedge \\ (\tau, Y) \in \llbracket \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge X \bowtie Y \neq \emptyset \\ \forall \tau'' \forall Z \left((\tau'', Z) \in \llbracket \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge (\tau' < \tau'' < \tau) \right) \\ \Rightarrow X \bowtie Z = \emptyset \end{array} \right\}$$

Example 7 Consider the GPM expressions (u_1, P_1) , (u_2, P_2) and the streams from Example 6. Then for the evaluation of the skip-till-next operator on these GPM expressions for the time boundaries $[10, 25]$, i.e. $\llbracket (u_1, P_1); (u_2, P_2) \rrbracket_{\Sigma}^{[10, 25]}$, we have the mappings as given below.

time	?h	?p	?l	?w	?v
20	:H1	:Pw1	L1	:W1	:V11
20	:H2	:Pw2	:L1	:W1	:V11

Due to the skip-till-next operator, there are only two matches in the aforementioned example. Both events (at $\tau = 10$ and $\tau = 15$) from the power-related stream match with just one event (at $\tau = 20$) from the weather-related stream.

We now define the semantics of the strict contiguity operator, where U is a set of stream names within a streamset Σ .

Definition 12 Given a sequence σ and a BOp expression Ψ , the **evaluation of the strict contiguity** (\bullet) sequence operator over a streamset Σ for the time boundaries $[\tau_b, \tau_e]$ is defined as follows:

$$\llbracket \sigma, \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \left\{ \begin{array}{l} (\tau, X \bowtie Y) \mid \exists \tau', (\tau', X) \in \llbracket \sigma \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \\ (\tau, Y) \in \llbracket \Psi \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \wedge \tau' < \tau \wedge X \bowtie Y \neq \emptyset \wedge \\ \forall u \forall \tau'' \forall G \left((\tau'', G) \in \Sigma(u) \wedge \tau'' > \tau' \Rightarrow \tau'' \geq \tau \right) \end{array} \right\}$$

The semantics of the strict contiguity operator follows the semantics of the skip-till-next operator, however with one important difference: the contiguity between the matched events. That is, an event is contiguous to another, only if there can be no other events between the two selected ones.

Example 8 Consider the GPM expressions and the named streams defined in Example 6. Then the evaluation of the strict contiguity operator described as $\llbracket (u_1, P_1), (u_2, P_2) \rrbracket_{\Sigma}^{[10, 25]}$, for time boundaries $[10, 25]$, will result in a single sequence match with the following mappings.

time	?h	?p	?l	?w	?v
20	:H2	:Pw2	:L1	:W1	:V11

This is due to the strict ordering of the strict contiguity operator. That is, for first event (at $\tau = 10$) in power-related stream, there is no immediately followed-by event in the weather-related stream.

5.4. Evaluation of Kleene+ Operator

We now move towards the definitions of the unary operator, namely Kleene+. We first define its semantics in a standalone manner and then recursively define it with the help of sequence σ .

Definition 13 The **evaluation of the standalone Kleene+** operator over the streamset Σ and for the time boundaries $[\tau_b, \tau_e]$ is defined using auxiliary constructs \cdot^k for integers $k > 0$ as follows:

$$\begin{aligned} \llbracket (u, P)^1 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \llbracket (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \\ \llbracket (u, P)^{k+1} \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \llbracket (u, P)^k; (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \\ \llbracket (u, P)^+ \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \bigcup_{k \in \mathbb{N}^*} \llbracket (u, P)^k \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \end{aligned}$$

The Kleene+ operator groups all the matched events with the defined GPM expression. Notice that the evaluation will not only match the longest sequence of matching patterns, but will also provide results for the shorter sequences (using the skip-till-next operator ($;$)). The case of the Kleene+ operator with sequence σ using additional skip-till-next and strict contiguity is defined as follows:

Definition 14 Let \bullet denote any of $;$, \bullet , or \cdot . Given a sequence expression σ , the **evaluation of the Kleene+** operator in a sequence over the streamset Σ and for the time boundaries $[\tau_b, \tau_e]$ are defined as follows:

$$\begin{aligned} \llbracket \sigma \bullet (u, P)^1 \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \llbracket \sigma \bullet (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \\ \llbracket \sigma \bullet (u, P)^{k+1} \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \llbracket \sigma \bullet (u, P)^k; (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \\ \llbracket \sigma \bullet (u, P)^+ \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} &= \bigcup_{k \in \mathbb{N}^*} \llbracket \sigma \bullet (u, P)^k \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \end{aligned}$$

Example 9 Consider the following, a GPM expression $(u_1, P_1) := (u_1, \{(?h, pow, ?p), (?h, loc, ?l)\})$ and a power-related named stream $(u_1, S_1) \in \Sigma$ as follows:

time	graph
10	:H1 :pow :Pw1 :H1 :loc :L1
25	:H2 :pow :Pw2 :H2 :loc :L2
...	...

A GPM expression

$$(u_2, P_2) := (u_2, \{(?w, value, ?v), (?w, loc, ?l)\})$$

and a weather-related named stream $(u_2, S_2) \in \Sigma$ as follows:

time	graph
15	:W1 :value :V11 :W1 :loc :L1
20	:W2 :value :V12 :W2 :loc :L1
...	...

The evaluation of the sequence $\llbracket (u_1, P_1); (u_2, P_2) \rrbracket_{\Sigma}^{[10,20]}$ with the Kleene+ and skip-till-next operators for the time boundaries $[10,20]$ is as follows:

time	?h	?p	?l	?w	?v
15	:H1	:Pw1	:L1	:W1	:V11
20	:H1	:Pw1	:L1	:W2	:V12

Notice that the Kleene+ operator collects one or more matches for (u_2, P_2) from the weather-related named stream.

5.5. The Graph Operator

The *Graph* operator within GPM expressions allows one to query both background KB and streams. In the previous sections, we omitted this construct because it needlessly makes the notations cumbersome: it would require adding an RDF dataset (in addition to a streamset) as a parameter of the evaluation function.

However, for completeness, we present the definition of the evaluation of the *Graph* operator. Now similarly to the Definition 7, we use a function $\Gamma(u)$ to select an RDF graph of name u from an RDF dataset D , such that:

$$\Gamma(u) = \begin{cases} \emptyset & \text{if } u \text{ is not a graph name in } D \\ G_D & \text{if } (u, G_D) \in D \end{cases}$$

Definition 15 Let D be an external RDF dataset and (v, P_D) be a graph pattern. Let (u, P) be the GPM expression defined over the streamset Σ , and $[\tau_b, \tau_e]$ be the time boundaries. The evaluation of the GPM expression and the Graph operator is defined as follows:

$$\llbracket (u, P) \text{ Graph } (v, P_D) \rrbracket_{(\Sigma, D)}^{[\tau_b, \tau_e]} = \left\{ (\tau, \llbracket P \rrbracket_G \bowtie \llbracket P_D \rrbracket_{G_D} \mid \exists \tau(\tau, G) \in \Sigma(u) \wedge \exists G_D \in \Gamma(v) \wedge \tau_b \leq \tau \leq \tau_e) \right\}$$

Example 10 Consider the same GPM expression (u_1, P_1) and power-related named stream $(u_1, S_1) \in \Sigma$ presented in Example 9. Now consider a graph pattern $(u_D, P_D) = (u_D, \{(?h, owner, ?n), (?h, address, ?a)\})$ defined over an external RDF graph D as follows:

G_D		
:H1	:owner	:john
:H1	:address	:paris
:H2	:owner	:smith
:H2	:address	:lyon

The evaluation of the GPM expression and the Graph operator $\llbracket (u_1, P_1) \text{ Graph } (u_D, P_D) \rrbracket_{(\Sigma, D)}^{[10,20]}$ is as follows:

time	?h	?p	?l	?n	?a
10	:H1	:Pw1	:L1	:john	:paris

5.6. Evaluation of SPASEQ Queries

In the previous sections, we outline the semantics of temporal operators and *Graph* operator of the SPASEQ query language. Herein, to sum it up, we present the evaluation of complete SPASEQ SELECT queries.

Definition 16 Let Ω be a mapping set, $\pi_{\mathcal{V}}$ be the standard SPARQL projection on the set of variables \mathcal{V} , and ω be the duration of the window. The evaluation of SPASEQ SELECT query $Q = (\mathcal{V}, \omega, \text{SeqExp})$ issued at time t , over the streamset Σ is defined as follows:

$$\llbracket Q \rrbracket_{\Sigma}^t = \bigcup_{k \in \mathbb{N}} \left\{ (\tau, \pi_{\mathcal{V}}(\Omega)) \mid (\tau, \Omega) \in \llbracket \text{SeqExp} \rrbracket_{\Sigma}^{[t+k \cdot \omega, t+(k+1) \cdot \omega]} \right\}$$

where

$$\pi_{\mathcal{V}}(\Omega) = \left\{ \mu_1 \mid \exists \mu_2 : \mu_1 \cup \mu_2 \in \Omega \wedge \text{dom}(\mu_1) \subseteq \mathcal{V} \wedge \text{dom}(\mu_2) \cap \mathcal{V} = \emptyset \right\}$$

The evaluation of the SPASEQ queries follows a push-based semantics, i.e. results are produced as soon as the sequence expression matches to the set of events within the streamset. Thus, the resulting set of mappings takes the shape of a stream of mappings, where the order within the mappings depends on the underlying executional framework. Note that the definition of $\llbracket Q \rrbracket_{\Sigma}^t$ is the intended one. It could be possible to define a continuous version of the query evaluation but we want to stay agnostic to how the solutions are provided. For instance, the evaluation could be performed on a static file with time series, possibly including future previsions; or the solutions could be provided in bulks every ω time units. Moreover, according to Definition 16, SPASEQ uses tumbling windows. This choice of windows is inspired from the classical CEP works and it is a specific design choice of SPASEQ. This enable us to focus on the other important aspects of the language and borrow existing techniques for windows. However, one could extend it for sliding windows without changing the core semantics of SPASEQ.

Example 11 Recall the two GPM expressions from Example 6, $(u_1, P_1) := (u_1, \{(?h, pow, ?p), (?h, loc, ?l)\})$ and $(u_2, P_2) := (u_2, \{(?w, value, ?v), (?w, loc, ?l)\})$. Now consider the power-related and weather-related named streams $(u_1, \mathcal{S}_1), (u_2, \mathcal{S}_2) \in \Sigma$ respectively as follows:

time	graph
10	:H1 :pow :Pw1 :H1 :loc :L1
25	:H2 :pow :Pw2 :H2 :loc :L2
...	...

time	graph
15	:W1 :value :V11 :W1 :loc :L1
40	:W2 :value :V12 :W2 :loc :L2
...	...

Then the evaluation of a SPASEQ query

$$Q = (\{?h, ?p, ?v\}, 50, ((u_1, P_1) ; (u_2, P_2)))$$

with the skip-till-next operator at time $\tau = 20$ over the streamset Σ , i.e. $\llbracket Q \rrbracket_{\Sigma}^{20}$, can be described as follows:

time	?h	?p	?v
15	:H1	:Pw1	:V11

In this section, we presented the detailed semantics of SPASEQ operators. Based on this, we also present some details how SPASEQ can be extended for operators such as negation and optional. Such details can be referred from Appendix B. The implementation details of SPASEQ operators are provided in the proceeding section.

6. Implementing the SPASEQ Query Engine

In this section, we move from the theory to the practical implementation of the SPASEQ query engine. We first present the NFA_{scep} model that is utilised to compile SPASEQ queries, and then provide details regarding the various system's blocks and optimisations used for the SPASEQ query engine.

In general, for a CEP system, the set of defined temporal operators are evaluated in a progressive way against the incoming events. That is, before a composite or complex event is detected through a full pattern match, partial matches of the query patterns emerge with time. These partial matches require to be taken into account, primarily within in-memory caches, since they express the potential for an imminent full match. There exists a wide spectrum of approaches to track the state of partial matches, and to determine the occurrence of a complex event. In summary, these approaches include rule-based techniques [13] that mostly represent a set of rules in a prolog like language, tree structures [51,5] where the pattern are parsed as join trees, graph-based representations [52,28] to merge all the rules within a single structure, and finally Finite State Machine representations, in particular Non-deterministic Finite Automata (NFA) [8,53]. The choice of these representations is motivated not only by their expressiveness measures, but also on the performance metrics that each approach tries to enhance. For instance, ETALIS [13], a rule-based engine, mostly focuses on how the complex rules are mapped and executed as Prolog objects and rules, while SASE+ [4,8], Zstream [5] and lazy evaluation [54] of NFAs focus on query-rewriting, predicate-related optimisations and memory management techniques.

Due to the separation of constructs, SPASEQ queries can be easily mapped to various different models. For instance, a query tree can be constructed from a SPASEQ query, where the GPM expression are mapped at the

leaves of the tree and the matched results are propagated to the root of the tree to extract a full match. However, considering the following points, we opt to use an NFA-based compilation and execution model for SPASEQ queries: (i) given the semantic similarity of SPASEQ's sequence expressions to the regular expressions, NFA would appear to be the natural choice; (ii) NFAs are expressive enough to capture all the complex patterns in SPASEQ; (iii) NFAs retain many attractive computational properties of Finite State Automata (FSA) on words, hence, by translating SPASEQ queries into NFAs, we can exploit several existing optimisation techniques [8,53].

In addition to the NFA-based model, we use various optimisation techniques to evaluate GPM expressions proposed by our system SPECTRA [45]: since SPASEQ employs an RDF graph model, it not only requires the efficient management of temporal operators, but also the efficient evaluation of graph patterns. In the following, we first describe the NFA_{scep} model for SPASEQ, and later present how SPASEQ queries are compiled and evaluated using NFA_{scep} . The optimisation techniques for the execution of SPASEQ queries are provided in Section 7.

6.1. The NFA_{scep} Model for SPASEQ

We extend the standard NFA model [41,55] in three ways. First, given that SPASEQ matches events with GPM expressions, we associate each automaton edge with a predicate, and for an incoming event, this edge is traversed iff the GPM expression is satisfied by this event. Second, in order to handle statefulness between GPM expressions (shared variables), we store in each automaton instance the mappings of those events that have contributed to the state transition of this instance. Third, we map all the SPASEQ operators (conjunction, disjunction, Kleene+, event selection strategies, etc.) on the NFA model: in the existing works [8,10], the mapping for the NFA model are only shown for a subsets of these operators and all the operators are not mapped onto a single model. We call such an automaton model as NFA_{scep} and it is defined as follows:

Definition 17 An NFA_{scep} automaton is a tuple $\mathcal{A} = \langle X, E, \Theta, \varphi, x_o, x_f \rangle$, where

- X : a set of states;
- E : a set of directed edges connecting states;
- Θ : is a set of state-transition predicates, where each $\theta \in \Theta$, $\theta = (U, op, P)$ is a tuple; U is a set of stream names, $op \in \{ \&, +, |, ; \} \cup \{ \emptyset \}$ is a temporal operator, and P is a graph pattern;

- φ : is a labelling function $\varphi : E \rightarrow \Theta$ that maps each edge to the corresponding state-transition predicate;
- x_o : $x_o \in X$ is an initial or starting state;
- x_f : $x_f \in X$ is a final state or acceptance state.

We define three types of states: *initial* (x_o), *ordinary* (x) and *final* (x_f) states. These state types are analogous to the ones used in the traditional NFA models to implement the operators such as sequence, Kleene+, etc. Each state, except the final state, has at least one forward edge. Note that, we use a structure $\mathcal{H} = (\theta_i, \mathcal{A}_i, \Omega)$ to store the set of mappings Ω corresponding to the state-transition predicate θ_i and the automaton instance \mathcal{A}_i . Hence, when an event makes an automaton instance traverse an edge, the mappings in that event are properly referenced.

Example 12 Figure 3 shows the compiled NFA_{scep} for the SPASEQ Query 1 with the sequence expression $SEQ(A, B+, C)$. It contains four states, each having a set of edges labelled with the state-transition predicates. The state-transition predicate (U, op, P) consists of three parameters: graph pattern P for the events with stream names (U); op describes the type of operator mapped to an edge, for instance edges of state x_1 contain the Kleene+ operator. The description of mapping from the SPASEQ Query 1 to the NFA_{scep} in Figure 3 is as follows:

- The SPASEQ Query 1 contains the sequence expression $SEQ(A, B+, C)$, which produces one initial state, two ordinary states and a final state.
- State x_0 has one edge with state-transition predicate (called as GPM A in Query 1) $(u_{s_1}, \emptyset, P_A)$, where $U_{s_1} = \{S1\}$ and P_A is the graph pattern. Since the sequence expression in Query 1 only contains the immediately followed-by operator, the NFA_{scep} can simply transit to the next state on matching the state-transition predicate.
- State x_1 maps GPM B with Kleene+. Therefore, it has two edges each with a state-transition predicate $(U_{s_2}, +, P_B)$, one with the destination state x_2 , and other with itself as destination (x_1) to consume one or more same kind of events; where $U_{s_2} = \{S2\}$ and P_B is the graph pattern.
- State x_2 has one edge, which is used to transit to the next state if an event matches the defined state-transition predicate $(U_{s_3}, \emptyset, P_C)$; where $U_{s_3} = \{S3\}$ and P_C is the graph pattern.




Figure 3. Compiled NFA_{scep} for SPASEQ Query 1 with SEQ(A, B+, C) expression

State-transition predicates are used to determine the action taken by a state to transit to another. For instance, in Figure 3 the state x_0 transits to x_1 , if (1) the incoming event is from the defined stream name U_{s_1} , (2) the evaluation of the graph pattern P_A does not produce an empty set. Furthermore, the event selection strategy also determines if there is a *followed-by* or *immediately followed-by* relation between the processed events. Note that, in the presence of the *Kleene+* operator, NFA_{scep} will exhibit a non-deterministic behaviour, since the state-transition predicates will not be mutually exclusive.

Considering the vocabulary from existing NFA works [8,53], we say that each instance of an NFA_{scep} or a partial match is called a *run*. A run depicts the partial matches of defined patterns, and contains the set of selected events. Each run has a current *active state*. A run whose final state has reached is a *matched run*, hence denoting that all the defined patterns are matched with the set of selected events. We call the output of the matched run as a *query match*.

6.2. Compilation of SPASEQ Queries

As discussed earlier, the two main components of the SPASEQ language are *sequence* and *GPM expressions*. Due to the separation of these components, one can provide multiple different techniques to compile and process them. The compilation process of graph patterns using the traditional relational operators (e.g., selection, projection, cartesian product, join, etc.) within each GPM expression is borrowed from our earlier work [45]. Herein, we focus on the sequence expression and show how the temporal operators are mapped onto the NFA_{scep} .

The sequence expression sorts the execution of GPM expressions according to its entries. Moreover, the temporal operators determine the occurrence criteria of such GPM matches and the event selection strategies are utilised to select the relevant events. These

constraints or properties are mapped on the NFA_{scep} through the compiled state-transition predicates, while the window constraints are computed during the evaluation of each automaton run.

Let (u_1, P_1) and (u_2, P_2) be two GPM expressions, the compilation of SPASEQ temporal operators onto an NFA_{scep} automaton is described as follows.

- *Simple GPM expression*: The NFA_{scep} for a simple GPM expression with a sequence expression forms a state which transits to the next one with the match of the GPM expression mapped at the state's edge. The NFA_{scep} automaton for a GPM expression (u_1, P_1) is presented in Figure 4.




Figure 4. Example of Compilation of the GPM Expression (u_1, P_1)

- *Kleene+*: The *Kleene+* operator selects a set of events if they match to the defined GPM expression. Its automaton is constructed using two edges with one edge having the same source and destination state. Thus, it can detect one or more consecutive events. The corresponding NFA_{scep} for $((u_1, P_1)+)$ is illustrated in Figure 5.




Figure 5. Example of Compilation of the *Kleene+* Operator $((u_1, P_1)+)$

- *Strict Contiguity Operator*: The construction of NFA_{scep} for this operator is similar to the compilation of a simple GPM expression, where a single edge for the corresponding state – having different source and destination states – is constructed. The corresponding NFA_{scep} automaton for $((u_1, P_1), (u_2, P_2))$ is illustrated in Figure 6.




Figure 6. Example of Compilation of the *Strict Contiguity Operator* $((u_1, P_1), (u_2, P_2))$

- *Skip-Till-Next and Skip-Till-Any Operators*: These operators require the irrelevant events to be skipped. Thus, two different edges emanate from the corresponding state. One has the same source and destination states: this transition matches any kind of event. The second edge is destined for the next state with the defined state-transition predicate. Note that the construction of both of these selection strategies is the same with the difference how they are evaluated at the run time. The corresponding NFA_{scep} automaton for $((u_1, P_1); (u_2, P_2))$ is presented in Figure 7, where $U = \{u_1, u_2\}$ is the set of stream names.




Figure 7. Example of Compilation of the *skip-till-next* (or *skip-till-any*) operator $((u_1, P_1); (u_2, P_2))$

- *Conjunction Operator*: This operator detects the simultaneous occurrence of two or more events. Thus, there are two edges for the conjunction state, each destined for the same destination state. The NFA_{scep} automaton for $((u_1, P_1) \& (u_2, P_2))$ is illustrated in Figure 8, where the conjunction state has multiple edges, each having different state-transition predicates.
- *Disjunction operator*: This operator forms a similar automaton structure as that of conjunction operator, however with the difference of how it is executed for an active run. That is, only one edge has to be matched with the incoming event. The NFA_{scep} automaton for $((u_1, P_1) | (u_2, P_2))$ is illustrated in Figure 9.

In the aforementioned discussion, we show the compilation of SPASEQ operators independently. For the




Figure 8. Example of Compilation of the *Conjunction Operator* $((u_1, P_1) \& (u_2, P_2))$




Figure 9. Example of Compilation of the *Disjunction Operator* $((u_1, P_1) | (u_2, P_2))$

full SPASEQ query, these operators can be combined together to form a complex NFA_{scep} automaton. In order to show this, let σ be a sequence with a set of GPM expression and binary/unary operators. The compilation of sequence expression $(\sigma; (u, P)^+)$ with the additional skip-till-next and *Kleene+* operators is presented in Figure 10. Notice from Figure 10 how the concatenating process is simply the mapping of the last state of sequence σ onto the initial state of the GPM expression $(u, P)^+$.




Figure 10. Example of Compilation of the *Sequence Expression* $(\sigma; (u, P)^+)$

To conclude, this section presented the mapping of SPASEQ queries onto equivalent NFA_{scep} automata. In the next section, we show how NFA_{scep} automata are executed while considering the window constraints defined within a query.

6.3. Evaluation of NFA_{scep} Automaton

The compiled NFA_{scep} automaton represents the model that a matched sequence should follow. Thus, in

Algorithm 1: Processing streamset with NFA_{scep}

Input: Σ : streamset, \mathcal{A} : NFA_{scep} Automaton, ω : time window

- 1 $\mathcal{R} \leftarrow \{\}$: list of active runs
- 2 $\mathcal{H} \leftarrow \{\}$: cache history
- 3 $\mathcal{D} \leftarrow \{\}$: conjunction edge-timestamp map
- 4 **foreach** event $G_e \in \Sigma$ **do**
- 5 get the initial state x_0 from \mathcal{A}
- 6 get the final state x_f from \mathcal{A}
- 7 get the stream name u from the event G_e
- 8 **PROCESSEVENT** ($G_e, u, x_0, x_f, \mathcal{H}, \mathcal{R}, \mathcal{D}, \mathcal{A}, \omega$)

order to match a set of events emanating from a streamset, a set of runs is initiated at run-time. This set of runs contains partially matched sequences and a run that reaches to its final state represents a matched sequence. When a new event enters the NFA_{scep} evaluator, it can result in several actions to be taken by the system. We describe them as follows:

- New runs can be created, or new runs are duplicated (cloned) from the existing ones in order to cater the *Kleene+* operator, thus registering multiple matches.
- If the newly arrived events match with state-transition predicates (θ) of the active states, existing runs transit from one active state to another.
- Existing runs can be deleted, either because the arrival of a new event invalidates the constraints defined in the NFA_{scep} model such as event selection strategies, conjunction, etc. or the selected events in those runs are outside the defined window.

These conditions can be generalised into an algorithm that (i) keeps track of the set of active runs (\mathcal{R}), (ii) starts a new run or deletes the obsolete ones, (iii) chooses the right event for the state-transition predicates ($\theta \in \Theta$), (iv) calls the GPM evaluator to match an event with the graph pattern (P) which is provided in the state-transition predicate (θ), and (v) keeps track of the mappings of matched events with a structure \mathcal{H} .

Algorithm 1 presents the initialisation process of various data structures and how a streamset Σ is processed against the NFA_{scep} automaton \mathcal{A} . The initialised data structures include (i) a list of currently active runs (\mathcal{R}), where each run stores partial matches; (ii) a history cache \mathcal{H} to store the mappings of matched events; (iii) an edge-timestamp map \mathcal{D} to store the mapping of events and their timestamps that are matched at a con-

Algorithm 2: PROCESSEVENT with NFA_{scep}

Input: G_e : Graph Event, u : stream name, x_0 : initial state, x_f : final state, \mathcal{H} : cache history, \mathcal{R} : list of active runs, \mathcal{D} : conjunction edge-timestamp map, \mathcal{A} : NFA_{scep} Automaton, ω : time window

1 Function

PROCESSEVENT ($G_e, u, x_0, x_f, \mathcal{H}, \mathcal{R}, \mathcal{D}, \mathcal{A}, \omega$)

- 2 **foreach** run $r \in \mathcal{R}$ **do**
- 3 **if** $\omega >$ *initialising time of* r **then**
- 4 remove r from the list of active runs \mathcal{R}
- 5 **continue**
- 6 $x_i \leftarrow$ **GETACTIVESTATE**(r)
- 7 $E \leftarrow$ **GETEDGECSET**(x_i)
- 8 get θ for an edge $e \in E$ s.t $\theta \neq \epsilon$ and $\theta = (U_\theta, op_\theta, P_\theta)$
- 9 **if** $op_\theta = '+'$ **then**
- 10 **KLEENEPLUS** ($G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$)
- 11 **else if** $op_\theta = ';' \text{ or } op_\theta = '\emptyset'$ **then**
- 12 **EVENTSELECTION** ($G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$)
- 13 **else if** $op_\theta = '&'$ **then**
- 14 **CONJUNCTION**($G_e, u, x_f, \mathcal{H}, \mathcal{R}, \mathcal{D}, r, x_i, E$)
- 15 **else if** $op_\theta = '|'$ **then**
- 16 **DISJUNCTION**($G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$)
- 17 // Check if an event G_e can create a new run from the initial state x_0
- 18 $E_0 \leftarrow$ **GETEDGECSET**(x_0)
- 19 get θ for the edge $e \in E_0$ s.t $\theta \neq \epsilon$ and $\theta = (U_\theta, op_\theta, P_\theta)$
- 20 **if** $u \in U_\theta$ **and** $GPM(G_e, P_\theta, \mathcal{H})$ **then**
- 21 initiate a new run r_i of \mathcal{A} with active state x_1
- 22 $\mathcal{R} \leftarrow \mathcal{R} \cup r_i$

junction operator's state (lines 1-3). The algorithm selects the initial state x_0 and final state x_f of automaton \mathcal{A} , and the stream name u of the event to be processed (lines 5-6). This information along-with the initialised structures is passed to the **PROCESSEVENT** function (see Algorithm 2), where each incoming event G_e is matched with the active automaton's runs.

In Algorithm 2, we present the general execution of SPASEQ operators with the arrival of an event. Algo-

rithm 2 begins by iterating over the list of active runs. The list of active runs \mathcal{R} is used to determine if (i) the existing runs are expired or not, i.e. their initialisation time is outside the window boundary, and hence to be deleted (lines 2-5); (ii) the active state of the active run can be matched with the newly arrived event (lines 9-16). The algorithm starts by comparing the initialising time of the run r under evaluation and the defined time window. The run r is deleted if its outside the defined window (line 3). The algorithm then extracts the current active state of the run, and according to the mapped operators it selects the appropriate function for the corresponding operator. In the end, the algorithm checks if the incoming event can start a new run or not (lines 19-22). That is, it matches the initial state's (x_0) edge of the automaton \mathcal{A} with the incoming event G_e using the GPM function (line 22)¹. In case of a match and if the event is from the same stream the edge is waiting for ($u \in U_\theta$), it initiates a new run r_i of the automaton \mathcal{A} with the active state x_1 , i.e. it proceeds to the next state (line 23). This new run is then added to the list of active runs \mathcal{R} (line 24).

The aforementioned general algorithm gives an overall view of how SPASEQ queries are executed. To keep the discussion brief, the evaluation details of SPASEQ temporal operators are provided in Appendix C. Note that, herein, we only provide the implementation of *strict contiguity* and *skip-till-next* event selection strategies, since *skip-till-any* operator results in a exponential time complexity [53] and may not be suitable for some real-world applications. The optimised implementation of the *skip-till-any* operator will be the subject of discussion for our future endeavours.

6.4. Design of the SPASEQ Query Engine

Having provided the details of compiling the SPASEQ queries onto NFA_{scep} automata and their evaluation strategies, herein we provide an overview of the system architecture of SPASEQ query engine.

Figure 11 shows the architecture of the SPASEQ query engine. Its main components are *input manager*, *queue manager*, *query optimiser*, *NFA evaluator*, and an RDF graph processor in *GPM evaluator*. In the following, we briefly discuss these components.

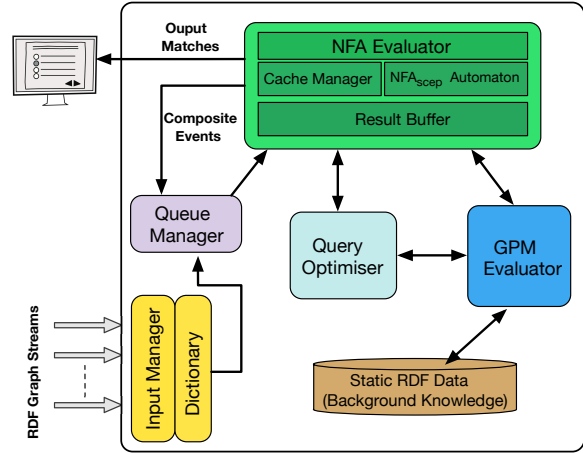


Figure 11. Architecture of the SPASEQ Query Engine

The queue and input managers do their usual job of feeding the required data into the NFA evaluator. Since our system employs streamsets, there are multiple buffers to queue the data from a set of streams. Moreover, the newly constructed events for SPASEQ queries can also be fed back to the queue manager for further processing. The incoming data from streams are first mapped to the numeric IDs using dictionary encoding². The input manager also utilises an efficient parser^{3,4} to parse the RDF formatted data into the internal format of the system.

At the heart of SPASEQ query engine, the GPM evaluator uses the GPM expressions, events, cache history and edge-timestamp mappings (see Appendix C) to match the defined graph pattern with the incoming events. We employ our SPECTRA GPM engine [45] for this task. SPECTRA is a main memory graph pattern processor but uses specialised data structures and indexing techniques suitable for the optimised evaluation of RDF graph events. We employ two main operators of SPECTRA for SPASEQ and they are described as follow:

²

Dictionary encoding is a usual process employed by a variety of RDF-based systems [56,57]. It reduces the memory footprints by replacing strings with short numeric IDs, and also increases the system performance by using numeric comparisons instead of costly string comparisons.

³

We employ the performance intensive NxParser, which is a non-validating parser for the Nx format, where x = Triples, Quads, or any other number.

⁴

NxParser: <https://github.com/nxparser/nxparser>, last accessed: November, 2017.

¹

Herein, for sake of brevity, we only show the simple option where the complex operators are not mapped at the initial state. However, in practice, it is implemented using the function for the defined operator at the initial state.

- The `SUMMARYGRAPH` operator enables the system to avoid storing all the triples within an event by exploiting the structure of the graph patterns. That is, the graph pattern P is used to prune all the triples within an event that do not match the subjects, predicates or objects defined in P . The pruned set of triples within an event, called *Summary Graph*, are materialised into a set of vertically partitioned tables called *views*. Each view, a two-column table (s, o) stores all the triples for each unique predicate in a summarised event.
- The `QUERYPROCESSOR` employs the set of views to implement multi-join operations using incremental indexing. That is, for each view a *sibling list* is used to incrementally determine the set of joined subject/objects that belong to a specific branch within a graph. This results in an incremental solution, where the creation and maintenance of the indexes are the by-product of the join executions between views, not of updates within a defined window.

The use of specialised indexing techniques and data structures enables SPECTRA to outperform existing RSP systems up to an order of magnitude [45]. Thus, employing such optimisations also aids the evaluation of GPM expressions in SPASEQ. Furthermore, the evaluated events can also be enriched using the static RDF data, where such data are loaded into the memory using the vertically partitioned tables.

Next comes the NFA evaluator of SPASEQ. It contains the compiled NFA_{scep} automata and employs the GPM evaluator to compute the GPM expressions mapped on the state-transition predicates. Its subcomponent, the cache manager stores the mappings of matched events and mappings for the conjunction operator, i.e. cache history (\mathcal{H}) and conjunction edge-timestamp map (\mathcal{D}). Finally, the query optimiser employs various techniques to reduce the load for GPM evaluator and the number of active runs. The detailed description of query optimiser is presented in the proceeding section.

7. Query Optimisations

The two main resources in question for processing CEP queries are CPU usage, and system memory. As discussed in Section 6, many different strategies have been proposed to find an optimal way of utilising CPU and memory in CEP systems. One of the main ben-

efits of using an NFA model as an underlying execution framework is that we can take advantage of the rich literature on such techniques. These optimisation techniques can be borrowed into the design of SCEP, while customising them for RDF graph streams. In this section, we describe how such techniques are applicable for SCEP, and also propose new ones considering the query processing over a streamset. First, we review the evaluation complexity for the main operators of the SPASEQ query language.

7.1. Evaluation Complexity of NFA_{scep}

The evaluation complexity of NFA_{scep} provides a quantitative measure to establish the cost of various SPASEQ operators. Herein, we first describe the cost of temporal operators in terms of the GPM evaluation function, and later provide the upper bound of time-complexity in terms of number of active runs.

Incoming events are matched to the GPM expressions mapped on the state's edges and such evaluation decides if a state can transit to the next one. Given n number of events in a streamset, the cost of comparing a graph pattern P with n number of events within a window for the unary operators and event selection strategies is described as follows:

$$cost_n = \sum_{i=1}^n \mathbf{Eval}(P, G_e^i),$$

where $\mathbf{Eval}(P, G_e^i)$ represents the cost of matching a graph pattern P with a RDF graph-based event G_e^i . For a BGP $\mathbf{Eval}(P, G_e^i) = |P| \cdot |G_e^i|$ from Theorem 1 in [58]. Now we extend this cost to include the cost of the binary operators (BOP). Given n events and k BOP operators, the cost function is extended as follows:

$$cost_n^k = \sum_{i=1}^n \sum_{j=1}^k \mathbf{Eval}(P_j, G_e^i),$$

Notice that due to the k number of BOP operator, each event in worst case scenario has to be matched with all the defined GPM expressions for the BOP operators.

Prior works analysing the complexity of NFA evaluation often consider the number of runs or partial matches created by an operator and employ upper bounds on its expected value [8,53]: since each event has to traverse the list of partial matches to find the complete matches. We adopt the same approach for analysing the complexity of NFA_{scep} evaluation. Given

an NFA_{scep} and streamset, for each incoming event the system has to check all the active runs to determine if the newly arrived event results in (i) state transition from the current active state to the next one, (ii) duplication of a new run, (iii) deletion of the active run. Query operators that result in creating new runs or those which increase the number of active runs are considered to be the most expensive ones. In order to simplify the analysis, we make the following assumptions:

1. We ignore the cost of evaluating a GPM expression over each event, i.e the **Eval** function.
2. We ignore the selectivity measures of the state-transition predicates, i.e. the events that are not matched are either skipped or result in deleting a run of an NFA_{scep} automaton. Hence, our focus is on the worst-case behaviour.

Based on this, let us consider that there are n number of events with a window and a new event arrives at a current active state of a run, where the active state may contain the following set of operators: *strict contiguity*, *skip-till-next*, *Kleene+*, conjunction and disjunction.

Theorem 1 *The upper bound of evaluation complexity of event selection strategies, i.e. strict contiguity and skip-till-next, is linear-time $\mathcal{O}(n)$, where n is the total number of runs generated for the n input events within a window.*

Proof Sketch. None of the operators described in Theorem 1 duplicate runs from the existing ones. Each has only one GPM expression to be matched with the incoming events. Let us consider the case of event selection operators. Given a sequence expression $((u_i, P_i) op (u_j, P_j))$, where $op = \{ ', ';\}$, mapped to states x_i and x_j . With the arrival of an event G_e at τ , where an event selection operator is mapped at state x_j , it can result in the following actions: (i) the run will transit to the next state, (ii) the event will be skipped due to *skip-till-next* operator, and (iii) the run will be deleted. Since we are considering the worst-case behaviour, let us dismiss the situations (ii) (iii). In situation (i) there will be no extra run created for the above mentioned operators, and each incoming event will be matched to only one GPM expression. Thus the evaluation cost remains linear and for n number of events there can be only n runs. \square

Although the upper bound of the operators described in Theorem 1 has the same evaluation complexity, there exists discrepancies when considering the real-world

scenarios [53]. Due to the skipping nature of *skip-till-next* operator, the life-span of its runs can be longer on a streamset. In particular, those runs that do not produce matches, and instead loop around a state by ignoring incoming events until the defined window expires. On the contrary, the expected duration of a run is shorter for the *strict contiguity* operator, since a run fails immediately when it reads an event that violates its requirements. Such difference in their evaluation cost is visible in our experimental analysis (Section 8).

The case of conjunction and disjunction operators is slightly different, and therefore has a different upper case bound. That is, for each conjunction/disjunction operator: (i) more than one edge has different source and destination states with distinct GPM expressions, (ii) the edges do not have an ϵ -transition. Thus, in worst case each incoming event has to match to the complete set of state-transition predicates. Hence, for k such edges of a conjunction/disjunction operator, and n events within a window, we can provide the upper bound on the complexity of these operators as follows:

Theorem 2 *The upper bound of evaluation complexity of conjunction and disjunction is $\mathcal{O}(n \cdot k)$, where n is the total number of events within a window and k is number of GPM expressions mapped to the edges of a state.*

Proof Sketch. For the conjunction and disjunction operators no new runs are initiated from old ones. If an event arrives at active state x , it either matches to the set of edges defined and moves to the next state, or it stays at the current active states and waits for new events (in case of the conjunction operator). However, for both operators, each incoming event can end up traversing the whole set of k mapped edges. Thus, even if the number of active runs remains the same, each event may have to be matched with a number of GPM expressions mapped at the state's edges. For k such edges and n events, in worst case the cost will be $\mathcal{O}(n \cdot k)$. \square

Theorem 3 *The upper bound of evaluation complexity of each Kleene+ operator is quadratic-time $\mathcal{O}(n^2)$ for n events within a window.*

Proof Sketch. Consider the following sequence expression $((u_j, P_j)+)$, which is mapped onto the state x_j with the *Kleene+* operator. Let us consider that x_j is an active state. If an event G_e arrives at time τ , and if it matches the GPM expression of the state-transition predicate, it will duplicate the current active run and append the duplicated one to the list of active runs.

Thus, for each newly matched event at a *Kleene+* state, a new run is added to the active list, and for n such events, there will be in total n^2 runs to be generated considering all the events are matched to the GPM expression of the state x_j , i.e. worst case behaviour. \square

The *Kleene+* operator is the most expensive in the lot, in terms of number of active runs. Based on the observations in Theorems 1, 2 and 3, we adopt some of the optimisation strategies previously proposed, and also propose some new ones. We divide these techniques into two classes from the view point of operators and system: *local*, and *global* levels. Local-level optimisation techniques are targeted at the specific operators considering their attributes, while the global-level optimisation are for all the operators, and are implemented at the system level. In the following section, we present these techniques in details.

7.2. Global Query Optimisations

The evaluation of an NFA_{scep} automaton is driven by the state-transition predicates being satisfied (or not) for an incoming event. The number of active runs of an NFA_{scep} automaton, and the number of state-transition predicates that each run could potentially traverse can be very large. Therefore, the aim of global optimisation is to reduce the total number of active runs by (i) deleting, as soon as possible, the runs that cannot produce matches, and (ii) indexing the runs to collect the smaller number of runs that are actually affected by an event.

7.2.1. Pushing Windows and Stateful Predicates

Pushing the defined windows and stateful predicates are usually employed by the CEP systems to evict, as soon as possible, the events that are outside the window or cannot produce matches for the defined predicates [8, 5]. Hence, we employ the same technique of window pushing to efficiently delete runs that are outside the time window. In Algorithm 2 (lines 3-5), we push the window check before iterating over the active runs list. For the stateful predicates, we customised our GPM evaluator. That is, we use the FILTER expression in the GPM construct to first determine if the incoming event can match the defined mapping without first evaluating the initiating the complete GPM process. This results in pruning the irrelevant events and reducing the load over the GPM evaluator. As an example of this, consider the SPASEQ Query 1. Its GPM expressions share the stateful variables of location (?l) and electricity fare

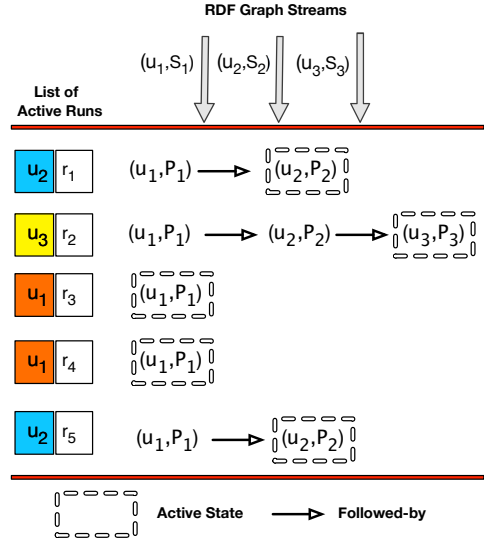


Figure 12. Partitioning Runs by Stream Names

(?fr). Pushing these two joins, we can easily ignore the events that would not contain the expected mappings of these variables, and consequently the system does not have to process the complete GPM expressions (GPM B and C) for such events.

7.2.2. Indexing Runs by Stream Names

SPASEQ queries are evaluated over a streamset, where the edges from each state contain the stream name which is used to match the graph patterns. Therefore, each active state waits for a specific type of event from a specific stream, and later invokes the GPM evaluator. This means we can use such property of streamset to extract the set of runs, from active run list, that can be affected by the incoming event. Based on this, we index each run by the stream name of its active state (see Figure 12). More precisely, the index takes the stream name as a key and the corresponding run as the value. These indexes are simple hash tables, and for each incoming event it essentially returns a set of runs that can be affected by the incoming event. These indexes proved to be a useful feature for processing events from a streamset. Note that the naive implementation using a single list of runs would be inefficient in this case: each incoming event would iterate over all the active runs, and initiate the matching process for each of them.

Example 13 Figure 12 shows a set of streams employed by a set of active runs. The active runs r_1 , r_5 are waiting for an event from stream with name u_2 ,

the active run r_2 is waiting for an event from stream with name u_3 , and active runs r_3 and r_4 are waiting for the events from the stream with name u_1 . Hence these runs are evaluated only for the events coming from the desired stream.

7.2.3. Memory Management

The memory requirements usually grow in proportion to the input stream size or the matched results [53]. For our system, the three main data structures that require tweaking are the *cache manager*, the *result buffer* and the *indexed active run list*. In this context, our first step is to use the *buy bulk* design principle. That is, we allocate memory at once or infrequently for resizing. This complies to the fact that the dynamic memory allocation is typically faster when allocation is performed in bulk, instead of multiple small requests of the same total size. Second, since the cache manager and the result buffer are indexed with the dynamically generated run ids, we use the expired runs – which either are complete or not – to locate the exact expired runs to be deleted. These runs are added to a pool: when a new run is created, we try to recycle the memory from such pool. This limits the initialisation of new runs and reduces the load over the garbage collector [8]. Note that we use hash-based indexing for all the data structures, which means the position of expired runs can be found in theoretically constant time.

7.3. Local Query Optimisation

Local query optimisation is devised for the conjunction and disjunction operators, where the chief problem is how to select the GPM expression from a set of edges and how to reduce the load on the GPM evaluator. The knowledge of the runs affected by an incoming event is not sufficient, we also have to determine which edge these runs will traverse. These issues are not catered by the existing NFA models since they (i) work on single stream mode; (ii) most of them do not provide the evaluation techniques for disjunction and conjunction operators.

To better illustrate the problem, let us start by examining the sequence expression $((u_1, P_1) \mid (u_1, P_2) \mid (u_1, P_3) \mid (u_1, P_4))$ for the disjunction operator. Figure 13 shows the related NFA_{scep} automaton. Now consider an input stream u_1 , and an event G_e^i at time τ_i arrives at the state x_1 . The basic sequential way of processing G_e^i would be to first match with P_1 then P_2 , P_3 and finally with P_4 : since all the graph patterns are waiting for an event from the stream u_1 . Now the ques-




Figure 13. Compilation of Disjunction Operator for $((u_1, P_1) \mid (u_1, P_2) \mid (u_1, P_3) \mid (u_1, P_4))$

tion is how to choose the less costly graph pattern to be selected first by the state, such that if it matches the automaton moves to the next state, hence complying to the optional operator.

The optimal way of processing the disjunction operator would be to sort the graph patterns according to their cost, and select the cheapest one for the first round of evaluation. That is, if $\mathbf{c}(P_i, G_e^i)$ is the cost of matching a graph pattern with an event G_e^i , then we require a sorted list such that $\mathbf{c}(P_j, G_e^i) < \mathbf{c}(P_k, G_e^i) < \dots < \mathbf{c}(P_m, G_e^i)$. The question is how to determine the cost of GPM evaluation. There can be two different ways to it.

1. Use the selective measures and structure of the graph patterns. That is, how much the GRAPH-SUMMARY operator be handy for them (see Section 6.4).
2. Adaptively gather the statistics about the cost of matching a specific graph pattern, and sort the graph pattern accordingly.

Let us focus on the first approach. As according to Theorem 1 in [59], the cost of matching a graph pattern and an event is directly proportional to each of their sizes. That is, if there are more triple patterns $tp \in P$, then there will be more join operations on different vertically-partitioned views: this can give us a fair bit of idea about the costly graph patterns. Furthermore, due to the presence of filters, the GRAPH-SUMMARY operator can prune most of unnecessary triples, and consequently reduce the cost of the GPM operation. Following this reasoning, we keep a sorted set of graph patterns $\langle P_1, P_2, \dots, P_k \rangle$ within the state which mapped the conjunction/disjunction operator, each associated with a stream name u_i . This set is sorted by checking the number of triple patterns, and the selec-




Figure 14. Compilation of Conjunction Operator for $((u_1, P_1) \& (u_1, P_2) \& (u_1, P_3))$

tivity of subjects, predicates and objects within a graph pattern during the query compilation [45]. The state can utilise this set to first inspect the less costly graph patterns for the incoming events. This can lead to a less costly disjunction operator with few calls to the GPM evaluator.

The second approach is based on the statistics measures. That is, the system during its life-span observes which graph pattern has been utilised successfully in the past and is less costly compared with others. This approach can be built on top of the technique discussed above. Herein, the implementation of such optimisation technique is considered as future work.

The conjunction operator, however, contains an additional challenge on top of the one discussed above. To illustrate this, let us consider a sequence expression $((u_1, P_1) \& (u_1, P_2) \& (u_1, P_3))$ for the conjunction operator. Figure 14 shows the NFA_{scep} automaton for it. Hence, for the state x_1 to proceed to the next state, it has to successfully match all the defined state-transition predicates, such that events satisfying them should occur at the same time. Thus, if an event G_e^i arrives and matches to one of the state-transition predicate, the automaton buffers its result, timestamp, and waits for the remaining events. Now consider a situation, where events G_e^i and G_e^j arrive at τ_1 and match with the GPM expressions (u_1, P_1) and (u_1, P_2) respectively. Then the automaton waits for an event to satisfy GPM expression (u_1, P_3) . Now consider an event G_e^m arrives at τ_2 . It results into two constraints to be examined (i) if $\tau_1 = \tau_2$, and (ii) if G_e^m matches with the GPM expression (u_3, P_3) . Here, if any of the above mentioned constraints would not match, then it means the run has to be deleted and all the previous GPM evaluations were useless: the process of matching an event with a graph pattern is expensive and it stresses the CPU utilisation.

Our approach to address this issue is to employ a *lazy evaluation* technique. Conceptually, it delays the eval-

uation of graph patterns until it gets enough evidence that these matches would not be useless. Its steps are described as follow:

1. Buffer the events from streams until the number of events with the same timestamps is equal to the number of edges (with distinct GPM expressions) going out from the state that maps conjunction operator.
2. After the conformity of the above constraint, choose graph patterns according to their costs (as discussed for disjunction operator).

The main idea underlying our lazy evaluation strategy is to avoid unnecessary high cost of the GPM, and to start the GPM process when it is probable enough that it would return the desired results. The idea of lazy batch-based processing is the reminiscent of work [5] on buffering the events and processing them as batches. To keep the discussion brief, the detailed algorithm for our lazy evaluation is described in Appendix D

In this section, we presented various strategies to optimise the evaluation of SPASEQ temporal operators. In the next section, we present the quantitative analysis of the SPASEQ operators and the effect of our optimisation strategies.

8. Experimental Evaluation

8.1. Experimental Setup

Datasets. We used four real-world datasets and their associated queries for our experimental evaluation.

The Stock Market Dataset (SMD) is a real-world dataset [60,28] from the New York Stock Exchange. It contains 225k transaction records of 10 companies in 1 sector during 12 hours. Each event in the dataset carries company, sector, and transaction identifiers, volume, price, time stamp, and type (sell or buy). We replicate this dataset 10 times with adjusted company, sector, and transaction identifiers such that the resulting data set contains transactions for 110 companies in 11 sectors. No other attributes except identifiers were changed in the replicas compared to the original. This dataset is then mapped to RDF N-Triples format, where each attribute of the stock event is mapped as a triple and a whole event (set of triples) refer to an RDF graph.

The *UMass Smart Home Dataset (SHD)* [61] is a real-world dataset and provides power measurements that include heterogeneous sensory information, i.e. power-related events for power sources, weather-related events from sensors (i.e. thermostat) and events for renewable energy sources. We use a smart grid ontology [30] to map the raw eventual data into N-Triples format for three different streams: the power stream (S_1), the power storage stream (S_2) and the weather stream (S_3). In total the dataset contains around 30 million triples, 8 million events.

The *Credit Card Dataset (CCD)* is a real dataset of credit card transactions [62,27]. In this dataset, each event is a transaction accompanied by several arguments, such as the time of the transaction, the card ID, the amount of money spent, etc. The total number of transactions in this summary dataset was around 1.5 million. We mapped the dataset to the RDF N-Triples format, where each attribute of the transaction is mapped as a triple and a whole transaction refer to an RDF graph.

The *Traffic Dataset (TD)* is a real traffic dataset gathered from sensors deployed within the city of Aarhus, Denmark [36]. Traffic data are collected by observing the vehicle count between two points over a duration of time. The dataset includes the average vehicle speed, vehicle count, estimated travel time and congestion level between the two points set over each segment of road. The dataset is then mapped to the sensor network ontology as provided by the CityBench [36].

Queries. We evaluate the workload of four main queries and their variations for the above mentioned datasets. That is, the queries for the **UC 1** (SHD dataset and Q1 from Section 4), **UC 3** (CCD dataset and Q2 from Appendix A) **UC 4** (SMD dataset and Q3 and Q4 from Appendix A) and **UC 5** (TD dataset and Q5 from Appendix A). The properties of these queries are described as follows: Q1 uses the multistreams and provides a sequence using three GPM clauses with stateless or constant filters; Q2 provides a sequence over two GPM clauses using stateful or variable filters; Q3 and Q4 employ three to seven GPM clauses for V-shaped and head-and-shoulder patterns with stateful filters; Q5 uses multistreams and provides a sequence with conjunction over three GPM clauses using stateless or constant filters.

Methodology. To show the efficiency of our approach, we compare our system with EP-SPARQL. EP-SPARQL is the only system which provide a SCEP

language and its implementation. SPASEQ and EP-SPARQL differ w.r.t each other in terms of semantics and data model. Hence, they may produce different results for the same query. Therefore, the aim of our comparative analysis is to employ the same use case, its queries and dataset to measure the performance differences between the two. This strategy is mostly utilised by the information retrieval systems [63]. Although EP-SPARQL does not provide any event selection strategies at the query level, its underlying system ETALIS provides various event selection strategies. We use the *recent* (EPSPARQL-R) and *unrestricted* (EPSPARQL-UR) strategy from ETALIS for our experiments, where the *recent* strategy can be mapped to the strict congruity (SPASEQ-SC) and *unrestricted* to skip-till-next (SPASEQ-STN). Note that we also use C-SPARQL for our set of experiments. Although C-SPARQL provides a time function to allow simple sequences over the RDF triple streams, it does not perform well on the selected datasets and queries, and takes several hours to process windows of moderate sizes. Hence, the results could not fit properly in our charts and are not included. Note that the GPM process of SPASEQ is based on our previous work called SPECTRA [45]. The readers are encouraged to refer to our previous work for the detailed comparison between the GPM of SPECTRA and other RSP systems. Moreover, for multiple streams, we use events from the streams in the order defined within the queries. This enables us to get the maximum number of matches and compare systems under heavy workloads. For all the sets of experiments, we use time and event-based windows. For the time window, we use the event's source timestamps. We measured two common metrics [8,28,53,64] for the CEP system, namely average *CPU time* over the windows and average throughput. Average CPU time is measured in seconds and milliseconds as the sum of total elapsed time in all windows divided by the number of windows. Average throughput is the number of triples processed per second averaged over all the processed windows. We execute each experiment three times and report their average results here.

8.2. Results and Analysis

We start our analysis by first providing the comparative analysis with the existing SCEP systems for the same datasets and use cases.

Comparative Analysis

Question 1. How does the SPASEQ engine perform w.r.t. the EP-SPARQL engine?

Fig. 15 and 16 showcase the CPU time and throughput respectively of both SPASEQ and EP-SPARQL for the two event selection strategies over the four selected datasets using various window sizes. From Fig. 15 and 16, SPASEQ using the SC and STN event selection strategies outperforms EP-SPARQL for R and UR event selection strategies in terms of CPU time and throughput measures. The details of these results are as follows.

EP-SPARQL. In general, EP-SPARQL uses a Prolog wrapper based on the *event driven backward chaining rules* (EDBC), and schedules the execution via a declarative language using backward reasoning. This first results in an overhead of object mappings. Second, each time a new consequent of a rule is matched, the entire rule set is searched to see what else can be concluded. Hence resulting in high computation costs. The CPU time of both EP-SPARQL-R and EP-SPARQL-UR grows with the size of the window. However, since EP-SPARQL-R provides a strict sequence over the incoming events, only the most recent rules are selected to produce the matches. This results in less computation cost compared with EP-SPARQL-UR: EP-SPARQL-UR skips the irrelevant events and a broad set of rules is selected and processed. Hence, EP-SPARQL-UR results in a considerable high computational cost for each incoming event. Hence the CPU time for EP-SPARQL-UR increases quadratically (subsequently throughput decreases) with the increase in the window size. The same effect can be seen in Fig. 16, when comparing the throughput measures of EP-SPARQL-R and EP-SPARQL-UR. Furthermore, the CPU costs for both EP-SPARQL-R and EP-SPARQL-UR have similar measures for all the datasets and queries. This is because of evaluating the backward chaining rules, where the defined constant filters in the queries are not taken into account beforehand and the system does not prune the irrelevant rules before starting the complete rule matching procedure. Moreover, EP-SPARQL uses a goal-based memory management technique, i.e. periodic pruning of expired goals using alarm predicates, which is expensive for large window sizes.

SPASEQ. The CPU time of SPASEQ-SC and SPASEQ-STN also increases (subsequently the throughput decreases) with the increase of the window size. However, it performs much better than its counterpart. The reasons are as follows. SPASEQ employs the NFA_{scep}

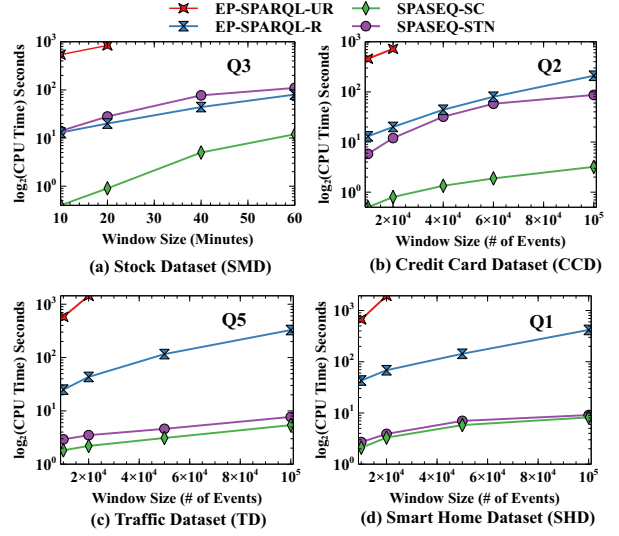


Figure 15. Average CPU Time comparison of SPASEQ (Strict Contiguity (SC) and Skip-Till-Next (STN)) with EP-SPARQL (Recent (R) and Unrestricted (UR)) over the four selected datasets (SMD, CCD, TD and SHD).

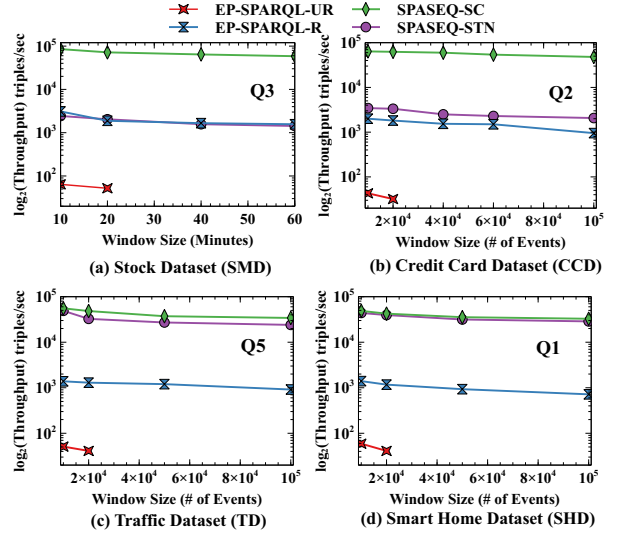


Figure 16. Average Throughput comparison (triples/second) of SPASEQ (Strict Contiguity (SC) and Skip-Till-Next (STN)) with EP-SPARQL (Recent (R) and Unrestricted (UR)) over the four selected datasets (SMD, CCD, TD and SHD).

model with various optimisation strategies to reduce the cost of evaluating GPM and the state-transition predicates. It utilises efficient “right-on-time” garbage collection for the deceased runs, and optimisations such as pushing temporal windows and stateful joins, and incremental indexing from SPECTRA – SPASEQ’s un-

derlying GPM engine – reduce the average computation overheads and life-span of an active run. Due to these optimisation strategies, SPASEQ has different CPU cost for different strategies and queries. For the CCD and SMD queries (Q2 and Q3), the initial NFA_{scep} state (GPM expression) has variable or stateful filters. Hence, each incoming event can potentially start a new NFA_{scep} run or a partial match. This results in a large number of runs to be evaluated for each incoming event. This effect, however, is not evident for SPASEQ-SC, since it expects a strict sequence of events and it deletes runs as soon as the incoming events violate it. For the SPASEQ-STN, events that are not matched are skipped and the runs life-times are much longer, hence the size of total active runs. This means, for each incoming event, the system has to go through a larger list of runs to be matched. This behaviour of event selection strategies is in line with our theoretical analysis in Section 7. For the SHD and TD queries (Q1 and Q5), the initial NFA_{scep} state (GPM expression) has static or stateless filters. Hence, only specific events can start a new run. This results in a small number of active runs to be produced and processed for an incoming event. For the same reason, for both TC and SMD datasets, both SPASEQ-SC and SPASEQ-STN have similar CPU costs and throughput. From Fig. 15 (c), we can also see that the conjunction operator of SPASEQ performs much better than EP-SPARQL for the two event selection strategies. This is due to the lazy evaluation of the conjunction operator, where the GPM process is started only when we have enough events that can produce a match. In summary, on all the datasets, SPASEQ-SC is two orders of magnitude less costly than EP-SPARQL-R, and SPASEQ-STN is two to three orders of magnitude less costly than EP-SPARQL-UR.

To further consolidate our comparison analysis, we showcase the performance of both systems by increasing the number of sequence clauses over windows of size 40 minutes. We use the SMD dataset and create a set of queries (extensions of Q3 and Q4), since SMD is much costly than other datasets and queries. From Fig. 17, the CPU cost of both EP-SPARQL-R and EP-SPARQL-UR increases with the number of sequence clauses. However, the cost of SPASEQ-SC and SPASEQ-STN provides similar measures. This is due to the fact that the CPU cost of SPASEQ depends on the number of active runs, since they result in large number of GPM evaluation operations. Increasing the number of sequences would not increase the number of active runs but the active life of runs. Hence, increasing the number of sequences does not increase the CPU cost dras-

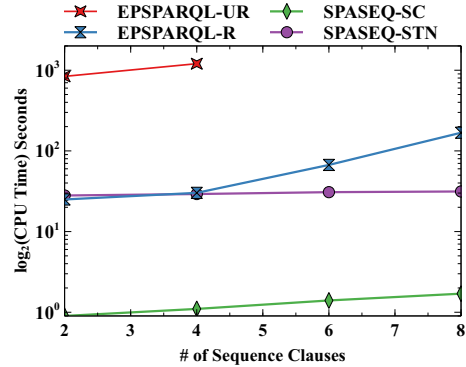


Figure 17. Comparison of SPASEQ (Strict Contiguity (SC) and Skip-Till-Next (STN)) with EP-SPARQL (Recent (R) and Unrestricted (UR)) by increasing the number of sequence clause for SMD dataset and queries over a window of 40 minutes.

tically compared with EP-SPARQL. For EP-SPARQL, increasing the number of sequence clauses results in more complex backward chain reasoning for incoming event; hence the increase in the CPU cost.

Question 2. How does EP-SPARQL and SPASEQ perform using the background KB?

To analyse the CPU cost by using the background KB, we use the SMD and CCD datasets and their respective queries. We generated the simulated background KB for these datasets. The background KB of SMD contains information about the company name, address, URL, phone number, while the background KB of CCD contains information about the cardholder name, issuing bank, bank address and type of the card. We increase the size of the background KB from 10K to 1M triples. We do not use show the CPU time for EP-SPARQL-UR for this set of experiments, since it does not perform well with the integration of background KB.

Fig. 18 shows the performance comparison of both systems. From Fig. 18, the CPU cost of EP-SPARQL-R increases quadratically with the increase in the number of triples for background KB. The reasons are as follows. EP-SPARQL bases its reasoning process on the ETALIS engine, where the background KB is used to map the complete set of inference rules before starting the event processing. With the arrival of a new event, not only the defined sequences are matched, but also all the inference rules are triggered at the same time. Hence, EP-SPARQL spends considerable amount of time on matching the events with a large set of rules, even for the events that would not produce matches. SPASEQ, on the other hand, employs an optimised way

of processing background information due to its separation of the query constructs. That is, the evaluation of sequence clauses in a SPASEQ query is separated from the evaluation of the background KB. Hence, triples from the background KB are joined only when the BGP defined in the query is matched with the incoming event. That is, the background KB is materialised only once by considering the patterns defined on the background KB, and only the events that match with the defined BGP are matched with the materialised background KB. Hence, unnecessary computation costs are saved. Furthermore, we employ the optimised indexing techniques from the SPECTRA [45] – our underlying RDF processing engine – for an efficient computation of the background KB and incoming events. For these reasons, the CPU cost of SPASEQ increases in a sub-linear manner by increasing the number of triples in the background KB. The results from this set of experiments also highlights the importance of providing explicit constructs for the background KB, rather than using it at the implementation level.

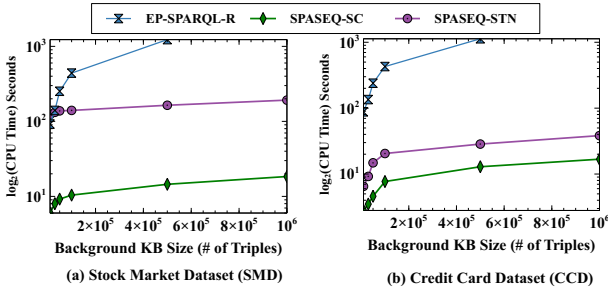


Figure 18. Comparison of EP-SPARQL and SPASEQ over the background KB using SMD and CCD datasets and queries

8.3. General Analysis of the SPASEQ Query Engine

In this section, we present the general analysis of SPASEQ’s query evaluation strategies and the cost of its temporal operators. We start by describing the cost of its internal operations.

Question 3. What is the cost of various operations for the SPASEQ engine?

The analysis of the cost of various SPASEQ operations would aid us in discovering the bottlenecks of the systems. For this set of experiments we use the SMD dataset and query Q3, since it is the most expensive from our previous analysis.

Fig. 19, shows the cost of three main operations of the SPASEQ query engine, namely parsing time, GPM

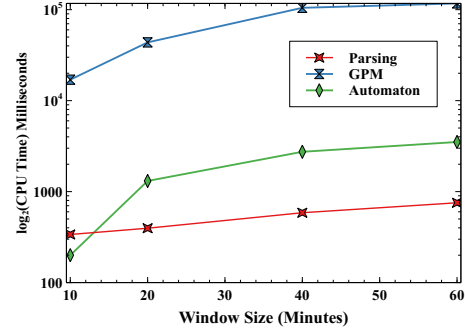


Figure 19. Cost analysis (in milliseconds) of various SPASEQ operations using the SMD and STN event selection strategy.

evaluation time and automaton evaluation time. The parsing time is the time to parse the incoming set of triples in N3 form to the internal representation of the SPASEQ engine. The GPM evaluation time is the total time to match each incoming event with the defined GPM expressions mapped at the automaton states. The automaton evaluation time determines the time to (i) create new automaton runs, and (ii) transit from one state to another. From Fig. 19, it is evident the GPM evaluation is the most expensive operation compared with parsing and automaton evaluation. However, this is expected since the RDF data model – which requires GPM – provides much more expressiveness than the traditional relational data model and is at the backbone of the Semantic Web. The cost of the automaton evaluation is the standard one since we need to create new runs to start a new partial match, and we proceed to the next state to produce the full matches. The parsing time is proportional to the number of parsed events and is quite less than the other operations. Note that the parsing time can be reduced by parsing the simple text-based triples – as used by many RSP systems and EP-SPARQL – than the standard full URI-based triples in N3 form.

The cost of SPASEQ operations shows us that optimising the GPM evaluation has a direct effect on the overall cost of a SCEP system. Since we use SPECTRA [45] as our underlying GPM engine – which outperforms existing RSP systems – the SPASEQ engine provides an optimised performance. Next, we determine the effect of our optimisation strategy on the performance of the SPASEQ engine.

Question 3. How does the cost of a simple sequence clause compare with the cost of the Kleene+ Operator?

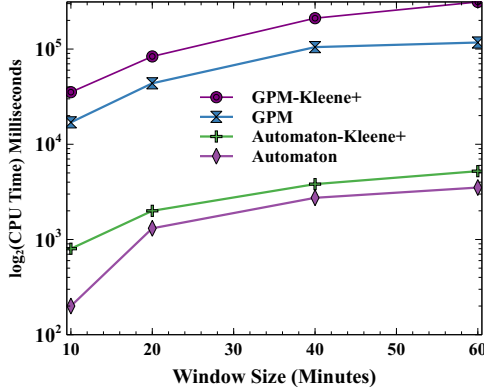


Figure 20. Cost analysis (in milliseconds) of various SPASeq operations for *Kleene+* operator and simple sequence operators using the SMD and STN event selection strategy.

In this set of experiments, we compare the CPU cost of GPM and automaton evaluation for the *Kleene+* operator and simple sequence clause using SMD and its respective queries. We do not compare the parsing time since it is the same for both queries. From Fig. 20, we can see that *Kleene+* operator results in an increase of the GPM and automaton evaluation costs'. This is because the number of runs generated for a query with a simple sequence clause is relatively proportional to the number of events that result in starting a new run from its initial state x_0 . However, the same is not the case with the *Kleene+* operator. If an event matches to the *Kleene+* operator, the system duplicates an additional run and adds it to the active runs list; hence to match the one or more relations over event streams. This means, following the same intuition from earlier, each newly arrived event has to process a large number of active runs. This results in an extra cost for the *Kleene+* operator to process an event. Nevertheless, our system provides good performance measures for the expressive operators such as the *Kleene+*. The experimental cost of the *Kleene+* operator is in line with our theoretical analysis in Section 7.

Question 4. How does the strategy for indexing runs by stream names affect the performance of the system?

In order to determine the effectiveness of indexing runs by stream names, we employ the SHD and Q1 for UC 1 with STN event selection strategy, since it is costly compared with the SC strategy. Recall from Section 7.2, we index runs by stream name, thus when an event arrives, only the runs whose active state is waiting

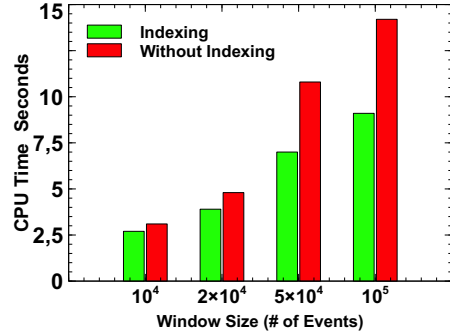


Figure 21. Analysis of the cost of indexing runs by stream names using SHD dataset and its query.

for such an event are used. Consequently, it reduces the overhead of going through the whole list of available active runs. Fig. 21 shows the results of our evaluation with variable window sizes. According to the results, the performance differences between the indexed and non-indexed approach is not evident at smaller windows with less number of events. This is due to the fact that small numbers of runs are produced/remain active for the smaller windows, hence indexing of runs does not results in a comparatively smaller set of runs to be probed for each event. However, the effectiveness of the indexing technique becomes quite clear with the increase in the window size. That is, a large number of runs is produced with a smaller set of them waiting for an event from a specific stream.

Question 4. How does the lazy evaluation affect the performance of the conjunction operator?

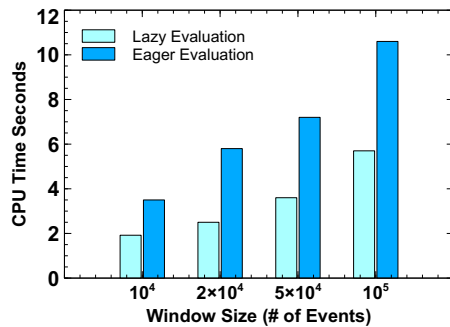


Figure 22. Analysis of eager and lazy evaluation strategies for conjunction operator using TD dataset and its query.

For this set of experiments, we employ the TC dataset and its respective query Q5 with conjunction operator and STN event selection strategy. Fig. 22 shows the results of the conjunction operator with lazy and

eager evaluation strategies. Recall from Section 7.3, lazy evaluation delays the computation of all the state-transition predicates until the number of the events with the same timestamp is equal to the number of state-transition predicates. As shown in Fig. 22, the lazy evaluation strategy performs much better on smaller windows and relatively better on larger ones: the eager evaluation results in a larger number of useless calls to the GPM evaluator, while lazy evaluation performs a batch-based call to the GPM evaluator. Thus, with lazy evaluation, a set of events is evaluated against a set of GPM expressions, only if all the buffered events (for a conjunction state) has the same timestamp. For the smaller window, if the number of buffered events is not equal to the number of edges from a conjunction state, the GPM evaluator is not invoked. Hence, with the expiration of the window, the runs are deleted without matching events and without using the additional resources. Contrary to this, the eager evaluation strategy calls the GPM evaluator for each incoming event and a large number of such calls proved to be useless for smaller windows.

9. Conclusion

In this paper, we presented the syntax, semantics and implementation of a SCEP query language called SPASEQ. We provided the motivation behind SPASEQ and pointed out various qualitative differences between other SCEP languages. Such analysis showcased the usability and expressivity of the SPASEQ query language. During the design phase of our language, we carefully consulted existing CEP techniques and the lessons learned. Thus, a suitable compromise between the expressiveness of a SCEP language and how it can be implemented in an effective way is made possible. We also proposed an NFA_{scep} model to map the SPASEQ operators and showed how they can be evaluated over a streamset. Furthermore, we also provided multiple optimisation techniques to evaluate SPASEQ operators in an optimised manner. Lastly, while utilising real-world datasets we showcased the usability and performance of SPASEQ query engine. Our future endeavours include: extension of the language with new operators, a through semantic comparative analysis with the EP-SPARQL, further optimisation strategies for the *Kleene+* operator and the evaluation of the SPASEQ operators in a distributed environment. We believe that SPASEQ can ignite the SCEP research community and will open the doors for the new insights and optimisation techniques in this field.

References

- [1] Lukasz Golab and M. Tamer Özsu. Issues in data stream management. *SIGMOD Rec.*, 32(2):5–14, June 2003.
- [2] Daniel J. Abadi, Don Carney, Ugur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. Aurora: A new model and architecture for data stream management. *The VLDB Journal*, 12(2):120–139, August 2003.
- [3] Gianpaolo Cugola and Alessandro Margara. Processing flows of information: From data stream to complex event processing. *ACM Comput. Surv.*, 44(3):15:1–15:62, June 2012.
- [4] Eugene Wu, Yanlei Diao, and Shariq Rizvi. High-performance complex event processing over streams. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*, SIGMOD '06, pages 407–418, New York, NY, USA, 2006. ACM.
- [5] Yuan Mei and Samuel Madden. Zstream: A cost-based query processor for adaptively detecting composite events. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, SIGMOD '09, pages 193–206, New York, NY, USA, 2009. ACM.
- [6] David Luckham. The power of events: An introduction to complex event processing in distributed enterprise systems. In *Rule Representation, Interchange and Reasoning on the Web*, pages 3–3, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [7] Barzan Mozafari, Kai Zeng, and Carlo Zaniolo. High-performance complex event processing over xml streams. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, SIGMOD '12, pages 253–264, New York, NY, USA, 2012. ACM.
- [8] Jagrati Agrawal, Yanlei Diao, Daniel Gyllstrom, and Neil Immerman. Efficient pattern matching over event streams. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 147–160, New York, NY, USA, 2008. ACM.
- [9] Barzan Mozafari, Kai Zeng, Loris D'antoni, and Carlo Zaniolo. High-performance complex event processing over hierarchical data. *ACM Trans. Database Syst.*, 38(4):21:1–21:39, December 2013.
- [10] Lars Brenna, Alan Demers, Johannes Gehrke, Mingsheng Hong, Joel Ossher, Biswanath Panda, Mirek Riedewald, Mohit Thatte, and Walker White. Cayuga: A high-performance event processing engine. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, SIGMOD '07, pages 1100–1102, New York, NY, USA, 2007. ACM.
- [11] Thomas BERNHARDT and Alexandre VASSEUR. ESPER-complex event processing. In *Online Article*, 2010.
- [12] Drool fusion. <http://www.drools.org/>. Accessed: 2016-06-03.
- [13] Darko Anicic, Sebastian Rudolph, Paul Fodor, and Nenad Stojanovic. Stream reasoning and complex event processing in ETALIS. *Semant. web*, 3(4):397–407, October 2012.
- [14] Apache Flink. <https://flink.apache.org/>.
- [15] Alessandro Margara, Jacopo Urbani, Frank van Harmelen, and Henri Bal. Streaming the web: Reasoning over dynamic data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 25(0), 2014.
- [16] Christian Bizer, Tom Heath, and Tim Berners-Lee. Linked data - the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.

- [17] Tom Heath and Christian Bizer. Linked data: Evolving the web into a global data space (1st edition). In *Synthesis Lectures on the Semantic Web: Theory and Technology*. Morgan and Claypool., 2011.
- [18] Danh Le-Phuoc, Minh Dao-Tran, Josiane Xavier Parreira, and Manfred Hauswirth. A native and adaptive approach for unified processing of linked streams and linked data. In *Proceedings of the 10th International Conference on The Semantic Web - Volume Part I, ISWC'11*, pages 370–388, Berlin, Heidelberg, 2011. Springer-Verlag.
- [19] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, and Michael Grossniklaus. Querying rdf streams with c-sparql. *SIGMOD Rec.*, 39(1):20–26, September 2010.
- [20] Jean-Paul Calbimonte, Oscar Corcho, and Alasdair J. G. Gray. Enabling ontology-based access to streaming data sources. In *Proceedings of the 9th International Semantic Web Conference on The Semantic Web - Volume Part I, ISWC'10*, pages 96–111, Berlin, Heidelberg, 2010. Springer-Verlag.
- [21] Darko Anicic, Paul Fodor, Sebastian Rudolph, and Nenad Stojanovic. Ep-sparql: A unified language for event processing and stream reasoning. In *Proceedings of the 20th International Conference on World Wide Web, WWW '11*, pages 635–644, New York, NY, USA, 2011. ACM.
- [22] Özgür L. Özcepe Veronika Thost, Jan Holste. On Implementing Temporal Query Answering in DL-Lite (extended abstract). In Diego Calvanese and Boris Konev, editors, *Proceedings of the 28th International Workshop on Description Logics, Athens, Greece, June 7-10, 2015.*, volume 1350 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2015.
- [23] Daniele Dell’Aglío, Minh Dao-Tran, Jean-Paul Calbimonte, Danh Le Phuoc, and Emanuele Della Valle. A query model to capture event pattern matching in RDF stream processing query languages. In *Knowledge Engineering and Knowledge Management - 20th International Conference, EKAW 2016, Bologna, Italy, November 19-23, 2016, Proceedings*, pages 145–162, 2016.
- [24] Jeremy J. Carroll, Christian Bizer, Pat Hayes, and Patrick Stickler. Named graphs, provenance and trust. In *Proceedings of the 14th International Conference on World Wide Web, WWW '05*, pages 613–622, New York, NY, USA, 2005. ACM.
- [25] Oscar Corcho Daniele Dell’Aglío Emanuele Della Valle Shen Gao Alasdair J G Gray Danh Le-Phuoc Robin Keskisärkkä Alejandro Llaves Alessandra Mileo Bernhard Ortner Adrian Paschke Monika Solanki Roland Stühmer Kia Teymourian Peter Wetz Darko Anicic, Jean-Paul Calbimonte. W3C Community Group. Latest version available as http://streamreasoning.github.io/RSP-QL/RSP_Requirements_Design_Document/#.
- [26] Ioannis Flouris, Nikos Giatrakos, Antonios Deligiannakis, Minos Garofalakis, Michael Kamp, and Michael Mock. Issues in complex event processing. *J. Syst. Softw.*, 127(C):217–236, May 2017.
- [27] Alexander Artikis, Nikos Katzouris, Ivo Correia, Chris Baber, Natan Morar, Inna Skarbovsky, Fabiana Fournier, and Georgios Paliouras. A prototype for credit card fraud management: Industry paper. In *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems, DEBS '17*, pages 249–260, New York, NY, USA, 2017. ACM.
- [28] Olga Poppe, Chuan Lei, Salah Ahmed, and Elke A. Rundensteiner. Complete event trend detection in high-rate event streams. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17*, pages 109–124, New York, NY, USA, 2017. ACM.
- [29] Barzan Mozafari, Kai Zeng, and Carlo Zaniolo. From regular expressions to nested words: Unifying languages and query execution for relational and XML sequences. *PVLDB*, 3(1):150–161, 2010.
- [30] Syed Gillani, Frédérique Laforest, and Gauthier Picard. A generic ontology for prosumer-oriented smart grid. In *EDBT/ICDT Workshops*, volume 1133 of *CEUR Workshop Proceedings*, pages 134–139. CEUR-WS.org, 2014.
- [31] Guangyi Liu, Wendong Zhu, Chris Saunders, Feng Gao, and Yang Yu. Real-time complex event processing and analytics for smart grid. *Procedia Computer Science*, 61(Supplement C):113 – 119, 2015. Complex Adaptive Systems San Jose, CA November 2-4, 2015.
- [32] Nijat Mehdiyev, Julian Krumeich, David Enke, Dirk Werth, and Peter Loos. Determination of rule patterns in complex event processing using machine learning techniques. *Procedia Computer Science*, 61(Supplement C):395 – 401, 2015. Complex Adaptive Systems San Jose, CA November 2-4, 2015.
- [33] Elias Alevizos, Alexander Artikis, and George Paliouras. Event forecasting with pattern markov chains. In *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems, DEBS '17*, pages 146–157, New York, NY, USA, 2017. ACM.
- [34] Cagri Balkesen, Nihal Dindar, Matthias Wetter, and Nesime Tatbul. Rip: Run-based intra-query parallelism for scalable complex event processing. In *Proceedings of the 7th ACM International Conference on Distributed Event-based Systems, DEBS '13*, pages 3–14, New York, NY, USA, 2013. ACM.
- [35] Yuan Mei and Samuel Madden. Zstream: A cost-based query processor for adaptively detecting composite events. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data, SIGMOD*, pages 193–206, New York, NY, USA, 2009. ACM.
- [36] Muhammad Intizar Ali, Feng Gao, and Alessandra Mileo. *City-Bench: A Configurable Benchmark to Evaluate RSP Engines Using Smart City Datasets*, pages 374–389. Springer International Publishing, Cham, 2015.
- [37] Chunyao Song, Tingjian Ge, Cindy Chen, and Jie Wang. Event pattern matching over graph streams. *Proc. VLDB Endow.*, 8(4):413–424, December 2014.
- [38] Jeremy J. Carroll, Christian Bizer, Pat Hayes, and Patrick Stickler. Named graphs. *Web Semantics: Science, Services and Agents on the World Wide Web*, 3(4):247 – 267, 2005. World Wide Web Conference 2005—Semantic Web Track.
- [39] Aidan Hogan, Marcelo Arenas, Alejandro Mallea, and Axel Polleres. Everything you always wanted to know about blank nodes. *Web Semantics: Science, Services and Agents on the World Wide Web*, 27–28:42 – 69, 2014. Semantic Web Challenge 2013.
- [40] Raman Adaikkalavan and Sharma Chakravarthy. Snooipib: Interval-based event specification and detection for active databases. *Data Knowl. Eng.*, 59(1):139–165, October 2006.
- [41] Yanlei Diao and Neil Immerman. Sase+: An agile language for kleene closure over event streams. *UMASS Technical Report*, 2007.
- [42] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, and Michael Grossniklaus. Querying rdf streams with c-sparql. *SIGMOD Rec.*, 39(1):20–26, September

- 2010.
- [43] Srdjan Komazec, Davide Cerri, and Dieter Fensel. Sparkwave: Continuous schema-enhanced pattern matching over rdf data streams. In *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems*, DEBS '12, pages 58–68, New York, NY, USA, 2012. ACM.
- [44] Arvind Arasu, Shivnath Babu, and Jennifer Widom. The cql continuous query language: Semantic foundations and query execution. *The VLDB Journal*, 15(2):121–142, June 2006.
- [45] Syed Gillani, Gauthier Picard, and Frédérique Laforest. Spectra: Continuous query processing for rdf graph streams over sliding windows. In *Proceedings of the 28th International Conference on Scientific and Statistical Database Management*, SSDBM '16, pages 17:1–17:12, New York, NY, USA, 2016. ACM.
- [46] Minh Dao-Tran and Danh Le Phuoc. Towards enriching CQELS with complex event processing and path navigation. In *Proceedings of the 1st Workshop on High-Level Declarative Stream Processing co-located with the 38th German AI conference (KI 2015)*, Dresden, Germany, September 22, 2015., pages 2–14, 2015.
- [47] Riccardo Tommasini, Pieter Bonte, Emanuele Della Valle, Erik Mannens, Filip De Turck, and Femke Ongenaë. *Towards Ontology-Based Event Processing*, pages 115–127. Springer International Publishing, Cham, 2017.
- [48] Özgür Lütfü Özçep, Ralf Möller, and Christian Neuenstadt. A stream-temporal query language for ontology based data access. In Carsten Lutz and Michael Thielscher, editors, *KI 2014: Advances in Artificial Intelligence*, pages 183–194, Cham, 2014. Springer International Publishing.
- [49] Richard Cyganiak, David Wood, and Markus Lanthaler. RDF 1.1 Concepts and Abstract Syntax. Technical report, W3C, January 2014.
- [50] Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of sparql. *ACM Trans. Database Syst.*, 34(3):16:1–16:45, September 2009.
- [51] Charles L. Forgy. Expert systems. chapter Rete: A Fast Algorithm for the Many Pattern/Many Object Pattern Match Problem, pages 324–341. IEEE Computer Society Press, Los Alamitos, CA, USA, 1990.
- [52] S. Gatzju and K. R. Dittrich. Detecting composite events in active database systems using petri nets. In *Research Issues in Data Engineering, 1994. Active Database Systems. Proceedings Fourth International Workshop on*, pages 2–9, Feb 1994.
- [53] Haopeng Zhang, Yanlei Diao, and Neil Immerman. On complexity and optimization of expensive queries in complex event processing. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 217–228, New York, NY, USA, 2014. ACM.
- [54] Ilya Kolchinsky, Izchak Sharfman, and Assaf Schuster. Lazy evaluation methods for detecting complex events. In *Proceedings of the 9th ACM International Conference on Distributed Event-Based Systems*, DEBS '15, pages 34–45, New York, NY, USA, 2015. ACM.
- [55] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation (3rd Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2006.
- [56] Thomas Neumann and Gerhard Weikum. The rdf-3x engine for scalable management of rdf data. *The VLDB Journal*, 19(1):91–113, February 2010.
- [57] Jeremy J. Carroll, Ian Dickinson, Chris Dollin, Dave Reynolds, Andy Seaborne, and Kevin Wilkinson. Jena: Implementing the semantic web recommendations. In *Proceedings of the 13th International World Wide Web Conference on Alternate Track Papers & Posters*, WWW Alt. '04, pages 74–83, New York, NY, USA, 2004. ACM.
- [58] Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of sparql. *ACM Trans. Database Syst.*, 34(3):16:1–16:45, September 2009.
- [59] Syed Gillani, Gauthier Picard, and Frédérique Laforest. Continuous graph pattern matching over knowledge graph streams. In *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems*, DEBS '16, pages 214–225, New York, NY, USA, 2016. ACM.
- [60] Stock Trace Dataset. http://davis.wpi.edu/datasets/Stock_Trace_Data/.
- [61] David Irwin Emmanuel Cecchet Prashant Shenoy Sean Barker, Aditya Mishra and Jeannie Albrecht. Smart: An open data set and tools for enabling research in sustainable homes. In *Proceedings of the 2012 Workshop on Data Mining Applications in Sustainability (SustKDD 2012)*, 2012.
- [62] Elias Alevizos, Alexander Artikis, and George Paliouras. Event forecasting with pattern markov chains. In *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems*, DEBS '17, pages 146–157, New York, NY, USA, 2017. ACM.
- [63] Mark Sanderson and Justin Zobel. Information retrieval system evaluation: Effort, sensitivity, and reliability. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '05, pages 162–169, New York, NY, USA, 2005. ACM.
- [64] Medhabi Ray, Chuan Lei, and Elke A. Rundensteiner. Scalable pattern sharing on event streams*. In *Proceedings of the 2016 International Conference on Management of Data*, SIGMOD '16, pages 495–510, New York, NY, USA, 2016. ACM.
- [65] Ken Thompson. Programming techniques: Regular expression search algorithm. *Commun. ACM*, 11(6):419–422, June 1968.

Appendix

A. SPASEQ By Examples

In this section, we provide the SPASEQ queries for the use cases described in Section 2

```

1 PREFIX pred: <http://example/>
2 SELECT ?vessel ?n ?cname
3 WITHIN 30 MINUTES
4 FROM STREAM S1 <http://creditCard.org/boats>
5
6 WHERE {
7
8   SEQ (A, B+)
9
10  DEFINE GPM A ON S1 {
11    ?card pred:id ?id1.
12    ?card pred:amount ?amount1.
13    ?card pred:loc ?loc.
14    FILTER (?s1 > 0)
15  }
16
17  DEFINE GPM B ON S1 {
18    ?card pred:id ?id1.

```

```

19 ?card pred:amount ?amount2.
20
21 FILTER (?amount2 *10 >= ?amount1)
22
23 GRAPH <http://creditcard.org/db> {
24   ?id :cardHolder ?name.
25   ?id :holderAddress ?address.
26   }
27 }
28
29
30 }

```

Query 2: Credit Card Fraud Detection, Big after Small: SPASEQ query

```

1 PREFIX pred: <http://example/>
2 SELECT ?company ?p1 ?p2 ?p3 ?p4 ?p5 ?p6 ?p7 ?vol1 ?vol2
   ?vol3 ?vol4 ?vol5 ?vol6 ?vol7
3 WITHIN 60 MINUTES
4 FROM STREAM S1 <http://stockmarket.org/stocks/google>
5
6 WHERE {
7
8   SEQ (A, B, C, D, E, F, G)
9
10  DEFINE GPM A ON S1 {
11    ?company pred:price ?p1.
12    ?company pred:volume ?vol1.
13  }
14
15  DEFINE GPM B ON S1 {
16    ?company pred:price ?p2.
17    ?company pred:volume ?vol2.
18    FILTER (?p2 > ?p1)
19  }
20
21  DEFINE GPM C ON S1 {
22    ?company pred:price ?p3.
23    ?company pred:volume ?vol3.
24    FILTER (?p3 < ?p2 && ?p3 > ?p1).
25  }
26
27  DEFINE GPM D ON S1 {
28    ?company pred:price ?p4.
29    ?company pred:volume ?vol4.
30    FILTER (?p4 > ?p2).
31  }
32
33  DEFINE GPM E ON S1 {
34    ?company pred:price ?p5.
35    ?company pred:volume ?vol5.
36    FILTER (?p5 < ?p4 && ?p5 > ?p3).
37  }
38
39  DEFINE GPM F ON S1 {
40    ?company pred:price ?p6.
41    ?company pred:volume ?vol6.
42    FILTER (?p6 > ?p5 && ?p6 < ?p4).
43  }
44
45  DEFINE GPM G ON S1 {
46    ?company pred:price ?p7.
47    ?company pred:volume ?vol7.
48    FILTER (?p7 < ?p6 && ?p7 > ?p5).
49  }
50 }

```

Query 3: Head and Shoulders Pattern: SPASEQ query

```

1 PREFIX pred: <http://example/>
2 SELECT ?company ?p1 ?p2 ?p3 ?vol1 ?vol2 ?vol3
3 WITHIN 60 MINUTES
4 FROM STREAM S1 <http://stockmarket.org/stocks/google>
5
6 WHERE {
7
8   SEQ (A, B, C,)
9
10  DEFINE GPM A ON S1 {
11    ?company pred:price ?p1.
12    ?company pred:volume ?vol1.
13  }
14
15  DEFINE GPM B ON S1 {
16    ?company pred:price ?p2.
17    ?company pred:volume ?vol2.
18    FILTER (?p2 > ?p1)
19  }
20
21  DEFINE GPM C ON S1 {
22    ?company pred:price ?p3.
23    ?company pred:volume ?vol3.
24    FILTER (?p3 < ?p2 ).
25  }
26
27 }

```

Query 4: V-Shaped Pattern: SPASEQ query

```

1 PREFIX ssn: <http://purl.oclc.org/NET/ssnx/ssn#>
2 PREFIX sao: <http://purl.oclc.org/NET/sao/ssn#>
3 SELECT ?vessel ?n ?cname
4 WITHIN 30 MINUTES
5 FROM STREAM S1 <http://www.insight-centre.org/dataset/
   SampleEventService#AarhusTrafficData182955>
6 FROM STREAM S2 <http://www.insight-centre.org/dataset/
   SampleEventService#AarhusTrafficData195578>
7 FROM STREAM S3 <http://www.insight-centre.org/dataset/
   SampleEventService#AarhusTrafficData195446>
8
9 WHERE {
10
11  SEQ (A; (B&C))
12
13  DEFINE GPM A ON S1 {
14    ?traffic ssn:observedProperty ?p1.
15    ?traffic sao:hasValue ?v1.
16    FILTER (?v1 > 3)
17  }
18
19  DEFINE GPM B ON S2 {
20    ?traffic ssn:observedProperty ?p1.
21    ?traffic sao:hasValue ?v2.
22
23    FILTER (?v2 >= ?v1)
24  }
25
26  DEFINE GPM C ON S3 {
27    ?traffic ssn:observedProperty ?p1.
28    ?traffic sao:hasValue ?v3.
29
30    FILTER (?v3 >= ?v2)
31  }
32
33 }
34 }

```

Query 5: Traffic Management: SPASEQ query

```

1 PREFIX pred: <http://example/>
2 SELECT ?vessel ?n ?cname
3 WITHIN 30 MINUTES
4 FROM STREAM S1 <http://harbour.org/boats>
5
6 WHERE {
7
8   SEQ (A, B+, C)
9
10  DEFINE GPM A ON S1 {
11    ?vessel pred:speed ?s1.
12    ?vessel pred:location "harbour".
13    ?vessel pred:direction ?dir1.
14    FILTER (?s1 > 0)
15  }
16
17  DEFINE GPM B ON S1 {
18    ?vessel pred:speed ?s2.
19    ?vessel pred:location ?l.
20    ?vessel pred:direction ?dir1.
21    FILTER (?s2 > ?s1)
22  }
23
24  DEFINE GPM C ON S1 {
25    ?vessel pred:speed 0.
26    ?vessel pred:location "fishingarea".
27    ?vessel pred:direction ?dir3.
28
29    GRAPH <http://harbour.org/db> {
30      ?vessel :name ?n.
31      ?vessel :operatedBy ?company.
32      ?company :name ?cname.
33    }
34  }
35 }

```

Query 6: Trajectory Classification: SPASEQ query

B. A Note on the Extension of SPASEQ

As discussed earlier, the design of SPASEQ encouraged the extensibility of the language with new operators. Herein, we present two operators and discuss how they can be integrated into the SPASEQ query model and their effects on the semantics of SPASEQ.

B.0.1. The case of Negation Operator

Negation is a unary operator and is used to describe the non-occurrence of certain events. The standalone semantics of the negation operator can easily be defined. However, discrepancies arise when it is used with the event selection strategies. For instance, the evaluation of the negation operator for a GPM expression (u, P) over a streamset Σ and the time boundaries $[\tau_b, \tau_e]$ can be described as follows:

$$\llbracket (u, P)! \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \{(\tau, \emptyset) \mid \exists (u, S) \in \Sigma \wedge \forall \tau (\tau, G) \in S, \llbracket P \rrbracket_G = \emptyset \wedge \tau_b \leq \tau \leq \tau_e\}$$

Thus, for each event that does not match with the GPM expression, the evaluation of the negation operator returns an empty set associated with a timestamp.

The use of the negation operator in conjunction with the *skip-till-next* or *skip-till-any* operator results in discrepancies, where it is difficult to differentiate between the results of a negation operator and the results of a simple GPM evaluation in case of non-occurrence of an event. Therefore, the negation operator requires a new structure such that we can differentiate between the empty set from the evaluation of GPM expressions and the GPM expressions with the negation operator. That is, contrary to the natural join of mappings with empty set $\emptyset \bowtie \Omega = \Omega \bowtie \emptyset = \emptyset$, for the negation operator we need a structure such that

$$? \bowtie \Omega = \Omega \bowtie ? = \Omega$$

An identity element from a commutative monoid family can be employed to showcase the aforementioned behaviour of the negation operator.

B.0.2. The case of Optional Operator

The optional operator selects an event if it matches to the defined GPM expression; otherwise it ignores the event. Since it corresponds to the zero or at-most one occurrence of an event, it suffers from the same issues as discussed for the negation operator. Hence, the remedies for the negation operator can directly be applied to define its semantics. For instance, let us denote the optional operator with ‘?’; then its evaluation can be described using the results of the negation operator as follows:

$$\llbracket (u, P)? \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} = \llbracket (u, P) \rrbracket_{\Sigma}^{[\tau_b, \tau_e]} \cup \llbracket (u, P)! \rrbracket_{\Sigma}^{[\tau_b, \tau_e]}$$

The above discussion highlights what kinds of issues arise with the integration of the negation and optional operators in SPASEQ and point-outs some of the remedies. In this paper, we do not present the complete semantics of these operator to keep the discussion focused on the core operators of SPASEQ. However, in Appendix C.4, we present some of the techniques to implement these operators. The description of their detailed semantics will be the topic of our future endeavours.

C. Evaluation of SPASEQ Operators

Herein, we present the details how SPASEQ operators are evaluated using the NFA_{scep} model.

C.1. Evaluation of the Kleene+ Operator

Previously we show the generic execution for NFA_{scpep} automaton for a sequence expression. Herein, we present how the unary operator, i.e. *Kleene+*, is evaluated. The evaluation of *Kleene+* operator is described in Algorithm 3 with details as follows.

Algorithm 3: Evaluation of the *Kleene+* Operator

```

1 Function KLEENEPLUS( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$ )
2   get  $\theta$  for an edge  $e \in E$  s.t  $\theta = (U_\theta, op_\theta, P_\theta)$ 
3   if  $u \in U_\theta$  and  $GPM(G_e, P_\theta, \mathcal{H})$  then
4     clone a new run  $r_c$  from  $r$  with active state
      as  $x_i$ 
5     SETACTIVESTATE( $r_c, x_{i+1}$ ) where
       $x_{i+1} \neq x_i$ 
      // If it is the final state
6     if  $x_i = x_f$  of  $r_c$  then
7       a query match has been found
8       remove  $r$  from the list of active runs  $\mathcal{R}$ 
9     else
10       $\mathcal{R} \leftarrow \mathcal{R} \cup r_c$ 
11  else if  $u \notin U_\theta$  or  $\neg GPM(G_e, P_\theta, \mathcal{H})$  then
12    remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

The evaluation of the *Kleene+* operator is an interesting one, since the state-transition predicates of the two edges are not mutually exclusive. Thus, to cater the non-determinism of the *Kleene+* operator, a new run is duplicated/cloned from the existing one in case of a match. Algorithm 3 presents the evaluation of *Kleene+* operator. It first uses an edge from the active state to employ the comparison of stream names to make sure the event is from the stream the edge is waiting for (line 3). The algorithm then uses event G_e , graph pattern P_θ and history cache \mathcal{H} to execute the GPM process (line 3). If the newly arrived event G_e is matched with the graph pattern P_θ : (i) a new run r_c is cloned from the run under evaluation with the same active state x_i (line 4); (ii) the cloned run transits to the next state x_{i+1} (line 5); (iii) if the new active state of the cloned run is the final state then a query match has been found (lines 6,7); (iv) otherwise the cloned run is added to the list of active runs \mathcal{R} (line 9). It then skips to the next run in \mathcal{R} , hence the run under the evaluation stays at the same state. Otherwise, in case of no match, it removes the run r from the list of active runs \mathcal{R} (lines 10-11). This way the system can keep track of one or more matched events of the same kind (see Figure 23).

Example 14 Consider Figure 23. In this example, r_i represents run i , x_0, x_1, x_2 , and x_f represent the states using sequence expression SEQ ($A, B+, C$) (From SPASEQ Query 1), and G_e^k represents an event that occurred at time k . The arrival of G_e^1 results in a new run r_1 and the automaton transits from state x_0 to x_1 if G_e^1 matches $(U_{s_1}, \emptyset, P_A)$. Now considering the next event G_e^2 matches $(U_{s_2}, +, P_B)$; the automaton results in a non-deterministic move due to the *Kleene+* operator at the state x_2 . Hence, the algorithm creates a new run r_2 with active state as x_2 , i.e. transiting from x_1 , while r_1 stays at the same state x_1 . When G_e^3 arrives and matches to $(U_{s_3}, \emptyset, P_C)$, r_2 moves to the final state and the match for r_2 is complete with events G_e^1, G_e^2 and G_e^3 , while r_1 remains active to consume for future events. Finally, after the arrival and match of events G_e^4 and G_e^5 with the corresponding GPM expressions, r_1 reaches the final state with a match using events G_e^1, G_e^4 and G_e^5 .

Algorithm 4: Evaluation of the Event Selection Strategies


```

1 Function EVENTSELECTION( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$ )
2   get  $\theta$  for an edge  $e \in E$  s.t  $\theta = (U_\theta, op_\theta, P_\theta)$ 
3   if  $u \in U_\theta$  and  $GPM(G_e, P_\theta, \mathcal{H})$  then
4     SETACTIVESTATE( $r, x_{i+1}$ )
      // If it is the final state
5     if  $x_i = x_f$  then
6       a query match has been found
7       remove  $r$  from the list of active runs  $\mathcal{R}$ 
8   else if  $op_\theta = \text{';'}$  and  $u \notin U_\theta$  or  $\neg GPM(G_e, P_\theta, \mathcal{H})$  then
9     skip the event for the followed-by operator
10  else
11    remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

C.2. Evaluation of the Event Selection Operators

This section presents the evaluation of event selection strategies. Algorithm 4 shows the evaluation of *strict contiguity* and *skip-till-next*. Herein, we only discuss the implementation of the skip-till-next operator. The optimised implementation of skip-till-any operator will be the topic of our future endeavours.

Figure 23. Execution of NFA_{scep} runs for the SPASEQ Query 1, as described in Example 14

C.2.1. Strict Contiguity Operator

The evaluation of the *strict contiguity* operator is rather simple due to the strictness of how an event should follow other. That is, the incoming event is compared with the state-transition predicate and if there is a match the run transits to the next state; otherwise it is deleted. Therefore, Algorithm 4 first gets the set of edges for the active state and selects the state-transition predicate (line 2). It then compares the stream names and the incoming event G_e with the graph pattern P_θ using the GPM function (line 3). If there is a match, the current active state x_i transits to the next state (line 4). In case the transited state is the last state a query match is found (lines 5-6). Otherwise the run under evaluation is deleted from \mathcal{R} (line 11).

C.2.2. Skip-Till-Next Operator

From our earlier discussion, the *skip-till-next* operator skips the irrelevant events until a query match is completed for the defined sequence expression. This extended functionality is described in Algorithm 4 (lines 8-9). That is, it first compares the stream names. If the event is from the same stream, the state's edge is waiting for, i.e. $u \in U_\theta$, it employs the same procedure as described for the *immediately followed-by* operator. In case the arriving event is from a different stream, i.e. $u \notin U_\theta$ (line 8) or the event is not matched, the algorithm skips the event (line 11). Hence, the run stays at the same state. Algorithm 4 can be extended for the skip-till-any configuration by initiating a new run each time an event is matched with the graph pattern P_θ . However, this will result in an exponential time complexity and will not be suitable for real-world applications.

C.3. Evaluation of the Binary Operators

Finally, we present how the conjunction and disjunction operators are evaluated in an NFA_{scep} automaton.

Algorithm 5: Evaluation of the Conjunction Operator

```

1 Function CONJUNCTION( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, \mathcal{D}, r, x_i, E$ )
2    $\tau \leftarrow \text{GETTIMESTAMP}(G_e)$ 
3    $AME \leftarrow \{\}$  // Already Matched Edges
4   foreach each edge  $e \in E$  do
5     if  $(r, x_i, e) \in \mathcal{D}$  then
6        $AME \leftarrow AME \cup e$ 
7   get timestamp  $\tau_{old}$  for  $a \in AME$  using  $\mathcal{D}$ 
8   if  $\tau \neq \tau_{old}$  then
9     remove  $r$  from the list of active runs  $\mathcal{R}$ 
10  else
11     $E \leftarrow E \setminus AME$ 
12    foreach edge  $e \in E$  do
13      get  $\theta$  from the edge  $e$  s.t.
14       $\theta = (U_\theta, op_\theta, P_\theta)$ 
15      if  $u \in U_\theta$  and  $GPM(G_e, P_\theta, \mathcal{H})$  then
16        insert  $(r, x_i, e)$  and  $\tau$  in  $\mathcal{D}$ 
17         $AME = AME \cup e$ 
18     $E \leftarrow \text{GETEDGEBYSET}(x_i)$ 
19    if  $|AME| = |E|$  then
20       $\text{SETACTIVESTATE}(r, x_{i+1})$ 
21      // If it is the final state
22      if  $x_i = x_f$  then
23        a query match has been found
24        remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

Algorithms 5 and 6 show the execution of these operators.

C.3.1. Conjunction Operator

The case of the conjunction operator is rather complicated: there are two or more outgoing edges – each with a distinct state-transition predicate – and the run should move to the next state if all the state-transition predicates are matched with the consecutive events having the same timestamp. This means we need to track the edges of a conjunction state that are already been matched and their timestamps. Therefore, we use a mapping structure (\mathcal{D}) to store the mapping between a tuple (r, x, e) and a timestamp τ of the matched events; where r is the run, x is state with conjunction operator and e is the edge of the state that is matched with an event having timestamp τ . The evaluation of the conjunction operator is presented in Algorithm 5. Its main evaluation starts by: (i) obtaining the timestamp τ of the newly arrived event (line 2) and (ii) initialising a set (AME) to store already matched edges (line 3). It then iterates over the set of edges E and checks if any of the edges have already been matched or not, using the mapping structure \mathcal{D} , while adding the matched edges to the AME set (lines 4-6). It then extracts the timestamp τ_{old} from an element in AME set using the mapping structure \mathcal{D} (line 7). This timestamp τ_{old} is used to check if the newly arrived event has the same timestamp τ . Otherwise, the algorithm removes the run r from \mathcal{R} , since it has violated the condition of conjunction operator (lines 8-9). If $\tau = \tau_{old}$, it extracts the set of already matched edges (AME). It then removes the edges from edge set E that are already matched with the previous events (line 11). This pruned set E is then used to match the incoming event with the selected edges. The algorithm iterates over E and for each $e \in E$ it takes the state-transition predicate θ to compare the stream name and graph pattern using the GPM function (line 12-14). If there is a match it marks the edge as matched by adding its mapping in \mathcal{D} and inserting the edge in the AME set (lines 15-16). In the end, the algorithm again extracts all the edges for the conjunction state and uses the AME set to check if all the edges are matched or not (line 17-18). In case all the edges have completed the matching procedure it transits the run r to the next state (line 19). If the transited state is the last state a query match is found for the conjunction operator.

C.3.2. Disjunction Operator

The disjunction operator resembles with the conjunction operator. However, in this case the run can transit to the next state if the incoming event matches to at-least one of the state's edges. Algorithm 6 presents

Algorithm 6: Evaluation of the Disjunction Operator

```

1 Function DISJUNCTION( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$ )
2    $Matched \leftarrow false$ 
3   foreach edge  $e \in E$  do
4     get  $\theta$  from the edge  $e$  s.t  $\theta = (U_\theta, op_\theta, P_\theta)$ 
5     if  $u \in U_\theta$  and  $GPM(G_e, P_\theta, \mathcal{H})$  then
6        $Matched \leftarrow true$ 
7   if  $Matched$  then
8     SETACTIVESTATE( $r, x_{i+1}$ )
9     // If it is the final state
10    if  $x_i = x_f$  then
11      a query match has been found
12      remove  $r$  from the list of active runs  $\mathcal{R}$ 
13  else
14    // If none of the edges  $e \in E$ 
15    match with the event  $G_e$ 
16    remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

the evaluation of the disjunction operator. It starts by initiating a variable $Matched$ to keep track if any of the state's edge is matched with the incoming event. It then iterates over each edge using its state-transition predicate θ (line 3-10), and compares the stream names U_θ and the graph pattern P_θ with the event G_e (line 5). If any of the edge matches to the event, it updates the value of the $Matched$ variable (line 6). Next if the value of $Matched$ variable is $true$, it transits the run to the next state (line 8). Similarly to the other algorithms, if such transition resulted in arriving at the final state of the automaton, a query match is found and the run is deleted (line 9-10). Otherwise, if none of the edges are matched with the event, it deletes the run under evaluation (line 13).

In the previous sections, we presented the execution of the SPASEQ temporal operators and how they use the NFA_{scep} automaton to implement the required functionalities. The discussion regarding the run-time optimisations is provided in Section 7.

C.4. Compilation and Evaluation of Negation and Optional Operators

In Section B, we discussed the case of integrating negation and optional operators in SPASEQ. We described the issues that can arise while integrating their semantics with the core SPASEQ operators. Herein, we

showcase the techniques to compile and implement these operators. To allow the expressivity of optional (“?”) and negation (“!”) operators, we extend op and φ in NFA_{scep} definition (Definition 17), such that $op \in \{ \text{'\&'}, \text{'+'}, \text{'|'}, \text{'\&'}, \text{'?'}, \text{'!'}, \text{'\&'} \} \cup \{\emptyset\}$ and $\varphi : E \rightarrow \Theta \cup \{\epsilon\}$, where ϵ denotes the instantaneous transition [65]. The compilation process of optional and negation operators is described as follows:

- *Optional*: The optional operator selects an event if it matches to the defined GPM expression; otherwise it ignores the event. Its compilation results in two edges: one with an ϵ -transition and the other with the defined state-transition predicate. The corresponding NFA_{scep} automaton for $((u_1, P_1)?)$ is illustrated in Figure 24. The ϵ transition allows the transition to the next state without a match.




Figure 24. Compilation of the Optional Operator for $((u_1, P_1)?)$

- *Negation*: This operator detects if either no match of an event occurs or there is no occurrence of the expected event. Thus, it behaves similarly to the optional operators, however the GPM process is opposite. That is, if an event matches to defined GPM expression, then it violates the condition of the sequence. The corresponding NFA_{scep} automata for $((u_1, P_1)!)$ is illustrated in Figure 25.




Figure 25. Compilation of the Negation Operator for $((u_1, P_1)!)$

We now present the evaluation of the negation and optional operators. The evaluation functions of these operators can be integrated in Algorithm 2.

Algorithm 7: Optional Operator’s Evaluation

```

1 Function OPTIONAL( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$ )
2   get  $\theta$  for an edge  $e \in E$  s.t  $\theta \neq \epsilon$  and
    $\theta = (U_\theta, op_\theta, P_\theta)$ 
3   if  $u \in U_\theta$  then
4     if  $G_{PM}(G_e, P_\theta, \mathcal{H})$  then
5       SETACTIVESTATE( $r, x_{i+1}$ )
6     else
7       if PROCEEDINGSTATE( $G_e, u, x_i, x_f$ )
8         then
9           SETACTIVESTATE( $r, x_{i+2}$ )
10        else
11          SETACTIVESTATE( $r, x_{i+1}$ )
12    else if  $u \notin U_\theta$  then
13      if PROCEEDINGSTATE( $G_e, u, x_i, x_f$ ) then
14        SETACTIVESTATE( $r, x_{i+2}$ )
15      else
16        SETACTIVESTATE( $r, x_{i+1}$ )
17    // If it is the final state
18    if  $x_i = x_f$  then
19      a query match has been found
20      remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

Algorithm 8: Comparing Proceeding State’s Transition Predicate

```

1 Function PROCEEDINGSTATE( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r$ )
2   select proceeding state  $x_{i+1}$ 
3   if  $x_{i+1} = x_f$  then
4     return false
5    $E \leftarrow \text{GETEDGESET}(x_{i+1})$ 
6   get  $\theta$  for an edge  $e \in E$  s.t  $\theta \neq \epsilon$  and
    $\theta = (U_\theta, sf_\theta, op_\theta, P_\theta)$ 
7   if  $u \in U_\theta$  and  $G_{PM}(G_e, P_\theta, \mathcal{H})$  then
8     return true
9   else
10    return false

```

C.4.1. Optional Operator

The main difference between the evaluation of the optional and negation operators is to differentiate between the occurrence and non-occurrence of an event. For instance, given a sequence expression $\text{SEQ}(A, B?, C)$ and events G_e^1 and G_e^2 . If G_e^1 matches A then we have to match G_e^2 with both B and C; since

Algorithm 9: Negation Operator's Evaluation

```

1 Function NEGATION( $G_e, u, x_f, \mathcal{H}, \mathcal{R}, r, x_i, E$ )
2   get  $\theta$  for an edge  $e \in E$  s.t  $\theta \neq \epsilon$  and
    $\theta = (U_\theta, sf_\theta, op_\theta, P_\theta)$ 
3   if  $u \in U_\theta$  then
4     if  $\neg G_{PM}(G_e, P_\theta, \mathcal{H})$  then
5       if PROCEEDINGSTATE( $G_e, u, x_i, x_f$ )
6         then
7           SETACTIVESTATE( $r, x_{i+2}$ )
8         else
9           SETACTIVESTATE( $r, x_{i+1}$ )
10      else
11        remove  $r$  from the list of active runs  $\mathcal{R}$ 
12      return
13  else if  $u \notin U_\theta$  then
14    if PROCEEDINGSTATE( $G_e, u, x_i, x_f$ ) then
15      SETACTIVESTATE( $r, x_{i+2}$ )
16    else
17      SETACTIVESTATE( $r, x_{i+1}$ )
18  // If it is the final state
19  if  $x_i = x_f$  then
20    a query match has been found
21  remove  $r$  from the list of active runs  $\mathcal{R}$ 

```

it is possible that the event G_e^2 does not match with B but with C , and in such a case we can have a query match. Therefore, the evaluation of the optional operator makes sure that the event that does not match with the optional state is compared with the next state's state-transition predicate. The evaluation of the optional operator is shown in Algorithm 7. It first employs the edge with state-transition predicate $\theta \neq \epsilon$ (lines 1) and uses the stream names to make sure the event is from the stream the edge is waiting for (line 3). The algorithm then uses event G_e , graph pattern P_θ and the history cache \mathcal{H} to execute the GPM process (line 4). If the event G_e matches with the selected edge, it transits to the next state (line 5). Otherwise it checks the event G_e with the next state-transition's predicate (line 7-10) using the CHECKNEXTSTATE function in Algorithm 8. If the event G_e matches to the next state-transition predicate – considering if it is not the final state – it transits the active state x_i to x_{i+2} (line 8). Furthermore, if the event is not from the desired stream, the run takes the ϵ -transition and compares the next state-transition predicate to either transit to the next state x_{i+1} or next

of the next state x_{i+2} (lines 12-15). During the evaluation of the operator, if the current state reaches the final state, the algorithm output the query match (line 16-18).

Algorithm 10: Optimised evaluation of the conjunction operator

```

Input:  $G_e$ : Graph Event,  $u$ : stream name,  $r$ : active run,  $x_f$ : final state,  $r$ : active run,  $\mathcal{H}$ : cache history,  $\mathcal{D}$ : conjunction edge-timestamp map,  $\mathcal{B}$ : event buffer
1  $x_i \leftarrow$  GETACTIVESTATE( $r$ )
2  $E \leftarrow$  GETEDGESET( $x_i$ )
3  $\tau \leftarrow$  GETTIMESTAMP( $G_e$ )
4 if  $|\mathcal{B}| \neq \emptyset$  then
5   get  $\tau_{old}$  from an event  $G_e^k \in \mathcal{B}$ 
6   if  $\tau \neq \tau_{old}$  then
7      $\mathcal{B} \leftarrow \emptyset$ 
8     remove  $r$  from the list of active runs  $\mathcal{R}$ 
9   else
10     $\mathcal{B} = \mathcal{B} \cup G_e$ 
11    remove edges from  $E$  using  $\mathcal{D}$  that are already matched
12    if  $|E| = |\mathcal{B}|$  then
13      sort  $E$  according to the graph patterns
14      foreach each edge  $e \in E$  do
15        get  $\theta$  from the edge  $e$ 
16        foreach each event  $G_e^i \in \mathcal{B}$  do
17          if  $u \in U_\theta$  and
18             $G_{PM}(G_e^i, P_\theta, sf_\theta, \mathcal{H})$  then
19              remove  $G_e^i$  from event buffer  $\mathcal{B}$ 
20              insert  $(r, x, e)$  and  $\tau_i$  in  $\mathcal{D}$ 
21     $E \leftarrow$  GETEDGESET( $x_i$ )
22    get the set of matched edges  $AME$  from  $\mathcal{D}$ 
23    if  $|E| = |AME|$  then
24      SETACTIVESTATE( $r, x_{i+1}$ )
25      // If it is the final state
26      if  $x_i = x_f$  then
27        a query match has been found

```

C.4.2. Negation Operator

The negation operator is evaluated in a similar fashion compared with the optional operator. However, in this case, the run transits to the next state (producing an identity element) if there is no match with the mapped

graph pattern P_θ . The evaluation of the negation operator is described in Algorithm 9. It begins by selecting the state-transition predicate and compares the stream names and graph pattern P_θ with the event G_e using GPM function (lines 4). If GPM returns *false*, i.e. the event is not matched with the graph pattern, the algorithm uses the CHECKNEXTSTATE function to determine if such an event can be matched with the next state-transition predicate x_{i+1} (line 5), as described for the optional operator. In case the current active state x_i is matched with the event G_e , it means the run has violated the negation operator and should be deleted (line 10-11). If the event is from a different stream, the evaluation of the negation operator is same as that of optional operator (lines 12-19).

D. Local Query Optimisations of Disjunction and Conjunction Operators

Algorithm 10 shows the lazy evaluation of the conjunction operator and it extends Algorithm 5. It employs an event buffer \mathcal{B} to cache the set of events having the same timestamps. The algorithm first takes the set of edges E for an active state and timestamp τ of the newly arrived event (lines 1-2). It then checks if there exists any previously buffered event in \mathcal{B} . If so it checks their timestamp τ_{old} and the timestamp τ of

the newly arrived event (line 5). If the timestamps do not match it deletes the run under evaluation while removing all the previously buffered events in \mathcal{B} (lines 4-7). Otherwise it adds the event G_e in the event buffer \mathcal{B} , which is later used to evaluate the unmatched edges (line 10). At this stage, the algorithm prunes the edge set E with the edges that have already been matched – using the mapping structure \mathcal{D} (line 11) (as described in Algorithm 5). If the number of unmatched edges in E is equal to the number of buffered event \mathcal{B} , it starts the matching process using the lazy evaluation strategy. Before starting the matching process, it first sorts the edges according to the selectivity of the graph patterns in state-transition predicates, hence using the low cost edges first (line 13). The algorithm then iterates over the sorted edges E and the buffered events set to match the selected edge (lines 12-22). If an edge $e \in E$ matches the buffered event $G_e^i \in \mathcal{B}$, its mapping is added into \mathcal{D} and the matched event is removed from the buffer \mathcal{B} (lines 18-19). In the end, the algorithm has to determine if all the edges of the active states are matched and should it transit to the next state or not. Therefore, it examines the number of actual edges of the state and the number of them that have already been matched (line 25). If all the edges are matched the run transits to the next state. Otherwise, it waits at the same state to receive more events having the same timestamps (line 28).