Stream reasoning: A survey and outlook

A summary of ten years of research and a vision for the next decade

Daniele Dell'Aglio [a,*], Emanuele Della Valle [b], Frank van Harmelen [c] and Abraham Bernstein [d]

[a] *DDIS, Department of Informatics, University of Zurich, Binzmühlestrasse 14, 8050 Zürich, Switzerland*
E-mail: dellaglio@ifi.uzh.ch; ORCID: https://orcid.org/0000-0003-4904-2511
[b] *DEIB, Politecnico di Milano, Via Ponzio, 34/5, 20133 Milano, Italy*
E-mail: emanuele.dellavalle@polimi.it; ORCID: https://orcid.org/0000-0002-5176-5885
[c] *Department of Computer Science, Vrije Universiteit Amsterdam, de Boelelaan 1081a, 1081HV Amsterdam, The Netherlands*
E-mail: frank.van.harmelen@vu.nl; ORCID: https://orcid.org/0000-0002-7913-0048
[d] *DDIS, Department of Informatics, University of Zurich, Binzmühlestrasse 14, 8050 Zürich, Switzerland*
E-mail: bernstein@ifi.uzh.ch; ORCID: https://orcid.org/0000-0002-0128-4602

Abstract. Stream reasoning studies the application of inference techniques to data characterised by being highly dynamic. It can find application in several settings, from Smart Cities to Industry 4.0, from Internet of Things to Social Media analytics. This year stream reasoning turns ten, and in this article we analyse its growth. In the first part, we trace the main results obtained so far, by presenting the most prominent studies. We start by an overview of the most relevant studies developed in the context of semantic web, and then we extend the analysis to include contributions from adjacent areas, such as database and artificial intelligence. Looking at the past is useful to prepare for the future: in the second part, we present a set of open challenges and issues that stream reasoning will face in the next future.

Keywords: Stream reasoning, stream processing

1. Introduction

Increasingly, applications require real-time processing of heterogeneous data streams together with large background knowledge bases. Consider the following examples. Which electricity-producing turbine has sensor readings similar (i.e., Pearson correlated by at least 0.75) to any turbine that subsequently

---

*Corresponding author. E-mail: dellaglio@ifi.uzh.ch.

had a critical failure in the past year [75]? When a sensor on a drill in an oil-rig indicates that it is about to get stuck, how long – according to historical records – can I keep drilling [74]? Where am I likely going to run into a traffic jam during my commute tonight and how long will it take, given current weather and traffic conditions [3,87,107]? Who are the current top influencer users that are driving the discussion about the top emerging topics across all the social networks [12,71]? Where shall I spend my evening given the presence of people and what their doing (predicted analysing the spatio-temporal correlation between privacy-preserving aggregates of Mobile Telecom Data and of geo-located Social Media posts) [13]? Who should be asked to go exercising, given people's past, possibly sedentary behaviour and allergies (accessed in a privacy-preserving manner) as well as current weather conditions and pollution/allergen levels [45]?

To answer these queries a system must be able to [37]:

R1. handle *volume*: a typical oil production platform is equipped with about 400.000 sensors; Facebook, as of February, 2017, has 1.86 billion of monthly active users,[1] etc.
R2. handle *velocity*: sensors on a power generation turbine can easily generate thousands of observations per minute; Instagram's users, as of February, 2017, like on average 2.92 million post per minute;[2] etc.
R3. handle *variety*: a large variety of static and streaming data sources and data management solutions exists in any domain. For instance, Milano has deployed some 600 traffic light systems equipped with inductive loops in the last 10 years: they use five different data formats, have different operational conditions, etc. Similarly, in social media, each network has its own data model and APIs.
R4. cope with *incompleteness*: sensors can run out of battery or networking links can break; in social media parts of the conversation may occur outside the social network [11], or the APIs that are used to access the social stream may have a limited sampling rate.
R5. cope with *noise*: sensors can be imperfect, faulty, or out of its ideal operational range; text can be worded in an ironic way and a sentiment mining solution may be unable to detect it with 100% correctness.
R6. provide answers in a *timely fashion*: answers should be generated within well-specified latency bounds, which depend on the application scenarios. The time available to answer depends on the application domain: in call centres, routing needs to be decided in sub-seconds (in real-time); in oil operations, the detection of dangerous situations must occur within minutes (in near real-time).
R7. support *fine-grained information access*: the issued query may require to locate exactly a turbine, a means of public transportation, an agent in a contact centre among thousands of similar ones.
R8. integrate *complex domain models*: social media analytics may require topic models to make sense of a conversation; oil production control systems may require to model operational and control processes; traffic monitors may require rich background knowledge about topology, planned and unplanned events to improve the accuracy of the analyses.
R9. capture *what users want*: the query should let users define analytics-aware tasks such as Pearson correlation as a mean of similarity, or complex concepts such as traffic jam and top influencer user.

Ten years ago, no system was able to address all these requirements simultaneously. The management of highly dynamic data (R2) in a timely fashion (R6) developed around the idea of stream processing, a computation paradigm where data is processed in motion, i.e., on the fly and as soon as it becomes

---

[1]See https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/. Accessed on February, 2017.
[2]See https://blog.hootsuite.com/instagram-statistics/. Accessed on February, 2017.

available. At that time, the most advanced stream processing solutions were developed in the context of Data Stream Management Systems (DSMSs) and Complex Event Processors (CEPs) [29]. DSMSs transform data streams in timestamped relations (usually through a so-called window operator) and process them with well known techniques such as relational algebras [6]. DSMSs allow the construction of systems able to compute aggregations and statistics (e.g. averages and Pearson correlation) over streaming data. CEPs, instead, look for patterns in the streams to identify when complex events occur [90]. CEPs focus on the derivation of (complex) events from the matching of patterns of events (e.g., sequences) in the input data. Both DSMS and CEP techniques were able to provide reactive fine-grained information access (R7) in the presence of noisy data (R5).

These system were, however, limited when data was heterogeneous (R3), with complex domain models (R8) and needed to combine rich background knowledge. In those cases, users, to solve the desired tasks, were required to put large manual effort in developing complex networks of queries. There was a high potential in finding systematic solutions to cope with these problems. The research of those years on Knowledge Representation (KR) and Semantic Web (SemWeb) were bringing relevant inputs in that direction. The study of complexity of description logics were bringing new results, opening the door to techniques able to perform reasoning tasks in polynomial and sub-polynomial time. Techniques like Ontology Based Data Access (OBDA) [25] were starting to show that complex domain models (R8) can be used to offer fine-grained information access (R7) to heterogeneous (R3) and incomplete datasets (R4). But these studies were focusing on static data, and solutions lacked the ability to provide reactive answers on data streams and to handle noise.

CEP, DSMS, KR and SemWeb supplied the ingredients to Stream Reasoning [37]. As the name suggests, this new research trend aims at studying how to perform online logical reasoning over highly dynamic data. Stream reasoners were envisioned as systems capable to address all the requirements above simultaneously (see Table 1).

Even if the goal can be easily stated, satisfying all of the above-stated requirements is challenging in multiple ways. Theoretically, it is difficult to create comprehensive data and processing models. Practically, it is non-trivial to guarantee reactiveness given the required functional requirements. Nonetheless, we have seen the emergence of various results [38,91]. The large majority of them where in the Semantic Web community, but also the broader Artificial Intelligence, Robotics, Databases, and Distributed Systems communities performed investigations in the Stream Reasoning area.

Focusing on the works of the semantic web community, different research groups proposed (i) data models and vocabularies to capture data streams through RDF and ontologies (e.g. [18,68]), (ii) continu-

Table 1

The requirements Stream Reasoning aims at covering and how DSMS, CEP and Semantic Web cover them (✓ indicates which area fits the requirement better)

Requirement	DSMS/CEP	SemWeb	Envisioned stream reasoning
R1: Volume	✓	✓	✓
R2: Velocity	✓		✓
R3: Variety		✓	✓
R4: Incompleteness		✓	✓
R5: Noise	✓		✓
R6: Timely fashion	✓		✓
R7: Fine-grained access	✓	✓	✓
R8: Complex domains		✓	✓
R9: What users want	✓		✓

ous query models, languages and prototypes (e.g. [5,15,24,82]), inspired by SPARQL [63] and collected under the *RDF Stream Processing* (RSP) label, (iii) extensions of the reasoning tasks over streams, as consistency check and closure (e.g. [16,84,108]), and (iv) applications built on top of the aforementioned results (e.g. [12,72,119]). However, a complete answer to the stream reasoning research question is still missing.

For this reason, in this article, we introduce a reference model to capture, organize and summarize the results that the semantic web community has obtained so far (Section 2). Then, we identify the challenges to address in the next years by analysing the research results of other communities and the current trends in big data processing (Section 3). In Section 4, we wrap up and draw some conclusions.

2. A reference model for stream reasoning

In the last ten years, several techniques have been proposed under the stream reasoning label. Such techniques are heterogeneous in terms of input, output and use cases. To present them, we present Fig. 1 as a reference model for stream reasoning [40].

In the remainder of this paper, we denote a data stream as a sequence of time-annotated items ordered according to a temporal criteria. Each item brings an information unit, e.g. a sensor observation, a social network action or a stock exchange. It is possible to represent the stream items in different formats. In the semantic web context, data is usually modelled according to RDF [31], a W3C recommendation where data is structured as directed labelled graph, with vertices and edges representing, respectively, resources and relations among them. We name RDF stream a stream where items are represented according to RDF, i.e., each item is represented through a directed labelled graph.
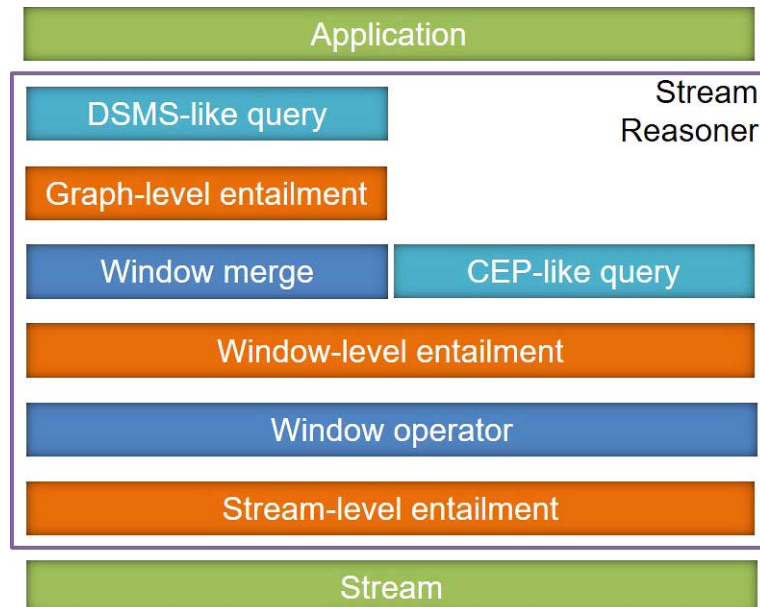

Fig. 1. A model to describe stream reasoners.

We now present a reference architecture for a generic continuous query answering system over streams, identified in Fig. 1 by the blue blocks. *Window operators* manage the access to the stream: they create time-dependent finite views over the streams, namely windows, over which processors perform the tasks. Window contains a portion of the input streams, i.e. a set of timestamped data items, that represents the data needed to solve the task at the current time instant. Several types of window operators exist, initially defined in CEP and DSMS research [29].

CEP aims at verifying if given sequences of events happen in the stream. In this sense, time annotations are key, since engines use them to determine if temporal constraints defined in the event patterns are satisfied. When not specified, event patterns are evaluated over a large portion of the stream, i.e. from the moment on which the engine starts to observe the stream up to the current moment. We can model this behaviour through *landmark windows*. Fixed an initial time instant, the window expands over time to capture portion of the stream that grows over time.

DSMSs perform operations that do not require to process time annotations after that windows have been computed, like aggregations and filters. That means, while the content of the stream items is important to solve the task, time annotations are not. In this sense, the *window merge* is an operation that moves from temporal data (i.e. time-annotated data items) to atemporal one (i.e. a collection of data items). *Sliding windows* are the typical window operators considered in DSMS. This operator creates a window with a fixed width in terms of time units or data items. The operator shifts (slides) the window over time, capturing the most recent part of the stream.

It is worth noting that time in CEP and DSMS is treated in different ways: in the former it has a predominant role while evaluating the event patterns, while in the latter it is used to identify the relevant portion of data over which an atemporal operation is performed.

How DSMS and CEP can be extended to perform reasoning tasks is one of the first problems studied in stream reasoning. In the semantic web community, this was addressed by adding DSMS and CEP features to SPARQL [63] – the query language for RDF-data recommended by W3C. Every SPARQL Query is built from Basic Graph Patterns (BGPs). A BGP is a direct labelled graph where vertices and edges can be either resources or variables. The evaluation of a BGP over an RDF graph consists of finding the set of mappings between variables in the BGP and the resources in the RDF graph such that the BGP matches the RDF graph. The presence of an inference process in SPARQL is modelled through entailment regimes [59]. Entailment regimes affect the basic graph pattern evaluation by extending the matching definition to take the edges and the nodes that can be logically inferred from those explicitly stated into account.

When we move to a model for processing streams, as the one in Fig. 1, we may observe that there is not one and only one correct moment to take the inference process into account. In other words, the key question we need to answer is: *while processing a stream, when is the correct moment to take the inference process into account?* The reference model described above allows different options, depicted through the orange boxes in Fig. 1. Inference can take place before or after the window (giving, respectively, stream- and window-level entailments), or still after that the window content is merged (giving graph-level entailment). This decision may affect the performance, the result and the behaviour of the engine. Taking into account this framework, group the main works developed in the state of the art of stream reasoning, in four different groups, that we present in the rest of this section.

2.1. Graph-level entailment

The last possible moment to take into account the inference processes is after the window merge. That means a window captures a portion of the stream and a merge operation creates a collection from the stream item contents.

We name *graph-level entailment* a continuous query evaluation where the inference process is considered at this moment. As explained above, this enables the execution of typical DSMS-like queries, e.g. aggregations and filters, but make not possible to evaluate CEP-like queries, since temporal annotations are lost and it is not possible to verify if temporal constraints are satisfied. This is represented in Fig. 1, where the graph-level entailment stacks on DSMS-like queries but not on CEP-like ones.

The graph-level entailment can be viewed as a direct application of SPARQL entailment regimes, since the inference process is taken into account in the context of the evaluation of (basic) graph patterns over graphs.

A large amount of research can be captured by this level. In the following, we group it in two main groups: systems that work under a simple entailment regime, i.e. RDF entailment, and the ones that considers more complex reasoning formalisms, e.g. Description Logics and ASP. They are presented respectively in Sections 2.1.1 and 2.1.3. To enable the comprehension of the latter, we briefly review some prominent work in incremental reasoning in Section 2.1.2.

2.1.1. Stream reasoning at graph level under simple entailment regime

The first case we consider is the one where the RDF entailment process is involved, which in SPARQL is named *simple entailment regime*. When a SPARQL query is evaluated under this regime, the engine verifies if a graph pattern matches the input graph, without considering ontological language or inference processes.

The first generation of continuous query answering systems for RDF data streams is represented in this category: engines where queries are registered and answers are continuously produced when new data is made available.

Several languages and prototypes were built by following the DSMS vision: sliding windows to capture finite portions of the streams to be queried through SPARQL as collections of RDF statements. Between 2010 and 2012, different groups proposed languages as the Continuous SPARQL (C-SPARQL) language [15], CQELS-QL [82] and SPARQL$_{stream}$ [24]. Such languages are similar and relies on the idea of extending SPARQL with sliding windows. For example, C-SPARQL extends the FROM clause in order to support sliding windows, while CQELS-QL pushes sliding windows in the GRAPH clause. There are differences on the prototyping sides as well. The C-SPARQL engine and CQELS, implementing respectively C-SPARQL and CQELS-QL, take as input RDF streams and process them as extended SPARQL query processors (e.g. ARQ[3]). morph-stream, an implementation of SPARQL$_{stream}$, adopts an OBDA-like approach: it processes relational streams by transforming the query from SPARQL$_{stream}$ to one to be registered to a DSMS engine such as Esper or GSN.

A different approach is adopted by the Incremental eNgine for STANding Sparql – INSTANS [110]. As query language, it adopts SPARQL 1.1, with an extension to its query evaluation model to continuously query data streams. The implementation of INSTANS relies on the RETE algorithm [54]: tasks are expressed as networks of queries and compiled in RETE-like structures to evaluate the results. More detailed comparison analyses of those solutions may be found in [41] and on the W3C RSP-CG wiki.[4]

---

[3]Cf. https://jena.apache.org/documentation/query/.
[4]Cf. https://www.w3.org/community/rsp/wiki/RDF_Stream_Processors_Implementation.

2.1.2. Incremental reasoning

Among the techniques to do reasoning, the closest to the stream reasoning idea are the incremental reasoning ones. They have been developed to cope with changes in knowledge bases. The intuition is to modify part of the materialisation of a knowledge base when updates happen, without re-materialising everything from scratch. As update operations occur these techniques identify the facts that have to be removed – since they were derived by deleted facts – and added – since they can be inferred by new facts. In this way, it is possible to avoid the recomputation of the whole materialisation upon changes.

DReD [126] has been proposed in 2005 by Volz et al. and it is inspired by techniques of view maintenance in databases. The idea is to compute two sets of axioms to be added and removed through a three-step process. In the first step, deletions are computed starting from the facts that should be deleted. This step produces an overestimation, due to the fact that some derived facts may still be inferred from other non-deleted facts. In the second step, the algorithm looks for these facts and remove them from the facts to be deleted. In the last step, new derivations are computed starting from the axioms that have been added to the knowledge base.

The first two steps of the DReD algorithm are critical for the performance of the technique: a considerable amount of computational effort may be required to identify the facts to be labelled as to be deleted in the first step and rectracted in the second. To overcome this limit, Motik and al. propose the Backward/Forward algorithm [97] in 2015. The idea is to use a combination of backward and forward inference to limit the number of overestimations in the first step. An implementation of this algorithm is available in RDFox [101], an in-memory datalog engine.

A different approach is the one considered by Ren and Pan in Truth Maintenance Systems in 2011 [108]. Differently from DReD, this technique builds and maintains a dependency graph. When deletions occur, this graph is used to decide if derived facts should be removed; when additions happen, the graph is updated by adding new edges and vertices. This approach is optimised for $\mathcal{EL}^{++}$ [9,10].

DynamiTE [125] is a framework proposed in 2013 to maintain the materialisation incrementally. One of the novelties introduced by Dynamite is the parallelisation of the inference process. The framework supports DReD and a counting algorithm proposed by the authors for the ρ DF fragment [100] of RDFS. On additions, DynamiTE recomputes the materialisation to add the new entailments through a parallel Datalog evaluation. On removals, it deletes the explicit and entailed axioms no longer valid. Several algorithms can perform this action: the authors considered DReD and a "counting" algorithm they defined, that exploits the idea of counting the number of justifications that entailed it.

2.1.3. Stream reasoning at graph level under other entailment regimes

Moving to non-simple entailment regimes opens a set of challenges given by the introduction of the ontological languages and the associated inference processes.

One of the first attempts of building a stream reasoner is Streaming Knowledge Bases [127] in 2008. The idea was to pipe a DSMS engine, TelegraphCQ, with a reasoning engine, the Jena Rule engine. The system is able to performs RDFS inference.

The Incremental Materialization for RDF Streams algorithm (IMaRS) [41], proposed in 2011 for the RDFS+ ontological language takes a step further. This technique focuses on the incremental maintenance of a window content as the window slides. The intuition behind this algorithm is that in this setting, deletions can be foreseen and are not random as in the incremental reasoning setting described above. An expiration time annotation is associated to all the axioms involved in the materialisation, and such information is exploited to identify only the facts to be deleted, avoiding the overestimation typical of DReD. An implementation of IMaRS is available in Sparkwave [79]. It implements the algorithm on the

top of the RETE algorithm and targets the RDF Schema entailment. RDF schema axioms are encoded as RETE rules and organised in a network. When new facts are added to the system, they are matched against the rules.

Another stream reasoner is StreamRule [95], proposed in 2013. As Streaming Knowledge Bases, the system is designed as a two-layer approach. The first layer is an RSP engine acting as a filter, that reduces the amount of data to be considered in the inference process. The second layer is a reasoning engine. What makes StreamRule unique w.r.t. the solution described above is the adoption of Answer Set Programming (ASP) [48], rather than a DL reasoner. ASP is a declarative problem-solving paradigm, characterized by rich modelling language maintaining very good performance, obtained by exploiting techniques from constraint solving. ASP works with static knowledge; Incremental ASP [58] overcomes such a limit and extends ASP to compute the solutions incrementally.

The principle of working on a graph captured with a time window was also applied to inductive reasoning. In [17], authors used graph-level inductive reasoning in developing a recommendation engine that suggest topics to users.

2.2. Window-level entailment

In the previous section, we explain that graph-level entailment is a direct application of atemporal reasoning techniques over streams. The main drawback of this approach is that the process uses only a subset of the information available in the stream: it considers the data item contents but not the relative temporal annotations. While DSMS queries can benefit if evaluated under graph-level entailment, CEP-like queries do not gain much advantages.

Window-level entailment overcomes these limitations by applying the inference process on the non-merged stream items captured by the window operators declared in the query. Differently from graph-level entailment, which works on graphs (ontologies), the window-level entailment applies the inference process to a window, i.e. a finite sequence of timestamped data items.

The Spatial and Temporal ontology Access with a Reasoning-based Query Language (STARQL) [102] proposes a framework to access and query heterogeneous sensor data through ontologies. STARQL is structured as a two-layer framework, composed of an Ontology Language, to model the data and its schema, and an Embedded Constraint Language, to compose the queries. STARQL offers window operators, clauses to express event matching and a layer to integrate static and streaming data. STARQL uses a sequence of time-annotated ontologies to make inference taking into account the temporal annotations of the streaming data. Recently, STARQL has been implemented in Exastream [73], a system that adopts an OBDA approach, similar to morph-stream.

2.3. Stream-level entailment

One of the main limitations of the window-level entailment is that it only considers a recent portion of the stream. For example, when a sliding window computes a new window, what happened in the past is forgotten and the inference process considers the data contained in the new fixed window.

Stream-level entailment overcomes this limitation, considering a larger portion of the stream than the one defined by the user through window operator. Even if the name suggests that this entailment regime considers the whole stream, in this case the reasoning is made on the top of a landmark window, which is a window that captures the stream from an initial time instant (e.g. when the source starts to supply the data when the engine starts to monitor the stream) up to now.

ETALIS (Event TrAnsaction Logic Inference System) [5] is a CEP-based stream reasoning engine. This query model processes streams where data items are annotated with two timestamps (i.e., time intervals). Users can specify event processing tasks in ETALIS using two declarative rule-based languages, ETALIS Language for Events (ELE) and Event Processing SPARQL [4] (EP-SPARQL). The former language is more expressive than the latter, even if it is less usable. A common point is that complex events are derived from simpler events using deductive prolog rules. EP-SPARQL supports backward temporal reasoning over RDFS, continuously evaluating the query over the whole stream received by the engine.

The Logic-based framework for Analyzing Reasoning over Streams [21] (LARS) defines a logic for modelling stream-related axioms. LARS models the notion of stream as a sequence of time-annotated formulas. In addition to the usual boolean operators (and, or, implies, not), the language introduces operators, such as $\Diamond$ and $\Box$ to express the fact that a formula holds respectively at some time in the future and every time in the future; @ to state that a formula holds at a specific time instant. LARS formulas can be evaluated against the whole stream, or the scope can be limited through the usage of the $\boxplus$ operator, that models a window operator. Through this operator, LARS is also able to capture reasoning at window-level entailment.

The principle of reasoning at stream-level was also applied to inductive reasoning. In [86], authors used stream-level inductive reasoning in predicting Knowledge in an Ontology Stream.

2.4. Measuring progresses

Ten years ago, we proposed [37] to initially assess progress by measuring the degree to which a stream reasoner meets the requirements we listed in Section 1. We left it as a community challenge to identify how to systematically evaluate stream-reasoning implementations against well-defined quality criteria, e.g. creating one or more domain specific benchmarks [60].

The first proposal in this direction was SRBench [131]. It provides seventeen continuous queries and a weather related dataset annotated with the Semantic Sensor Network ontology. Its goal is to assess the ability of implementations to cover different stream reasoning features even if little emphasis is dedicated to reasoning.

A second proposal was LSBench [83]. It proposes twelve continuous queries, a social network ontology, and a data generator. Its goal is to move from coverage of query language features to systematic measures of well-defined quality criteria. It does not include any reasoning, but raises the problem of assessing correctness.

The theme of correctness helped the community in opening the discussion on what a domain specific benchmark on stream reasoning should look for [113] and in deepening the understanding of the semantics of RDF continuous processing. The investigations that eventually led to RSP-QL semantics [44] also delivered CSRBench [42] as a side result. This benchmark is an extension of SRbench that includes an oracle able to tell which result shall be expected from a given implementation.

Two more mature proposals, CityBench [3] and YAbench [77], emerged from those pioneering works. They provide real and synthetic workloads together with software environments to ease the execution of the experiments and the measurement of the quality criteria.

Heaven [122] is the most recent work in this area. It takes a step in the direction of stressing expressive reasoning and offers an open-source test-stand to run reproducible and repeatable experiments. The goal of Heaven is to offer the community an environment where researchers can (i) take any dataset that contains events originated over time, an ontology that describes the dataset, a query workload that requires a certain entailment regime, and (ii) run a comparative experiment against all the stream reasoners included in the environment.

The next step in this area should be the integration of the benchmarking efforts that are maturing in parallel in the semantic web, e.g., the Social Network benchmark of the Linked Data Benchmark Council [51], the results of the Hobbit project [112], but also in the Big Data, e.g., the streaming workloads of HiBench [67].

3. Open problems and challenges for the next years

Stream reasoning research progressed and expanded its initial community to a growing number of practitioners. In this section, we review the requirements presented in Section 1 in the light of the current state of the art and outline open questions.

Table 2 summarizes the current state and serves as an indication towards possible directions for future stream reasoning research. DSMS, CEP and semantic web let users express and use complex notions in queries, such as trends, skylines and aggregations. However, users are demanding for even more sophisticated features (R9), as discussed in Section 3.1.

Almost every study in stream reasoning addressed the timely fashion requirement (R6). It is possible to observe it by considering that time performance is usually the most relevant axis when performing evaluations. Similarly, the focus has been on data streams with rich data models. In this sense, the requirement related to complex domains (R8) has been partially treated. However, we should consider this requirements in a broader sense. As we explain in Section 3.2 stream reasoning is expanding by considering (i) more expressive ontological languages and (ii) non-deductive reasoning approaches.

Velocity and fine-grained access (R2 and R7) have been the two requirements at the centre of the research so far. Indeed, all studies on stream reasoning so far addressed the velocity dimension of the problem, considering streaming data the main object of investigation. In future, as we present in Section 3.3, we expect to assist to the rise of *smart* streams that leverage semantics for processing.

Volume and variety (R1 and R3) have not been properly addressed, but they are key requirements to enable stream reasoning in real-world big data environments. We discuss this in Section 3.4.

Finally, massive data is common, but real-world data is often characterized also by being imperfect, i.e. incomplete and noisy (R4 and R5). In this direction, stream reasoning research is still at the starting point. In Section 3.5 we analyze these research opportunities.

Table 2

A review of the stream reasoning requirements w.r.t. the current state of the art ($\star$ = not specifically treated so far, $\star\star$ = treated but not resolved, $\star\star\star$ = universally addressed by all studies)

Requirement	Current Stream Reasoning
R1: Volume	$\star$
R2: Velocity	$\star\star\star$
R3: Variety	$\star\star$
R4: Incompleteness	$\star$
R5: Noise	$\star$
R6: Timely fashion	$\star\star$
R7: Fine-grained access	$\star\star\star$
R8: Complex domains	$\star\star$
R9: What users need	$\star\star$

3.1. Towards queries that capture user needs

Queries connect users to stream reasoners: they encode the needs of the former in a format readable and processable by the latter. Research on query languages has been one of the leading topics of the stream reasoning trend up to now. There has been the first wave of languages, e.g. [4,15,24,82], refined and improved by new proposals to increase the expressiveness and the supported operations, e.g. [57, 102]. In parallel, research has worked on comparing and contrasting such languages, e.g. [32,43,44,70].

Despite the results obtained so far, the research on query languages for stream reasoning is far from being complete. In this section, we present three directions that, in our opinion, are needed to reduce of the distance between stream reasoners and users (and consequently, real applications).

3.1.1. Capturing a wider set of tasks in queries

So far, language development focused primarily on SPARQL-like queries, where the main goal is matching of graph patterns over the streaming data under some entailment regimes. This kind of query sets the basis to filter, aggregate, as well as detect trends and complex event patterns in the data. However, there are other important tasks to be performed over streams. One example are spatial operations, that allow to define geographical and trajectory queries. They have been largely studied in semantic web as well as in DSMSs and CEPs, e.g. [81,104,105], but only recently stream reasoning started to move in this direction (e.g. [34]).

In several domains, there is the need to perform complex analyses of streaming data that make heavy use of aggregation, correlation functions and arithmetic operations as well as of inductive methods. Examples of this trend are big data processors, where solutions such as Apache Spark [129], Apache Flink [27] and Twitter Heron [80] offer an extensive set of operations, from counting algorithms up to machine learning methods [94]. Solutions in this direction are starting to appear, e.g. [75], and we expect a large interest in investigating this area in the near future.

Analytics provide insights, but users often also want fine grain access to the information that supports those insights. This calls for optimizing a special type of continuous queries that include preferences in the information need, e.g. continuous top-k queries [99]. Some initial works exist on top-k query answering in SPARQL, but to the best of our knowledge there is only a vision paper [39] that casts some light in this direction. No one has attempted, yet, to investigate this area.

Moreover, the graph structure of the RDF data model opens the doors to the application of the classical algorithms related to graph theory, e.g. path finding, connectivity, bipartiteness and clique search. How to apply such techniques to graph streams has been studied in the database literature [93].

Whilst there are examples for supporting these kinds of operations on static data (e.g., [76]), to the best of our knowledge, none of the stream reasoning-related query languages designed up to now support the above operations. Given the importance these non-deductive techniques have gained over the past few years (just think about the rise of machine learning), we strongly believe that the next years will see the addition of new operations and the design of languages that support them. It may still remain important to offer declarative languages (in addition to programmable APIs) to describe tasks, since they separate the expected output (the goal to the computation) from the way it should be computed and may – in some cases – be easier to handle for domain experts.

3.1.2. Tightening CEP and ontological reasoning

It is worth considering the relation between CEP and DL reasoning. CEP queries are *de-facto* production rules: if a set of temporal constraints over events are satisfied (condition), then a complex event is

triggered (action). That means a system that involved both CEP rules and DL axioms has two inference components: one based on CEP rules and another one based on ontological axioms.

This consideration opens the door to several questions, such as: which part of the knowledge better fits the CEP rules and which one better fits ontology axioms? We can imagine that while some parts of the knowledge can be modelled as CEP rules or ontological axioms only, others may be captured by both formalisms. Such modelling decisions have an impact on the evaluation engine, its performance and the results it computes.

Another question is: how can the two components be combined? Following the framework presented in Section 2, the CEP evaluation is on top of the entailment (both window- and stream-level one). A well-marked separation between the two components is a safeguard with regards to complexity, but may not fit some applications. For example, a use case may require that the output of CEP reasoning will result in changes to the ontological entailment, which in turn may trigger CEP rules. In such cases, the interaction of the two deduction sub-systems may affect the performance, with the risk to end up in endless processes. One possible solution to the problem is to identify the conditions under which CEP rules maintain the decidability of the stream reasoning process, similarly to the research on DL-safe rules in [98].

3.1.3. Forgetting knowledge

While processing streams, the identification of data that is useful for the current computation can be problematic. Several solutions have been proposed up to now, from sliding window mechanisms in DSMSs, to consumption and selection policies in CEPs. Such methods are key since they allow to discard data and keep under control the resource usage, and are currently adapted by stream reasoning solutions, e.g. sliding windows in CQELS and consumption policies in Etalis. However, they may lead to unexpected situations, e.g. a fact may fall out of the window while it is still relevant in the inference process.

In the context of stream reasoning, there is the opportunity to develop more sophisticated forgetting mechanisms. The notion of consumption is slightly different from the one of expiration [23]: consumed facts are not useful for processing, whilst expired facts are not true anymore. Incremental reasoning (Section 2.1.2) relies on the idea that data is removed because it expires. Similarly, initial work on stream reasoning like IMaRS assumes that consumption and expiration overlaps. However, this is not always the case: while some input data can be consumed (e.g. by a sliding window), some derivations may still be useful to solve the task. It follows that such derived data should not be consumed at the same time. We need models and techniques to manage consumption and expiration separately. In this way, the semantics of forgetting data becomes more precise, improving the quality of the engine answers.

Another approach is to exploit the semantics and knowledge about the data content to identify the relevant information. Ongoing research in DSMS is studying how to use knowledge about the streaming data to define windows, e.g. session windows in Google Dataflow [2] and frames [61]. The idea behind such techniques is to create windows by using specific information in the data (e.g. a session identifier in a server access log), rather than by using generic information such as time or number or elements. Moving from DSMS to stream reasoning, TEF-SPARQL [57] allows users to define facts as time-annotated elements by declaring a set of conditions on the input stream items. The approach presented in [128] introduces the notion of semantic importance, as a set of metrics assigned to the stream items, such as query contribution and provenance. These values lead the process of deciding which information should be consumed.

3.2. Towards sophisticated stream reasoning

As the stream reasoning name suggests, reasoning plays a crucial role. In this section, we describe two directions related to this topic. The first is the study of more expressive formalisms for deductive reasoning, such as temporal logics and, more in general, alternatives to description logics. The second is the integration of deductive reasoning with other types of reasoning, such as the inductive one.

3.2.1. Extending the range of logical formalisms

After ten years of stream reasoning investigations, it appears that logical languages are most popular for stream reasoning. Most of the works reported in this paper use or slightly extend OWL 2 DL and its fragments.

In the next years, we think it is important to investigate other inference approaches and how they can be combined with OWL. First steps in this direction were taken using ASP [48,58]. Important contribution may arise by the study of temporal logics [8] in the streaming context. The processing of data streams with Metric Temporal Logic was pioneered in [66] and it is now attracting again interests [26,121]. However, as already noticed in [37], several other logics also appear to be valid starting points, e.g. temporal action logic [47], step logic [50], active logic [49] and event calculus [116].

3.2.2. Integrating other types of reasoning

One element usually found in stream reasoning solutions is the presence of a conceptual model, possibly in combination with rich background knowledge. They can usually be described through an ontological language, which enables the derivation of implicit information. When queries extend the set of operators as described above, an interesting challenge will be to investigate how several reasoning techniques can coexist.

For example, let's consider the system that integrates machine learning and deductive reasoning algorithms described in the end of Section 2.1.3. The authors of [17] built a system that pipes the results of an RSP engine into a machine learning system. However, this is just one possibility. The latter may also feed the former (as in [87]). Moreover, it is possible to exploit more interactive paradigms, where results of machine learning and reasoning techniques are continuously exchanged to achieve a given goal.

In addition to the machine learning mentioned above, we hope to observe an increasing number of explorations that study how to combine (deductive) stream reasoning with other techniques, such as probabilistic reasoning, planning, natural language processing, sentiment analysis.

3.3. Towards semantic streams

One of the ways to introduce semantics is through annotations: they can describe data in a machine-readable fashion, and can consequently be read and processed by systems. Even if this is a typical semantic web use case, this direction has been investigated by few studies so far, e.g. [65,92].

There is a potential value in annotating streams: it enables engines to access a description of the stream and to use it to take decisions, e.g. dynamic discovery and selection of data sources.

The stream descriptor can provide quantitative and qualitative information about the content, e.g. statistical data about frequency and size, information about the vocabularies adopted in the stream items and provenance. Such information may help the stream reasoner in taking some decisions on how to process the stream, even before starting to receive it, e.g. relevancy of the data w.r.t. the registered queries, query optimisation or need for data eviction techniques.

The description of the stream may provide knowledge about its content. We can indeed observe that stream content is heterogeneous in nature. A stream may bring states (e.g. the temperature in a room),

producing data items in a periodic way (e.g. every 2 seconds) or when the state change (e.g. when the temperature increases of one degree). Streams may also describe sequences of actions (e.g. the log of the user interaction in a Web site).

When the query developer aims at describing a task for a stream reasoner, it is up to him to know what the stream carries and consequently to take proper decisions. By exploiting annotations about the stream it would be possible to improve the interoperability at application level: tools may assist the development of queries over the stream and help domain experts with limited technical skills. As usual, when talking about vocabularies and annotations, it will be important to compare the existing proposals as well as the upcoming ones, to find agreements and to come up with shared standards.

3.4. Towards scalable stream reasoners

The progress on stream reasoning foundations sets the basis to build a new generation of more sophisticated stream reasoning frameworks. Researchers integrate reasoning processes in a gradual way, from the application of reasoning over the window content as an ontology, e.g. Streaming Knowledge Bases and Sparkwave, to more sophisticated solutions that take into account also the time dimension and the transient nature of the data stream items in the reasoning process, e.g. ExaStream. In the following, we will present the main challenges that stream reasoner researchers and engineers are going to cope with in order to build new engines.

Scalability is an open and exciting challenge: which order of throughput will stream reasoners be able to support?

As known, there are theoretical results that set some constraints to the velocity that stream reasoners may reach. Reasoning computational complexity is strictly related to the adopted ontological language: the query answering task of the three OWL 2 dialects can vary from AC^0 and polynomial (w.r.t. ABox size) up to NP and EXP (w.r.t. ABox, TBox and query sizes). In this sense, it is not possible to expect that a stream reasoner is going to reach the performance of DSMS and CEP solutions.

However, here we observe a trade-off similar to the one between memory size and access time in computer systems, which is solved using a memory hierarchy. As proposed in [118], a hierarchy of processing steps of increasing complexity can tackle scalability. Technically, this is doable because reasoning can speed up by pushing down processing steps in the hierarchy (e.g., query rewriting) and by post-processing the results coming up from the layer underneath.

3.4.1. Approximating the results

A typical approach to scale DSMS and CEP systems is to move from complete and exact outputs to approximated ones. Such answers are acceptable in a wide range of scenarios, in particular when tasks require aggregations and small errors are tolerable, e.g. counting the number of people entering a square or calculating the average temperature in a building.

There are several ways to achieve approximation. Load shedding [120] techniques capture the idea that the system can produce the output by processing a portion of the stream and throwing away the remaining part. Over the years, several load shedding techniques have been proposed, to select (with the minimal effort) the data to be evicted to minimise the error of the answer. [22,56] introduce data eviction in stream reasoning. The main difference of applying load shedding in stream reasoning is the more complex nature of the data items and the reasoning process itself. Removing data in aggregation operations can introduce errors that can be estimated and controlled. In stream reasoning, the eviction of a single fact may drastically affect the inference process, with high impact on the correctness of the answer.

Besides data eviction, that discards data before processing it, summarization exploits the idea that output can be computed starting with a summary of the input data rather than from the whole dataset. DSMS, CEP and DL reasoning have extensively used these techniques. Summaries in DSMS and CEP, also named sketches, are used to reduce the memory consumption of the engine and to approximate results of aggregations [46]. Several sketch methods are available, usually tailored to specific kinds of aggregation query, e.g. [28]. In the DL-reasoning context, summaries follow a similar idea [53,130]: the ontology (or part of it) is transformed in a *smaller* representation, over which it is possible to perform reasoning tasks.

While the above approaches focus on data, other techniques work to simplify the processing, gaining in performance while introducing some degree of approximation. An example is [103]: authors propose methods to reason over ontologies represented in OWL 2 DL through inference processes of OWL 2 EL, i.e. a tractable fragment of OWL 2 DL. Axioms that cannot be treated in OWL 2 EL (e.g. inverse properties) are managed through ad-hoc rules, applied before and after the reasoning process. In this way, it is usually possible to apply faster algorithms to perform the reasoning task, moving from a situation of certain answer to approximated ones (under some correctness constraints).

3.4.2. Parallelizing and distributing the stream reasoners

Parallelization and distribution can be seen as an opportunity or as a challenge. So far stream reasoning was addressed bringing data streams and contextual (or background) knowledge in one single point. If this point is a cluster, parallelization and distribution (in the cluster) is an approach to engineer scalable and elastic stream reasoners. The first part of this section discusses this opportunity. However, data streams are parallel and distributed sources in nature. The same applies to many Web data sources to join data streams with. Pushing computation to those sources (see also Fog Computing as a broader research field) is the challenge presented in the second half of this section.

When talking about parallelization and distribution, the intuitive idea is that the processing can be split and executed at the same time in multiple locations, e.g. multiple processors in a machine or different nodes in a cluster or a cloud platform.

Looking at stream reasoning, we can find only some attempts in this direction, such as DynamiTE. However, in adjacent areas several investigations are available: in stream processing, e.g. [1,78], in big data processing, e.g. [27,80,129], and in SPARQL query evaluation, e.g. [62,106,114].

Engineering a distributed stream reasoner is a challenging task that touches several scientific and technical problems. Ideally, such a system should maximise the throughput, finding a perfect balance between network load (i.e. how data route through the nodes) and machine load (i.e. the computation loads assigned to the nodes). It is, therefore, important to understand which are the best topologies, operators and data distributions to perform the stream reasoning task.

One initial study in this direction is [52]: authors propose to apply graph partitioning over linked data streams in the context of continuous query answering. The goal is to reduce the network load and consequently improve the performance of the system. Further work is needed to understand the problem of how to cope with the presence of inference processes in the context of reasoning.

Several contributions are available on parallel and distributed reasoning in non-streaming settings. One possible way to achieve this is to treat the data as a set of interconnected ontologies: first reasoning over each ontology is locally performed and the inference completes by exchanging messages, e.g. [14,115]. More recent work exploits new parallelization paradigms to perform the reasoning process, e.g. [96,124]. The usual problems of distributed reasoners are related to termination, i.e. decide when nodes can stop the computation, and to duplicates, i.e. the less duplicated derivations, the higher the performance. The problem in stream reasoning is exacerbated by the need to provide reactive answers.

The last direction we highlight is related to the data distribution problem. When considering scenarios like the Web, it often happens that the data is distributed, controlled by several actors and exposed through services, e.g. SPARQL endpoints or Web APIs, working with either pull or push mechanisms. In such contexts, it often happens that data cannot be centralised and permanently stored in local memory of the stream engine. For example, data cannot fit in the engine memory, it may change over time (and services may not publish notifications), or can only be stored for limited amounts of time for legal reasons. This setting poses interesting challenges to stream reasoning, where responsiveness is one of the most critical requirements. One of the possible ways to see the problem is the following: given the data that can be pushed to the processor, which is the needed contextual remote data to be pulled to solve the given task? In other words, the challenge is how to achieve the integration of the local and remote data without losing responsiveness. The initial effort [35,55] works in the setting of linked data integration of streaming and contextual data for query evaluation purposes. The idea is to adopt caches where to store a portion of the remote data, updating it depending on the recent stream content. Another relevant work is the one in [69], where the authors study the problem of integrating distributed dynamic data and process it through a set of rules. Further techniques are required, since moving and replicating data in the processing nodes impacts the performance.

3.4.3. Reasoning outside the window

The stream-level entailment offers an additional opportunity for stream reasoning. The main challenge is on the resource usage. Given the absence of a sliding window, which introduces a consumption mechanism for the formula validities, a reasoner operating under stream-level entailment may require an unrealistic amount of memory and processing capabilities. In other words, is it possible to build a generic framework to perform stream reasoning under stream-level entailment? Under which conditions would it work without exceeding the assigned amount of resources? Neither LARS nor EP-SPARQL, described in Section 2.3, is targeting such problems. LARS allows defining formulas over the whole stream without limiting the reasoning effect to the window content. However, it is a theoretical framework and no implementations are available at the moment.

Looking at the problem from a database perspective, we can observe that a DSMS can compute the average of those numbers incrementally with finite memory. This relates to the database notion of non-blocking operators. Our intuition is that reasoning outside the window is feasible for *non-blocking reasoning tasks*. For instance, it is possible to compute the materialisation of a stream under the DL-fragment of RDFS when the TBox is fixed. The inference can be applied to each stream item independently, avoiding the storage of streaming data to compute future derivations. That is the case of EP-SPARQL and Etalis: given a fixed schema, all the inference rules are triggered by one data item. In other words, the reasoning process does not need to access the past to compute new derivations. Etalis can still have problems with the memory management: the registered queries may require an infinite amount of memory to compute all the solutions. However, this is a problem related to the CEP-nature of the system rather than to the reasoning one. Researchers investigating online monitoring (i.e. the assessment of a temporal logic formula w.r.t. a stream) are developing solid foundations on this topic. In 2013, [20] introduces the notion of trace-length independence, to indicate monitors able to operate with space resources independent from the number of stream events. Recently, [19] refined the trace-length independence by proposing the notion of *event-rate independence*, to indicate monitors able to work with an amount of space independent from the number of events in a time unit.

Building systems that perform window-level entailment requires coping with the problem of introducing controls on the resource usage, in particularly on the memory usage. A task can usually be modelled through a conceptual model, one or more data streams and a set of operations (represented by a query).

There are classes of queries that guarantee an upper bound on resource usages, given an underlying logical language. EP-SPARQL and the studies mentioned above supply evidence about this since it is possible to execute its query with a limited amount of memory usage.

Answers to the above question can enable the constructions of new algorithms to perform stream- and window-level entailment, able to analyse the current scenario and decide which strategy to adopt to execute the query in a *safe* way avoiding to exceed the assigned resources. For example, a system may decide to compute the correct answer when the conditions allow it, and move to approximated results in the other cases, by adopting item consumption and summary techniques.

3.5. Towards robustness to imperfect data

The road to the usage of stream reasoning in real environments goes through the ability to process imperfect data, that can be either noisy or incomplete. Deductive reasoning is sensible to noise and incomplete data: one single error may lead a system to an overall inconsistent state. In this section, we discuss the open problems related to stream reasoning with imperfect data, analyzing first the ones related to heterogeneity and then the ones related to noise.

3.5.1. Overcoming the heterogeneity

As we depicted above, reasoning offers a set of methods and solutions to cope with the heterogeneity. In particular, such techniques focus on the problem of heterogeneity at schema-level: when models are different, OBDA is a solution to access such data through a conceptual shared model. For example, [33,65] propose to annotate the streams through an ontology, and to reason about those metadata to retrieve and integrate the streaming data needed for the processing. Tackling this problem is important, but heterogeneity issues affect the stream reasoning scenario in other ways.

Data streams can be heterogeneous because they are not synchronised. For instance, imagine two cameras monitoring the same street that report every minute the number of cars they counted. They report the same number only if they are in sync. If one camera starts the counting 20 seconds before the other one, the two counting will differ, but this second situation is normal in an open world. Similarly, a continuous query may require to join a data stream with background knowledge served by a pull API, e.g. to monitor the changes in the number of followers the users mentioned in a stream of microposts, because user profiles normally are only available via pull API. There is no guarantee that the API is returning values that are temporally aligned with the data stream. The solution of this problem requires both a rich semantic description of the data streams (see Section 3.3) and an extension of stream reasoning methods. A possible research direction is to offer a synchronisation service to be used to perform the stream reasoning tasks. [64] proposes to define synchronization policies, built on the top of a notion of state, to perform the streams alignment is these kind of services. Such services may also offer the opportunity to homogenise access to stored time-series and continuously computed predictions.

Another service that such a layer can provide is the on demand discretisation abstraction of quantities as facts. For instance, a stream reasoner may prefer the cameras of the example above to report the level of congestion in the street rather than the counting. However, different applications (or even the same application in different moments) may require different discretisations. A first step in this direction is reported in [123], where the authors report on a system able to answer continuous queries over data stream applying a rewriting method for query answering over temporal fuzzy DL-Lite ontologies.

Finally, heterogeneity can go beyond the data and affect the stream reasoner as well. Existing stream reasoning techniques differ from each other. It is evident when the goal is different, but it happens even when they perform the same task and user may expect the same output. There is an ongoing effort on

studying heterogeneity in stream reasoners. RSEP-QL [43,44] is a reference model to explain heterogeneity in stream reasoners under simple entailment, while LARS [32] has been introduced to capture the behaviour of stream reasoners under more complex regimes. Studying and understanding heterogeneity of stream reasoners is important to achieve interoperability, standardization and comparison.

3.5.2. Coping with noise

DSMS and CEP have always coped with noise [30]. We can distinguish two different types of noise, given that the streaming item is composed of some content and a time annotation.

The first one affects the stream item content. Sensors may break and stop to work. Or worse, they may degrade their precision and provide observations with some degree of error, leading the processing to wrong results. The problem is not limited to the sensor-generated data: streams generated from human interactions may contain syntactical or semantical errors in the data items. In a stream reasoning scenario, this may lead to wrong conclusions and consequently to wrong decisions and actions.

When the stream has a very simple schema (e.g. a time series), statistical methods can supply solutions to manage the noise. However, when we consider more complex schemata, more sophisticated methods may be required. Recently, techniques to cope with noise in stream reasoning emerged, e.g. [85]. The idea is to adopt machine learning to process noisy data and to learn models over which to apply deductive reasoning processes. Inductive reasoning is a powerful tool to cope with noise, but there are other solutions to explore in the next years. Looking at deductive reasoning techniques, inconsistency repair [7,88] and belief revision [109,111] offer solid foundations to build framework to identify noise and to decide the proper actions to be taken. The challenge for researchers is to find algorithms able to use them in the context of stream processing.

The other kind of noise is the one that affects the temporal annotations on the data item. When considering a single stream, the noise manifests in the out-of-order phenomena: for some reasons (e.g. during stream production or transmission), the stream engine is receiving stream items in a wrong order. However, the problem may become more complicated when considering multiple streams. Different sources are producing and publishing them, and this can lead to the introduction of noise, since they may adopt different, not perfectly aligned clocks to generate the temporal annotations. Moreover, one stream can be received sensibly before another one, since the transmission time between two points of a network can require different time.

Several solutions have been proposed to cope with noise in time in the context of stream and event processing, e.g. [89,117], and stream reasoning can get an advantage of such techniques as well. However, we believe that semantics can offer the opportunity to enhance the existing methods to manage such problems. Engines can use the rich and machine-readable descriptions of the data streams to monitor if the received data stream has the correct order. When not, the system would be aware of issues in the stream, with the possibility to take actions.

4. Conclusions

We are observing an impressive increase of the speed of data production and consumption. In this paper, we explained how stream reasoning aims at providing methods and tools to perform sophisticated analyses of such data.

In the beginning, stream reasoning grew with the idea of building such analyses on top of logical and deductive inference. DSMS, CEP and SemWeb offered solid starting points to kick off the research.

Through the years, we have observed the creation of languages, techniques and frameworks. Those studies pushed stream reasoning in a broader area, introducing reasoning techniques beyond the deductive ones. Semantic Scholar and Google Scholar count more than a 1000 articles containing "stream reasoning",[5] published in different areas, from semantic web to artificial intelligence. However, there is still a lot of research to be done.

In Section 3, we presented the main directions over which stream reasoning research can continue. Stream reasoners should offer richer query languages, which include a wider set of operators to encode user needs, and the engine to evaluate them. Reasoning took a more generic connotation, and now it includes inductive reasoning techniques in addition to deductive ones. This trend will grow, combining different techniques to overcome their respective limits. Solutions need to be engineered in scalable frameworks, i.e., they must be able to integrate and reason over huge amounts of heterogeneous data while guaranteeing time requirements. And it will be important to fill the gaps between theoretical models and reality, making stream reasoning solutions robust and able to cope with issues such as noise and heterogeneity. In parallel, it will be important to identify real problems and scenarios where stream reasoning may be a solution. Internet of Things and Industry 4.0 are examples of areas where to apply stream reasoning results. Moreover, it is necessary to develop benchmarking and evaluation activities, to compare and contrast the current solutions.

Results obtained up to now are important. In addition to the publications, some of the mature solutions were exploited in real scenarios, such as social media analytics and turbine monitoring. We should get inspired by such results, and see them as the foundations to build new research and to reach new ambitious achievements, to reach the goal of:

> making sense in real time of multiple, heterogeneous, gigantic and inevitably noisy and incomplete data streams in order to support the decision process of extremely large numbers of concurrent users [36].

Acknowledgements

References

[1] D.J. Abadi, Y. Ahmad, M. Balazinska, U. Çetintemel, M. Cherniack, J. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryvkina, N. Tatbul, Y. Xing and S.B. Zdonik, The design of the borealis stream processing engine, in: *CIDR*, 2005, pp. 277–289. Available online at http://cidrdb.org/cidr2005/papers/P23.pdf.

[2] T. Akidau, R. Bradshaw, C. Chambers, S. Chernyak, R. Fernández-Moctezuma, R. Lax, S. McVeety, D. Mills, F. Perry, E. Schmidt and S. Whittle, The dataflow model: A practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing, *PVLDB* **8**(12) (2015), 1792–1803. Available online at https://dl.acm.org/citation.cfm?id=2824076.

[3] M.I. Ali, F. Gao and A. Mileo, CityBench: A configurable benchmark to evaluate RSP engines using smart city datasets, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 9367, Springer, 2015, pp. 374–389. doi:10.1007/978-3-319-25010-6_25.

---

[5]Checked at February 2017.

[4] D. Anicic, P. Fodor, S. Rudolph and N. Stojanovic, EP-SPARQL: a unified language for event processing and stream reasoning, in: *WWW*, ACM, 2011, pp. 635–644. doi:10.1145/1963405.1963495.

[5] D. Anicic, S. Rudolph, P. Fodor and N. Stojanovic, Stream reasoning and complex event processing in ETALIS, *Semantic Web* **3**(4) (2012), 397–407. doi:10.3233/SW-2011-0053.

[6] A. Arasu, S. Babu and J. Widom, CQL: A language for continuous queries over streams and relations, in: *DBPL*, Lecture Notes in Computer Science, Vol. 2921, Springer, 2003, pp. 1–19. doi:10.1007/978-3-540-24607-7_1.

[7] M. Arenas, L.E. Bertossi and J. Chomicki, Consistent query answers in inconsistent databases, in: *PODS*, ACM Press, 1999, pp. 68–79. doi:10.1145/303976.303983.

[8] A. Artale and E. Franconi, A survey of temporal extensions of description logics, *Ann. Math. Artif. Intell.* **30**(1–4) (2000), 171–210. doi:10.1023/A:1016636131405.

[9] F. Baader, S. Brandt and C. Lutz, Pushing the EL envelope, in: *IJCAI*, Professional Book Center, 2005, pp. 364–369. Available online at https://dl.acm.org/citation.cfm?id=1642351.

[10] F. Baader, C. Lutz and S. Brandt, Pushing the EL envelope further, in: *OWLED (Spring)*, CEUR Workshop Proceedings, Vol. 496, CEUR-WS.org, 2008. Available online at http://ceur-ws.org/Vol-496/owled2008dc_paper_3.pdf.

[11] A. Bachmann, C. Bird, F. Rahman, P.T. Devanbu and A. Bernstein, The missing links: Bugs and bug-fix commits, in: *SIGSOFT FSE*, ACM, 2010, pp. 97–106. doi:10.1145/1882291.1882308.

[12] M. Balduini, I. Celino, D. Dell'Aglio, E. Della Valle, Y. Huang, T.K. Lee, S. Kim and V. Tresp, Reality mining on micropost streams – deductive and inductive reasoning for personalized and location-based recommendations, *Semantic Web* **5**(5) (2014), 341–356. doi:10.3233/SW-130107.

[13] M. Balduini, E. Della Valle, M. Azzi, R. Larcher, F. Antonelli and P. Ciuccarelli, CitySensing: Fusing city data for visual storytelling, *IEEE MultiMedia* **22**(3) (2015), 44–53. doi:10.1109/MMUL.2015.54.

[14] J. Bao, D. Caragea and V. Honavar, A tableau-based federated reasoning algorithm for modular ontologies, in: *Web Intelligence*, IEEE Computer Society, 2006, pp. 404–410. doi:10.1109/WI.2006.28.

[15] D.F. Barbieri, D. Braga, S. Ceri, E. Della Valle and M. Grossniklaus, C-SPARQL: a continuous query language for RDF data streams, *Int. J. Semantic Computing* **4**(1) (2010), 3–25. doi:10.1142/S1793351X10000936.

[16] D.F. Barbieri, D. Braga, S. Ceri, E. Della Valle and M. Grossniklaus, Incremental reasoning on streams and rich background knowledge, in: *ESWC (1)*, Lecture Notes in Computer Science, Vol. 6088, Springer, 2010, pp. 1–15. doi:10.1007/978-3-642-13486-9_1.

[17] D.F. Barbieri, D. Braga, S. Ceri, E. Della Valle, Y. Huang, V. Tresp, A. Rettinger and H. Wermser, Deductive and inductive stream reasoning for semantic social media analytics, *IEEE Intelligent Systems* **25**(6) (2010), 32–41. doi:10.1109/MIS.2010.142.

[18] D.F. Barbieri and E. Della Valle, A proposal for publishing data streams as linked data – A position paper, in: *LDOW*, CEUR Workshop Proceedings, Vol. 628, CEUR-WS.org, 2010. Available online at http://ceur-ws.org/Vol-628/ldow2010_paper11.pdf.

[19] D.A. Basin, B.N. Bhatt and D. Traytel, Almost event-rate independent monitoring of metric temporal logic, in: *TACAS (2)*, Lecture Notes in Computer Science, Vol. 10206, 2017, pp. 94–112. doi:10.1007/978-3-662-54580-5_6.

[20] A. Bauer, J.-C. Küster and G. Vegliach, From propositional to first-order monitoring, in: *RV*, Lecture Notes in Computer Science, Vol. 8174, Springer, 2013, pp. 59–75. doi:10.1007/978-3-642-40787-1_4.

[21] H. Beck, M. Dao-Tran, T. Eiter and M. Fink, LARS: A logic-based framework for analyzing reasoning over streams, in: *AAAI*, AAAI Press, 2015, pp. 1431–1438. Available online at https://dl.acm.org/citation.cfm?id=2887205.

[22] F. Belghaouti, A. Bouzeghoub, Z. Kazi-Aoul and R. Chiky, POL: A pattern oriented load-shedding for semantic data stream processing, in: *WISE (2)*, Lecture Notes in Computer Science, Vol. 10042, 2016, pp. 157–171. doi:10.1007/978-3-319-48743-4_13.

[23] I. Botan, G. Alonso, P.M. Fischer, D. Kossmann and N. Tatbul, Flexible and scalable storage management for data-intensive stream processing, in: *EDBT*, ACM International Conference Proceeding Series, Vol. 360, ACM, 2009, pp. 934–945. doi:10.1145/1516360.1516467.

[24] J. Calbimonte, H. Jeung, Ó. Corcho and K. Aberer, Enabling query technologies for the semantic sensor web, *Int. J. Semantic Web Inf. Syst.* **8**(1) (2012), 43–63. doi:10.4018/jswis.2012010103.

[25] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, M. Rodriguez-Muro, R. Rosati, M. Ruzzi and D.F. Savo, The MASTRO system for ontology-based data access, *Semantic Web* **2**(1) (2011), 43–53. doi:10.3233/SW-2011-0029.

[26] D. Calvanese, E.G. Kalayci, V. Ryzhikov and G. Xiao, Towards practical OBDA with temporal ontologies – (position paper), in: *RR*, Lecture Notes in Computer Science, Vol. 9898, Springer, 2016, pp. 18–24. doi:10.1007/978-3-319-45276-0_2.

[27] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi and K. Tzoumas, Apache flink™: Stream and batch processing in a single engine, *IEEE Data Eng. Bull.* **38**(4) (2015), 28–38. Available online at http://sites.computer.org/debull/A15dec/p28.pdf.

[28] G. Cormode and S. Muthukrishnan, An improved data stream summary: The count-min sketch and its applications, *J. Algorithms* **55**(1) (2005), 58–75. doi:10.1016/j.jalgor.2003.12.001.

[29] G. Cugola and A. Margara, Processing flows of information: From data stream to complex event processing, *ACM Comput. Surv.* **44**(3) (2012), 15. doi:10.1145/2187671.2187677.

[30] G. Cugola, A. Margara, M. Matteucci and G. Tamburrelli, Introducing uncertainty in complex event processing: Model, implementation, and validation, *Computing* **97**(2) (2015), 103–144. doi:10.1007/s00607-014-0404-y.

[31] R. Cyganiak, D. Wood and M. Lanthaler, RDF 1.1 Concepts and Abstract Syntax, W3C Recommendation, W3C, 2014. Available online at https://www.w3.org/TR/rdf11-concepts/.

[32] M. Dao-Tran, H. Beck and T. Eiter, Contrasting RDF stream processing semantics, in: *JIST*, Lecture Notes in Computer Science, Vol. 9544, Springer, 2015, pp. 289–298. doi:10.1007/978-3-319-31676-5_21.

[33] D. de Leng and F. Heintz, Ontology-based introspection in support of stream reasoning, in: *SCAI*, Frontiers in Artificial Intelligence and Applications, Vol. 278, IOS Press, 2015, pp. 78–87. doi:10.3233/978-1-61499-589-0-78.

[34] D. de Leng and F. Heintz, Qualitative spatio-temporal stream reasoning with unobservable intertemporal spatial relations using landmarks, in: *AAAI*, AAAI Press, 2016, pp. 957–963. Available online at https://dl.acm.org/citation.cfm?id=3015955.

[35] S. Dehghanzadeh, D. Dell'Aglio, S. Gao, E. Della Valle, A. Mileo and A. Bernstein, Approximate continuous query answering over streams and dynamic linked data sets, in: *ICWE*, Lecture Notes in Computer Science, Vol. 9114, Springer, 2015, pp. 307–325. doi:10.1007/978-3-319-19890-3_20.

[36] E. Della Valle, On Stream Reasoning, PhD thesis, Vrije Universiteit Amsterdam, 2015. Available online at http://dare.ubvu.vu.nl/handle/1871/53293.

[37] E. Della Valle, S. Ceri, F. van Harmelen and D. Fensel, It's a streaming world! Reasoning upon rapidly changing information, *IEEE Intelligent Systems* **24**(6) (2009), 83–89. doi:10.1109/MIS.2009.125.

[38] E. Della Valle, D. Dell'Aglio and A. Margara, Taming velocity and variety simultaneously in big data with stream reasoning: Tutorial, in: *DEBS*, ACM, 2016, pp. 394–401. doi:10.1145/2933267.2933539.

[39] E. Della Valle, S. Schlobach, M. Krötzsch, A. Bozzon, S. Ceri and I. Horrocks, Order matters! Harnessing a world of orderings for reasoning over massive data, *Semantic Web* **4**(2) (2013), 219–231. doi:10.3233/SW-2012-0085.

[40] D. Dell'Aglio, On Unified Stream Reasoning, PhD thesis, Politecnico di Milano, 2016. Availble online at http://hdl.handle.net/10589/122892.

[41] D. Dell'Aglio, M. Balduini and E. Della Valle, Applying semantic interoperability principles to data stream management, in: *Data Management in Pervasive Systems*, Data-Centric Systems and Applications, Springer, 2015, pp. 135–166. doi:10.1007/978-3-319-20062-0_7.

[42] D. Dell'Aglio, J. Calbimonte, M. Balduini, Ó. Corcho and E. Della Valle, On correctness in RDF stream processor benchmarking, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 8219, Springer, 2013, pp. 326–342. doi:10.1007/978-3-642-41338-4_21.

[43] D. Dell'Aglio, M. Dao-Tran, J. Calbimonte, D. Le Phuoc and E. Della Valle, A query model to capture event pattern matching in RDF stream processing query languages, in: *EKAW*, Lecture Notes in Computer Science, Vol. 10024, 2016, pp. 145–162. doi:10.1007/978-3-319-49004-5_10.

[44] D. Dell'Aglio, E. Della Valle, J. Calbimonte and Ó. Corcho, RSP-QL semantics: A unifying query model to explain heterogeneity of RDF stream processing systems, *Int. J. Semantic Web Inf. Syst.* **10**(4) (2014), 17–44. doi:10.4018/ijswis.2014100102.

[45] C. Dobbins, P. Fergus, M. Merabti and D. Llewellyn-Jones, Monitoring and measuring sedentary behaviour with the aid of human digital memories, in: *CCNC*, IEEE, 2012, pp. 395–398. doi:10.1109/CCNC.2012.6181016.

[46] A. Dobra, M.N. Garofalakis, J. Gehrke and R. Rastogi, Processing complex aggregate queries over data streams, in: *SIGMOD Conference*, ACM, 2002, pp. 61–72. doi:10.1145/564691.564699.

[47] P. Doherty, J. Gustafsson, L. Karlsson and J. Kvarnström, TAL: Temporal action logics language specification and tutorial, *Electron. Trans. Artif. Intell.* **2** (1998), 273–306. Availble online at http://www.ep.liu.se/ej/etai/1998/009/.

[48] T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer and H. Tompits, Combining answer set programming with description logics for the semantic web, *Artif. Intell.* **172**(12–13) (2008), 1495–1539. doi:10.1016/j.artint.2008.04.002.

[49] J. Elgot-Drapkin, S. Kraus, M. Miller, M. Nirkhe and D. Perlis, Active logics: A unified formal approach to episodic reasoning, Technical report, 1999. Availble online at http://hdl.handle.net/1903/1039.

[50] J.J. Elgot-Drapkin, Step-logic: Reasoning situated in time, PhD thesis, University of Maryland at College Park, 1988. Availble online at https://dl.acm.org/citation.cfm?id=915324.

[51] O. Erling, A. Averbuch, J. Larriba-Pey, H. Chafi, A. Gubichev, A. Prat-Pérez, M. Pham and P.A. Boncz, The LDBC social network benchmark: Interactive workload, in: *SIGMOD Conference*, ACM, 2015, pp. 619–630. doi:10.1145/2723372.2742786.

[52] L. Fischer, T. Scharrenbach and A. Bernstein, Scalable linked data stream processing via network-aware workload scheduling, in: *SSWS@ISWC*, CEUR Workshop Proceedings, Vol. 1046, CEUR-WS.org, 2013, pp. 81–96. Available online at http://ceur-ws.org/Vol-1046/SSWS2013_paper6.pdf.

[53] A. Fokoue, A. Kershenbaum, L. Ma, E. Schonberg and K. Srinivas, The summary abox: Cutting ontologies down to size, in: *International Semantic Web Conference*, Lecture Notes in Computer Science, Vol. 4273, Springer, 2006, pp. 343–356. doi:10.1007/11926078_25.

[54] C. Forgy, Rete: A fast algorithm for the many patterns/many objects match problem, *Artif. Intell.* **19**(1) (1982), 17–37. doi:10.1016/0004-3702(82)90020-0.

[55] S. Gao, D. Dell'Aglio, S. Dehghanzadeh, A. Bernstein, E. Della Valle and A. Mileo, Planning ahead: Stream-driven linked-data access under update-budget constraints, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 9981, 2016, pp. 252–270. doi:10.1007/978-3-319-46523-4_16.

[56] S. Gao, T. Scharrenbach and A. Bernstein, The CLOCK data-aware eviction approach: Towards processing linked data streams with limited resources, in: *ESWC*, Lecture Notes in Computer Science, Vol. 8465, Springer, 2014, pp. 6–20. doi:10.1007/978-3-319-07443-6_2.

[57] S. Gao, T. Scharrenbach, J. Kietz and A. Bernstein, Running out of bindings? Integrating facts and events in linked data stream processing, in: *SSN-TC/OrdRing@ISWC*, CEUR Workshop Proceedings, Vol. 1488, CEUR-WS.org, 2015, pp. 63–74. Available online at http://ceur-ws.org/Vol-1488/paper-07.pdf.

[58] M. Gebser, O. Sabuncu and T. Schaub, An incremental answer set programming based system for finite model computation, *AI Commun.* **24**(2) (2011), 195–212. doi:10.3233/AIC-2011-0496.

[59] B. Glimm and C. Ogbuji, SPARQL 1.1 Entailment Regimes, W3C Recommendation, W3C, 2013. Available online at https://www.w3.org/TR/sparql11-entailment/.

[60] J. Gray (ed.), *The Benchmark Handbook for Database and Transaction Systems (2nd Edition)*, Morgan Kaufmann, 1993. ISBN 1-55860-292-5.

[61] M. Grossniklaus, D. Maier, J. Miller, S. Moorthy and K. Tufte, Frames: Data-driven windows, in: *DEBS*, ACM, 2016, pp. 13–24. doi:10.1145/2933267.2933304.

[62] S. Gurajada, S. Seufert, I. Miliaraki and M. Theobald, TriAD: A distributed shared-nothing RDF engine based on asynchronous message passing, in: *SIGMOD Conference*, ACM, 2014, pp. 289–300. doi:10.1145/2588555.2610511.

[63] S. Harris and A. Seaborne, SPARQL 1.1 Query Language, W3C Recommendation, W3C, 2013. Available online at https://www.w3.org/TR/sparql11-query/.

[64] F. Heintz, DyKnow: A Stream-Based Knowledge Processing Middleware Framework, PhD thesis, Linköping University, Sweden, 2009. Available online at www.diva-portal.org/smash/record.jsf?pid=diva2:159661.

[65] F. Heintz and D. de Leng, Semantic information integration with transformations for stream reasoning, in: *FUSION*, IEEE, 2013, pp. 445–452. Available online at ieeexplore.ieee.org/document/6641314/.

[66] F. Heintz and P. Doherty, DyKnow: An approach to middleware for knowledge processing, *Journal of Intelligent and Fuzzy Systems* **15**(1) (2004), 3–13. Available online at http://content.iospress.com/articles/journal-of-intelligent-and-fuzzy-systems/ifs00220.

[67] S. Huang, J. Huang, J. Dai, T. Xie and B. Huang, The HiBench benchmark suite: Characterization of the MapReduce-based data analysis, in: *2010 IEEE 26th International Conference on Data Engineering Workshops (ICDEW)*, 2010, pp. 41–51. doi:10.1109/ICDEW.2010.5452747.

[68] Z. Huang and H. Stuckenschmidt, Reasoning with multi-version ontologies: A temporal logic approach, in: *International Semantic Web Conference*, Lecture Notes in Computer Science, Vol. 3729, Springer, 2005, pp. 398–412. doi:10.1007/11574620_30.

[69] T. Käfer, A. Harth and S. Mamessier, Towards declarative programming and querying in a distributed cyber-physical system: The I-VISION case, in: *CPS Data*, IEEE, 2016, pp. 1–6. doi:10.1109/CPSData.2016.7496418.

[70] R. Keskisärkkä, Query templates for RDF stream processing, in: *SR+SWIT@ISWC*, CEUR Workshop Proceedings, Vol. 1783, CEUR-WS.org, 2016, pp. 25–36. Available online at http://ceur-ws.org/Vol-1783/paper-03.pdf.

[71] R. Keskisärkkä and E. Blomqvist, Semantic complex event processing for social media monitoring-a survey, in: *Proceedings of Social Media and Linked Data for Emergency Response (SMILE) Co-Located with the 10th Extended Semantic Web Conference, Montpellier, France. CEUR Workshop Proceedings (May 2013)*, 2013. Available online at http://ceur-ws.org/Vol-1191/paper1.pdf.

[72] E. Kharlamov, S. Brandt, M. Giese, E. Jiménez-Ruiz, S. Lamparter, C. Neuenstadt, Ö.L. Özçep, C. Pinkel, A. Soylu, D. Zheleznyakov, M. Roshchin, S. Watson and I. Horrocks, Semantic access to siemens streaming data: The optique way, in: *International Semantic Web Conference (Posters & Demos)*, CEUR Workshop Proceedings, Vol. 1486, CEUR-WS.org, 2015. Available online at http://ceur-ws.org/Vol-1486/paper_23.pdf.

[73] E. Kharlamov, S. Brandt, E. Jiménez-Ruiz, Y. Kotidis, S. Lamparter, T. Mailis, C. Neuenstadt, Ö.L. Özçep, C. Pinkel, C. Svingos, D. Zheleznyakov, I. Horrocks, Y.E. Ioannidis and R. Möller, Ontology-based integration of streaming and static relational data with optique, in: *SIGMOD Conference*, ACM, 2016, pp. 2109–2112. doi:10.1145/2882903.2899385.

[74] E. Kharlamov, D. Hovland, E. Jiménez-Ruiz, D. Lanti, H. Lie, C. Pinkel, M. Rezk, M.G. Skjæveland, E. Thorstensen, G. Xiao, D. Zheleznyakov and I. Horrocks, Ontology based access to exploration data at statoil, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 9367, Springer, 2015, pp. 93–112. doi:10.1007/978-3-319-25010-6_6.

[75] E. Kharlamov, Y. Kotidis, T. Mailis, C. Neuenstadt, C. Nikolaou, Ö.L. Özçep, C. Svingos, D. Zheleznyakov, S. Brandt, I. Horrocks, Y.E. Ioannidis, S. Lamparter and R. Möller, Towards analytics aware ontology based access to static and streaming data, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 9982, 2016, pp. 344–362. doi:10.1007/978-3-319-46547-0_31.

[76] C. Kiefer, A. Bernstein and A. Locher, Adding data mining support to SPARQL via statistical relational learning methods, in: *ESWC*, Lecture Notes in Computer Science, Vol. 5021, Springer, 2008, pp. 478–492. doi:10.1007/978-3-540-68234-9_36.

[77] M. Kolchin, P. Wetz, E. Kiesling and A.M. Tjoa, YABench: A comprehensive framework for RDF stream processor correctness and performance assessment, in: *ICWE*, Lecture Notes in Computer Science, Vol. 9671, Springer, 2016, pp. 280–298. doi:10.1007/978-3-319-38791-8_16.

[78] A. Koliousis, M. Weidlich, R.C. Fernandez, A.L. Wolf, P. Costa and P.R. Pietzuch, SABER: Window-based hybrid stream processing for heterogeneous architectures, in: *SIGMOD Conference*, ACM, 2016, pp. 555–569. doi:10.1145/2882903.2882906.

[79] S. Komazec, D. Cerri and D. Fensel, Sparkwave: Continuous schema-enhanced pattern matching over RDF data streams, in: *DEBS*, ACM, 2012, pp. 58–68. doi:10.1145/2335484.2335491.

[80] S. Kulkarni, N. Bhagat, M. Fu, V. Kedigehalli, C. Kellogg, S. Mittal, J.M. Patel, K. Ramasamy and S. Taneja, Twitter heron: Stream processing at scale, in: *SIGMOD Conference*, ACM, 2015, pp. 239–250. doi:10.1145/2723372.2742788.

[81] K. Kyzirakos, M. Karpathiotakis, K. Bereta, G. Garbis, C. Nikolaou, P. Smeros, S. Giannakopoulou, K. Dogani and M. Koubarakis, The spatiotemporal RDF store strabon, in: *SSTD*, Lecture Notes in Computer Science, Vol. 8098, Springer, 2013, pp. 496–500. doi:10.1007/978-3-642-40235-7_35.

[82] D. Le Phuoc, M. Dao-Tran, J.X. Parreira and M. Hauswirth, A native and adaptive approach for unified processing of linked streams and linked data, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 7031, Springer, 2011, pp. 370–388. doi:10.1007/978-3-642-25073-6_24.

[83] D. Le Phuoc, M. Dao-Tran, M. Pham, P.A. Boncz, T. Eiter and M. Fink, Linked stream data processing engines: Facts and figures, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 7650, Springer, 2012, pp. 300–312. doi:10.1007/978-3-642-35173-0_20.

[84] F. Lécué, Diagnosing changes in an ontology stream: A DL reasoning approach, in: *AAAI*, AAAI Press, 2012. Available online at https://dl.acm.org/citation.cfm?id=2900740.

[85] F. Lécué, Towards scalable exploration of diagnoses in an ontology stream, in: *AAAI*, AAAI Press, 2014, pp. 87–93. Available online at https://dl.acm.org/citation.cfm?id=2893889.

[86] F. Lécué and J.Z. Pan, Predicting knowledge in an ontology stream, in: *IJCAI*, IJCAI/AAAI, 2013, pp. 2662–2669. Available online at https://dl.acm.org/citation.cfm?id=2540512.

[87] F. Lécué, S. Tallevi-Diotallevi, J. Hayes, R. Tucker, V. Bicer, M.L. Sbodio and P. Tommasi, Smart traffic analytics in the semantic web with STAR-CITY: Scenarios, system and lessons learned in Dublin city, *J. Web Sem.* **27** (2014), 26–33. doi:10.1016/j.websem.2014.07.002.

[88] D. Lembo and M. Ruzzi, Consistent query answering over description logic ontologies, in: *Description Logics*, Vol. 250, CEUR-WS.org, 2007. Available online at http://ceur-ws.org/Vol-250/paper_26.pdf.

[89] M. Liu, M. Li, D. Golovnya, E.A. Rundensteiner and K.T. Claypool, Sequence pattern query processing over out-of-order event streams, in: *ICDE*, IEEE Computer Society, 2009, pp. 784–795. doi:10.1109/ICDE.2009.95.

[90] D. Luckham, The power of events: An introduction to complex event processing in distributed enterprise systems, in: *RuleML*, Lecture Notes in Computer Science, Vol. 5321, Springer, 2008, p. 3. doi:10.1007/978-3-540-88808-6_2.

[91] A. Margara, J. Urbani, F. van Harmelen and H.E. Bal, Streaming the web: Reasoning over dynamic data, *J. Web Sem.* **25** (2014), 24–44. doi:10.1016/j.websem.2014.02.001.

[92] A. Mauri, J.-P. Calbimonte, D. Dell'Aglio, M. Balduini, M. Brambilla, E. Della Valle and K. Aberer, TripleWave: Spreading RDF streams on the web, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 9982, Springer, 2016, pp. 140–149. doi:10.1007/978-3-319-46547-0_15.

[93] A. McGregor, Graph stream algorithms: A survey, *SIGMOD Record* **43**(1) (2014), 9–20. doi:10.1145/2627692.2627694.

[94] X. Meng, J.K. Bradley, B. Yavuz, E.R. Sparks, S. Venkataraman, D. Liu, J. Freeman, D.B. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M.J. Franklin, R. Zadeh, M. Zaharia and A. Talwalkar, MLlib: Machine Learning in Apache Spark, *CoRR*, abs/1505.06807 (2015).

[95] A. Mileo, A. Abdelrahman, S. Policarpio and M. Hauswirth, StreamRule: A nonmonotonic stream reasoning system for the semantic web, in: *RR*, Lecture Notes in Computer Science, Vol. 7994, Springer, 2013, pp. 247–252. doi:10.1007/978-3-642-39666-3_23.

[96] B. Motik, Y. Nenov, R. Piro, I. Horrocks and D. Olteanu, Parallel materialisation of datalog programs in centralised, main-memory RDF systems, in: *AAAI*, AAAI Press, 2014, pp. 129–137. Available online at https://dl.acm.org/citation. cfm?id=2893896.

[97] B. Motik, Y. Nenov, R.E.F. Piro and I. Horrocks, Incremental update of datalog materialisation: The backward/forward algorithm, in: *AAAI*, AAAI Press, 2015, pp. 1560–1568. Available online at https://dl.acm.org/citation.cfm?id=2886537.

[98] B. Motik, U. Sattler and R. Studer, Query answering for OWL-DL with rules, *J. Web Sem.* **3**(1) (2005), 41–60. doi:10. 1016/j.websem.2005.05.001.

[99] K. Mouratidis, S. Bakiras and D. Papadias, Continuous monitoring of top-k queries over sliding windows, in: *SIGMOD Conference*, ACM, 2006, pp. 635–646. doi:10.1145/1142473.1142544.

[100] S. Muñoz, J. Pérez and C. Gutierrez, Simple and efficient minimal RDFS, *J. Web Sem.* **7**(3) (2009), 220–234. doi:10. 1016/j.websem.2009.07.003.

[101] Y. Nenov, R. Piro, B. Motik, I. Horrocks, Z. Wu and J. Banerjee, RDFox: A highly-scalable RDF store, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 9367, Springer, 2015, pp. 3–20. doi:10.1007/ 978-3-319-25010-6_1.

[102] Ö.L. Özçep, R. Möller and C. Neuenstadt, A stream-temporal query language for ontology based data access, in: *Description Logics*, CEUR Workshop Proceedings, Vol. 1193, CEUR-WS.org, 2014, pp. 696–708. Available online at http://ceur-ws.org/Vol-1193/paper_47.pdf.

[103] J.Z. Pan, Y. Ren and Y. Zhao, Tractable approximate deduction for OWL, *Artif. Intell.* **235** (2016), 95–155. doi:10.1016/ j.artint.2015.10.004.

[104] M. Perry and J. Herring, OGC GeoSPARQL – A Geographic Query Language for RDF Data, Technical report, Open Geospatial Consortium, 2012. Available online at http://www.opengeospatial.org/standards/geosparql.

[105] M. Potamias, K. Patroumpas and T.K. Sellis, Sampling trajectory streams with spatiotemporal criteria, in: *SSDBM*, IEEE Computer Society, 2006, pp. 275–284. doi:10.1109/SSDBM.2006.45.

[106] A. Potter, B. Motik, Y. Nenov and I. Horrocks, Distributed RDF query answering with dynamic data exchange, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 9981, 2016, pp. 480–497. doi:10. 1007/978-3-319-46523-4_29.

[107] D. Puiu, P.M. Barnaghi, R. Toenjes, D. Kuemper, M.I. Ali, A. Mileo, J.X. Parreira, M. Fischer, S. Kolozali, N. FarajiDavar, F. Gao, T. Iggena, Thu-Le Pham, C. Nechifor, D. Puschmann and J. Fernandes, CityPulse: Large scale data analytics framework for smart cities, *IEEE Access* **4** (2016), 1086–1108. doi:10.1109/ACCESS.2016.2541999.

[108] Y. Ren and J.Z. Pan, Optimising ontology stream reasoning with truth maintenance system, in: *CIKM*, ACM, 2011, pp. 831–836. doi:10.1145/2063576.2063696.

[109] M.M. Ribeiro, *Belief Revision in Non-Classical Logics*, Springer Briefs in Computer Science, Springer, 2013. doi:10. 1007/978-1-4471-4186-0.

[110] M. Rinne, M. Solanki and E. Nuutila, RFID-based logistics monitoring with semantics-driven event processing, in: *DEBS*, ACM, 2016, pp. 238–245. doi:10.1145/2933267.2933300.

[111] H. Rott, *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*, Vol. 42, Oxford University Press, 2001. ISBN 9780198503064.

[112] T. Saveta, E. Daskalaki, G. Flouris, I. Fundulaki, M. Herschel and A.N. Ngomo, LANCE: Piercing to the heart of instance matching tools, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 9366, Springer, 2015, pp. 375–391. doi:10.1007/978-3-319-25007-6_22.

[113] T. Scharrenbach, J. Urbani, A. Margara, E. Della Valle and A. Bernstein, Seven commandments for benchmarking semantic flow processing systems, in: *ESWC*, Lecture Notes in Computer Science, Vol. 7882, Springer, 2013, pp. 305–319. doi:10.1007/978-3-642-38288-8_21.

[114] A. Schätzle, M. Przyjaciel-Zablocki, S. Skilevic and G. Lausen, S2RDF: RDF querying with SPARQL on spark, *PVLDB* **9**(10) (2016), 804–815. Available online at https://dl.acm.org/citation.cfm?id=2977806.

[115] L. Serafini and A. Tamilin, DRAGO: Distributed reasoning architecture for the semantic web, in: *ESWC*, Lecture Notes in Computer Science, Vol. 3532, Springer, 2005, pp. 361–376. doi:10.1007/11431053_25.

[116] A. Skarlatidis, G. Paliouras, A. Artikis and G.A. Vouros, Probabilistic event calculus for event recognition, *ACM Trans. Comput. Log.* **16**(2) (2015), 11. doi:10.1145/2699916.

[117] U. Srivastava and J. Widom, Flexible time management in data stream systems, in: *PODS*, ACM, 2004, pp. 263–274. doi:10.1145/1055558.1055596.

[118] H. Stuckenschmidt, S. Ceri, E. Della Valle and F. van Harmelen, Towards expressive stream reasoning, in: *Semantic Challenges in Sensor Networks*, Dagstuhl Seminar Proceedings, Vol. 10042, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010. Available online at http://drops.dagstuhl.de/opus/volltexte/2010/2555/.

[119] S. Tallevi-Diotallevi, S. Kotoulas, L. Foschini, F. Lécué and A. Corradi, Real-time urban monitoring in Dublin using semantic and stream technologies, in: *International Semantic Web Conference (2)*, Lecture Notes in Computer Science, Vol. 8219, Springer, 2013, pp. 178–194. doi:10.1007/978-3-642-41338-4_12.

[120] N. Tatbul, U. Çetintemel, S.B. Zdonik, M. Cherniack and M. Stonebraker, Load shedding in a data stream manager, in: *VLDB*, 2003, pp. 309–320. Available online at https://dl.acm.org/citation.cfm?id=1315479.

[121] M. Tiger and F. Heintz, Stream reasoning using temporal logic and predictive probabilistic state models, in: *TIME*, IEEE Computer Society, 2016, pp. 196–205. doi:10.1109/TIME.2016.28.

[122] R. Tommasini, E. Della Valle, M. Balduini and D. Dell'Aglio, Heaven: A framework for systematic comparative research approach for RSP engines, in: *ESWC*, Lecture Notes in Computer Science, Vol. 9678, Springer, 2016, pp. 250–265. doi:10.1007/978-3-319-34129-3_16.

[123] A. Turhan and E. Zenker, Towards temporal fuzzy query answering on stream-based data, in: *HiDeSt@KI*, CEUR Workshop Proceedings, Vol. 1447, CEUR-WS.org, 2015, pp. 56–69. Available online at http://ceur-ws.org/Vol-1447/paper5.pdf.

[124] J. Urbani, S. Kotoulas, J. Maassen, F. van Harmelen and H.E. Bal, WebPIE: A web-scale parallel inference engine using MapReduce, *J. Web Sem.* **10** (2012), 59–75. doi:10.1016/j.websem.2011.05.004.

[125] J. Urbani, A. Margara, C.J.H. Jacobs, F. van Harmelen and H.E. Bal, DynamiTE: Parallel materialization of dynamic RDF data, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 8218, Springer, 2013, pp. 657–672. doi:10.1007/978-3-642-41335-3_41.

[126] R. Volz, S. Staab and B. Motik, Incrementally maintaining materializations of ontologies stored in logic databases, *J. Data Semantics* **2** (2005), 1–34. doi:10.1007/978-3-540-30567-5_1.

[127] O. Walavalkar, A. Joshi, T. Finin and Y. Yesha, Streaming knowledge bases, in: *In International Workshop on Scalable Semantic Web Knowledge Base Systems*, 2008. Available online at http://ebiquity.umbc.edu/paper/html/id/416/Streaming-Knowledge-Bases.

[128] R. Yan, M.T. Greaves, W.P. Smith and D.L. McGuinness, Remembering the important things: Semantic importance in stream reasoning, in: *SR+SWIT@ISWC*, CEUR Workshop Proceedings, Vol. 1783, CEUR-WS.org, 2016, pp. 49–54. Available online at http://ceur-ws.org/Vol-1783/paper-05.pdf.

[129] M. Zaharia, R.S. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M.J. Franklin, A. Ghodsi, J. Gonzalez, S. Shenker and I. Stoica, Apache spark: A unified engine for big data processing, *Commun. ACM* **59**(11) (2016), 56–65. doi:10.1145/2934664.

[130] X. Zhang, G. Cheng and Y. Qu, Ontology summarization based on rdf sentence graph, in: *WWW*, ACM, 2007, pp. 707–716. doi:10.1145/1242572.1242668.

[131] Y. Zhang, M. Pham, Ó. Corcho and J. Calbimonte, SRBench: A streaming RDF/SPARQL benchmark, in: *International Semantic Web Conference (1)*, Lecture Notes in Computer Science, Vol. 7649, Springer, 2012, pp. 641–657. doi:10.1007/978-3-642-35176-1.