

Deploying Spatial-Stream Query Answering in C-ITS Scenarios¹

Thomas Eiter^a, Ryutaro Ichise^{b,d}, Josiane Xavier Parreira^c, Patrik Schneider^{a,c,*} and Lihua Zhao^d

^a *Vienna University of Technology, Vienna, Austria*

E-mails: eiter@kr.tuwien.ac.at, patrik@kr.tuwien.ac.at

^b *National Institute of Informatics, Tokyo, Japan*

E-mail: ichise@nii.ac.jp

^c *Siemens AG Österreich, Vienna, Austria*

E-mail: josiane.parreira@siemens.com

^d *National Institute of Advanced Industrial Science and Technology, Tokyo, Japan*

E-mail: lihua.zhao@aist.go.jp

Editors: Catherine Faron, Université Côte d'Azur, CNRS, Inria, I3S, France; Chiara Ghidini, Fondazione Bruno Kessler, Trento, Italy

Solicited reviews: Maxime Lefrançois, École des Mines de Saint-Étienne, France; Manolis Koubarakis, National and Kapodistrian University of Athens, Greece; Three anonymous reviewers

Abstract. Cooperative Intelligent Transport Systems (C-ITS) play an important role for providing the means to collect and exchange spatio-temporal data via V2X-based communication between vehicles and the infrastructure, which will become a central enabler for road safety of (semi)-autonomous vehicles. The Local Dynamic Map (LDM) is a key concept for integrating static and streamed data in a spatial context. The LDM has been semantically enhanced to allow for an elaborate domain model that is captured by a mobility ontology, and for queries over data streams that cater for semantic concepts and spatial relationships. Our approach for semantic enhancement is in the context of ontology-mediated query answering (OQA) and features conjunctive queries over DL-LiteA ontologies that support window operators over streams and spatial relations between spatial objects. In this paper, we show how this approach can be extended to address a wider range of use cases in the three C-ITS scenarios *traffic statistics*, *traffic events detection*, and *advanced driving assistance systems*. We define for the mentioned use cases requirements derived from necessary domain-specific features and report, based on them, on extensions of our query language and ontology model. The extensions include temporal relations, numeric predictions and trajectory predictions as well as optimization strategies such as caching. An experimental evaluation of queries that reflect the requirements has been conducted using the real-world traffic simulation tool PTV Vissim. It provides evidence for the feasibility/efficiency of our approach in the new scenarios.

Keywords: Mobility, C-ITS, Query Answering, Ontology-based Data Access, Stream Reasoning, Temporal Relations

1. Introduction

The development in (semi)-autonomous vehicles leads to an extensive communication between vehicles and the infrastructure, which is covered by Cooperative Intelligent Transport Systems (C-ITS). These systems

produce temporal data (e.g., traffic light signal phases) and geospatial data (e.g., GPS positions), which are exchanged in vehicle-to-vehicle, vehicle-to-infrastructure, and combined communications (V2X). This aids to improve road safety by analyzing traffic scenes that could lead to accidents (e.g., red light violations), and to reduce emissions by optimizing traffic flow (e.g., dissolve traffic jams). A key technology for this is the Local Dynamic Map (LDM) [2] as an integration platform for

¹This article is a revised and extended version of a paper presented at EKAW 2018 [1].

*Corresponding author. E-mail: patrik@kr.tuwien.ac.at.

static, semi-static, and dynamic information in a spatial context.

In previous work, we have semantically enhanced the LDM to allow for an elaborate domain model that is captured by a mobility ontology, and for queries over data streams that cater for semantic concepts and spatial relationships [3]. Our approach is based on ontology-mediated query answering (OQA) and features conjunctive queries (CQs) over DL-Lite_A [4] ontologies that support window operators over streams and spatial relations between objects. We believe that OQA and the related ontology-based data access (OBDA) [5] are well suited for C-ITS applications, as an ontology can be used to model vehicles, traffic, and infrastructure details, and map to scalable stream database technology adding dynamicity to the model. For example, the definition of a hazardous situation is complex, ranging from bad road conditions to traffic jams [2]. Therefore, an expressive query language is crucial to cover C-ITS specific requirements needed for retrieving dynamic data and expressing complex patterns regarding, event detection for example. Furthermore, scalability and swift response time are crucial since fast changing traffic demands a quick response time of ranging below 1s to avoid accidents [2].

In this paper, we continue the work in [1, 3, 6] with the goal of showing how spatial-stream OQA can be used to address a wider set of C-ITS scenarios. For achieving this, the approach in [3] is extended with new domain-specific features beyond “generic” spatial-stream OQA. In cooperation with ITS domain experts from Siemens, the C-ITS scenarios – *traffic statistics*, *events detection*, and *advanced driving assistance systems* (ADAS) – were defined and used to single out requirements derived from a domain-specific list of features. We then formulate, for each use case, requirements that should be covered by our approach. The focus of the new, more specific features will be on *temporal relations*, e.g., *during*, as well as numerical and *trajectory predictions* on a query window. For the qualitative assessment of the use cases, requirements, and domain-specific features, we conducted several interviews with ITS experts, which supported our assumptions that temporal relations and predictions are important extensions, but also gave raise to important extensions for future work such as capturing uncertainty in the ontology and the query language. For the quantitative assessment, we provide a detailed report on the improved implementation. The improvements includes new temporal relations and predictions, but also improvements in the performance based on caching static

elements of the query, and the parallel execution of stream atoms. The implementation is evaluated in an experimental setting using queries that match with the features, where a real-world traffic simulation is used to generate the data. The results provide evidence for the potential feasibility and efficiency of our approach in these scenarios. Our contributions are briefly summarized as follows:

- we outline the field of V2X integration using LDMs and provide details on our ontology-based LDM (Section 2);
- we define three scenarios, use cases, desired features, and requirements (Section 3);
- we conducted expert interviews to obtain feedback on the scenarios, use cases, and query features (Section 4);
- we present our current approach including data model, query language, and evaluation strategy (Section 5);
- we report on the implementation of our approach in a prototype and comment on the implementation details (Section 6);
- we evaluate our work regarding the set of features and requirements based on a traffic simulation and assess the results (Section 7);
- we list related work, and evaluate existing stream reasoning systems, wherein we compare the performance of our prototype to the systems C-SPARQL [7] and CQELS [8] (Section 8).

In Section 9, we discuss lessons learned and conclude with ongoing and future work.

This article is a revised and extended version of our preliminary work [1], which we have enriched by the following extensions:

1. a report on interviews of C-ITS experts regarding the spatial-QA approach for real-world ITS applications, whose suggestions have been incorporated to improve the approach (Section 4);
2. a further version of the LDM ontology (Section 2) that builds on the Sensor-Observation-Sampling-Actuator (SOSA) vocabulary [9], following an expert suggestion;
3. a complexity assessment of the query answering approach (Section 5);
4. an outline of optimization strategies for the query rewriting and evaluation covering related parameters for the strategies (Section 5);
5. an increased set of experiments taking the optimization strategies into account (Section 6);

6. a quantitative and qualitative comparison with existing stream reasoning systems (Section 8); and
7. a discussion of the lessons learned (Section 9).

Furthermore, we have enhanced Sections 5 and 6 with the purpose of adding clarity and self-containment to the paper. For this, we have added definitions and details on the semantics and the initial algorithm for the query answering approach from our initial work [3].

2. C-ITS Data Integration and Query Answering

Our setting is the ongoing efforts in data integration and querying in the C-ITS domain. The base technologies for C-ITS are already available and experimentally deployed in infrastructure projects as in [2]. The communication technology is based on the IEEE 802.11p standard, and the data integration effort is the *Local Dynamic Map* (LDM); there are starting points for this work. IEEE 802.11p allows wireless access in vehicular environments, called V2X communications, which enables messaging between vehicles and the infrastructure. The messages are broadcast every 100ms by traffic participants, i.e., vehicles and roadside ITS stations, to update other participants about their current states [2]. The main standardized message types are [10–12]:

- *CAMs* (Cooperative Awareness Messages) provide high frequency status updates of a vehicle’s position, speed, and might include vehicle type, model, and turn signals;
- *MAPs* (Map Data Messages) describe the detailed topology of an intersection, including its lanes, their connections, and assigned traffic light (TL) signal groups;
- *SPaTs* (Signal Phase and Timing Messages) contain the projected TL signal phases (e.g., green, yellow, and red) for each lane;
- *DENMs* (Decentralized Environmental Notification Messages) informing whether specific events like road works or a traffic jam occur in a designated area.
- *CPMs* (Collective Perception Messages) are aimed to complement CAMs with information from the surroundings that is collected by the LIDAR and RADAR sensors of a vehicle or by an ITS station and partially shared with other traffic participants. For instance, a temporary obstacle such as a parked car can be detected by a vehicle and forwarded to the other participants;

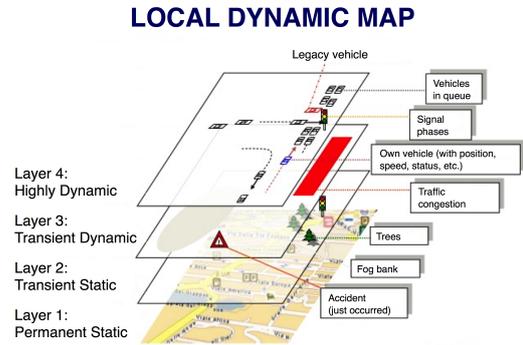


Figure 1. The four Layers of a LDM [2]

- *PAMs* (Precise Awareness Messages) are lightweight correction messages, which are used to correct a vehicle’s positions regarding a fixed “virtual anchor”, since in urban areas the GPS position could be imprecise.

2.1. Local Dynamic Map

The V2X technology does not yet consider the integration of the different types of messages. As a comprehensive integration effort, the EU SAFESPOT project [2] introduced the concept of an LDM, which acts as an integration platform to combine static geographic information system (GIS) maps, with dynamic environmental objects (e.g., vehicles or pedestrians) [13, 14]. The integration is motivated by advanced safety applications, which need an “overall” understanding of a traffic environment. The LDM consists of the following four layers (see Figure 1):

1. *Permanent static*: the first layer contains static information obtained from GIS maps and includes roads, intersections, and points-of-interest (POIs);
2. *Transient static*: the second layer extends the static map by detailed local traffic informations such as fixed ITS stations, landmarks, and intersection features like lanes;
3. *Transient dynamic*: the third layer contains temporary regional information like weather, road or traffic conditions (e.g., traffic jams), and traffic light signal phases;
4. *Highly dynamic*: the fourth layer contains dynamic information of road users taken from V2X messages or in-vehicle sensors like the GPS module.

Recent research by Netten et al. [15], and Shimada et al. [16] suggested that an LDM can be built on top

of a spatial relational RDBMS enhanced with streaming capabilities. Netten et al. recognize that an LDM should be represented by a world model, world objects, and data sinks on the streamed input [15]. However, an elaborate domain model captured by an LDM ontology and extended query processing or rule evaluation methods over spatial data streams were still missing in the current approaches. An ontology-based LDM has advantages regarding the maintainability and understandability of the model, since dependencies between the concepts are clearly defined and easy extendable without altering the underlying database (DB).

2.2. Ontology-based LDM

With the support of Siemens and AIST domain experts, we defined two ontologies capturing the main elements of an LDM. In our second modelling, we took the recommendations of the interviewed experts (see Section 4) into account and extended the initial ontology with the “standard” vocabulary of the Sensor-Observation-Sampling-Actuator (SOSA) ontology [9]. This includes concepts such as *Observation*, *Sensor*, and *Platform*, which are combined with the LDM specific vocabulary such as the four levels of the LDM. The resulting ontologies are partially shown in Figure 2a and Figure 2b, where asterisks indicate elements that are commented in detail below.²

Lightweight LDM Ontology. Our initial ontology is intended to make querying ITS streams simple, but still capturing the main concepts of a LDM. It follows a layered approach starting with a separation between the top concepts as follows:

- *V2XFeature* is the representation of V2X objects, such as the details of an intersection topology including lanes denoted as *V2XLane* and traffic lights denoted as *V2XSignalGroup*;
- *GeoFeature* represents the GIS aspects of the LDM including POIs, areas like parks, and road networks with *Geometry* as the geometrical representation of them;
- *LDMLayer* is the representation of the four layers of an LDM, where each feature can be assigned to one layer by the role *isOnLDMLayer*;

²available at <http://www.kr.tuwien.ac.at/research/projects/loctrafflog/LocalDynamicMapITS-v0.5-Lite.owl> and <http://www.kr.tuwien.ac.at/research/projects/loctrafflog/LocalDynamicMapSOSA-v0.1-Lite.owl>

- *Actor* is the concept that includes persons, vehicles, as well as roadside ITS stations, which are autonomous agents and are the main generator of streamed data;
- *Event* captures prototypical events that happen in the ITS domain. An important subclass of *Event* is *Hazard* that captures the different types of dangers, e.g., accidents that might occur;
- *CategoricalValues* specify the different categories such as signal phases, or vehicle roles used in the domain.

We also have defined “domain specific” roles and attributes such as *isPartOf*, *connected*, *intersects*, which deal with the different aspects of intersection and road topologies. In particular, the roles *spatialRelation*, resp., *temporalRelation* are of interest, since we have introduced several sub-roles that define possible spatial, resp. temporal, relations between features. Examples are *intersects*, resp., *during*, which are needed for querying vehicles that are passing a specific intersection during a specific period.

Extended SOSA Ontology. In the following, we give an outline of the main concepts of the extension of our initial ontology, aligning it with the SOSA ontology:

- *V2XFeature* is the representation of V2X objects and defined as before;
- *LDMLayer* is the representation of the four layers of an LDM and defined as before;
- *Observation* is taken from the SOSA ontology and is stated as “the act of carrying out a procedure to estimate a value of a *Feature* involving a *Sensor* and yielding a result” [9], which we call *Value*;
- *Value* is the result of an *Observation* and can either be a *Literal*, a *Geometry*, or a *CategoricalValue* from different static categories such as signal phases;
- *Sensor* defines (mostly) physical devices that are the implementation of the observation procedure by acting on changes in the environment, and is the main generator of streams of observations and the assigned values;
- *Platform* is the host or mounting of one or several sensors. A platform can be represented by mobile phones, vehicles, or roadside ITS stations.
- *Event* is a sub-concept and of *Feature* (as suggested in [9]) and captures prototypical events that occur in the ITS domain.

We have incorporated the “standard” roles and attributes of the SOSA ontology. This includes the central relations between *Platform*, *Sensor*, *Observation*, and *Value* using the roles *hosts*, *madeObservation*, and *hasResult* for connecting them to each other.

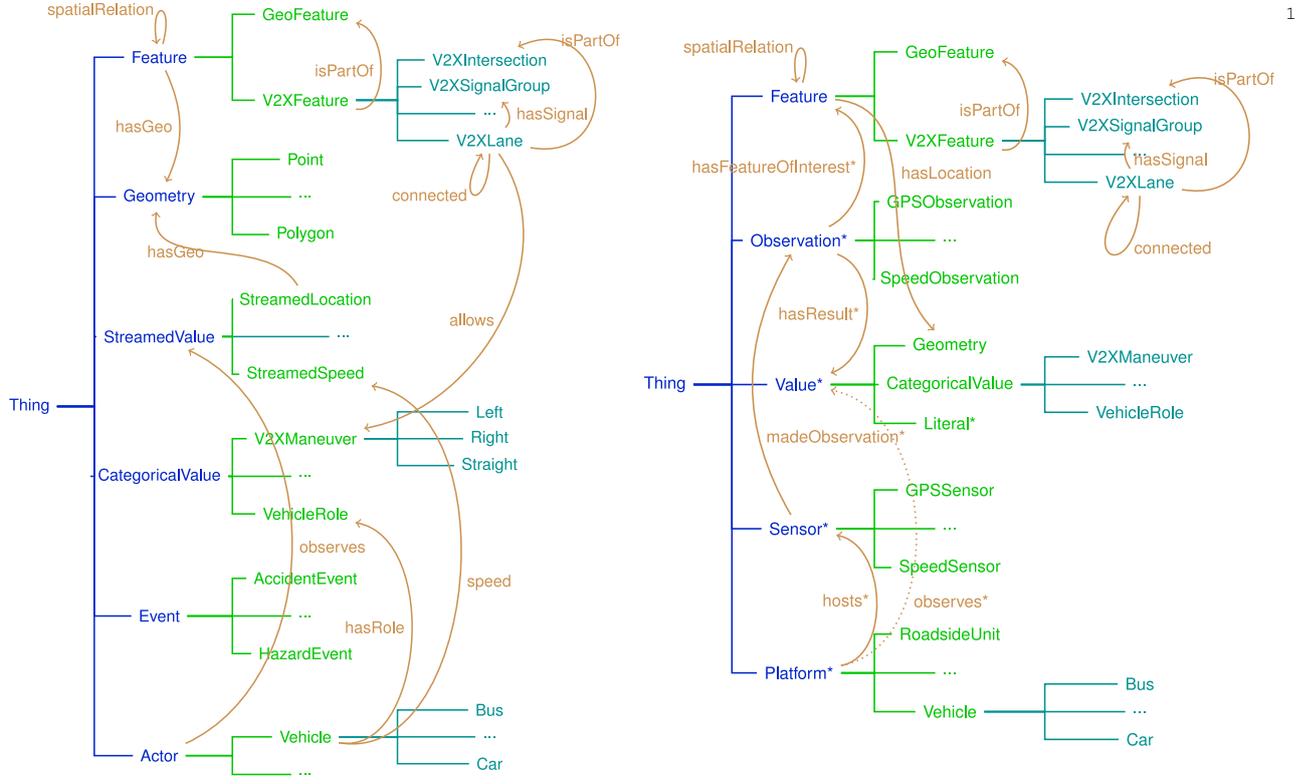


Figure 2. (a) The lightweight and (b) SOSA-based LDM Ontology (partial renderings)

The previously defined “domain specific” roles and attributes such as *isPartOf*, *connected*, *intersects*, as well as *spatialRelation* and *temporalRelation* are defined as before. An important help for simpler query formulation are the roles *speed*, *pos*, *acceleration* and *hasState*. These roles are sub-roles of the lightweight ontology role *ldm:observes* (shown by the dotted arrow in Figure 2b) and represent role-chains to simplify the more complex relations between the platform, e.g., vehicles, its sensors, and the measured values of observations, e.g., measured speed. Note that the lightweight ontology role *ldm:observes* and the SOSA role *sosa:observes* have different meanings, as the latter describes the relation between a sensor and observable properties. The top role *ldm:observes* can be defined by the following role chains:

$$\begin{aligned} \text{hosts} \circ \text{madeObservation} &\sqsubseteq \text{hostsObs}, \text{ and} \\ \text{hostsObs} \circ \text{hasResult} &\sqsubseteq \text{observes}, \end{aligned}$$

where sub-roles specify the particular range that is observed; hence we distinguish between *hasResult* for composed values and *hasSimpleResult* for atomic values represented by literals. However, adding role chains

to DL-Lite_A [4] would change the expressiveness of the language. We thus leave role chaining out, but may allow *query templates* that replace role chains.

Example 2.1. We illustrate according to Figure 2b a possible query template for *observes*, which is defined as follows:

$$\text{observes}(x, y) ::= \text{hosts}(x, u) \wedge \text{madeObservation}(u, v) \wedge \text{hasResult}(v, y),$$

where $::=$ defines that the atom *observes* in any query is replaced by the atoms on the right side of the argument.

OWL 2 QL. Both of the above LDM ontologies are represented in DL-Lite_A [4], which is the logical underpinning for the W3C standard OWL 2 QL. Apart from the restriction to DL-Lite_A, our methods are ontology-agnostic; hence the mentioned ontologies, but also other mobility ontologies could be used. In the rest of the paper, we will use the lightweight ontology for the examples, the experiments, and the expert evaluation. The queries and the experiments are transferable to the SOSA-based ontology, since the restriction to DL-Lite_A is retained. Furthermore, the expert evaluation induced the development of the SOSA-based ontology.

2.3. Spatial-Stream Query Answering

The OQA component is central part regarding the usage of a semantically enhanced LDM, since it allows us to access the streamed data in the LDM.

Example 2.2. The following query detects red-light violations on intersections by searching for vehicles (in y) with an *aggregated* trajectory and speed above 30km/h in a 8 secs window, projecting 3 secs into the future (represented as a negative distance), which move on lanes (in x) *during* the time intervals, where the signal phases of these lanes will turn to “Stop”, i.e., red, in a 10 secs window with a 5 secs look into the future to capture the current/next changes in the signal phases:

$$q_1(x, y) : \text{LaneIn}(x) \wedge \text{hasLoc}(x, u) \wedge \text{intersects}(u, p) \wedge \\ \text{Vehicle}(y) \wedge \text{pos}(y, p@i_p)[\text{traject_line}, 5s, -3s] \\ \wedge \text{speed}(y, v)[\text{mov_avg}, 5s, -3s] \wedge (v > 30) \wedge \\ \text{during}(p@i_p, s@i_s) \wedge \text{isManagedBy}(x, z) \wedge \\ \text{SignalGroup}(z) \wedge \text{hasState}(z, s@i_s)[\text{last}, 5s, -5s] \\ \wedge (s = 'Stop')$$

Query q_1 exhibits the different dimensions that need to be combined:

- (a) $\text{Vehicle}(y)$, $\text{isManagedBy}(x, z)$, $\text{SignalGroup}(z)$ and $\text{LaneIn}(x)$ are ontology atoms, where isManagedBy assigns traffic lights z to incoming lanes x . These atoms have to be unfolded with respect to the concept/role hierarchies of the LDM ontology;
- (b) $\text{intersects}(u, v)$ and $\text{hasLoc}(x, u)$ are spatial atoms, where the first checks spatial intersection and the second returns the object geometries;
- (c) $\text{speed}(y, v)[\text{traject_line}, 5s, -3s]$ and $\text{pos}(y, p@i_p)[\text{mov_avg}, 5s, -3s]$ define window operators that aggregate and *predict* the moving average of speed and positions (represented by a path) of the vehicle y over the streams speed and pos , respectively, and $\text{hasState}(z, s@i_s)[\text{last}, 5s, -5s]$ returns the traffic lights that have their last phase on “Stop”;
- (d) the relation $\text{during}(p@i_p, s@i_s)$ checks if “ p happens during s ”, where p is all the occurrences of trajectories on the set of time intervals of i_p , and s are the traffic light phases that are on “Stop” in the set of time intervals i_s , where i_p and i_s are derived from the trajectory aggregations and the phase duration of the traffic lights, respectively.

Note that we added for readability the implicit definitions of $p@i_p$ and $s@i_s$ to the original query syntax, where we would use solely p and s . The implicit variables i_p and i_s represents the time interval annotated to the aggregated values of variable p and s . For instance,

the average speed of a vehicle in the interval $[1, 6]$ as $\langle \text{car}, \text{speed}, 20 \rangle @ [1, 6]$. A user formulating a query can ignore this notation and use p and s .

3. Development of C-ITS Scenarios

In this section, we present three application scenarios that are used to define requirements and features split into three complexity levels. On the infrastructure side, we have C-ITS (roadside) stations that receive nearby V2X messages and send messages to inform other participants on their current state, i.e., the traffic light phases. Other participants such as vehicles share their states such as their current speed, acceleration, and position. On the vehicle side, ADAS perceive driving environments and make safe driving decisions to improve safety of autonomous vehicles. The ADAS use sensors such as Lidar/Radar or cameras, and process the sensor data to avoid accidents by detecting pedestrians, vehicles, or other obstacles [17]. The sensor data can be linked to our ontology-based LDM and enables the system to represent the driving environments.

3.1. Scenario Description

First, we give an overview of the three scenarios.

S1: Traffic Statistics. The focus of this scenario is on the collection of statistical data that concerns stops, throughput, traffic distribution, or types of participants by aggregating the streaming data on specific intersections. Regarding this scenario, we have identified the following use cases and related challenges:

1. *Object level:* for a single vehicle or station, the average speed, acceleration, number of stops, as well as the temperature could be collected;
2. *Road/Intersection level:* on this level, besides calculating a summary of road/lane level indicators such as average throughput, waiting time, the amount of stops, also matrices regarding transfers (e.g. how many cars head straight on), modality, and type mix, (e.g. which vehicle classes are present) could be determined;
3. *Network level:* on the network level, intersections are represented by nodes connected by roads. We could collect statistical summaries of indicators on intersections. For instance, estimating the transfer times and traffic flow between intersections.

S2: Hazardous Events Detection. An important C-ITS application is *road safety* [2], where a reliable event detection is central to find unexpected, hazardous events. This is a more challenging case, since it requires the combination of the topology, vehicle maneuvers, and temporal relations that might be evaluated over longer and shorter periods. We identified the following events as possibly hazardous:

1. *Simple vehicle maneuvers*: the following maneuvers are relevant for this case and are directly extractable from trajectories: (1) quick slow down/speed up; (2) drive straight on, turn left, turn right; (3) stop, unload, park;
2. *Complex vehicle maneuvers*: the aim is to detect lane changes, overtakes, and u-turns, which are complex maneuvers, composed of simpler maneuvers;
3. *Red-light violation*: in Example 2.2, a description of the detection of red-light violations is given;
4. *Vehicle breakdown/accident*: this event is based on the stop maneuvers, where we identify vehicles that are not moving and are inside a dangerous area of an intersection. This case can be extended to several vehicles;
5. *Traffic congestions*: this is a more complex event, where short and long term observations must be combined. Queuing cars could indicate a congestion and be detected by checking the stop maneuvers of several vehicles that are behind each other, but not stopped by a longer red light phase.

S3: ADAS. ADAS features are an important step towards fully automated driving by enabling the vehicle to take control of speed or breaking, where drivers still have the “full” control over the vehicle. The following challenges come for ADAS:

1. *Self monitoring*: self-monitoring is a central requirement of ADAS, where intelligent speed adaptation is an important feature to improve roadway safety;
2. *Obstructed view*: this concerns dangerous situations where a vehicle might collide with another vehicle, since they have no visual contact due to an obscured view (e.g., buildings). The crossing of the predicted trajectories of two vehicles has to be verified to provide a simple collision detection.
3. *Traffic rules*: the embedding of traffic rules like checking of traffic rules such as right-of-way rules could become an important requirement for autonomous driving.

3.2. Features for Spatial-Stream QA

The eight “standard” requirements for stream processing identified by [18], namely volume, velocity, variety, incompleteness, noise, timely fashion, fine-grained information access, complex domain models, and user intention, as well as the three entailment levels identified by [19] for stream reasoning systems, namely stream-, window-, and graph-level entailment, are not discussed here; they should hold for C-ITS stream systems as well. Besides the generic features *F1*, *F2*, *F3*, and *F9*, we also focus on domain specific features that are mapped to requirements crucial for enabling the above scenarios. For this, we distinguish for each feature three levels of fulfillment: *basic* (L1), *enhanced* (L2), and *advanced* (L3). We have identified the following feature sets:

F1 - Time model: possible time models are *point-based* (L1), and *interval-based* (L2), where L1 is the “simplest” representation. Also belonging to L2, on point-based data, applying aggregations can be represented by intervals based on point-based data items. If we apply an interval-based model, temporal relations (L3) such as Allen’s Time Interval Algebra [20] with operators like *before* that can be used for querying and inference.

F2 - Process paradigm: queries that are processed in a pull-based (L1) manner should be the baseline. Push-based processing (L2) in particular with sliding windows is already more challenging. If we allow a combined (L3) processing, we could treat high velocity, resp. low velocity, atoms by push-based, resp. as pull-based queries.

F3 - Query features: we consider “basic” query features that are include in a standard query language such as SPARQL or SQL. These features are part of L1 and include selection, projection, join, and filter, as well as time- and triple-based window operators. Combining the basic features by query nesting, unions of CQs, and allowing inter-stream joins belongs to L2.

F4 - Numerical aggregations: aggregations can be “simple” functions such as *sum* or *average* on either a set or multiset (bag) of data items (L1). L2 extends L1 by statistical functions such applying the *mean*, *median*, or *mode*, as well as *standard deviation*, *variance*, and *range*. Note that the aggregations are mainly over multisets, since we often have data items of different objects in a single stream.

F5 - Spatial aggregations: a wide range of spatial aggregations can be applied to geometric objects like points and lines (L1) and the aggregation functions need

to take the peculiarities of geometries into account, e.g., convex vs. concave objects. L2 adds the computation of spatial relations using a simple point-set model [21] or the more detailed *9-Intersection model* (L2) based on the aggregated objects. Smoothing and simplification of complex objects could also be included, which leads to L3.

F6 - Numerical predictions: predictions allow the generation of unknown data items projecting from the past into the future. Several prediction functions such as moving average (L1) or exponential smoothing (L2) regression should be available. Depending on the task, also more complex machine learning methods could be envisioned (L3).

F7 - Trajectory predictions: we predict a vehicle's movement, by linearly projecting the trajectory into the future (L1). More accurate results could be achieved by (1) a "point-to-curve" aggregation, and (2) calculating possible paths using a road graph (L2), and the usage of machine learning for trajectory predictions (L3).

F8 - Spatial matching: basic spatial matching is the extraction of specific features such as angles from the objects (L1). Advanced features include the matching of complex geometries such as road graphs (L2).

F9 - Advanced reasoning: include the different forms of reasoning ranging from rule-based to ontological reasoning that reaches beyond the expressivity of query answering in OWL 2 QL [22]. The underlying language can include "simple" implications as $b(x, y) \wedge c(y, z) \rightarrow a(x, z)$, but also "advanced" features such as recursion as in OWL 2 RL [22], and all combined with strong constraints and negation as failure (NAF) as present in Answer Set Programming (ASP) [23]. Note that in a limited form simple implications are already feasible in OWL 2 RL by allowing role-chaining such as $b \circ c \sqsubseteq a$.

3.3. Requirements

In Table 1, we show the requirements that are derived by analyzing each scenario and use case with respect to the needed features. The requirements build the base line for the implementation and a later experimental assessment. In case of single features, we only distinguish between L1 to L3 for required, blank for not required, and "P" for possibly required. For instance, in S2.2 for F1, a point-based time model (L1) suffices for detecting left/right turns; however, if we want to detect u-turns, an interval-based time model in combination with temporal relations (L2) will be needed. Further-

more, push-based queries are desired for swift reaction on changes.

The requirements of advanced reasoning (F9) start with QA using OWL2 QL as the baseline. This allows one to access the streams and to enrich them with a domain model. However, in some use cases the expressive power of QA is too limited, since more "advanced" features are needed. For instance, if a use case requires to take the road network and reachability into account, OWL2 RL has the expressive power for capturing it. If a use case requires to model traffic rules and traffic regulations, one needs, besides reachability, also constraints and NAF (as provided by ASP). For instance, a warning has to be generated if two traffic lights of crossing lanes are green at the same time.

In our previous work of [3], we already cover level L1 for the features *F1* to *F5*, but aim in this work to introduce new features such as time intervals, temporal relations, and unions of CQs. *F6*, *F7*, and *F8* are entirely new features.

4. Expert Interviews

We conducted four long interviews with ITS experts, who play different roles in the field. The first two experts work in industry and the other two experts come from academia. The interviews were conducted as guideline-based expert interview [24]. Following the guidelines of [24], we left the interviewee the freedom to choose the topic he/she prefers to discuss. However, we had prepared a set of questions, which we asked if the interviewee heads into the direction of a particular topic, e.g., the query language.

Importantly, guideline-based expert interviews are a standard practice of qualitative research in social sciences [25]. Maccoby et al. [26] define an interview as an "interchange in which one person attempts to elicit information or expressions of opinion or belief from another person or persons." If complex topics are to be investigated, long interviews can be a better choice compared to other (quantitative) methods in filling a knowledge gap; they allow to obtain a more profound insight into the positions of the experts [27]. Other approaches such as questionnaire-based interviews are more appropriate if many interviewees are participating and a quantitative but not qualitative analysis has to be conducted. Furthermore, we encountered limited availability of experts (for time and business reasons) that have also a good understanding of the methods and techniques surroundings.

Use Case	F1	F2	F3	F4	F5	F6	F7	F8	F9
<i>S1.1</i> (Object statistics)	L1	L1	L1	L1	L1				OWL2 QL
<i>S1.2</i> (Road/Intersection statistics)	L2	L1	L2	L2	L2	L1	L1		OWL2 QL
<i>S1.3</i> (Network statistics)	L2	L1	L2	L2	L2	L2	L1	P	OWL2 RL
<i>S2.1</i> (Simple maneuvers)	L1	L1	L1	L1	L2	P	P		OWL2 QL
<i>S2.2</i> (Complex maneuvers)	L2	L2	L2	L1	L2	L1	L1	L1	OWL2 QL
<i>S2.3</i> (Red-light violation)	L2	L2	L2	L1	L2	L1	L1	L1	OWL2 QL
<i>S2.4</i> (Vehicle breakdown)	L2	L2	L2	L1	L2	L1		P	OWL2 QL
<i>S2.5</i> (Traffic congestion)	L2	L3	L2	L2	L3	L2		P	OWL2 RL / ASP
<i>S3.1</i> (Self monitoring)	L1	L2	L1	L1	L1	P	P		OWL2 QL
<i>S3.2</i> (Obstructed view)	L1	L2	L2	L1	L2	L1	L1	L1	OWL2 QL
<i>S3.3</i> (Traffic rules)	L2	L2	L2	L2	L3	L1	L1	L1	OWL2 RL / ASP

Table 1

Requirement Matrix (Level L1/L2/L3 implies required, blank is not required, P is possibly required)

The main goal of the interviews is to answer the following general question:

“How suitable is our spatial-stream QA for real-world ITS applications?”

The above question can be split into more specific parts:

- *Q1*: “What are important technologies/developments for future ITS systems?”
- *Q2*: “How well is the LDM suited for the integration of vehicle/traffic control sensor data such as V2X messages?”
- *Q3*: “Is the presented ontology and the approach of ontology-based data access suitable to realise an LDM?”
- *Q4*: “We present three scenarios with several use cases, how relevant are they? Are different use cases needed?”
- *Q5*: “We present different query features, such as trajectory predictions, how important are these features in your opinion? Should they be extended?”
- *Q6*: “We present you an example for our queries for detecting red-light violations; how well is this query comprehensible for you? Do you believe another query language is better suited for this?”

In the remainder of this section, we give a summary of the interviews that were conducted with each expert. We do not transcribe the full interview, which range from 30 minutes to 1.5 hours, but only present a summarized version of each interview. In the Conclusion, we shall come back to recommendations and suggestions from the experts consulted for future work.

Expert 1. The first expert is an initiator of the V2X technology development in Europe. He was responsible, in a large ITS company, for the standardization of V2X messages in different committees of standardization organizations such as ETSI, ISO, and SAE. Additionally, he was participating in several large research projects such as *DRIVE C2X*.³

First, the expert pointed out that there are not only CAM, SPaT, MAP, and DENM messages, but also messages designed for public transport signal requests called SRM/SSN, as well as position correction messages for autonomous vehicles. The latter are important since in inner-cities, the exact position in combination with high-resolution maps are crucial for safety. An important task arising from this challenge is the continuous matching of the vehicle position to the high-resolution map. He stated that autonomous vehicles, even with advanced sensors, will in the near future *not* be able to drive autonomous and safely in inner-cities with complex intersections, in particular when the weather is unpredictable. He then said that messages like the CAM or DENM need to be sent near real-time to the surroundings, hence the existing 4G and upcoming 5G mobile standards are not suitable for this purpose, and the standard IEEE 802.11p (also called ITS-G5) is better suited, since it is based on WLAN technology and already available for low latency (in ms) communication. However, existing WLAN technology can not be used, since these protocols require sessions and authentication with a base station, which

³<http://www.drive-c2x.eu/project>

1 would cause a long delay for fast moving vehicles. He
 2 gave more details on DENM messages, which inform
 3 the surrounding on dangerous events, and noted that
 4 DENM messages are fired based on triggering conditions
 5 that have specific probabilities assigned. The triggering
 6 conditions are defined by the standardization bodies,
 7 and vehicle manufacturer have to implement them ac-
 8 cordingly. Furthermore, he highlighted that more de-
 9 velopment is needed to protect vulnerable road users,
 10 which requires the integration of bluetooth-based com-
 11 munication such as ZigBee, and also the sharing of
 12 data with infrastructure sensors such as mounted Li-
 13 dar/Radar stations.

14 Second, we discussed the LDM, and he identified the
 15 LDM as an integration platform for the V2X messages,
 16 where each vehicle has its own proprietary representa-
 17 tion of the LDM, since the ETSI standard defines only
 18 the interfaces to access it, but not its internal structure.
 19 He doubts though that a common definition of the LDM
 20 is needed, since every manufacturer should implement
 21 it on their own, where only the dynamic elements could
 22 be encoded as a snapshot (in a standardized data model)
 23 and exchanged with the surrounding. One discussed
 24 use case is the exchange of snapshots, where legacy
 25 vehicles and/or vulnerable road users are on the road,
 26 hence only fixed Lidar/Radar stations could detect them
 27 sending their snapshots to other V2X-based vehicles,
 28 since the quality of fixed Lidar/Radar stations should
 29 be more accurate than of the moving ones. Currently
 30 several companies develop Collective Perception Mes-
 31 sages (CPMs) [28], which figure as an exchange of
 32 the mentioned Lidar/Radar data. A future development
 33 could foster the exchange of sensor data by vehicles
 34 using the CPMs.

35 Third, we discussed the usage of ontologies and OQA
 36 for realizing the LDM, which he regarded initially scepti-
 37 cal since scalability might be an issue. After explain-
 38 ing the intention of OQA this skepticism was in parts
 39 redressed. Regarding the modelling of the ontology,
 40 he stated that users do not care how the (ontological)
 41 model is shaped, more important is the query language,
 42 since users will directly be confronted with it. He iden-
 43 tified that the query engine and the stream DB could be
 44 migrated and deployed on roadside ITS stations, where
 45 the user could cover custom use cases by writing and
 46 applying queries.

47 Fourth, he reviewed the query features, where he
 48 explained that an interval-based time model and push-
 49 based processing is favourable since more information
 50 can be extracted, and in C-ITS applications due to low
 51 latency push-based queries should be supported, which

1 could include some buffering techniques so not every
 2 change will be computed. He pointed out that “native”
 3 query languages such as SQL might be too complicated
 4 for writing queries, and a simplified, object-oriented
 5 representation (similar to CQ) could be favourable.
 6 While discussing the example query, we observed that
 7 he understands and reads queries as rules. He noted also
 8 that, if the query get complex as in the second scenario,
 9 rules might be easier usable to express problems, since
 10 they are more capturing “the way people think”. He
 11 also stated that the aggregation feature can be also used
 12 to guarantee privacy, since on aggregated values single
 13 vehicles are not distinguishable anymore. Furthermore
 14 aggregations can be used to calculate journey times in
 15 a road network. Finally, he agreed that predictions are a
 16 crucial feature for accident prevention, hence it should
 17 be more elaborated.

18 **Expert 2.** The second interviewed expert is the head of
 19 a traffic engineering department in a large ITS company
 20 and is responsible for managing R&D projects.

21 The expert described that historically traffic manage-
 22 ment was a closed, autarkic system, where traffic data
 23 was collected in a slow process using radar, road loops,
 24 and cameras. She identified V2X as an important step
 25 towards collecting real-time data on traffic and vehicles,
 26 where besides the mentioned CAM, SPaT, and DENM
 27 messages, Collective Perception Messages (CPMs) will
 28 play an important role, since they allow one to exchange
 29 locally perceived objects by a vehicles sensors, and ex-
 30 change the object data with other V2X vehicles. She
 31 believed that the LDM could be used in combination
 32 with CPMs, where a CPM could be aligned with a ve-
 33 hicle’s own LDM. However, she saw no immediate de-
 34 mand for the use of an ontology-based extension of a
 35 LDM and the use of spatial-stream queries to access the
 36 (streaming) data, since they use their own tools and lan-
 37 guages for processing V2X messages. She commented
 38 though, that an ontology-enhanced LDM could be used
 39 as a data integration platform, which could be used for
 40 future data analytics tasks.

42 **Expert 3.** The interviewed expert is an associate pro-
 43 fessor in an European university, where he works in
 44 the fields where reasoning and learning over streams is
 45 applied to safe autonomous systems including robots,
 46 boats, and drones.

47 First, we discussed the LDM, and he noted that in
 48 robotics similar techniques have been used for long,
 49 but with the additional challenge that these dynamic
 50 maps have to be built on-the-fly and cooperatively. An
 51 interesting challenge arises if different agents have a

1 local representation of an LDM, and for coordination
2 a global perspective has to be constructed based upon
3 them. He identified another important task in this con-
4 text, which is the matching of identical entities (in the
5 sense of a physical grounding of the same entity) de-
6 tected by different agents/sensors, which they solved
7 by using bridge-rules.

8 Second, we discussed the LDM ontology, where he
9 identified that meta-data is important, and a possibility
10 of capturing uncertainty should be part of an ontology.
11 Uncertainty introduces the additional challenge that
12 measured observations, e.g., speed, are not crisp any-
13 more, but are inside confidence intervals were the ob-
14 servation holds. He also stated that by using an average
15 or most-likely values instead of intervals, we encounter
16 a loss of information. This uncertainty also can occur
17 in the classification of objects, wherein an instance of a
18 vehicle is detected, but it turns out to be a bike. After
19 discussing our query rewriting technique, he described
20 their approach in robotics, which aims at matching the
21 data to the query (or formula), instead of rewriting
22 queries and checking for instances in the DB. The data
23 is processed/transformed until it matches the queries.
24 If new streamed data appears, they continuously try to
25 match the streamed data to the queries.
26

27 Third, he reviewed the query and could understand
28 most under certain assumptions (i.e., that aggregations
29 are grouped), and he believed that including predictions
30 are an important extension, which they consider for
31 their approach as well. Since predictions need an under-
32 lying model, he sees that it is important to re-evaluate
33 and monitor the predictions to detect concept drifts. He
34 also commented on the semantics of query language
35 and windows, where it is important to define when to
36 forget data in the stream, and how long a data item
37 holds into the future. Regarding the language features,
38 he noted that intervals are good way to also express
39 temporal uncertainty. Further it would be crucial to re-
40 spect the underlying time model, which could be dense
41 continuous or discrete. If someone assumes the finite
42 number of changes assumption, it is possible to map
43 the observation from the continuous into the discrete
44 space. Finally, he believed that numerical aggregations
45 and predictions could be enriched by different types of
46 filters, whereby he mentioned that Kalman or particle
47 filters [29] are often used in robotics.
48

49 **Expert 4.** This interviewed expert is a senior researcher
50 in the fields of stream reasoning, graph DBs, and au-
51 tonomous driving in an European university, and has

1 also been involved in the development of stream rea-
2 soning tools.

3 First, we discussed the LDM, which he saw as a
4 suitable approach to integrate map and streaming data.
5 He pointed out though that classical DB technology
6 as B+ trees [30] are not suited for spatial-stream data,
7 but advances in the field of moving object DBs [31]
8 result in efficient query languages and spatio-temporal
9 indexing methods. Furthermore, the LDM does not
10 consider a 3D representation of maps, which is crucial
11 if the map is used in the context of autonomous driving,
12 since detected objects of a image recognition step (e.g.,
13 building), need to be anchored in a 3D map.
14

15 Second, he evaluated the LDM ontology, where he
16 agreed on the modelling regarding map features, but
17 criticized that the LDM ontology misses to concept of
18 sensors, resp., observations, which is crucial in C-ITS
19 applications. He noted that elements such as *Stimulus*,
20 *Sensor*, and *Observation* pattern of the Semantic Sensor
21 Network (SSN) ontology [32] could be incorporated in
22 the LDM ontology.
23

24 Third, he reviewed the query language, and easily
25 understood the presented query (for red-light viola-
26 tions). He noted that the definition of the ontology and
27 spatial atoms are clear, but the stream atoms need a
28 formal grammar to capture an atoms parameterization
29 such as *speed[avg, 10s]*. Subsequently, the definition
30 of *speed[last, 5, -5s]*, i.e., the last element in a 10 secs
31 window projecting 5 secs to the future, was not clear
32 to him. He further suggested that using a SPARQL-
33 dialect, either C-SPARQL [7] or CQELS [8] could be
34 suitable languages. Evaluating the language features,
35 he believed that all presented features are important,
36 but more focus in terms of research and development
37 should be given numerical- (F6) and trajectory predic-
38 tion (F7), as he saw them as crucial for autonomous
39 driving, since these are the main techniques used, for
40 instance in object recognition and motion planning. He
41 also believed that the combination of (stream) reason-
42 ing and machine learning is in its early stages and needs
43 more attention from the community. He further sug-
44 gested that the aggregates (F4) could be extended with
45 top-*k* aggregate functions [33].
46

47 Fourth, he criticized that the scenario S3 is too
48 generic, and should be narrowed to more specific tasks
49 in the domain of autonomous driving; he suggested to
50 include motion planing as a replacement for the generic
51 scenario.

5. Approach for Spatial-Stream Query Answering

First, we introduce the data model and the definition of a spatial-stream knowledge base, which leads to our main focus of spatial-stream queries. Then, we introduce the “standard” query rewriting and extend it with temporal relations. Finally, we describe how the queries are evaluated putting the focus on stream aggregation and predictions. We start from previous work in [3], which introduced spatial ontology-mediated query answering over Mobility Streams using DL-Lite_A [4]. We focus on pull-based queries that are evaluated at one single time point called the *query time* \mathbb{T}_i .

5.1. Data Model and Knowledge Base

Our data model is *point-based* and captures the *valid time*, extracted from the V2X messages, saying that some *data item* is valid at that time point. Importantly, while evaluating a query, the model can change (temporarily) to an *interval-based* model that results from the window and aggregation functions. To capture streaming data, we introduce the *timeline* \mathbb{T} , which is a *closed interval* of (\mathbb{N}, \leq) . A *data stream* is a triple $D = (\mathbb{T}, \nu, P)$, where \mathbb{T} is a timeline, $\nu : \mathbb{T} \rightarrow \langle \mathcal{F}, \mathcal{S}_{\mathcal{F}} \rangle$ is a function that assigns to each element of \mathbb{T} (called a *timestamp*) data items of $\langle \mathcal{F}, \mathcal{S}_{\mathcal{F}} \rangle$, where \mathcal{F} (resp. $\mathcal{S}_{\mathcal{F}}$) is a *stream* (resp. *spatial-stream*) *DB*; the integer P is called *pulse* defining the general interval of consecutive data items on the timeline (cf. [34]), which naturally induces a stream of data items. We always have a *main pulse* with a fixed interval length that defines the highest granularity of the validity of data points. Larger pulses for streams with lower frequency can be defined, which can be exploited to perform optimizations such as caching. The pulse also aligns the data items that arrive asynchronously in the DB to the timeline. A spatial or spatial-stream DB is a Data Stream Management System that support spatial and geospatial objects and operators (e.g., ODYSSEUS, PIPELINEDB or SQLSTREAM).⁴

Example 5.1. For the timeline $\mathbb{T} = [0, 100]$, we have the stream $F_{CAM} = (\mathbb{T}, \nu, 1)$ of vehicle positions and speed at the assigned time points for the individuals c_1 , c_2 and b_1 :

$$\begin{aligned} \nu(0) &= \{speed(c_1, 30), pos(c_1, (5, 5)), speed(c_2, 10), \\ &\quad pos(c_2, (4, 4)), speed(b_1, 10), \dots\}, \\ \nu(1) &= \{speed(c_1, 29), pos(c_1, (6, 5)), speed(c_2, 0), \end{aligned}$$

⁴ <http://odysseus.informatik.uni-oldenburg.de/>, <https://www.pipelinedb.com/>, and <https://sqlstream.com/>

$$pos(c_2, (5, 4)), speed(b_1, 5), \dots\}, \dots$$

A “slower” stream $F_{SPaT} = (\mathbb{T}, \nu, 5)$ captures the next signal state of a traffic light: $\nu(0) = \{hasState(t_1, Stop)\}$ and $\nu(5) = \{hasState(t_1, Go)\}$. The static ABox contains the assertions $Car(c_1)$, $Car(c_2)$, $Bike(b_1)$, and $SignalGroup(t_1)$. A different “annotated” representation by applying the function ν on F_{CAM} yields $\{speed(c_1, 30)@t0, \dots, speed(c_1, 29)@t1\}$, which is better suited for an interval-based time model.

We consider a vocabulary of individual names Γ_I , domain values Γ_V (e.g., \mathbb{N}), and spatial objects Γ_S . Given atomic concepts A , atomic roles P , and atomic attributes U_C , we define (a) *basic concepts* B , *basic roles* Q , and *basic value-domains* E (attribute ranges); (b) *complex concepts* C , *complex role expressions* R , and *complex attributes* V_C ; and (c) value-domain expressions D :

$$\begin{aligned} Q &::= P \mid P^- \\ B &::= A \mid \exists Q \mid \delta(U_C) \\ E &::= \rho(U_C) \\ C &::= \top_C \mid B \mid \neg B \mid \exists Q.C' \\ D &::= \top_D \mid D_1 \mid \dots \mid D_n \\ R &::= Q \mid \neg Q \\ V_C &::= U_C \mid \neg U_C \end{aligned} \quad (1)$$

where P^- is the *inverse* of P , \top_D is the *universal value-domain* and \top_C is the *universal concept*; furthermore, U_C is a given attribute with domain $\delta(U_C)$ (resp. range $\rho(U_C)$). A DL-Lite_A *knowledge base* (KB) is a pair $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ where the TBox \mathcal{T} and the ABox \mathcal{A} consist of finite sets of axioms as follows:

- *inclusion assertions* of the form $B \sqsubseteq C$, $Q \sqsubseteq R$, $E \sqsubseteq D$, and $U_C \sqsubseteq V_C$; respectively
- *functionality assertions* of the form *funct* Q and *funct* U_C ;
- *membership assertions* of the form $A(a)$, $D(c)$, $P(a, b)$, and $U_C(a, c)$, where a, b are individual names in Γ_I and c is a value in Γ_V .

We extend DL-Lite_A and introduce the possibility to specify the *localization* of atomic concepts and roles. For this, we extend the standard DL-Lite_A syntax as follows:

$$\begin{aligned} C &::= \top_C \mid B \mid \neg B \mid \exists Q.C' \mid (loc A) \mid (loc_s A) \\ R &::= Q \mid \neg Q \mid (loc Q) \mid (loc_s Q), \end{aligned} \quad (2)$$

where $s \in \Gamma_S$ and the concept and roles are as defined before. Intuitively, $(loc A)$ is the set of indi-

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

viduals in A that can have a spatial extension (e.g., $(loc\ Parks)$), and $(loc_s\ A)$ is the subset where it is s (e.g., $(loc_{(48.20,16.37)}\ Vienna)$).

The extension with streaming consists of the following axiom schemes

$$(stream_F\ C) \text{ and } (stream_F\ R), \quad (3)$$

where F is a particular stream over either complex concepts C or roles R of \mathcal{T} . Intuitively, $(stream_F\ C)$ and $(stream_F\ R)$ are axioms defining that concept C and role R are based on streamed items from the data stream F . In a static setting, F is always empty and hence C/R can be treated like an ordinary concept/role. In a hypothetical case where all data items in F are known, the concept/role is interpreted over all time points of \mathbb{T} . The dynamic case with changing data items is discussed in Subsection 5.3.

Example 5.2. A TBox may contain $(stream_{CAM}\ speed)$, $(stream_{CAM}\ loc\ pos)$, $(stream_{CAM}\ Vehicle)$, and $(stream_{SPAT}\ hasState)$; we have the axioms $Car \sqsubseteq Vehicle$, $Bike \sqsubseteq Vehicle$, $Ambulance \sqsubseteq Vehicle$, and $Ambulance \sqsubseteq \exists hasRole.Emergency$ describing vehicle classes.

Finally, a *spatial-stream knowledge base* is a tuple $\mathcal{K} = \langle \mathcal{T}, \mathcal{A}, \mathcal{S}_A, \langle \mathcal{F}, \mathcal{S}_F \rangle, \mathcal{B} \rangle$,

where \mathcal{T} (\mathcal{A} , resp.) is a DL-Lite_A TBox (ABox, resp.), \mathcal{S}_A is a spatial DB, and $\langle \mathcal{F}, \mathcal{S}_F \rangle$ is a spatial-stream DB. Furthermore, $\mathcal{B} \subseteq \Gamma_I \times \Gamma_S$ is a partial function called the *spatial binding* from \mathcal{A} to \mathcal{S}_A and $\langle \mathcal{F}, \mathcal{S}_F \rangle$. The binding function \mathcal{B} does not relate to the *mapping function* known from OBDA mapping of concepts and roles to an underlying DB (as used in [35]), but *guarantees* that spatial objects in \mathcal{A} have a spatial extension in \mathcal{S}_A and $\langle \mathcal{F}, \mathcal{S}_F \rangle$.

Semantics. Without a spatial-stream extension, we keep the semantics of DL-Lite_A defined as usual (see [4]), based on interpretations $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ of a DL-Lite_A KB $\langle \mathcal{T}, \mathcal{A} \rangle$, where $\Delta^{\mathcal{I}}$ is a non-empty domain and $\cdot^{\mathcal{I}}$ an interpretation function of the vocabulary.

First, we give a semantics to $(loc\ Q)$ and $(loc_s\ Q)$ for individuals of Q with *some spatial extension* respectively *located at s* , such that a KB $\mathcal{K}_S = \langle \mathcal{T}, \mathcal{A}, \mathcal{S}, \mathcal{B} \rangle$ can be transformed into an ordinary DL-Lite_A KB $\mathcal{K}_O = \langle \mathcal{T}', \mathcal{A}' \rangle$, using the fresh spatial top concept C_{S_T} and spatial concepts C_s . An interpretation of \mathcal{K}_S is a structure $\mathcal{I}_S = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}}, b^{\mathcal{I}} \rangle$, where $\langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is an interpretation of $\langle \mathcal{T}, \mathcal{A} \rangle$, and $b^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Gamma_S$ is a partial function that assigns some individuals a location, such

that for every $a \in \Gamma_I$, $(a, s) \in \mathcal{B}$ implies $b^{\mathcal{I}}(a^{\mathcal{I}}) = s$. We extend the semantics with $(loc\ Q)$ and $(loc_s\ Q)$, where Q is an atomic role in \mathcal{T} :

$$(loc\ Q)^{\mathcal{I}_S} \supseteq \bigcup_{s \in \Gamma_S} (loc_s\ Q)^{\mathcal{I}_S}, \text{ where} \\ (loc_s\ Q)^{\mathcal{I}_S} = \{(a_1, a_2) \in Q^{\mathcal{I}} \mid (a_2, s) \in b^{\mathcal{I}}\}. \quad (4)$$

The transformation of \mathcal{K}_S to an ordinary DL-Lite_A KB \mathcal{K}_O is described in [36] and guarantees that we can transform \mathcal{K}_S into an ordinary DL-Lite_A KB for query rewriting.

Second, we give an initial streaming semantics by interpreting the streamed KB over the full timeline. The timeline is captured by a finite sequence $\mathcal{F}_A = (\mathcal{F}_i)_{\mathbb{T}_{\min} \leq i \leq \mathbb{T}_{\max}}$ of temporal ABoxes, which is obtained via the evaluation function v on \mathcal{F} and \mathbb{T} (similar to [37]). Then, we define the interpretation of the point-based model over \mathbb{T} as a sequence $\mathcal{I}_F = (\mathcal{I}_i)_{\mathbb{T}_{\min} \leq i \leq \mathbb{T}_{\max}}$ of interpretations $\mathcal{I}_i = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}_i} \rangle$; then \mathcal{I}_F is a model of \mathcal{K}_F , denoted $\mathcal{I}_F \models \mathcal{K}_F$, iff $\mathcal{I}_i \models \mathcal{F}_i$ and $\mathcal{I}_i \models \mathcal{T}$, for all $i \in \mathbb{T}$.

The semantics of the $(stream_F\ C)$ and $(stream_F\ R)$ axioms is along the same lines. A stream axiom is satisfied, if a complex concept C (resp. role R) holds over all the time points of stream $F = (\mathbb{T}, v, P)$; thus we restrict our models such that:

$$(stream_F\ C)^{\mathcal{I}} = \bigcap_{i \in tp(\mathbb{T}, P)} C^{\mathcal{I}_i} \text{ and} \\ (stream_F\ R)^{\mathcal{I}} = \bigcap_{i \in tp(\mathbb{T}, P)} R^{\mathcal{I}_i}, \quad (5)$$

where $tp(\mathbb{T}, P)$ is a set of time points determined by the segmentation of \mathbb{T} by P . This allows us to check *satisfiability* of a KB and gives us a notion of global consistency over the full timeline, which is expected for log data of theoretical nature.

5.2. Query Language

Our query language is based on conjunctive queries (CQs) and adds spatial-stream capabilities (as shown in Example 2.2). A spatial-stream CQ $q(\mathbf{x})$ is a formula:

$$\bigwedge_{i=1}^m Q_{O_i}(\mathbf{x}, \mathbf{y}) \wedge \bigwedge_{j=1}^n Q_{S_j}(\mathbf{x}, \mathbf{y}) \wedge \\ \bigwedge_{k=1}^o Q_{D_k}(\mathbf{x}, \mathbf{y}) \wedge \bigwedge_{l=1}^p Q_{T_l}(\mathbf{x}, \mathbf{y}) \quad (6)$$

where \mathbf{x} are the *distinguished (answer) variables*, \mathbf{y} consists of *non-distinguished (existentially quantified) variables, objects, and constant values*:

- each atom $Q_{O_i}(\mathbf{x}, \mathbf{y})$ has the form $A(z)$ or $P(z, z')$, where A is a class name, P is a property name of the LDM ontology, and z, z' are from \mathbf{x} or \mathbf{y} ;

- each atom $Q_{S_j}(\mathbf{x}, \mathbf{y})$ is from the vocabulary of *spatial* relations and of the form $S(z, z')$, where z, z' represents geometries matched by S , where S is one of the following relations: $S = \{\textit{intersects}, \textit{contains}, \textit{next}, \textit{equals}, \textit{within}, \textit{disjoint}, \textit{outside}\}$;
- each atom $Q_{D_k}(\mathbf{x}, \mathbf{y})$ is similar to $Q_{O_i}(\mathbf{x}, \mathbf{y})$ but adds stream operators that relate to Continuous Query Language operators [38]. We have a window $[agr, b, e]$ over a stream D_k , where b and e are the bounds of the window in time units (positive for past, negative for future) and an *aggregate function* agr applied to the data items in the window:
 - $[agr, b]$ represents the aggregate of last (if positive) or next (if negative) b time units of stream D_k ;
 - $[b]$ represents the single tuple of stream D_k at index b with $b = 0$ if it is the current tuple;
 - $[agr, b, e]$: represents the aggregate of a window $[b, e]$ in the past/future of D_k .
- each atom $Q_{T_i}(\mathbf{x}, \mathbf{y}) = (T_1(z_1, z'_1), \dots, T_q(z_q, z'_q))$ represents a disjunction of *temporal* relations, where the variables z_1 to z'_q represent matches, i.e., individuals annotated with time points or intervals, which are filtered by the temporal relation T_i . For time points, $T_i = T_i^P$ is a binary arithmetic relation, namely $\{<, \leq, =, \geq, >\}$. For intervals, we choose the 13 relations of Allen’s Time Interval Algebra [20], where $T_i = T_i^I$ is taken from $\{\textit{before}, \textit{equal}, \textit{meets}, \textit{overlaps}, \textit{during}, \textit{starts}, \textit{finishes}\}$ and the set of their inverses (e.g., \textit{during}^-), which filter variable matches according to the start/end points of the intervals.

The “historic” window operator $[agr, b, e]$ is derived from Brandt et al. [34] and allows us to query logs represented by data streams. Details on handling the temporal relations and aggregate functions are given below. We also have added a limited form of disjunction in our temporal relations.

5.3. Query Rewriting by Stream Aggregation

For the evaluation of spatial-stream CQs, we have to extend OQA to handle spatial and streaming data, which is not considered in the standard approach as [4]. In detail, we aim at answering pull-based queries at a *single* time point T_i with stream atoms that define *aggregate functions* on different windows sizes relative to T_i . For this, we consider a semantics based on *epistemic aggregate queries* (EAQ) [39] over ontologies by

dropping the order of time points inside a window and handle the streamed data items as *bags* (multi-sets).

In the rest of this subsection, we give a detailed outline of evaluating EAQs following [3].

Epistemic aggregate queries (EAQs). Introduced by [39], EAQs are defined over bags of numeric and symbolic values, called *groups*, denoted as $\{\mid \cdot \mid\}$. Aggregates cannot be directly transferred to DL-Lite_A, due to the certain answer semantics, where each model has different groups, which result from the grouping of unknown individuals; this leads to (undesired) empty answers. In [39], the authors extended the semantics for DL-Lite with aggregates by an epistemic operator \mathbf{K} and a two-layer evaluation including the completion of the query atoms w.r.t \mathcal{T} . The basic idea is to close the atoms, so only known individuals are grouped and aggregated. More formally, a simplified EAQ is defined as⁵

$$q_a(\mathbf{x}, agr(\mathbf{y})) : \mathbf{K} \mathbf{x}, \mathbf{y}, \mathbf{z}, \phi, \quad (7)$$

where \mathbf{x} are the *grouping variables*, $agr(\mathbf{y})$ is the *aggregate function and variable*, and ϕ is a CQ called *main conditions*; \mathbf{z} are the disjoint existential variables of ϕ . We call $\mathbf{w} := \mathbf{x} \cup \mathbf{y} \cup \mathbf{z}$ the \mathbf{K} -variables of ϕ . Values for a tuple \mathbf{d} are grouped in multisets, denoted as $H_{\mathbf{d}}$, which are defined as:

$$H_{\mathbf{d}} := \{\mid \pi(\mathbf{y}) \mid \pi \in KS_{at_{\mathcal{I}, \mathcal{K}}}(\mathbf{w}; \phi) \text{ such that } \pi(\mathbf{x}) = \mathbf{d} \mid\}, \quad (8)$$

where $KS_{at_{\mathcal{I}, \mathcal{K}}}(\mathbf{w}; \phi)$ is the set of satisfying \mathbf{K} -matches of ϕ for the model \mathcal{I} (of \mathcal{K}). A \mathbf{K} -match for a query q_a in an interpretation $\mathcal{I} = \langle \Delta^{\mathcal{I}}, \cdot^{\mathcal{I}} \rangle$ is defined by a function π that maps \mathbf{w} to the domain $\Delta^{\mathcal{I}}$. KS_{at} is defined as follows:

$$KS_{at_{\mathcal{I}, \mathcal{K}}}(\mathbf{w}; \phi) := \{\pi \in Eval(\phi, \mathcal{I}) \mid \pi(\mathbf{w}) \in Cert(aux_{q_a}, \mathcal{K})\}, \quad (9)$$

where $aux_{q_a}(\mathbf{w}) \leftarrow \phi$ is an auxiliary atom used to map \mathbf{w} only to *known solutions*. $Cert(aux_{q_a}, \mathcal{K})$ are the certain answers of aux_{q_a} over \mathcal{K} .

Then, the set of epistemic \mathbf{K} -answers for an EAQ query over a model \mathcal{I} of \mathcal{K} can be defined as:

$$q_a^{\mathcal{I}} := \{(\mathbf{d}, agr(H_{\mathbf{d}})) \mid \mathbf{d} = \pi(\mathbf{x}), \text{ for some } \pi \in KS_{at_{\mathcal{I}, \mathcal{K}}}(\mathbf{w}; \phi)\} \quad (10)$$

⁵We simplified EAQs of [39] by omitting ψ and consider only aggregates with a single variable.

where H_d and KS are defined as above. The *epistemic certain answers* $ECert(q_a, \mathcal{K})$ for a query q_a over \mathcal{K} is the set of \mathbf{K} -answers that are answers in every model \mathcal{I} of \mathcal{K} . Calvanese et al. [39] showed that $ECert(q_a, \mathcal{K})$ can be computed by a “general algorithm” GA, which (1) computes the certain answers, (2) projects them on the \mathbf{K} -variables, and (3) aggregates the resulting tuples. Importantly, evaluating EAQs reduces to standard CQ evaluation over DL-Lite_A with LOGSPACE data complexity.

Lifting EAQs to Streams. Our approach is to evaluate the EAQs over *filtered* and *merged* temporal ABoxes, which are created according to the window operator and \mathbb{T}_i . We introduce a function that filters and merges (streamed) data items, relative to the window size and \mathbb{T}_i , and creates several *windowed* ABoxes \mathcal{A}_{\boxplus} , which are the unions of the static ABox \mathcal{A} and the filtered stream ABoxes from \mathcal{F} . The EAQ aggregates are applied on each *windowed* ABox \mathcal{A}_{\boxplus} by aggregating normal objects, concrete values, and spatial objects. More formally, a stream atom $\phi \boxplus_T^L agr$ is evaluated as EAQ

$$q_\phi(\mathbf{x}, agr(y)) : \mathbf{K} \mathbf{x}, y, \mathbf{z}. \phi \boxplus_T^L \quad (11)$$

over ontologies, where \mathbf{x} are the grouping variables and y is the aggregate variable, \mathbf{z} are the disjoint existential variables, and ϕ is a query atom in the same *scope* of the window operator \boxplus_T^L and aggregate functions agr .

In a stream setting, $KS_{at_{\mathcal{I}_{\boxplus}, \mathcal{K}_{\boxplus}}}(\mathbf{w}; \phi)$ is now the set of \mathbf{K} -matches of ϕ for a model \mathcal{I}_{\boxplus} of \mathcal{K}_{\boxplus} , where the *windowed* ABox \mathcal{A}_{\boxplus} is defined as $\mathcal{A}_{\boxplus} = \mathcal{A} \cup \bigcup \{ \mathcal{A}_i \mid w_b \leq i \leq w_e \}$, where w_b , resp., w_e is the absolute time point for the begin, resp., end of a window. We have four cases for the calculation of w_b and w_e based on the window size L and a pulse P , where P enlarges L according to its interval length:

- a current window with $L = 0$, i.e. $w_b = w_e = \mathbb{T}_i$;
- a past window with $L > 0$ leading to $w_b = (\mathbb{T}_i - L)$ and $w_e = \mathbb{T}_i$;
- a future window with $L < 0$ that is $w_b = \mathbb{T}_i$ and $w_e = (\mathbb{T}_i + |L|)$; and
- the entire history of a stream resulting in $w_b = 0$ and $w_e = \mathbb{T}_i$.

We obtain KB $\mathcal{K}_{\boxplus} = \langle \mathcal{T}, \mathcal{A}_{\boxplus} \rangle$ as above; the *epistemic (certain) answers* for q_ϕ over \mathcal{K}_{\boxplus} are naturally defined as

$$ECert_{\boxplus}(q_\phi, \mathcal{K}_{\boxplus}) = \bigcap_{\mathcal{I}_{\boxplus} \models \mathcal{K}_{\boxplus}} q_\phi^{\mathcal{I}_{\boxplus}}, \text{ where}$$

$$q_\phi^{\mathcal{I}_{\boxplus}} = \{ (\mathbf{d}, agr(H_d)) \mid \mathbf{d} = \pi(\mathbf{x}), \text{ for some } \pi \in KS_{at_{\mathcal{I}_{\boxplus}, \mathcal{K}_{\boxplus}}}(\mathbf{w}; \phi) \} \quad (12)$$

Algorithm 1: NSQ - Naive Stream Query Answering

Input: A stream conjunctive query q , time point \mathbb{T}_i , and a KB $\mathcal{K}_{\mathcal{F}}$

Output: Set of tuples \mathcal{O}

/ Step 1: Detemporalize */*

foreach Q_{F_i} with $\phi_i \boxplus_T^L agr$ of q **do**

$\mathcal{A}_{\boxplus_i} \leftarrow \mathcal{A} \cup_{w_b \leq i \leq w_e} \mathcal{A}_j$ according to \boxplus_T^L and \mathbb{T}_i ;

$\mathcal{K}_{\boxplus_i} \leftarrow \langle \mathcal{T}, \mathcal{A}_{\boxplus_i} \rangle$;

build $aux_i(\mathbf{x}, y, \mathbf{z})$ from $\phi_i \boxplus_T^L agr$;

build $q_{i,1}(\mathbf{x}, y, \mathbf{z}^{\phi_i})$ from $\text{PerfectRef}(aux_i, \mathcal{T})(\mathbf{x}, y, \mathbf{z})$;

build $q_{i,2}(x, agr(y))$ from $q_{i,1}(\mathbf{x}, y, \mathbf{z}^{\phi_i})$ and $\phi_i \boxplus_T^L agr$;

$R_{i,1} := \text{evaluate Answer}(aux_i, \mathcal{K}_{\boxplus_i})$ /* certain answers */;

$R_{i,2} := \text{evaluate } q_{i,2}$ over $R_{i,1}$ /* \mathbf{K} projection */;

$R_i := \text{evaluate } q_{i,2}$ over $R_{i,2}$ /* aggregation */;

Add R_i to \mathcal{A}_{aux} and replace Q_{F_i} in q with $R_i(x, y)$;

/ Step 2: Standard evaluation */*

$\mathcal{O} := \text{evaluate Answer}(q, \langle \mathcal{T}, \mathcal{A} \cup \mathcal{A}_{aux} \rangle)$;

are the \mathbf{K} -matches that are answers in the model \mathcal{I}_{\boxplus} of \mathcal{K}_{\boxplus} . In $ECert_{\boxplus}$, we did not yet address the *validity* of an assertion, say in \mathcal{A}_{\boxplus_1} , until the next assertion in \mathcal{A}_{\boxplus_3} . Two semantics are suggestive: an *initial semantics* ignores intermediate time points, and thus \mathcal{A}_{\boxplus_2} will be unknown. A second, an *inertia-based semantics* fills the missing gaps with the previous assertions.

$ECert_{\boxplus}$ gives the certain answers for a single EAQ including the ontology atoms in the same scope as the stream atoms. A full spatial-stream CQ q can be answered by answering each EAQ q_{ϕ_k} separately and joining the answers.

Naive Algorithm. Finally, we introduce the algorithm NSQ (see Algorithm 1), where \mathbf{z}^{ϕ} are the non-distinguished variables in ϕ and PerfectRef (resp. Answer) is the “standard” query rewriting (resp. evaluation) as in [4]; NSQ extends the GA algorithm of [39] to compute the answers for stream CQs as follows:

- (1) calculate the epistemic answer for each stream atom over the different windowed ABoxes and store the result in an *auxiliary ABox* using fresh atoms. Furthermore, replace each stream atom with a new auxiliary atom;
- (2) calculate the certain answers over \mathcal{A} and the auxiliary ABox, using standard DL-Lite_A query evaluation.

We omit the details of different aggregate functions used in $ECert_{\boxplus}$ and refer to [39] and [3] for an in-depth discussion.

5.4. Query Rewriting with Temporal Relations

At first sight, spatial and temporal relations could be treated similarly. As shown in [3], we are able to evaluate spatial relations regarding their *Point-Set Topologi-*

cal Relations. This amounts to pure set theoretic operations on point sets using the function $points(p)$, which defines the (infinite) set of points of a geometry p that is a sequence $p = (p_1, \dots, p_n)$ of (defined) points. For instance, the relation $inside(x, y)$ between geometries is defined as $\{(x, y) : points(y) \subseteq points(x)\}$. However, for temporal relations, we distinguished point-based relations that can be encoded as simple arithmetic filters, from interval-based relations, where the 13 relations of Allen’s Time Interval Algebra (IA) [20] can hold between two intervals. The domain of IA relations is the set of intervals $\mathbb{I} = \{[p_1], \dots, [p_k]\}$ over the linear order of \mathbb{T} defined as $[p_i] = [p_i, \bar{p}_i]$ with $p_i < \bar{p}_i$. The binary *basic IA* relations are defined according to their start/end points as follows [20]:

$$\begin{aligned}
before(x, y) &= \{(x, y) : \underline{x} < \bar{x} < \underline{y} < \bar{y}\} \\
meets(x, y) &= \{(x, y) : \underline{x} < \bar{x} = \underline{y} < \bar{y}\} \\
overlaps(x, y) &= \{(x, y) : \underline{x} < \underline{y} < \bar{x} < \bar{y}\} \\
starts(x, y) &= \{(x, y) : \underline{x} = \underline{y} < \bar{x} < \bar{y}\} \\
finishes(x, y) &= \{(x, y) : \underline{y} < \underline{x} < \bar{x} = \bar{y}\} \\
during(x, y) &= \{(x, y) : \underline{y} < \underline{x} < \bar{x} < \bar{y}\} \\
equal(x, y) &= \{(x, y) : \underline{y} = \underline{x} < \bar{x} = \bar{y}\}
\end{aligned} \tag{13}$$

Interpretation of IA relations. IA relations can be interpreted over the sets of intervals \mathbb{I}_A and \mathbb{I}_B in two ways: (a) *IA filtering*, where each relation is treated as a single binary constraint. In that sense, the temporal relation acts as a filter on all intervals in $\mathbb{I}_A \times \mathbb{I}_B$ that match the relations regarding their start/end points; (b) *IA reasoning*, which requires the computation of the path consistency of all temporal relations over the intervals in $\mathbb{I}_A \cup \mathbb{I}_B$ using the predefined composition table of [20]. The composition table is defined as a set of transitive rules on basic relations, which are applied until no new *general* relations can be inferred. For instance, if we have the edges $during(I_1, I_2)$ and $during(I_2, I_3)$, we can infer a new relation $during(I_1, I_3)$. Note that only with approach (b) all possible (chained) relations between intervals are derivable. A well-known representation for IA relations are IA graphs (also called IA networks), which are directed graphs, where the vertices are the intervals of \mathbb{I}_A and \mathbb{I}_B and the edges represent the IA relations that hold between two intervals. Hence, an IA graph (closed by transitive rules) is a materialization of all relations that can hold between intervals, and can be used to check the relations if a directed edge exists.

From timestamps to intervals. Time intervals are not directly represented in our streamed data, but are

an intermediate product of the EAQ evaluation and are used to annotate the resulting objects. After the evaluation of an EAQ, the results (also called answers) are *aggregated* data items, hence each single timestamp is not “meaningful” anymore; we need to assign time intervals to the answers, which are taken either from the window itself, or are extracted from the aggregated data items.

We already have introduced the function v that assigns to each element of \mathbb{T} the data items of $\langle \mathcal{F}, \mathcal{S}_{\mathcal{F}} \rangle$. Now, we define the function $v^* : \mathcal{E}_{\phi} \rightarrow \mathbb{I}$, where \mathcal{E}_{ϕ} is the set of epistemic (certain) answers that result from evaluating the EAQ q_{ϕ} (as described before) and \mathbb{I} is the set of intervals $\{[p_1], \dots, [p_k]\}$ over the linear order of \mathbb{T} ; q_{ϕ} represents a single query with the atom $\phi[agr, b, e]$, where agr , b , and e are defined as before. The function v^* assigns to each answer of q_{ϕ} in \mathcal{E}_{ϕ} one interval of \mathbb{I} . For instance, we would assign the window size of $[2, 6]$ to an answer as follows: $v^*(speed(c_1, 50)) = [2, 6]$. Note that we also use a shorter notation using $@$, which would be in our example: $speed(c_1, 50)@[2, 6]$.

The function v^* can be defined in different ways. In a first approach, we use \mathbb{T}_i and the window sizes b and e for generating the interval for the annotations. For instance, having \mathbb{T}_5 and $speed[avg, 3, -1]$, we would annotate to each grouped/aggregated match the interval $[2, 6]$. In a second approach, we extract for each grouped/aggregated answer of an EAQ the upper and lower bounds of the timestamps of each data item that is in the group/aggregation.

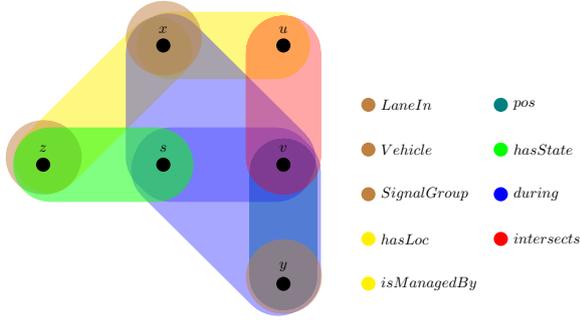
More sophisticated approaches could include a segmentation of the data items, thus creating different fragmented sub-intervals.

5.5. Query Evaluation by Hypertree Decomposition

The four types of query atoms need different evaluation techniques over separate DB entities. Ontology atoms are evaluated over the static ABox A using a “standard” DL-Lite_A query rewriting, i.e., PerfectRef [4]. For spatial atoms, we need to dereference the bindings to the spatial ABox S_A and evaluate the spatial relations to filter spatial objects. Stream atoms are computed as EAQ to group and summarized over the temporary ABoxes based on the streams as described before.

For temporal atoms, we consider three techniques, where the first is only suitable for time points and the others are suitable for time intervals:

- (1) Adding the filter conditions to the rewritten query for delimiting possible time points;

Figure 3. Hypergraph of q_2 (hyper-edge for *intersects* is simplified)

- (2) Rewriting each IA relation of Q_{T_i} into a filter that encodes the equation with the start/end points (as defined before);
- (3) Full IA-based reasoning.

In the third technique, we would have to provide full IA reasoning. This includes the closure of the IA graph by applying the (predefined) transitive rules on all intervals derived from an EAQ. Then, the derived intervals with the annotated objects from the IA graph are extracted, which hold according to the queried relations in Q_{T_i} .

In [3], we introduced two spatial query evaluation strategies assuming two restrictions:

- *no bounded variables* occur in spatial atoms, and
- the CQ is *acyclic*.

Queries with bound variables in spatial atoms are unlikely to occur, since we introduced a separation between spatial objects and their spatial extension, where the former is shown in the results and the later is used for filtering. Acyclicity, roughly speaking, is given if a CQ has no proper cycle between join variables.

The main strategy is based on the decomposition of the query hypergraph and the derived join plan, is well-suited for implementing spatial-stream CQs, as it gives us fine-grained caching, full control over the evaluation, and possibly the handling of different DB entities. Details are given in the standard DB literature such as [40].

The main steps of our query evaluation strategy are as follows. First, we construct the acyclic hypergraph H_q from q and label each hyperedge in H_q with

- l_O for an ontology edge,
- l_S for a spatial edge, and
- l_F for a stream edge with the window size added too.

Then, we build the join tree J_q of H_q and extract the subtrees J_{ϕ_i} in H_q , such that each node is covered by the same labels. Thus, we have sub-CQs that share the same aggregation/prediction functions and the same window size l . For each subtree J_{ϕ_i} , we perform the detemporalization of the stream CQ q_{ϕ_i} by extracting

and computing the results, which are stored in a virtual relation (a temporary table) R_{ϕ_i} . Finally, we traverse J_q bottom up, left-to-right, to evaluate q_{ϕ_i} for each subtree J_{ϕ_i} (without stream atoms) and cache the results in memory for future queries.

Example 5.3. The following example is a simplified version of q_1 , where the layers distinguish between ontology (first, second), stream/temporal (third, fourth), and spatial (fifth line) atoms:

$$\begin{aligned}
 q_2(x, y) : & \text{LaneIn}(x) \wedge \text{isManagedBy}(x, z) \wedge \\
 & \text{SignalGroup}(z) \wedge \text{Vehicle}(y) \wedge \\
 & \text{pos}(y, v)[\text{line}, 10s] \wedge \text{during}(v, s) \wedge \\
 & \text{hasState}(z, s)[\text{last}, 5s, -5s] \wedge (s = 'Stop') \wedge \\
 & \text{hasLoc}(x, u) \wedge \text{intersects}(u, v)
 \end{aligned}$$

Based on the hypergraph and its decomposition shown in Fig. 3, we have the following evaluation order:

- (1) $q_{2_{F1}}(y, v@i_v) : \text{Vehicle}(y) \wedge \text{pos}(y, v@i_v)[\text{line}, 10s]$;
- (2) $q_{2_{N1}}(x, u) : \text{LaneIn}(x) \wedge \text{hasLoc}(x, u)$;
- (3) $q_{2_{F2}}(x, s@i_s) : \text{LaneIn}(x) \wedge \text{isManagedBy}(x, z) \wedge \text{SignalGroup}(z) \wedge \text{hasState}(z, s@i_s)[\text{last}, 5s, -5s] \wedge (s = 'Stop')$;
- (4) $q_{2_{T1}}(y, v) : q_{2_{F1}}(y, v@i_v) \wedge \text{during}(v@i_v, s@i_s) \wedge q_{2_{F2}}(x, s@i_s)$;
- (5) $q_2(x, y) : q_{2_{T1}}(y, v) \wedge \text{intersects}(u, v) \wedge q_{2_{N1}}(x, u)$.

5.6. Complexity of Query Answering

In this section, we address the computational cost of our spatial-stream QA approach, where we refer to the following complexity classes:

$AC^0 \subsetneq LOGSPACE \subseteq NLOGSPACE \subseteq P \subseteq NP$; their definitions can be found in any standard textbook (e.g., [41]). As already stated, we are mainly interested in the reasoning service of QA with databases and DL-Lite_A. For analyzing the computational complexity of QA, we distinguish between *data*, *query*, and *combined* complexity, which is the complexity with respect to the size of the database (i.e., the ABox), the size of the query, and the combined size of both, respectively [42]. For our work, the smallest class of relevance is AC^0 , which is captured by the family of polynomial sized, constant depth, arbitrary Boolean fan-in circuits. A central result of complexity theory is that the evaluation of FO-queries (of which CQs are a fragment) is complete for (logtime uniform) AC^0 with respect to data complexity.

DL-Lite. The DL-Lite family of languages such as DL-Lite_R and DL-Lite_A, extends QA with the idea of rewriting the ontology \mathcal{O} into the CQ q resulting in a union of CQs q' that is evaluated over a DB (called ABox)

1 \mathcal{A} . This is based on the idea that computing an answer
 2 for the rewritten query $ans(q', \mathcal{A})$ is the same as com-
 3 puting the answer over the full KB $ans(q, (\mathcal{O}, \mathcal{A}))$; this
 4 property is called *FO-rewritability*. Note that in the DL-
 5 Lite family, a QA technique needs to handle existen-
 6 tially quantified roles in the head of inclusion assertions,
 7 which can be done in different ways. The languages
 8 above have combined complexity in NP and data com-
 9 plexity in LOGSPACE in [4]; hence, the evaluation of
 10 the rewritten CQ q can be delegated to an RDBMS.
 11 The only drawback is that q' can possibly be exponen-
 12 tially larger than q (and thus it takes exponential time
 13 to construct q' in the worst case)

14 The motivation for our work was to extend DL-Lite_A
 15 with spatial and stream capabilities by *keeping* FO-
 16 rewritability and the LOGSPACE data complexity. This
 17 allows us to rewrite spatial-stream CQ and evaluate
 18 them over a Data Stream Management System. In the
 19 following, we introduce the already covered extensions
 20 and highlight the restrictions, which we imposed on the
 21 extended query language, so we stay within LOGSPACE
 22 data complexity.

23 **Spatial Extensions.** Dealing with spatial queries, three
 24 aspects have to be considered regarding data complex-
 25 ity. First, we restrict the input to a finite (active) do-
 26 main. This restriction is due to the structure of Boolean
 27 fan-in circuits, where each relation is encoded by input
 28 nodes concerning all domain elements. Finiteness is
 29 guaranteed, if we encode geometries as sets of points,
 30 which we call *admissible geometries*. The binary (spat-
 31 ial) relation between two admissible geometries can be
 32 defined according to pure set operations of these points,
 33 e.g., the *contains* relation is defined by the subset op-
 34 eration. Second, spatial atoms of the form $Q_S(z, z')$
 35 need to be rewritten into unions of CQs of the form
 36 $\bigvee_{s, s' \in \Gamma_S} (C_s(z) \wedge C_{s'}(z') \wedge S(s, s'))$, where C_s and $C_{s'}$
 37 are fresh spatial concepts for the geometries s and s' ,
 38 and the spatial relation $S(s, s')$ is calculated according
 39 to set operations mentioned. Rewriting the spatial rela-
 40 tions into the query generates large unions of CQs, but
 41 FO-rewritability is retained. However, a query blowup
 42 may be avoided by rewriting the query to a Datalog
 43 program (see query rewriting techniques in Section 5.8).
 44 Third, an interaction with the existentially quantified
 45 parts of the ontology is not given, since we introduce
 46 fresh concepts that are not connected to the rest of the
 47 ontology, where existentially quantified roles might oc-
 48 cur (more details are given in [36]).

49 **Stream Relations.** As is with spatial data, the first
 50 concern regards the finiteness of the domain, since our
 51

1 timeline \mathbb{T} might be infinite. However, by applying
 2 the “timestamp” function v^* (as shown in the previous
 3 section), we limit the possible time points by the up-
 4 per/lower bound of the windows sizes. As discussed in
 5 Subsection 5.3, aggregates can not directly be lifted to
 6 DL-Lite_A, since groups, namely bags of values, lead
 7 to empty answers having the certain answer semantics
 8 of DL-Lite_A. A suitable technique was developed by
 9 Kostylev and Reutter [43], who defined a *range seman-*
 10 *tics* for aggregate queries that is based on intervals of
 11 aggregated values for each group over all models. Their
 12 approach has the drawbacks that only COUNT queries
 13 are feasible and that the data complexity increases to
 14 coNP-completeness for Boolean queries, where a tuple
 15 and a (counting) number is checked whether it is in the
 16 certain answers. An alternative semantics based on epis-
 17 temic aggregate queries is in our case favourable as the
 18 data complexity is lower, and besides COUNT also other
 19 aggregation function such as SUM are usable. However,
 20 this approach has the drawback that all known values
 21 have to be collected before the query is answered (as
 22 performed in line 10 of Algorithm 1); this is a costly
 23 operation, which can be partially mitigated by caching
 24 the static values and tracking the added/deleted value
 25 in a particular window. An extensive evaluation of miti-
 26 gation is beyond the scope of this article but considered
 27 for future work.

28 **Temporal Relations.** Earlier in this section, we intro-
 29 duced two possible strategies for rewriting temporal IA
 30 relations of the form $Q_T(z, z')$: (a) *IA filtering*, where
 31 each *basic* relation is treated as single binary constraint,
 32 and (b) *IA reasoning*, which includes the inference of
 33 possibly 2^{13} *general* relations from the 13 basic rela-
 34 tions. The IA reasoning based strategy is not feasible,
 35 as we aim to retain LOGSPACE data complexity, since
 36 already checking *Satisfiability* of a given IA network is
 37 NP-complete [44]. The IA filtering based strategy does
 38 not affect FO-rewritability, since every temporal rela-
 39 tions acts as a *single binary constraint* that fullfils the
 40 filter criteria as shown in Table 13. Furthermore, the
 41 filtering is applied only to named individuals, due the
 42 nature of EAQs.

43 **Query Structure and Decompositions.** A large body
 44 of works have been dedicated to connecting hyper-
 45 graphs, (acyclic) DB schemes, and join trees, an
 46 overview is given in [40, 45]. For decomposing a query
 47 q , the query hypergraph $H(q) = (V, E)$ represents the
 48 variables in q by vertices V and the atoms in q with
 49 shared variables by hyperedges in E . In case of an
 50 *acyclic conjunctive query* (ACQ), which is defined in
 51

terms of the acyclicity of $H(q)$, a *join tree* can be generated from $H(q)$ that yields a plan for computing the query q . Gottlob et al. show that the relation between ACQs, hypergraphs, and join trees strengthening previous results by showing that answering Boolean ACQ has a combined complexity of LOGCFL-completeness [46, 47]. We point out that the construction of the join tree and the derived evaluation order is a preprocessing step, therefore the computational complexity is of lower importance.

Note that there are different notions of acyclicity, i.e., α , β , and γ -acyclicity; we focus in this work on α -acyclicity, which can be efficiently tested by the *GYO-reduction* (cf. [40, 45]). We say that a CQ is *acyclic*, if its hypergraph is α -acyclic. A specific join tree J_H can be found via the *maximum-weight spanning tree* T_S of the *intersection graph* I_H of H , where edge weights of T_S are edge counts of V in I_H (cf. [40]).

5.7. Aggregations and Predictions

In this section, we give more details on the aggregation/prediction functionalities, where normal objects, spatial objects, and constant values need different types of aggregate functions.

For *normal objects* and *constant values*, we allow the aggregate functions *count*, *first*, and *last* on the streamed data items. For *last* and *first*, we need to search the bag of data items, as the sequence of time is lost in our representation. This is achieved by iteratively checking if we have a match at one of the points in time. In the implementation, the first and last match can be simply cached while processing the stream. For *individuals* and constant values that are *numerical*, we allow a wider range of aggregate and prediction functions on the streamed data items:

- *Order of items*: *first*, *last*, where they give the first or last element in the stream, respectively;
- *Simple aggregations*: *count*, *min*, *max*, *sum*, and *avg*;
- *Descriptive statistics (DS)*: *mean*, *sd*, *var*, and *median* are standard statistical functions that calculate the mean, standard deviation, variance, and median as expected;
- *Predictions*: we apply predefined regression methods to predict values from existing (time-series) data items inside a window. Model building (i.e., the training) and prediction should be fast, hence we support the following lightweight methods:
 - (a) *lin_reg* calculates the log-linear regression model;

- (b) *mov_avg* calculates the moving average of the past values;
- (c) *exp_smooth* applies simple exponential smoothing; and
- (d) *grad_boost* uses gradient boosting with regression trees.

Since the order of items is lost due to the bag semantics, the temporal annotations (e.g., $speed(c_1, 50)@10$) are needed in the prediction functions as the second dimension. We allow different regression methods with increasing complex models. On small windows with a required fast response time, *mov_avg* and *exp_smooth* is preferable, while on larger windows, e.g., for traffic predictions, *grad_boost* could be applied.

For *spatial objects*, geometric aggregate functions are applied to the bag of data items p_b that represent geometries. As with *first* or *last*, we must rearrange them to create a valid geometry, i.e., a sequence $p_o = (p_1, \dots, p_n)$ of points. We allow these functions to derive new geometries (among others):

- *point*: we evaluate the function *last* to get the last data item p_n of the sequence p_o ;
- *line*: we create a sequence of points p_o representing a path by calculating a total order on the bag of points p_b , such that we have a starting point using *last* and iterate backwards finding the next point by *Euclidean distance*;
- *line_angle*: the angle (in degrees) of *line* regarding a reference system is calculated by
 - (1) applying the function *line*,
 - (2) obtaining a simplified geometry using smoothing, and
 - (3) calculating the angles between the lines of the simplified geometry;
- *polygon*: similar to *line*, but we create a polygon (p_1, \dots, p_n) , where the start- and endpoints are the same, i.e., $p_1 = p_n$, by (1) determining the *convex hull* of the bag of points, and (2) extracting all pairs of points representing the convex hull;
- *traject_line* and *traject_heading* are simple techniques to project possible trajectories from past points. The former is linearly projecting the trajectory based on the previous points and the current speed. The latter calculates the trajectory based on the last point and the last heading of the vehicle.

For the trajectory computation, besides a simple linear, also a curvature-based model could be applied. To improve the accuracy of the model, we could use the speed of the last data points, so a speed-up or slow

down would be taken into account. Since this extension needs further investigation, it will be considered for future work.

5.8. Optimization Techniques

We outline three suitable optimization techniques for the query processing, namely by (a) optimizing the query rewriting itself with the aim of generating small rewritings, while keeping the number of generated rewritings low; (b) caching of static sub-queries for faster re-evaluation of the full query; and (c) the parallelization of the stream sub-queries for improving the performance for streamed, no-cachable data. A fourth optimization techniques focuses on improving the query decomposition with the aim of finding larger sub-queries for reducing in-memory joins is out of the scope of this work and needs further investigation.

Query Rewriting. Optimizations for DL-Lite query rewriting is a well studied topic. As recognized in [48], the original *PerfectRef* algorithm suffers from producing large rewritings, since an exhaustive factorization to guarantee completeness, leads to unnecessarily large unions of CQ. Significant improvements regarding *PerfectRef* were presented with *REQUIEM* [48], where the factorization was replaced with functional terms, and in *Presto* [49] where unnecessary existential joins are eliminated by computing the most-general subsumees and producing a non-recursive Datalog program instead of an union of CQ as the rewriting. Finally, the most recent optimization technique, called *tree-witness*-based rewriting [50], which boils down to introducing fresh constant symbols (nulls) for non-ABOX elements of the canonical model to allow for matches of the query in the extended (canonical) model. The *tree-witness*-based rewriting was implemented in [35], where also additional optimization techniques such as pushing joins into unions, eliminating sub-queries, and self-join elimination were added. Our approach is based on the *PerfectRef* algorithm, which is implemented with minor optimizations in the query rewriter of OWLGRES 0.1 [51]. A more efficient implementation based on *tree-witness*-based rewriting is planned for future work. For possible optimizations in the query rewriting, we define the following parameters:

- (1) structure of the ontology;
- (2) structure of the query; and
- (3) number of existentially-quantified query atoms.

Caching. For RDBMS, several approaches for caching were developed over the years. Traditionally, caching

between client- and server systems (with a DB present) was page-based (i.e., groups of related tuples), where full pages are cached on the client and requested from the server if missing. Tuple-based caching introduces a smaller data granularity, where individual tuples are cached on the client. Both caching approaches can implement different cache replacement policies such as First in First Out (FIFO), Least Recently Used (LRU), or Most Recently Used (MRU) [40]. More sophisticated caching approaches, called “semantic caching”, introduce “semantic spaces” used for caching, which allow a segmentation of the domain space of the attributes of a (DB) relation [52].

Our approach for caching the sub-queries relates to “semantic caching”. However, the segmentation of the attributes/tuples is given (a) by the separation between static and stream query atoms, and (b) by the structure of the TBox, i.e., the concept/role hierarchies as well as the domain range of roles of a query atom. Note that our segmentation does not consider spatial data, which also could be used for caching using the Euclidean distance for segmentation. For a cache-based optimization, we define the following parameters:

- (1) caching strategies as mentioned above;
- (2) cache size and location, i.e., in-memory caching or integrated caching as used by PostgreSQL;
- (3) cache invalidation based on the stream update frequency;
- (4) number and size of static/spatial sub-queries; and
- (5) ratio between static/spatial and streamed sub-queries.

Query Parallelization. Similar to caching techniques, parallel query execution is a standard RDBMS query evaluation technique. We focus in our work on *intra-query* parallelism and do not cover *inter-query* parallelism, since we restrict ourselves to execute a single query at the time for evaluating our approach. *Intra-query* parallelism can be broken down to *inter-operator* and *intra-operator* parallelism, where we omit *intra-operator* executions, which would allow for the execution of multiple processes of a single operator such as merge-join [53]. *Inter-operator* execution of different operators is, besides caching, the central optimization method for our evaluation of queries. *Inter-operator* execution can be implemented by *vertical*, i.e., pipeline-based execution of the operators, or by *horizontal* (also called bushy) execution. The horizontal, *inter-operator* execution fits well to our hypertree-based approach, since it can be used for a horizontal separation into sub-queries if there are no dependencies, i.e., hyper-edges,

The *standard query evaluator* performs the DL-Lite_A query rewriting using OWLGRES 0.1 [51], which is a prototypical implementation of the *PerfectRef* algorithm. More efficient implementation as in ONTOP [35] are planned to replace the OWLGRES-based rewriter.

The *stream aggregator* detemporalizes for each stream sub-CQ q_i the streams by grouping/aggregating the data items and performing the following steps:

- (1) extracting the data items according to the defined window size;
- (2) evaluating q_i without rewriting and storing the “known solutions” in memory as $R_{i,1}$;
- (3) evaluating q_i' with rewriting over $R_{i,1}$ and storing them in memory as $R_{i,2}$;
- (4) applying the prediction function on $R_{i,2}$ and adding the predicted data items;
- (5) applying the grouping/aggregation function on $R_{i,2}$ and producing the outcome $R_{i,3}$.

The *stream predictor* is integrated in the stream aggregator, wherein predictions are applied on the collected data items after they are grouped. We provide standard implementations for the functions *mov_avg* and *exp_smooth*. For *grad_boost*, we use the state-of-art library XGBOOST.¹⁰

The *spatial evaluator* evaluates the different spatial relations based on the grouped/aggregated data items. For performance reasons, we do not compile the spatial relation to SQL, but evaluate them in-memory using the functions of the JTS TOPOLOGY SUITE.¹¹ This is feasible as we have a clear separation in the join tree J_q by sub-query type.

The *temporal evaluator* supports the mentioned *IA filtering* technique, since temporal relations can be directly rewritten into SQL by encoding the relations as joins, where each relation is encoded as a filter on the start/end points of the aggregated data items. The second technique, *IA reasoning*, is planned for future work.

Note that we designed the prediction function as an integrated part of the stream evaluator. However, predictions could also be treated as external data streams, generating data items continuously. This would be an appealing extension, but would change the query evaluation process and needs further investigation. The same considerations would apply to the trajectory predictions. The aggregate function *traject* is designed with the same intention; we take the existing points (as coordinates)

and project a single path into the future. Currently, we apply a simple straight-line projection to create new points. However, taking the curvature into account is a desired extension.

Optimizations. In Subsection 5.8, we outlined suitable optimization techniques, where the techniques regarding query rewriting come at no cost, since it is part of the rewriter component. For instance, the query rewriter of OWLGRES 0.1 only rewrites the concepts/roles of the ontology \mathcal{O} that have assertions in the ABox. Using components of ONTOP would lead to smaller unions of CQ and faster evaluation times, since tree-witness-based rewriting was implemented therein [35].

We have implemented an initial caching technique, which is based on the separation of static/spatial and streamed sub-queries. Our implementation applies in-memory storage and mixes a LRU policy for the static sub-queries and a FIFO policy for streamed sub-queries, where data streams naturally capture a FIFO policy.

As already mentioned before, our approach favours a vertical and inter-operator parallelism, since our hypertree-based query decomposition results in a join tree, which allows us to extract stream sub-queries that are independent of each other. After the query is parsed, the query evaluator collects all independent stream sub-queries, execute them in-parallel on PIPELINEDB, and collects all the results of these queries. After the last stream sub-query is finished, the query evaluator continues to evaluate further operators by traversing the join tree to the top.

7. Evaluation

We evaluated our platform regarding the requirements/features (cf. Table 1) derived from the use cases. The requirements are encoded into a set of queries that include the introduces features of Subsection 3.2. The lightweight LDM ontology, queries, the experimental setup, log files, and collected results, as well as the implementation are available on the evaluation website.¹²

7.1. Scenario Data

For having realistic traffic data, we generated our streaming data with the microscopic traffic simulation tool PTV VISSIM,¹³ which allows us to simulate real-

¹⁰<https://xgboost.readthedocs.io/en/latest/>

¹¹<https://github.com/locationtech/jts>

¹²<http://www.kr.tuwien.ac.at/research/projects/loctrafflog/ekaw2018>

¹³<http://vision-traffic.ptvgroup.com/en-us/products/ptv-vissim/>

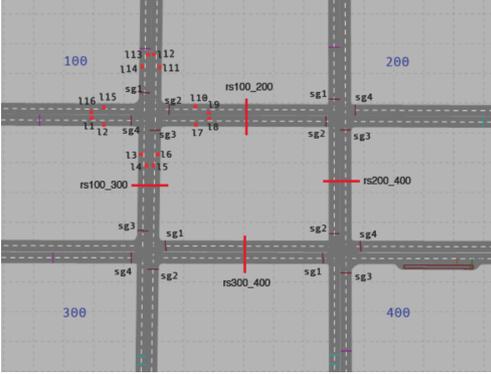


Figure 5. Four Intersection Scenario

istic driving and traffic light behavior, as well as the possibility to create unexpected events like accidents. We extract the actual state of each Vissim simulation step, and store the result as JSON in log files. A log file player is provided to replay the different simulations by feeding the log data to PIPELINEDB. For varying the data throughput, we adjusted the following parameters: (a) replayed with 5ms, 10ms, 50ms, 100ms delay, where 5ms are the fastest updates (i.e., simulating sensors) and 100ms is the real-time speed of the Vissim simulation; (b) we simulated light, medium, and heavy traffic in our scenario, where we have approx. 20, 50, and 150 vehicles, respectively, simultaneously on the road network. We modeled a real-world scenario shown in Fig. 5, which is based on a grid layout with four intersections of four roads crossing, and two incoming and outgoing lanes per street. The two incoming lanes of each side have traffic light controllers assigned; all maneuvers (turn left/right, straight on) to outgoing lanes are allowed. The main traffic flow is from north to south and west to east. We encode the structure of the full intersections into static ABox instances as follows:

- (a) intersections, roads, lanes, signal groups, and vehicles as concept assertions;
- (b) geometries for each lane, road, etc. as attribute assertions; and
- (c) lane connectivity, signal group assignments, etc. as role assertions.

7.2. Queries for Experiments

Based on the requirements, we derived a set of queries to assess each scenario, where each query aims at answering a specific problem of the use case taking the set of features into account. We use a more compact representation, where the commas between atoms are

conjunctions and disjunctions are explicitly declared using the *or* statement.

For the use case *SI.1* (object statistics), query $q_{1.1}$ determines the average and maximum speed of BMWs and VWs in the last 10 secs.

$$q_{1.1}(x, u, v) : \text{Vehicle}(x), \text{vehicleMaker}(x, z), \text{speed}(x, u)[\text{avg}, 10s], \text{speed}(x, v)[\text{max}, 10s], (z='BMW' \text{ or } z='VW')$$

For the use case *SI.2* (intersection statistics), we count vehicles according to their engine type. Sub-queries $q_{1.2a}$ and $q_{1.2b}$ select cars with either diesel or petrol engine that pass intersection $i100$. Query $q_{1.2}$ aggregates the sub-queries and returns the count of diesel in y and petrol vehicles in z , respectively:

$$q_{1.2a}(x, y) : \text{Vehicle}(y), \text{pos}(y, z)[\text{line}, 10s], \text{vehicleEngine}(y, m), (m='Petrol'), \text{intersects}(z, u), \text{hasLoc}(x, u), \text{Intersection}(x), x = 'i100'$$

$$q_{1.2b}(x, y) : \text{Vehicle}(y), \text{pos}(y, z)[\text{line}, 10s], \text{vehicleEngine}(y, m), (m='Diesel'), \text{intersects}(z, u), \text{hasLoc}(x, u), \text{Intersection}(x), x = 'i100'$$

$$q_{1.2}(x, y, z) : q_{1.2a}(x, y)[\text{count}, 10s], q_{1.2b}(x, z)[\text{count}, 10s]$$

For the use case *SI.3* (network statistics), we have two linked intersections $i100$ and $i200$. Query $q_{1.3}$ traces the vehicles that start at $i100$ and counts those passing through $i200$. A delay allows us to check the vehicle's position 7s later, and the temporal relation *before* ensures that a vehicle first passes $i100$ and then $i200$.

$$q_{1.3a}(x, v) : \text{Vehicle}(x), \text{pos}(x, v)[\text{line}, 6s], \text{intersects}(v, u), \text{Intersection}(r), \text{hasLoc}(r, u), (r = 'i100')$$

delay(7s)

$$q_{1.3b}(x, z) : \text{Vehicle}(x), \text{pos}(x, z)[\text{line}, 6s], \text{intersects}(z, w), \text{Intersection}(r), \text{hasLoc}(r, w), (r = 'i200')$$

$$q_{1.3c}(x) : q_{1.3a}(x, v), \text{before}(v, z), q_{1.3b}(x, z)$$

For the use case *S2.1* (simple maneuvers), query $q_{2.1}$ returns all vehicles x that turned left or right in the last 6s, where the function *match* is an extension of the aggregate function *line_angle* that incorporates a filter with a predefined interval on the results of *line_angle*, e.g., they have to be between -175 and -15 degrees heading. As a result, both queries are combined by an union resulting in the count of all vehicles performing the two maneuvers.

$$q_{2.1l}(x) : \text{Vehicle}(x), \text{pos}(x, y)[\text{line}, 6s], \text{Intersection}(r), \text{intersects}(y, u), \text{hasLoc}(r, u), (r = 'i100'), \text{match}(y)[\text{angle}, -175, -15]$$

$$q_{2.1r}(x) : \text{Vehicle}(x), \text{pos}(x, y)[\text{line}, 6s], \text{Intersection}(r), \text{intersects}(y, u), \text{hasLoc}(r, u), (r = 'i100'), \text{match}(y)[\text{angle}, 15, 175]$$

$$q_{2.1}(x) : q_{2.1l}(x) \text{ or } q_{2.1r}(x)$$

In use case *S2.2* (complex maneuvers), query $q_{2.2}$ detects illicit lane changes in terms of crossing the middle marker (i.e., a white line). This is detected by evaluat-

ing whether a vehicle passed for a certain period from an in-lane to an out-lane or vice versa.

$$q_{2.2}(x, y) : \text{LaneIn}(z), \text{hasLoc}(z, u), \text{intersects}(u, v), \text{Vehicle}(x), \\ \text{pos}(x, v)[\text{line}, 6s, 3s], \text{pos}(x, w)[\text{line}, 3s, 0s], \\ \text{intersects}(t, w), \text{hasLoc}(y, t), \text{LaneOut}(y)$$

For the use case S2.3 (red-light violation), we modified Ex. 1 by taking trajectory and speed prediction into account, which allows us a more precise detection of violations, since we can rule out vehicles that are slowing down or are about to change lanes.

$$q_{2.3}(x, y) : \text{LaneIn}(x), \text{hasLoc}(x, u), \text{intersects}(u, v), \text{Vehicle}(y), \\ \text{pos}(y, v)[\text{traject_line}, 5s, -3s], (r > 10), \\ \text{speed}(y, r)[\text{mov_avg}, 5s, -3s], \text{hasSignalGroup}(x, z), \\ \text{SignalGroup}(z), \text{hasState}(z, \text{Stop})[\text{last}, 5, -5]$$

For the use case S2.4 (vehicle breakdown), we check with $q_{2.4}$, if a car has stopped for longer than 30s, while (using the *during* relation) it is located inside our intersections, but not on one of the park lanes (using the *disjoint* relation).

$$q_{2.4}(x, y) : \text{Vehicle}(x), \text{speed}(x, r)[\text{avg}, 30s], (r < 1), \text{inside}(v, u), \\ \text{pos}(x, v)[\text{line}, 15s], \text{hasLoc}(y, u), \text{Intersection}(y), \\ \text{during}(v, r), \text{disjoint}(v, z), \text{hasLoc}(p, z), \text{ParkLane}(p)$$

The use case S2.5 (traffic congestion) can be evaluated by a query similar to S2.4, but with the extension that stop-and-go traffic can be excluded by checking if there is no movement while the traffic light phases are on “Go”, hence the stopping of the traffic is not caused by the traffic light.

$$q_{2.5}(x, y) : \text{Vehicle}(x), \text{speed}(x, r)[\text{avg}, 30s], \text{pos}(x, v)[\text{line}, 30s], \\ (r < 1), \text{intersects}(v, u), \text{hasLoc}(y, u), \text{LaneIn}(y), \\ \text{hasSignalGroup}(y, z), \text{hasState}(z, s)[\text{last}, 30s], \\ (s = 'Go'), \text{during}(s, r), \text{SignalGroup}(z)$$

For the use case S3.1 (self monitoring), we aim to detect with $q_{3.1}$, if our ego vehicle defined by $isEgo(x)$ is exceeding a speed limit, which is assigned to each lane and checking, if the vehicle is currently located on such a lane.

$$q_{3.1}(x, y) : \text{LaneIn}(y), \text{hasLocation}(y, u), \text{intersects}(u, v), \\ \text{pos}(x, v)[\text{line}, 5s], \text{Vehicle}(x), \text{speed}(x, r)[\text{max}, 5s], \\ isEgo(x), \text{speedLimit}(y, s), (r > s)$$

In use case S3.2 (obstructed view), we compute query $q_{3.2}$, where our prototype (as part of the ego vehicle) aims to detect vehicles that very likely will collide in 2s on an intersection by checking whether their predicted trajectories will cross another vehicle’s trajectory.

$$q_{3.2}(x, y) : \text{Vehicle}(y), isEgo(y), \text{pos}(y, v)[\text{traject_line}, 2s, -1s], \\ \text{intersects}(v, w), (r > 10), \text{speed}(x, r)[\text{mov_avg}, 5s, -2s], \\ \text{Vehicle}(x), \text{pos}(x, w)[\text{traject_line}, 2s, -1s]$$

In S3.3 (traffic rules), our ego vehicle approaches an uncontrolled intersection at the same time with other

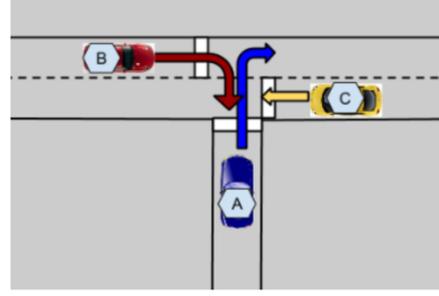


Figure 6. Use Case Traffic Rules

vehicles. According to (local) traffic regulations, preference (shown Fig. 6) is given to (1) the vehicle on the main road, (2) the one that is not changing lanes, and (3) the vehicle approaching from the right. Vehicle C has preference over A and B, where B would have preference over A, but the preference can not be given, since A is on a single-lane road. We can express the traffic regulations with the following Datalog rules:

$$\text{willCross}(x, y) \wedge \text{straightOn}(x) \wedge \text{turnRight}(y) \wedge \text{onMainRoad}(x) \\ \wedge \text{onMainRoad}(y) \rightarrow \text{giveWay}(y, x) \quad (1)$$

$$\text{willCross}(x, y) \wedge \text{onMainRoad}(x) \wedge \text{onTributRoad}(y) \\ \rightarrow \text{giveWay}(y, x) \quad (2)$$

$$\text{willCross}(x, y) \wedge \text{giveWay}(x, y) \wedge \text{onSingleLaneRoad}(x) \\ \rightarrow \text{giveWay}(y, x) \quad (3)$$

$$\text{vehicle}(x) \wedge \text{vehicle}(y) \wedge \text{giveWay}(y, x) \rightarrow \text{stop}(y) \quad (4)$$

The atom $\text{willCross}(x, y)$ matches all vehicles that might collide and can be evaluated by $q_{3.2}(x, y)$ (modified without $isEgo(x)$). The atoms $\text{turnRight}(x)$ and $\text{straightOn}(x)$ can be evaluated by the queries $q_{2.1r}(x, y)$ assuming the queries are treated as atomic rules with $q(x, y)$ as the head. Atoms $\text{onMainRoad}(x)$, $\text{onTributRoad}(y)$, and $\text{onSingleLaneRoad}(x)$ can be evaluated in the spirit of $q_{1.2a}(x, y)$ checking spatial containment. Then, the rules of S3.3 can be expressed as unions of CQs, but with the difficulties that the order of the query evaluation effects the completeness of the results, i.e., Rule (2) has to be evaluated before (3).

7.3. Results

We conducted our experiments on a Mac OS X 10.14.4 system with an Intel Core i7 2.9GHz, 8GB of RAM, and a 250GB SSD. The average of 21 runs for query rewriting time and evaluation time was calculated. The results are shown in Table 2, where the average evaluation time (AET) t is measured in seconds for three traffic densities and four update delays in ms. The columns represent the number of sub-queries $\#Q$ with stream queries in brackets, the size of rewritten atoms $\#A$, Opt denoting if optimizations are active or not, the AET, and d (%) is the average of deviations between the

1 optimized and non-optimized runs for a single traffic
 2 density with four delays. The column, resp., row *Avg.*
 3 *of d (%)* holds the calculated average over the full rows,
 4 resp., columns.

5 The new experiments confirm results of [3] with
 6 closer to “real-world” queries and simulation data. The
 7 AET ranges between 0.86s and 2.06s with the exception
 8 of use case *S3.3*, which emulates rules using unions
 9 of CQs. Query $q_{3.1}$ shows the highest delay of 2.06s,
 10 since the join condition of ($r > s$) is evaluated inline
 11 and not on the DB, which adds a delay of 0.4s with
 12 larger windows. Our baseline query is $q_{1.1}$ tested with
 13 100ms delay and low traffic. It has an AET of 0.86s,
 14 where 0.23s is the time-to-load (TOL), 0.63s is needed
 15 for query evaluation of two stream atoms, where we
 16 artificially delaying the next stream atom evaluation by
 17 150ms. The artificial delay is empirically determined
 18 and needed for PIPELINEDB to set up the *continuous*
 19 *views* (CVs) on-the-fly; ignoring this would lead to
 20 missing results.

21 The added functions for statistics, matching, and pre-
 22 dictions, i.e., *mov_avg* and *traject_line* do not affect
 23 performance, since they are applied on small windows
 24 with few data items. If we would apply the function
 25 *grad_boost* for predictions on large windows for long-
 26 term traffic predictions, our query time could rise con-
 27 siderably, since prediction time (without a preprocessed
 28 training step) can be above 20s.

29 **Discussion on Optimizations.** A second set of exper-
 30 iments highlights the effect of the optimizations based
 31 on caching and sub-query parallelization. The results
 32 are shown in the second row of each query in Table
 33 2. The optimizations improve the original AET ranging
 34 between 0.49s and 1.05s, to the new results with a min-
 35 imal AET of 0.37s (in $q_{1.1}$ with 100ms delay on low
 36 traffic density), and with the worst result of a maximal
 37 AET of 0.74s (in query $q_{2.5}$ with 5ms delay and high
 38 traffic density). Note that we do not take query $q_{3.3}$ into
 39 account, since it encodes rules and is discussed below.

40 The performance gain in AET is shown in percentage
 41 terms for each traffic density in the columns 9, 14, 19,
 42 and 20, which includes the calculated average of the
 43 other three columns. We observe that the gain for each
 44 query is almost uniform over the traffic densities, and
 45 the variation is between the queries (shown in column
 46 20), and smaller between the four delays (shown in row
 47 23). Regarding the results in column 20, the queries $q_{1.3}$
 48 and $q_{3.1}$ show the highest, and the queries $q_{2.5}$ and $q_{3.3}$
 49 the lowest gains. The improvement in $q_{1.3}$ is related to
 50 the effect that the same query is executed twice with a
 51

1 7s delay, hence the caching has a strong impact. The
 2 improvement of $q_{3.1}$ is mainly related to the fact that
 3 a single ego-vehicle is observed, and in the case that
 4 the vehicle is not present no result are calculated. The
 5 (undesirable) lowest gain in query $q_{2.5}$ is related to an
 6 implementation detail of the sub-query parallelization,
 7 which concurrently executes only pairs of sub-queries;
 8 hence the third stream query is run after the initial two
 9 stream queries. The average performance gain over all
 10 queries is 58.53%, which is encouraging, however more
 11 optimizations are feasible. Finally, we observed that
 12 the optimizations mitigate two negative effects on the
 13 overall performance:

- 14 (1) the rewriting of the ontology of sub-queries into
 15 views is only done in the initial run, since after-
 16 wards the results are taken from the cache;
- 17 (2) the artificial delay of 150ms for evaluating con-
 18 secutive stream sub-queries is dropped, since this
 19 evaluation is newly conducted in parallel.

20 Note that query $q_{3.3}$ is a special case and shows the
 21 least performance gain, since four rules are evaluated
 22 sequentially as queries starting with query $q_{2.1}$ twice,
 23 $q_{3.2}$, and finally merging the results of the queries at
 24 the end. This has the effect that the caching is local,
 25 performed for each query separately, and not global.

27 7.4. Feature Coverage

28 As shown with the queries, we covered in the im-
 29 plementation all initial levels (L1) of features that are
 30 defined in the scenarios/uses cases. We support tem-
 31 poral relations based on a (partial) interval-based data
 32 model (*F1*) that are evaluated by pull-based queries
 33 (*F2*). Then, we allow temporal relations and nested
 34 queries that include unions of CQs (*F3*). However, we
 35 have not yet implemented the *IA reasoning* for tem-
 36 poral relations, since an in-memory evaluation of the
 37 transitive rules completing the IA graph needs further
 38 investigation. Regarding *F4* and *F5*, we have imple-
 39 mented the initial set of numerical, descriptive statisti-
 40 cal, and spatial aggregation functions. For *F6*, we cov-
 41 ered *mov_avg* and *exp_smooth* for fast, simple predic-
 42 tions, and support *grad_boost* for long-term traffic fore-
 43 casting. For trajectory prediction (*F7*), we have imple-
 44 mented a method based on a simple linear path calcu-
 45 lation. However, more accurate trajectory predictions
 46 would be desired. Feature *F8* is covered by the atom
 47 *match(y, z)[angle, 0, 15]*, and *F9* is partially covered by
 48 unions of CQs, but transitive rules are out of scope for
 49 this work.
 50
 51

	#Q	#A	Opt	(l) with ms delay					(m) with ms delay					(h) with ms delay					Avg. of d (%)
				5	10	50	100	d (%)	5	10	50	100	d (%)	5	10	50	100	d (%)	
$q_{1.1}$	3(2)	42	N	1.35	1.18	0.95	0.86		1.45	1.30	0.99	0.88		1.46	1.35	1.14	0.99		
			Y	0.45	0.43	0.38	0.37	61.80	0.47	0.46	0.41	0.38	61.90	0.48	0.48	0.43	0.39	63.61	62.44
$q_{1.2}$	6(2)	43	N	1.30	1.20	1.01	0.96		1.33	1.24	1.04	1.00		1.41	1.38	1.07	1.01		
			Y	0.47	0.45	0.41	0.40	61.02	0.49	0.48	0.43	0.39	61.03	0.52	0.53	0.46	0.42	60.03	60.69
$q_{1.3}$	8(2)	44	N	1.44	1.35	1.15	1.08		1.47	1.37	1.23	1.09		1.45	1.44	1.30	1.20		
			Y	0.49	0.45	0.41	0.39	65.22	0.49	0.46	0.41	0.41	65.54	0.52	0.52	0.45	0.43	64.39	65.05
$q_{2.1}$	6(2)	43	N	1.31	1.20	1.01	0.98		1.43	1.29	1.09	0.99		1.48	1.40	1.13	1.02		
			Y	0.50	0.47	0.41	0.39	60.57	0.52	0.50	0.44	0.40	61.03	0.57	0.54	0.46	0.43	60.01	60.54
$q_{2.2}$	7(2)	45	N	1.36	1.26	1.05	1.00		1.47	1.29	1.08	1.03		1.51	1.43	1.13	1.06		
			Y	0.63	0.54	0.43	0.40	57.47	0.66	0.65	0.52	0.45	53.22	0.61	0.58	0.50	0.48	57.38	56.02
$q_{2.3}$	7(3)	50	N	1.57	1.50	1.27	1.21		1.63	1.53	1.30	1.22		1.72	1.65	1.37	1.27		
			Y	0.66	0.63	0.56	0.54	56.81	0.69	0.65	0.57	0.55	56.56	0.67	0.66	0.59	0.55	58.67	57.35
$q_{2.4}$	5(2)	46	N	1.24	1.21	0.98	0.92		1.28	1.24	1.06	0.97		1.28	1.29	1.13	0.99		
			Y	0.47	0.44	0.42	0.39	60.12	0.48	0.47	0.43	0.39	60.96	0.54	0.48	0.46	0.41	59.62	60.23
$q_{2.5}$	7(3)	43	N	1.44	1.38	1.16	1.08		1.50	1.41	1.20	1.11		1.55	1.47	1.26	1.17		
			Y	0.69	0.68	0.59	0.56	50.02	0.74	0.68	0.62	0.56	50.08	0.74	0.74	0.66	0.59	49.78	49.96
$q_{3.1}$	5(2)	43	N	1.85	1.72	1.40	1.32		1.89	1.79	1.48	1.35		2.06	2.04	1.57	1.38		
			Y	0.51	0.44	0.41	0.38	72.19	0.47	0.48	0.43	0.41	72.22	0.49	0.52	0.44	0.40	73.43	72.62
$q_{3.2}$	5(3)	63	N	1.41	1.34	1.23	1.17		1.48	1.43	1.27	1.20		1.56	1.51	1.31	1.21		
			Y	0.62	0.62	0.57	0.55	54.10	0.66	0.62	0.59	0.54	55.15	0.67	0.66	0.60	0.57	55.11	54.79
$q_{3.3}$	12(5)	43	N	3.02	2.80	2.42	2.39		3.26	2.98	2.58	2.38		3.36	3.20	2.66	2.44		
			Y	1.63	1.57	1.40	1.34	44.01	1.71	1.63	1.48	1.35	44.69	1.82	1.75	1.53	1.44	43.65	44.12
Avg. of d(%)				59.87	59.56	57.36	57.15	58.49	60.46	59.06	56.95	57.12	58.40	60.52	59.94	57.47	56.86	58.70	58.53

Table 2

Results (t in secs) for scenario with (l)ow, (m)edium, and (h)heavy traffic, where Y means results with caching and parallelization, and N means results non-optimizations; d (%) is the average of deviations between optimized/non-optimized runs.

7.5. Summary of Expert Evaluation

Overall, the experts confirmed that (a) the extension to time intervals and temporal relations, and (b) the inclusion of prediction capabilities are important extensions of the initial spatial-stream OQA approach. However, it turned out that the experts identified several limitations and interesting extensions.

The first limitation regards the *LDM ontology*, which should be aligned with the *SSN ontology* [32] to include patterns such as *Stimulus*, *Sensor*, and *Observation*. Furthermore, it should capture a spatial model based 3D maps. The second limitation regards the unclear definition of the *grammar* of stream atoms, which should be clarified in this work. The third limitation is more generic and captures the assumption that rules would be better suited to capture the complex use cases, which is most apparent in the traffic rules use case.

The first extension regards Collective Perception Messages [28], which could be added as a new type of message used to share local sensor data. The second extension addresses Scenario 3, which could be extended

with use cases that are taken from motion planning tasks in autonomous driving. The third extension identifies Kalman filters [29] and top-k aggregates [33] as more powerful aggregation functions. The fourth extension is likely the furthest reaching, since the ontology and query language (with the underlying semantics) could support uncertainty on the level of (a) data items and (b) of TBox assertions, e.g., inclusion assertions. The fifth extension describes the handling of data items inside windows, where adaptably forgetting data items and extending their validity into future could be added to have a more flexible query language. The last extension is again more generic and captures the use of LDMs in an autonomous agent-based scenario, which would lead to the challenge of aligning different LDMs to a global dynamic map.

8. Related Work and System Comparison

First, we give a general overview on related work, and will discuss in-depth related systems in the sec-

tions afterwards, where we provide an qualitative and quantitative comparison with selected systems.

8.1. Overview

In a more formal setting, stream processing started with Data stream management systems (DSMSs) such as STREAM [38], Aurora [55], and TelegraphCQ [56]. Notably in STREAM, the authors introduced the Continuous Query Language (CQL) [38], which provides an explicit operational semantics based on stream-to-relation, relation-to-relation, and relation-to-stream operators. Furthermore, three kinds of window operators, namely partition-based, time-based, and tuple-based windows were applied in CQL.

With the need for more complex domain models, such as provided by background KBs, streams were lifted to a “semantic” level, leading to the development of RDF stream processing engines, such as Streaming SPARQL [57], C-SPARQL [7], SPARQLstream [58], and CQELS [8]. These systems proposed processing of RDF streams integrated with other Linked Data sources and background KBs. The C-SPARQL framework also has been extended to support deductive and inductive reasoning [59]. A system for efficient spatio-temporal query execution was proposed in [60], which supports spatial operators as well as aggregate functions over temporal features. EP-SPARQL [61] and LARS [62] propose languages that extend SPARQL respectively CQs with stream reasoning, but translate KBs into more expressive and less efficient logic programs. Regarding spatial RDF stream processing, a few SPARQL extensions were proposed, such as SPARQL-ST [63] and st-SPARQL [64].

In [65], DL-Lite_A was extended with temporal operators such as used in LTL [66], introducing a semantics based on a two-sorted model separating the object from the temporal domain. Besides [65], similar temporal QA was investigated in [37] and [67], which are all on the theoretical side and for which no implementation has been yet provided. Finally, our framework was build on the results for EAQs provided by [39], but we investigated EAQs in a streaming setting and introduced more complex queries supporting also spatial and temporal relations.

8.2. Comparison with Existing Systems

There is a wide range of systems for stream processing, stream reasoning, and event detection available. For a comparison with our approach, we focus on systems that fulfill the following criteria: (1) abil-

ity to handle streaming and/or temporal data, either by having a window operator, or supporting incremental updates; (2) ability to provide reasoning, either based on ontology-, rule-, or query rewriting-based methods; and (3) a maintained implementation of the system has to be available.

As commented in Section 3, we will not re-evaluate the eight “standard” requirements of [18], and the three entailment levels for stream reasoning systems of [19]. Still, this work overlaps on some features with the comprehensive study of Margara et al. [68], where the authors compare systems according to the criteria:

- continuous queries (F2/F3),
- background data (F9),
- time model (F1),
- reasoning (F9), and
- temporal operators (F1).

The other criteria in [68], namely data transformation, uncertainty management, historical data, quality of service, and parallel/distributed processing are not investigated in the context of this work.

The comparison in Table 3 shows the evaluation of the selected systems on the generic features *F1* (Time model), *F2* (Process paradigm), *F3* (Query features), and *F9* (Advanced reasoning), as well as the ITS domain specific features *F4* (Numerical aggregations), *F5* (Spatial aggregations), *F6* (Numerical predictions), *F7* (Trajectory predictions), and *F8* (Spatial matching). We separate the table into two groups, where the first group represents query-based systems, and the second group rule-based systems. In the following, we give details on the systems that are compared.

8.2.1. C-SPARQL

C-SPARQL was introduced by Della Valle et al. [7] and was one of the first contributions in the area of stream processing with reasoning extensions. The main intention of this work lies in bridging the gap between the world of stream processing systems, i.e., stream RDBMS, and the Semantic Web taking RDF and SPARQL into account. C-SPARQL includes (a) a language for continuous queries over streams of RDF data, (b) an evaluation engine for this language, whereby C-SPARQL has the distinguishing features of (i) supporting timestamped RDF triples, (ii) supporting continuous queries over streams, and (iii) defining ad-hoc, explicit operators for data aggregation. The results of C-SPARQL queries are continuously updated as new data items appear on the stream, hence an efficient evaluation of sliding windows is possible.

System	F1	F2	F3	F4	F5	F6	F7	F8	F9
<i>C-SPARQL</i>	Point	Pull	SPARQL + windows	Yes	Pre	Pre	Pre	No	RDFS
<i>CQELS</i>	Point	Push	SPARQL + windows	Yes	Pre	Pre	Pre	No	RDF
<i>INSTANS</i>	Point	Push	SPARQL + event patterns	Yes	Pre	Pre	Pre	No	RDFS + Rete
<i>SPARQLstream / Morph-streams</i>	Point	Pull	SPARQL + windows	Yes	Pre	Pre	Pre	No	OWL2 QL
<i>ONTOP (DatalogMTL)*</i>	Point + MTL	Pull	SPARQL + windows	Yes	Yes	Yes	Pre	No	DatalogMTL
<i>ONTOP (STARQL)*</i>	Point	Pull	SPARQL + windows	Yes	Yes	Yes	Pre	No	OWL2 QL
<i>TrOWL</i>	Point	Pull	CQ	Lmt	Pre	Pre	Pre	No	OWL2 EL or approx. OWL2 DL
<i>Clingo (Multi-shot ASP)*</i>	Point	Pull	Rules + ext. atoms	Yes	Ext	Ext	Ext	Ext	ASP
<i>ETALIS (EP-SPARQL)</i>	Interval + full AIA	Push	Rules or SPARQL + windows	Yes	Lmt	Pre	Pre	No	Prolog
<i>RDFox</i>	Point	Pull	Rules	Yes	Pre	Pre	Pre	No	Datalog / OWL2 RL (incremental)
<i>Laser (LARS)*</i>	Point + LTL	Pull	Rules + windows	Yes	Pre	Pre	Pre	No	Stratified LARS
<i>Ticker (LARS)*</i>	Point + LTL	Pull	Rules + windows	Yes	Pre	Pre	Pre	No	LARS
<i>Spatial-stream OQA</i>	Point/Interval with limited AIA	Pull	CQ + windows	Yes	Yes	Yes	Lmt	Lmt	OWL2 QL

Table 3

Qualitative comparison of F1 - F9 on selected systems, where * indicates comments in the system descriptions; *Ext*, *Pre*, and *Lmt* means evaluation by using of external atoms, a preprocessing step needed, and with limited coverage

8.2.2. CQELS

Similar to C-SPARQL, CQELS [8] also offers a query language and processing engine for answering queries over a combination of static and stream RDF data. While C-SPARQL adopts a “black box” approach, i.e. static and stream query elements are first divided and sent to the respective underlying stream and RDF query engines, CQELS, on the other hand, takes a “white box” approach by natively implementing the required query operators for the RDF-based data model, both for streams and static data. This native approach enables better performance and can dynamically adapt to changes in the input data. Another difference to C-SPARQL is that CQELS takes an eager query execution strategy: input data is processed as soon as it arrives in the system, contrary to the periodic evaluation of C-SPARQL which is triggered periodically, regardless of the input data throughput.

8.2.3. ETALIS with EP-SPARQL

The ETALIS system [61, 69] was applied to the ITS domain and offers the combination of Datalog-style rules with a background KB. In ETALIS a Prolog-based language is used to express complex event patterns with predicates like *during(event1, 5)* or *begins(event1, event2)*. The background KB is also encoded into the rule language, which in combination with the temporal and causal parts can be used for query answering. The standard Prolog query evaluation, which is based on a request-response paradigm, is altered to an event-driven backward chaining (EDBC) of rules. A standard Prolog system is then triggered to evaluate a query and additional EDBC rules when new data is arriving at the system. EP-SPARQL [61] is an approach that extends ETALIS and introduces windows and the handling of RDF streams for lifting the rule-engine to a Semantic Web/Linked Data settings.

8.2.4. INSTANS

INSTANS [70] is an event processing engine based on handling multiple interconnected SPARQL queries with updates. It supports continuous evaluation of incoming RDF data using an encoding of SPARQL queries into Rete-like structures. INSTANS supports stateless/stateful filters using its internal memory, enrichment, (de-)composition, aggregation and pattern matching on the streamed events. The authors have implemented their approach, where they provide a conversion of the SPARQL queries to Rete rules, which then are evaluated on the Rete [71] rule engine JESS [72].

8.2.5. SPARQLstream/Morph-streams

SPARQLstream [58] extends standard SPARQL with time windows over streams similar to C-SPARQL, but adds the relation-to-stream operator for dealing with relational streams. SPARQLstream was further extended to Morph-streams [73], where OBDA techniques are applied to access the underlying streams, which then are stored in a stream processing system (SPS). R2RML mappings are used to create virtual streams on-the-fly, which can be accessed in their SPARQL extension. The authors implemented and tested their query rewriting techniques for different SPS, namely for SNEE, Esper, and Global Sensor Networks.

8.2.6. Clingo with Multi-shot ASP

Multi-shot Answer Set Programming (ASP) is an extension of existing ASP solving techniques, which deals in a reactive fashion if new information arrives at the logic program, instead of solving the program from scratch. Clingo 4 [74] supports natively multi-shot solving by offering high-level constructs and control capacities via the scripting languages Lua and Python. The authors introduced the *#external* directive, which allows a flexible handling of undefined atoms. Additionally, the solver needs to keep track of the sequence of system states, which is defined using a simple operational semantics, where operations such as *create*, *add*, or *assignExternal* can be used to modify the states.

8.2.7. LARS with Ticker/Laser

The LARS framework [75] is a recent development in stream reasoning that considers lifting ASP to streams with generic windows for capturing data snapshots with the aim of generalizing time-, tuple- and partition-based window functions. To this end, the ASP syntax is extended with window operators and with temporal operators for evaluating truth at every, some, and a specific (exact) time point in a stream; variables ranging over domain constants or time points are al-

lowed in rules. Semantically, answer sets of ASP programs are naturally generalized to answer streams of LARS programs. Fragments of LARS programs have been implemented in Ticker [76] and Laser [77], based on plain LARS rules.

In Ticker,¹⁴ full negation is supported, with sliding time-based and tuple-based windows. It comes with two evaluation modes, viz. a static ASP encoding, which employs Clingo for repeated solving, and an incremental ASP encoding, which performs model update by truth maintenance techniques; the latter tends to be faster.

Laser¹⁵ is geared towards high performance evaluation and thus focuses on positive and stratified programs with sliding windows. It aims at fast model update by an efficient substitution management, for which it extends semi-naive evaluation of Datalog by incorporating the temporal dimension and tracks intervals of guaranteed formula validity. This avoids redundant re-derivations and allows for efficient removal of expired derivations. Laser was shown to outperform Ticker, C-SPARQL and CQELS in micro-benchmarks.

8.2.8. ONTOP with STARQL and DatalogMTL

The STARQL framework of Özçep et al. [78] is an effort of streamifying ODBA by introducing an extensible query language, hence it is closely related to our approach. It uses a first-order logic fragment for temporal reasoning over sequences of ABoxes. The framework extends the first-order query rewriting of DL-Lite with intra-ABox reasoning. In a second extension of ONTOP, the authors added Metric Temporal Logic (MTL) to allow querying of log data using LTL operators that are extended with time intervals [34, 79]. For this purpose, they introduced datalogMTL, which combines non-recursive Datalog with the MTL operators.

8.2.9. RDFox

RDFox [80] is the combination of a scalable main-memory RDF store that supports materialisation and parallel Datalog reasoning, which also includes SPARQL query answering. The Datalog materialisation is based on a novel parallel reasoning algorithm extending the well-known DRed algorithm, computing incremental updates very efficiently on its internal triple store and thus is well suited to handle data streams.

¹⁴<https://github.com/hbeck/ticker>

¹⁵<https://github.com/karmaresearch/laser>

8.2.10. TrOWL

TrOWL [81] is an incremental DL reasoner over the expressive OWL2 DL language. It handles streams of KBs instead of using fixed time windows over streams, hence allowing to add and remove axioms from the KB on-the-fly. The authors applied syntactic approximation to reduce the reasoning complexity, which guarantees soundness but loses completeness in certain cases. TrOWL provides justifications for the following entailments: atomic concept subsumption, atomic class assertion and atomic object property assertion.

8.3. Feature Comparison

Based on the literature review and discussions with the authors of the systems, we conducted a feature evaluation of the above systems. In Table 3, we summarized the reviews and discussions, where we use the underlying specifications for the three levels of fulfillment: *basic*, *enhanced*, and *advanced*. For F1, the basic level matches to a point-based time model, the enhanced level would relate to an interval-based model, and the advanced level would include AIA over an interval-based model. For F2, the basic level includes pull-based, the enhanced level push-based queries, and the advanced level the combination of pull- and push-based queries. In F3 and F9, we list the query, rule, or ontology language and the possibility to allow windows, but do not classify the fulfillment level. For F4 to F8, we evaluate how a specific feature is covered by the *basic* fulfillment levels, since most systems are generic reasoners, and are not intended to support ITS-specific features.

As shown in Table 3, for F1 the ONTOP- and LARS-based systems offer similar or richer query- and ontology languages, since these systems allow the use of LTL and MTL operators. For F2, our work is on a par with most presented systems, except the two push-based systems CQELS and INSTANS, which support a higher reactivity, since they support an eager query execution strategy. For F4 and F5, our work covers the widest range of numerical aggregation and prediction functions, with the exception of STARQL, which covers similar functionalities. One of the motivations of this work are the coverage of F6 to F8; hence our work is the only approach that supports spatial aggregations and predictions.

8.4. Performance Comparison

We conducted our experiments again on a Mac OS X 10.14.4 system with an Intel Core i7 2.9GHz, 8GB of

RAM, and a 250GB SSD. We calculated the average of 50 runs for query evaluation time with warm starts, hence, we did not restart the systems over 50 runs in each experiment.

For the experiments, we compared our prototype on two selected queries with the state-of-the-art systems C-SPARQL and CQELS, which support limit reasoning but are designed to deal with high velocity and volume streams. The comparison of all presented queries is not feasible, since C-SPARQL and CQELS do not support natively the features spatial/temporal relations, spatial aggregation, and inline predictions. Hence, we selected the two queries $q_{1.1}$ and $q_{2.3}$, where the first is our baseline comparison and the second is our running example. We pre-calculated the missing features such as the spatial relations in the log player and materialized the outcome as streamed data items. Furthermore, we adapted our CQs to the SPARQL dialects of each system. In Figure 7, we give the encoding of $q_{1.1}$ as an example.

The results of our experiments are shown in Table 4 where t is the AET in seconds for different traffic densities and update delays in ms. The baseline systems C-SPARQL and CQELS outperform our prototype in the range between 70ms in $q_{1.1}$ (C-SPARQL on heavy traffic with 100ms delay) and 657ms in $q_{2.3}$ (CQELS on light traffic with 5ms delay). The results are an important indicator for a lower bound of QA over streams. Yet, the results are not fully comparable since CQELS or C-SPARQL respectively, do not support RDFS or OWL2 QL, respectively; there is a trade-off between performance and expressivity, where in CQELS the full TBox is omitted, and in C-SPARQL axioms with existentially-quantified variables on the right-hand side of inclusion assertions are ignored leading to incomplete results. Furthermore, both systems do not support directly spatial relations, aggregates, and predictions, which amount to approximately 100ms evaluation time, since these functions are precomputed by the log player in the experiments, which is not reflected in the AET. The results also indicate that by increasing traffic density, the performance seems to align in $q_{1.1}$, where the grouping/aggregation might be the most demanding operation for C-SPARQL and CQELS.

After profiling the runtime of our prototype, we noticed that approximately 200ms are lost by establishing a connection to PIPELINEDB, which could be mitigated by connection pooling using a persistent connection pool such as provided by PG BOUNCER.¹⁶

¹⁶<http://pgfoundry.org/projects/pgbouncer>

```

1 PREFIX : <http://www.kr.tuwien.ac.at/its/ldm/items#
2 SELECT ?vehicle (AVG(?vall) AS ?as) (MAX(?vall) AS ?ms)
3 FROM STREAM <http://www.kr.tuwien.ac.at/its/streams/vehicles>
4 WHERE
5   {
6     [RANGE 10s STEP 1s]
7     { ?vehicle :observed ?obs . ?obs :speed ?vall .
8       ?vehicle :vehicleMaker ?maker . FILTER
9         ( regex(str(?maker), 'BMW', 'i') ||
10           regex(str(?maker), 'VW', 'i') ) }
11   }
12 GROUP BY ?vehicle
13
14 PREFIX : <http://www.kr.tuwien.ac.at/its/ldm/items#
15 SELECT ?vehicle (AVG(?vall) AS ?as) (MAX(?vall) AS ?ms)
16 WHERE
17   {
18     [RANGE 10s]
19     { ?vehicle :observed ?obs . ?obs :speed ?vall .
20       ?vehicle :vehicleMaker ?maker .
21       FILTER ( regex(str(?maker), 'BMW', 'i') ||
22                 regex(str(?maker), 'VW', 'i') ) } }
23 GROUP BY ?vehicle

```

Figure 7. C-SPARQL encoding (left) and CQELS encoding (right) of query $q_{1.1}$

Query System	(l) with ms delay				(m) with ms delay				(h) with ms delay			
	5	10	50	100	5	10	50	100	5	10	50	100
$q_{1.1}$												
C-SPARQL	0.019	0.032	0.031	0.027	0.132	0.134	0.113	0.110	0.252	0.347	0.324	0.314
CQELS	0.034	0.044	0.043	0.038	0.029	0.038	0.043	0.041	0.056	0.043	0.035	0.043
Spatial-stream QA	0.452	0.430	0.383	0.370	0.470	0.463	0.411	0.379	0.484	0.478	0.434	0.389
$q_{2.3}$												
C-SPARQL	0.045	0.025	0.058	0.070	0.050	0.090	0.048	0.077	0.307	0.253	0.265	0.353
CQELS	0.001	0.003	0.013	0.034	0.001	0.002	0.006	0.017	0.001	0.003	0.009	0.020
Spatial-stream QA	0.658	0.634	0.561	0.542	0.692	0.653	0.574	0.548	0.670	0.660	0.585	0.554

Table 4

Results (t in secs) for query $q_{1.1}$ and $q_{2.3}$ with the scenarios (l)ow, (m)edium, and (h)easy traffic

9. Conclusion and Future Work

This work was sparked by the need of applying spatial-stream query answering as an effort to integrate and access streamed mobility data, e.g., vehicle movements, in a spatial context over the complex C-ITS domain. In [3] we have introduced simple aggregate queries over streams, which often do not suffice to capture more complex C-ITS use cases. In this paper, we presented an extension with temporal relations and numerical/trajectory predictions, which allows us to query complex mobility patterns such as traffic statistics, or complex events such as detecting (potential) accidents. Based on the newly developed scenarios of *traffic statistics*, *event detection*, and *advanced driving assistance systems* (ADAS), we have defined a set of domain-specific features such as trajectory computation, which are matched with the scenarios/use cases to define key requirements. Given the new features, we extended the LDM ontology, the spatial-stream query language, and extended the methods used for query answering accordingly. We also redesigned the architecture and optimized the system by pre-compiling the static query elements and executing stream atoms in parallel. The experimental evaluation provides evidence for an improved performance of approx. 40% and an evaluation time below 700ms. This indicates that poten-

tially the feasibility and efficiency of our approach in the mentioned scenarios is given.

Lessons Learned. The presented approach of spatial-stream query answering is well-suited for data integration and query answering in the C-ITS domain. The concept of a LDM was a good starting point, since it has been developed and standardized by the C-ITS community and was already extended by Netten et al. to an ontology-like model [15]. In particular, Semantic Web Technologies play to their strengths in easily modelling a complex domain such as C-ITS, and allowing the (expert) user to formulate powerful queries on top of the streams that are integrated by the ontology. This can be seen by the new scenario ADAS, where small modifications of the ontology and new queries open up a new application field. However, our approach of using OQA revealed some limitations that are discussed in the expert summary.

Using spatial-stream CQs for capturing the scenarios worked out to our satisfaction in most of the use cases. But as illustrated with use cases S2.4, S2.5, and S3.3, our language reached its limits regarding “usability” and also “expressivity”. If we have a larger set of rules as in S3.3 (even without transitivity), the conversion to unions of CQs becomes cumbersome and inefficient (AET is between 2.46s and 3.34s), hence rule engines such as in Ticker [76] or Laser [77] might be better

1 suited. In *S1.3* and *S2.5*, we can see that qualitative
 2 temporal relations like *before* are convenient, since we
 3 are able to avoid the implicit encoding of temporal rel-
 4 ation using shifted windows of different sizes, which
 5 needs an a-priori knowledge of appropriate window
 6 sizes and the relative positions to each other. This makes
 7 a windows-based handling of temporal relations inflexi-
 8 ble and prone to errors. Furthermore, the usage of CQs
 9 with sub-queries can be directly transferred to SPARQL
 10 queries; the underlying evaluation system would have to
 11 be adjusted to deal with the mobility-specific features.

12 **Outlook.** We believe that stream processing/reasoning
 13 methods could well be applied to the mentioned C-ITS
 14 scenarios, which was confirmed by the experts; they
 15 also acknowledged that the new features such as time
 16 intervals with temporal relations and prediction capa-
 17 bilities are important extensions of the initial spatial-
 18 stream OQA approach.

19 Nevertheless, the experts identified practical and the-
 20 oretical extensions that should be addressed in future
 21 research. On the practical side, they suggested that the
 22 LDM ontology could be (and has been) combined with
 23 the SSN ontology [32], where Collective Perception
 24 Messages [28] could be used to integrate local sensor
 25 data of other vehicles. This could be further elaborated
 26 such that the different LDMs (of each vehicle) could be
 27 aligned to a single global dynamic map. Furthermore,
 28 the experts suggested that we could integrate Kalman
 29 filters [29] and top-k aggregates [33] to provide more
 30 powerful aggregates. On the theoretical side, our meth-
 31 ods could be extended to capture uncertainty on the
 32 level of data items and also of TBox assertions, which
 33 would lead to a change in the underlying semantics and
 34 the computational properties. Furthermore, our meth-
 35 ods could be extended to handle windows in a more
 36 flexible manner by flexibly forgetting data items or ex-
 37 tending their validity into the future. Finally, one expert
 38 suggested that the focus of Scenario 3 could more on
 39 motion planning instead of ADAS.

40 In addition to the suggestions of the experts, we be-
 41 lieve that efforts on following issues would be benefi-
 42 cial: (a) allowing the integration of external (domain-
 43 specific) modules such as functions for advanced tra-
 44 jectory prediction, which would be similar to exter-
 45 nal atoms in ASP [23]; (b) the possibility of analytic
 46 queries over longer periods, hence an extension with
 47 a transient cache and variable window sizes would be
 48 needed; (c) the full integration of OQA with IA rela-
 49 tions, which would include the representation of IA net-
 50 works and the rewriting of subsets of the composition
 51

1 table; and (d) handling more complex queries and rules
 2 while still maintaining scalability, which would bring
 3 our approach closer to Ticker [76] and Laser [77].

4 As discussed in the section above, ongoing and future
 5 research should be directed to extend the languages,
 6 methods, and the platform to fulfill the defined require-
 7 ments, which will enable us to apply our approach and
 8 prototype to more complex scenarios such as advanced
 9 traffic monitoring and motion planing.

10 **Acknowledgements.** This work has been supported by
 11 the Austrian Research Promotion Agency project Loc-
 12 TraffLog (FFG 5886550) and DynaCon (FFG 861263),
 13 as well as by the European Commission through
 14 IoTCrawler (H2020 contract 779852).
 15

16 References

- 17
- 18
- 19
- 20
- 21
- 22
- 23
- 24
- 25
- 26
- 27
- 28
- 29
- 30
- 31
- 32
- 33
- 34
- 35
- 36
- 37
- 38
- 39
- 40
- 41
- 42
- 43
- 44
- 45
- 46
- 47
- 48
- 49
- 50
- 51
- [1] T. Eiter, R. Ichise, J.X. Parreira, P. Schneider and L. Zhao, Deploying Spatial-Stream Query Answering in C-ITS Scenarios, in: *Proc. of EKAW 2018*, 2018, pp. 386–406. https://doi.org/10.1007/978-3-030-03667-6_25.
 - [2] L. Andreone, R. Brignolo, S. Damiani, F. Sommariva, G. Vivo and S. Marco, SAFESPOT Final Report, Technical Report, D8.1.1, 2010, Available at http://www.safespot-eu.org/documents/D8.1.1_Final_Report_-_Public_v1.0.pdf.
 - [3] T. Eiter, J.X. Parreira and P. Schneider, Spatial Ontology-Mediated Query Answering over Mobility Streams, in: *Proc. of ESWC 2017*, 2017, pp. 219–237. https://doi.org/10.1007/978-3-319-58068-5_14.
 - [4] D. Calvanese, G.D. Giacomo, D. Lembo, M. Lenzerini and R. Rosati, Tractable Reasoning and Efficient Query Answering in Description Logics: The *DL-Lite* Family, *Journal of Automated Reasoning* **39**(3) (2007), 385–429. <https://doi.org/10.1007/s10817-007-9078-x>.
 - [5] R. Kontchakov, M. Rodriguez-Muro and M. Zakharyashev, Ontology-Based Data Access with Databases: A Short Course, in: *Reasoning Web. Semantic Technologies for Intelligent Data Access - 9th International Summer School 2013, Mannheim, Germany. Proceedings*, 2013, pp. 194–229. https://doi.org/10.1007/978-3-642-39784-4_5.
 - [6] T. Eiter, J.X. Parreira and P. Schneider, Detecting Mobility Patterns using Spatial Query Answering over Streams, in: *Proc. of Stream Reasoning Workshop 2017*, 2017. <http://ceur-ws.org/Vol-1936/paper-02.pdf>.
 - [7] D.F. Barbieri, D. Braga, S. Ceri, E.D. Valle and M. Grossniklaus, C-SPARQL: a Continuous Query Language for RDF Data Streams, *International Journal of Semantic Computing* **4**(1) (2010), 3–25. <https://doi.org/10.1142/S1793351X10000936>.
 - [8] 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: *ISWC 2011*, 2011, pp. 370–388. https://doi.org/10.1007/978-3-642-25073-6_24.
 - [9] K. Janowicz, A. Haller, S.J.D. Cox, D.L. Phuoc and M. Lefrançois, SOSA: A lightweight ontology for sensors, observations, samples, and actuators, *Journal of Web Semantics* **56** (2019), 1–10. <https://doi.org/10.1016/j.websem.2018.06.003>.

- [10] ETSI EN 302 637-2 (V1.3.2), Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service, Technical Report, ETSI, 2014.
- [11] ETSI EN 302 637-3 (V1.2.2), Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service, Technical Report, ETSI, 2014.
- [12] ETSI TS 103 191-3 (V1.1.1), Intelligent Transport Systems (ITS); Testing; Conformance test specifications for Signal Phase And Timing (SPAT) and Map (MAP); Part 3: Abstract Test Suite (ATS) and Protocol Implementation eXtra Information for Testing (PIXIT), Technical Report, ETSI, 2015.
- [13] ETSI TR 102 863 (V1.1.1), Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Local Dynamic Map (LDM); Rationale for and Guidance on Standardization, Technical Report, ETSI, 2011.
- [14] ETSI EN 302 895 (V1.1.0), Intelligent transport systems - Extension of map database specifications for Local Dynamic Map for applications of Cooperative ITS, Technical Report, ETSI, 2014.
- [15] B. Netten, L. Kester, H. Wedemeijer, I. Passchier and B. Driessen, DynaMap: A Dynamic Map for road side ITS stations, in: *Proc. of ITS World Congress 2013*, 2013. <https://trid.trb.org/view.aspx?id=1322235>.
- [16] H. Shimada, A. Yamaguchi, H. Takada and K. Sato, Implementation and Evaluation of Local Dynamic Map in Safety Driving Systems, *J. Transportation Technologies* **5**(2) (2015), 102–112. <http://dx.doi.org/10.4236/jtts.2015.52010>.
- [17] L. Zhao, R. Ichise, Z. Liu, S. Mita and Y. Sasaki, Ontology-Based Driving Decision Making: A Feasibility Study at Uncontrolled Intersections, *IEICE Trans.* **100-D**(7) (2017), 1425–1439. <https://doi.org/10.1587/transinf.2016EDP7337>.
- [18] M. Stonebraker, U. Çetintemel and S.B. Zdonik, The 8 requirements of real-time stream processing, *SIGMOD Record* **34**(4) (2005), 42–47. <https://doi.org/10.1145/1107499.1107504>.
- [19] D. Dell’Aglío, E.D. Valle, F. van Harmelen and A. Bernstein, Stream reasoning: A survey and outlook, *Data Science, IOS Press* **1**(1–2) (2017), 59–83. <https://doi.org/10.3233/DS-170006>.
- [20] J.F. Allen, Maintaining Knowledge about Temporal Intervals, *Com. ACM* **26**(11) (1983), 832–843. <https://doi.org/10.1145/182.358434>.
- [21] R.H. Güting, Geo-Relational Algebra: A Model and Query Language for Geometric Database Systems, in: *EDBT 1988*, 1988, pp. 506–527. https://doi.org/10.1007/3-540-19074-0_70.
- [22] R. Kontchakov and M. Zakharyashev, An Introduction to Description Logics and Query Rewriting, in: *Reasoning Web. Reasoning on the Web in the Big Data Era - 10th International Summer School 2014, Athens, Greece. Proceedings*, 2014, pp. 195–244. https://doi.org/10.1007/978-3-319-10587-1_5.
- [23] T. Eiter, G. Ianni and T. Krennwallner, Answer Set Programming: A Primer, in: *Reasoning Web. Semantic Technologies for Information Systems, 5th International Summer School 2009, Brixen-Bressanone, Italy, Tutorial Lectures*, 2009, pp. 40–110. https://doi.org/10.1007/978-3-642-03754-2_2.
- [24] A. Bogner, B. Littig and W. Menz, *Das Experteninterview*, VS Verlag für Sozialwissenschaften, 2009. <https://doi.org/10.1007/978-3-322-93270-9>.
- [25] J. Mason, *Qualitative Researching*, SAGE Publications Ltd, 2002.
- [26] E.E. Maccoby and N. A.Maccoby, The interview: A tool of social science, in: *Handbook of social psychology*, Vol. 1, Addison-Wesley, 1954.
- [27] V. Minichiello, R. Aroni and T. Hays, *In-depth Interviewing: Principles, Techniques, Analysis*, 3rd edn, Pearson Australia Group, 2008.
- [28] A. Rauch, F. Klanner, R.H. Rasshofer and K. Dietmayer, Car2X-based perception in a high-level fusion architecture for cooperative perception systems, in: *2012 IEEE Intelligent Vehicles Symposium, IV 2012, Alcal de Henares, Madrid, Spain*, 2012, pp. 270–275. <https://doi.org/10.1109/IVS.2012.6232130>.
- [29] G. Welch and G. Bishop, An Introduction to the Kalman Filter, Technical Report, 95-041, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 1995. <http://www.cs.unc.edu/~welch/kalman/kalmanIntro.html>.
- [30] D. Comer, Ubiquitous B-Tree, *ACM Comput. Surv.* **11**(2) (1979), 121–137. <https://doi.org/10.1145/356770.356776>.
- [31] R.H. Güting and M. Schneider, *Moving Objects Databases*, Morgan Kaufmann, 2005. <https://doi.org/10.1016/B978-0-12-088799-6.X5000-2>.
- [32] A. Haller, K. Janowicz, S.J.D. Cox, M. Lefrançois, K. Taylor, D.L. Phuoc, J. Lieberman, R. García-Castro, R. Atkinson and C. Stadler, The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation, *Semantic Web* **10**(1) (2019), 9–32. <https://doi.org/10.3233/SW-180320>.
- [33] I.F. Ilyas, W.G. Aref and A.K. Elmagarmid, Supporting top-k join queries in relational databases, *The VLDB Journal* **13**(3) (2004), 207–221. <https://doi.org/10.1007/s00778-004-0128-2>.
- [34] S. Brandt, E.G. Kalayci, R. Kontchakov, V. Ryzhikov, G. Xiao and M. Zakharyashev, Ontology-Based Data Access with a Horn Fragment of Metric Temporal Logic, in: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, San Francisco, California, USA.*, 2017, pp. 1070–1076. <https://dl.acm.org/doi/10.5555/3298239.3298397>.
- [35] M. Rodríguez-Muro, R. Kontchakov and M. Zakharyashev, Ontology-Based Data Access: Ontop of Databases, in: *Proc. of ISWC 2013*, 2013, pp. 558–573. https://doi.org/10.1007/978-3-642-41335-3_35.
- [36] T. Eiter, T. Krennwallner and P. Schneider, Lightweight Spatial Conjunctive Query Answering Using Keywords, in: *Proc. of ESWC 2013*, 2013, pp. 243–258. https://doi.org/10.1007/978-3-642-38288-8_17.
- [37] S. Borgwardt, M. Lippmann and V. Thost, Temporalizing rewritable query languages over knowledge bases, *Journal of Web Semantics* **33** (2015), 50–70. <https://doi.org/10.1016/j.websem.2014.11.007>.
- [38] A. Arasu, S. Babu and J. Widom, The CQL continuous query language: semantic foundations and query execution, *The VLDB Journal* **15**(2) (2006), 121–142. <https://doi.org/10.1007/s00778-004-0147-z>.
- [39] D. Calvanese, E. Kharlamov, W. Nutt and C. Thorne, Aggregate queries over ontologies, in: *Proc. of ONISW 2008*, 2008, pp. 97–104. <https://doi.org/10.1145/1458484.1458500>.
- [40] D. Maier, *The Theory of Relational Databases*, Computer Science Press, 1983.
- [41] C.H. Papadimitriou, *Computational complexity*, Academic Internet Publ., 2007.
- [42] M.Y. Vardi, The Complexity of Relational Query Languages (Extended Abstract), in: *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, STOC 1982, ACM, 1982.

- New York, NY, USA, 1982, pp. 137–146. <https://doi.org/10.1145/800070.802186>.
- [43] E.V. Kostylev and J.L. Reutter, Answering Counting Aggregate Queries over Ontologies of the DL-lite Family, in: *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI 2013, AAAI Press, 2013, pp. 534–540. <https://dl.acm.org/doi/10.5555/2891460.2891534>.
- [44] J.F. Allen, Maintaining Knowledge about Temporal Intervals, *Communications of the ACM* **26**(11) (1983), 832–843. <https://doi.org/10.1145/182.358434>.
- [45] R. Dechter, Tree Decomposition Methods (Chapter 9), in: *Constraint Processing*, R. Dechter, ed., The Morgan Kaufmann Series in Artificial Intelligence, Morgan Kaufmann, San Francisco, 2003, pp. 245–269.
- [46] G. Gottlob, N. Leone and F. Scarcello, The complexity of acyclic conjunctive queries, *Journal of the ACM* **48**(3) (2001), 431–498. <https://doi.org/10.1145/382780.382783>.
- [47] G. Gottlob, N. Leone and F. Scarcello, Hypertree Decompositions: A Survey, in: *Mathematical Foundations of Computer Science 2001*, J. Sgall, A. Pultr and P. Kolman, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 37–57. https://doi.org/10.1007/3-540-44683-4_5.
- [48] H. Pérez-Urbina, I. Horrocks and B. Motik, Practical Aspects of Query Rewriting for OWL 2, in: *Proceedings of the 6th International Conference on OWL: Experiences and Directions - Volume 529*, OWLED 2009, CEUR-WS.org, Aachen, Germany, Germany, 2009, pp. 152–159. http://ceur-ws.org/Vol-529/owlled2009_submission_17.pdf.
- [49] R. Rosati and A. Almatelli, Improving Query Answering over DL-Lite Ontologies, in: *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning*, KR 2010, AAAI Press, 2010, pp. 290–300. <https://dl.acm.org/doi/10.5555/3031748.3031786>.
- [50] S. Kikot, R. Kontchakov and M. Zakharyashev, Conjunctive Query Answering with OWL 2 QL, in: *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, KR 2012, AAAI Press, 2012, pp. 275–285. <https://dl.acm.org/doi/10.5555/3031843.3031876>.
- [51] M. Stocker and M. Smith, Owlgres: A Scalable OWL Reasoner, in: *Proc. of OWLED 2008*, 2008. http://ceur-ws.org/Vol-432/owlled2008eu_submission_25.pdf.
- [52] Q. Ren, M.H. Dunham and V. Kumar, Semantic Caching and Query Processing, *IEEE Trans. on Knowl. and Data Eng.* **15**(1) (2003), 192–210. <https://doi.org/10.1109/TKDE.2003.1161590>.
- [53] G. Graefe, Query Evaluation Techniques for Large Databases, *ACM Computing Surveys* **25**(2) (1993), 73–169. <https://doi.org/10.1145/152610.152611>.
- [54] A. Dermaku, T. Ganzow, G. Gottlob, B.J. McMahan, N. Musliu and M. Samer, Heuristic Methods for Hypertree Decomposition, in: *Proc. of MICAI 2008: Advances in Artificial Intelligence*, 2008, pp. 1–11. https://doi.org/10.1007/978-3-540-88636-5_1.
- [55] D. Carney, U. Çetintemel, M. Cherniack, C. Convey, S. Lee, G. Seidman, M. Stonebraker, N. Tatbul and S. Zdonik, Monitoring streams: a new class of data management applications, in: *Proc. of VLDB 2002*, 2002, pp. 215–226. <https://doi.org/10.1016/B978-155860869-6/50027-5>.
- [56] S. Madden, M. Shah, J.M. Hellerstein and V. Raman, Continuously adaptive continuous queries over streams, in: *2002 ACM SIGMOD International Conference on Management of Data*, 2002, pp. 49–60. <https://doi.org/10.1145/564691.564698>.
- [57] A. Bolles, M. Grawunder and J. Jacobi, Streaming SPARQL - Extending SPARQL to Process Data Streams, in: *Proc. of ESWC 2008*, 2008, pp. 448–462. https://doi.org/10.1007/978-3-540-68234-9_34.
- [58] J. Calbimonte, J. Mora and Ó. Corcho, Query Rewriting in RDF Stream Processing, in: *Proc. of ESWC 2016*, 2016, pp. 486–502. https://doi.org/10.1007/978-3-319-34129-3_30.
- [59] D. Barbieri, D. Braga, S. Ceri, E.D. 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. <https://doi.org/10.1109/MIS.2010.142>.
- [60] H.N.M. Quoc and D. Le Phuoc, An Elastic and Scalable Spatiotemporal Query Processing for Linked Sensor Data, in: *Proc. of SEMANTICS 2015*, ACM, 2015, pp. 17–24. <https://doi.org/10.1145/2814864.2814869>.
- [61] D. Anicic, P. Fodor, S. Rudolph and N. Stojanovic, EP-SPARQL: a unified language for event processing and stream reasoning, in: *Proc. of WWW 2011*, 2011, pp. 635–644. <https://dl.acm.org/doi/10.1145/1963405.1963495>.
- [62] H. Beck, M. Dao-Tran, T. Eiter and M. Fink, LARS: A Logic-Based Framework for Analyzing Reasoning over Streams, in: *Proc. of AAAI 2015*, 2015, pp. 1431–1438.
- [63] M. Perry, P. Jain and A.P. Sheth, SPARQL-ST: Extending SPARQL to Support Spatiotemporal Queries, *Geospatial Semantics and the Semantic Web* **12** (2011), 61–86. https://doi.org/10.1007/978-1-4419-9446-2_3.
- [64] M. Koubarakis and K. Kyzirakos, Modeling and Querying Metadata in the Semantic Sensor Web: The Model stRDF and the Query Language stSPARQL, in: *ESWC 2010*, 2010, pp. 425–439. https://doi.org/10.1007/978-3-642-13486-9_29.
- [65] A. Artale, R. Kontchakov, A. Kovtunova, V. Ryzhikov, F. Wolter and M. Zakharyashev, First-Order Rewritability of Temporal Ontology-Mediated Queries, in: *Proc. of IJCAI 2015*, 2015, pp. 2706–2712. <https://dl.acm.org/doi/10.5555/2832581.2832627>.
- [66] A. Pnueli, The Temporal Logic of Programs, in: *Proc. of Annual Symposium on Foundations of Computer Science 1977*, 1977, pp. 46–57. <https://doi.org/10.1109/SFCS.1977.32>.
- [67] S. Klarman and T. Meyer, Querying Temporal Databases via OWL 2 QL, in: *Proc. of RR 2014*, 2014, pp. 92–107. https://doi.org/10.1007/978-3-319-11113-1_7.
- [68] A. Margara, J. Urbani, F. van Harmelen and H.E. Bal, Streaming the Web: Reasoning over dynamic data, *J. Web Semant.* **25** (2014), 24–44. <https://doi.org/10.1016/j.websem.2014.02.001>.
- [69] D. Anicic, S. Rudolph, P. Fodor and N. Stojanovic, Stream reasoning and complex event processing in ETALIS, *Semantic Web* **3**(4) (2012), 397–407. <https://doi.org/10.3233/SW-2011-0053>.
- [70] M. Rinne and E. Nuutila, Constructing Event Processing Systems of Layered and Heterogeneous Events with SPARQL, in: *On the Move to Meaningful Internet Systems: OTM 2014 Conferences - Confederated International Conferences: CoopIS, and ODBASE 2014, Amantea, Italy, Proceedings*, 2014, pp. 682–699. https://doi.org/10.1007/978-3-662-45563-0_42.
- [71] C.L. Forgy, Rete: A fast algorithm for the many pattern/many object pattern match problem, *Artificial Intelligence* **19**(1) (1982), 17–37. [https://doi.org/10.1016/0004-3702\(82\)90020-0s](https://doi.org/10.1016/0004-3702(82)90020-0s).
- [72] E. Friedman-Hill, *Jess in Action: Rule-Based Systems in Java*, Manning Publications, 2003. ISBN 978-1-930-11089-2.

- [73] J.-P. 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. <https://doi.org/10.4018/jswis.2012010103>.
- [74] M. Gebser, R. Kaminski, B. Kaufmann and T. Schaub, Clingo = ASP + Control: Preliminary Report, *CoRR* **1405.3694** (2014). <https://arxiv.org/abs/1405.3694>.
- [75] H. Beck, M. Dao-Tran and T. Eiter, LARS: A Logic-based framework for Analytic Reasoning over Streams, *Artif. Intell.* **261** (2018), 16–70. <https://dl.acm.org/doi/10.5555/2887007.2887205>.
- [76] H. Beck, T. Eiter and C. Folie, Ticker: A system for incremental ASP-based stream reasoning, *Theory and Practice of Logic Programming* **17**(5–6) (2017), 744–763. <https://doi.org/10.1017/S1471068417000370>.
- [77] H.R. Bazoobandi, H. Beck and J. Urbani, Expressive Stream Reasoning with Laser, in: *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, Proceedings, Part I*, 2017, pp. 87–103. https://doi.org/10.1007/978-3-319-68288-4_6.
- [78] Ö.L. Özçep, R. Möller and C. Neuenstadt, A Stream-Temporal Query Language for Ontology Based Data Access, in: *KI 2014: Advances in Artificial Intelligence - 37th Annual German Conference on AI, Stuttgart, Proceedings*, 2014, pp. 183–194. https://doi.org/10.1007/978-3-319-11206-0_18.
- [79] S. Brandt, E.G. Kalayci, V. Ryzhikov, G. Xiao and M. Zakharyashev, Querying Log Data with Metric Temporal Logic, *Journal of Artificial Intelligence Research* **62** (2018), 829–877. <https://doi.org/10.1613/jair.1.11229>.
- [80] Y. Nenov, R. Piro, B. Motik, I. Horrocks, Z. Wu and J. Banerjee, RDFox: A Highly-Scalable RDF Store, in: *The Semantic Web - ISWC 2015 - 14th International Semantic Web Conference, Bethlehem, PA, USA, Proceedings, Part II*, 2015, pp. 3–20. https://doi.org/10.1007/978-3-319-25010-6_1.
- [81] Y. Ren and J.Z. Pan, Optimising ontology stream reasoning with truth maintenance system, in: *Proceedings of the 20th ACM Conference on Information and Knowledge Management, CIKM 2011, Glasgow, United Kingdom, October 24-28, 2011*, 2011, pp. 831–836. <https://doi.org/10.1145/2063576.2063696>.