

Context-aware query derivation for IoT data streams with DIVIDE enabling privacy by design

Mathias De Brouwer^{*}, Bram Steenwinckel, Ziyi Fang, Marija Stojchevska, Pieter Bonte, Filip De Turck, Sofie Van Hoecke and Femke Ongenaë

IDLab, Ghent University – imec, Belgium

E-mails: mrdbrouw.DeBrouwer@UGent.be, Bram.Steenwinckel@UGent.be, Ziye.Fang@UGent.be, Marija.Stojchevska@UGent.be, Pieter.Bonte@UGent.be, Filip.DeTurck@UGent.be, Sofie.VanHoecke@UGent.be, Femke.Ongenaë@UGent.be

Editors: Haridimos Kondylakis, FORTH-ICS, Greece; Praveen Rao, University of Missouri, USA; Kostas Stefanidis, Tampere University, Finland

Solicited reviews: Fotis Aisopos, National Centre of Scientific Research Demokritos, Greece; three anonymous reviewers

Abstract. Integrating Internet of Things (IoT) sensor data from heterogeneous sources with domain knowledge and context information in real-time is a challenging task in IoT healthcare data management applications that can be solved with semantics. Existing IoT platforms often have issues with preserving the privacy of patient data. Moreover, configuring and managing context-aware stream processing queries in semantic IoT platforms requires much manual, labor-intensive effort. Generic queries can deal with context changes but often lead to performance issues caused by the need for expressive real-time semantic reasoning. In addition, query window parameters are part of the manual configuration and cannot be made context-dependent. To tackle these problems, this paper presents DIVIDE, a component for a semantic IoT platform that adaptively derives and manages the queries of the platform's stream processing components in a context-aware and scalable manner, and that enables privacy by design. By performing semantic reasoning to derive the queries when context changes are observed, their real-time evaluation does not require any reasoning. The results of an evaluation on a homecare monitoring use case demonstrate how activity detection queries derived with DIVIDE can be evaluated in on average less than 3.7 seconds and can therefore successfully run on low-end IoT devices.

Keywords: Context-aware query derivation, Internet of Things, cascading reasoning, semantic reasoning, homecare monitoring

1. Introduction

1.1. Background

In the healthcare domain, many applications involve a large collection of Internet of Things (IoT) devices and sensors [40]. Many of those systems typically focus on the real-time monitoring of patients in hospitals, nursing homes, homecare or elsewhere. In such systems, patients and their environment are being equipped with different

^{*}Corresponding author. E-mail: mrdbrouw.DeBrouwer@UGent.be.

devices and sensors for following up on the patients' conditions, diseases and treatments in a personalized, context-aware way. This is achieved by integrating the data collected by the IoT devices with existing domain knowledge and context information. As such, analyzing this combination of data sources jointly allows a system to extract meaningful insights and actuate on them [66].

Integrating and analyzing the IoT data with domain knowledge and context information in a real-time context is a challenging task. This is due to the typically high volume, variety and velocity of the different data sources [2]. To deal with these challenges, semantic IoT platforms can be deployed [11]. They generally contain stream processing components that integrate and analyze the different data sources by continuously evaluating semantic queries. To deploy this, Semantic Web technologies are typically employed: ontologies are designed to integrate and model the data from different heterogeneous sources and its relationships and properties in a common, machine-interpretable format, and existing stream reasoning techniques are used by the data stream processing components [30].

Currently, the configuration and management of queries that run on the stream processing components of a semantic IoT platform are manual tasks that require a lot of effort from the end user. In the typical IoT applications in healthcare, those queries should be context-aware: the context information determines which sensors and devices should be monitored by the query, for example to filter specific events to send to other components for further analysis. For example, a patient's diagnosis in the Electronic Health Record (EHR) determines the monitoring tasks that should be performed in the patient's hospital room, while the indoor location (room) of the patient in a homecare monitoring environment determines which in-home activities can be monitored. Changes in this context information regularly occur. For example, the profile information of patients in their EHR can be updated, or the in-home location of the patient can evolve over time. Hence, the management of the queries should be able to deal with such context changes. Currently, no semantic IoT platform component exists that allows to configure, derive and manage the platform's queries in an automated, adaptive way. Therefore, platforms typically apply one of two existing approaches to achieve this.

The first approach to introduce context-awareness into semantic queries is by defining them in a generic fashion. A generic query uses generic ontology concepts in its definitions to perform multiple contextually relevant tasks. This way, semantic reasoners will reason in real-time on all available streaming, context and domain knowledge data to determine the contextually relevant sensors and devices to which the query is applicable. The advantage of this approach is that such queries are prepared to deal with contextual changes: due to their generic nature, they should not be updated often. However, the disadvantage of highly generic queries is the high computational complexity of the semantic reasoning during their evaluation. This is caused by complex ontologies in IoT domains such as healthcare that require expressive reasoning [16]. In healthcare applications that involve a large number of sensors, it is practically challenging to do this in real-time: queries take longer to evaluate, causing lower performance and difficulty to keep up with the required query execution frequencies. Typically, central components in an IoT platform have more resources and are therefore more likely to overcome this challenge. However, running all queries on central components would require all generated IoT streaming data to be sent over the network, causing the network to be highly congested all the time. In addition, the central server resources would be constantly in use, and local decision making would no longer be possible. Importantly, this would also imply no flexibility in preserving the patient's privacy by keeping sensitive data locally. Looking at local and edge IoT devices to run those generic queries instead, resources are typically lower, making the performance challenges an even bigger issue of the generic query approach.

An alternative approach that can be adopted is installing multiple specific queries on the stream processing components that filter the contextually relevant sensors for one specific task. Evaluating such non-generic queries reduces the required semantic reasoning effort, solving the performance issues of the generic approach. However, this approach even further increases the required manual query configuration and management effort for the end user: whenever the context changes, the queries should be manually updated. This is infeasible to do in practice. As a consequence, current platforms do not apply this approach often and mostly work with generic queries instead.

Moreover, the definition of generic stream processing queries does not contain any means to make the window parameters of the query dependent on the application context and domain knowledge. Currently, an end user should configure these query parameters, and cannot let the system define them based on the data. This can be a problem in some specific monitoring cases. For example, the size of the data window on which a monitoring task such as in-home activity detection should be executed, may depend on the type of task, and therefore be defined in the domain

knowledge. Another example is when the execution frequency of a monitoring task depends on certain contextual events happening in the patient's environment.

In addition, preserving the privacy of the patients is of utmost importance in healthcare systems [1]. In IoT platforms, lots of the data generated by the IoT devices can contain privacy-sensitive information. Depending on where the data processing components are being hosted, this privacy-sensitive data may have to be sent over the IoT network, potentially exposing it to the outside world. Therefore, the IoT data is ideally processed close to where it is generated to reduce the amount of information sent over the network as much as possible. With regards to this, a semantic IoT platform should enable privacy by design [21]: it should allow an end user to build privacy by design into an application by precisely controlling which data is kept locally, and which data is sent over the network.

Finally, a semantic IoT platform component that would solve the aforementioned issues, should also be practically usable. Currently, existing semantic IoT healthcare platforms use semantic reasoners or stream reasoners that are configured with existing sets of generic semantic queries [66]. Defining such queries and ensuring their correctness is a delicate and time-consuming task. Hence, a new component should not introduce a completely different means of defining generic queries, but instead reduce the required changes to these definitions to a minimum. This implies that it should start from the generic definition of stream processing queries. Moreover, the other configuration tasks of the component should also be as minimal as possible to increase overall usability.

1.2. Research objectives and paper contribution

In summary, there is a need for a semantic IoT platform component that fulfills the different requirements tackled in the previous subsection, so that it can be applied in a healthcare data management system. Hence, we set the following research objectives for the design of such an additional semantic IoT platform component:

1. The component should reduce the manual, labor-intensive query configuration effort by managing the queries on the platform's stream processing components in an automated, adaptive and context-aware way.
2. The evaluation of queries managed by the component should be performant, also on low-end IoT edge devices with fewer resources. Network congestion and overuse of central resources should be avoided.
3. The component should allow for the query window parameters to be context-dependent.
4. The component should enable privacy by design: it should allow end users to integrate privacy by design into an application by defining, on different levels of abstraction, which data is kept locally and which parts of the data can be sent over the network.
5. The component should be practically usable, minimizing the effort to integrate it into an existing system.

This paper presents DIVIDE, a semantic IoT platform component that we have designed to achieve the presented research objectives. DIVIDE automatically and adaptively derives and manages the contextually relevant specific queries for the platform's stream processing components, by performing semantic reasoning with a generic query definition whenever contextual changes occur. As a result, the derived queries will efficiently monitor the relevant IoT sensors and devices in real-time, and still do not require any real-time reasoning during their evaluation.

The contribution of this paper is the methodological design and proof-of-concept of the DIVIDE component, fulfilling the requirements associated with the above research objectives. In the paper, DIVIDE is applied and evaluated on a realistic homecare monitoring use case, to demonstrate how it can be used in a practical IoT application context that works with privacy-sensitive information.

1.3. Paper organization

The remainder of this paper is structured as follows. Section 2 discusses some related work. In Section 3, the eHealth use case scenario is further explained, translated into the technical system set-up, and semantically described with an ontology. Section 4 presents a general overview of the DIVIDE system. Further functional and algorithmic details of DIVIDE are provided in Section 5 and Section 6 using the running use case example, while Section 7 zooms in on the technical implementation of DIVIDE. Section 8 describes the evaluation set-up with the different evaluation scenarios and hardware set-up. Results of the evaluations are presented in Section 9, and further discussed in Section 10. Finally, Section 11 concludes the main findings of the paper and highlights future work.

2. Related work

2.1. Semantic Web, stream processing and stream reasoning

Using Semantic Web technologies such as the Resource Description Framework (RDF) and the Web Ontology Language (OWL), heterogeneous data sources can be consolidated and semantically enriched into a machine-interpretable representation using ontologies [11]. An ontology is a model that semantically describes all domain-specific knowledge by defining domain concepts and their relations and attributes. Within RDF, an Internationalized Resource Identifier (IRI) is used to refer to every resource defined in an ontology [31]. Semantic reasoners can interpret semantic data to derive new knowledge based on the definitions in the ontologies. The complexity of the semantic reasoning depends on the expressivity of the underlying ontology [53]. Different ontology languages exist. They range from RDFS, which has the lowest expressivity, to OWL 2 DL, which has the highest expressivity.

RDFox [55] and VLog [72] are state-of-the-art OWL 2 RL reasoners. OWL 2 RL contains all constructs that can be evaluated by a rule engine. These constructs can be expressed by simple Datalog rules. By design, these engines are not able to handle streaming data. However, RDFox can also run on a Raspberry Pi, and any ARM-based IoT edge device in general. In addition, previous research has shown it can also successfully run on a smartphone [48].

Notation3 Logic (N_3) [14] is a rule-based logic that is often used to write down RDF. N_3 is a superset of RDF/Turtle [25], which implies that any valid RDF/Turtle definitions are valid N_3 as well.

Stream Reasoning (SR) [30] state-of-the-art contains three main approaches: Continuous Processing (CP) engines, Reasoning Over Time (ROT) frameworks and Reasoning About Time (RAT) frameworks. CP engines have continuous semantics, high throughput, and low latency but do not perform reasoning. ROT frameworks solve reasoning tasks continuously with high throughput and low latency, but do not consider time. RAT frameworks do consider time in the reasoning task, but may lack reactivity due to the high latency. These various approaches each investigate the trade-off between the expressiveness of reasoning and the efficiency of processing [30].

RDF Stream Processing (RSP) identifies a family of CP engines that solve information needs over heterogeneous streaming data, which is typical in IoT applications. It addresses data variety by adopting RDF streams as data model, and solves data velocity by extending SPARQL with the continuous semantics [65]. Different RSP engines exist, such as C-SPARQL [9], CQELS [46], Yasper [70] and RSP4J [69]. Queries can be registered to these engines that are used to continuously filter the defined data streams. A data window is placed on top of the data stream. Parameters of the window definition include the size of the data window that is added to the query's data model, and the window's sliding step which directly influences the query's evaluation frequency.

RSP-QL [29] is a reference model that unifies the semantics of the existing RSP approaches. RSP has been extended to support ROT in various ways: (i) solutions incorporating efficient incremental maintenance of materializations of the windowed ontology streams [10,44,54,73], (ii) solutions for expressive Description Logics (DL) [47,68], and (iii) a solution for Answer Set Programming (ASP) [52]. More central to ROT is the logic-based framework for analyzing reasoning over streams (LARS) [13] that extends ASP for analytical reasoning over data streams. LASER [12] is a system, based on LARS, that employs a tractable fragment of LARS that ensures uniqueness of models. BigSR [59] employs Big Data technologies (e.g., Apache Spark and Flink) to evaluate the positive fragment of LARS. C-Sprite [17] focuses on efficient hierarchical reasoning to improve the throughput and application on edge devices by efficiently filtering out unnecessary data in the stream. A similar approach to filter out unnecessary streaming data in ASP exists, by investigating the dependency graph of the input data [57]. RDF Event Processing (RSEP) identifies a family of approaches that extend CP over RDF Streams with event pattern matching [4]. RSEP extends RSP with a reactive RAT formalism with limited expressiveness [49]. RSEP-QL [28] is an extension of RSP-QL that incorporates the language features from Complex Event Processing (CEP) [50]. StreamQR [20] rewrites continuous RSP queries to multiple parallel queries, allowing for the support of ontologies that are expressed in the \mathcal{ELHI} logic. The CityPulse project [58] presents the combination of RSP, CEP and expressive reasoning through ASP.

The most advanced attempts to develop expressive Stream Reasoning increased the reasoning expressiveness, but at the cost of limited efficiency. DyKnow [38] and ETALIS [5] combine RAT and ROT reasoning, but perform CP at an extremely slow speed. STARQL [56] is a first step in the right direction because it mixes RAT, and ROT reasoning utilizing a Virtual Knowledge Graph (VKG) approach [75] to obtain CP. Cascading Reasoning [64] was

proposed to solve the problem of expressive reasoning over high-frequency streams by introducing a hierarchical approach consisting of multiple layers. Although several of the presented approaches adopt a hierarchical approach [5,52,56], only a recent attempt has laid the first fundamentals on realizing the full vision of cascading reasoning with Streaming MASSIF [18].

2.2. Semantic IoT platforms and privacy preservation

Today, different IoT platforms exist that extend big data platforms with IoT integrators [22,43]. FIWARE [24] is a platform that offers different APIs that can be used to deploy IoT applications. Sofia2 [61] is a semantic middleware platform that allows different systems and devices to become interoperable for smart IoT applications. SymbIoTe [62] goes a step further and abstracts existing IoT platforms by providing a virtual IoT environment provisioned over various cloud-based IoT platforms. The Agile [35] and BIG IoT [19] platforms focus on flexible IoT APIs and gateway architectures, such as VICINITY [23] and INTER-IoT [36] which also provide an interoperability platform. bIoTope [42] addresses the requirement for open platforms within IoT systems development.

Zooming in on IoT-based healthcare systems, a large number of solutions have risen in the last few years [15,40,41,51]. Jaiswal et al. surveyed 146 healthcare for IoT solutions in recent years, and classified them in five categories: sensor-based, resource-based, communication-based, application-based, and security-based approaches. They identified scalability and interoperability as two big challenges that are yet to be solved by many systems. Especially the latter is a challenge with the heterogeneity of data originating from different sources. This challenge can be solved with Semantic Web technologies.

Focusing on IoT healthcare systems that involve semantic technologies, multiple solutions already exist. For example, in the topic of homecare monitoring, Zgheib et al. [76] proposed a scalable semantic framework to monitor activities of daily living in elderly, to detect diseases and epidemics. The proposed framework is based on several semantic reasoning techniques that are distributed over a semantic middleware layer. It makes use of CEP to extract symptom indicators, which are fed to a SPARQL engine that detects individual diseases. C-SPARQL is then employed on a stream of diseases to detect possible epidemics. While this approach zooms in largely on scalability for this specific use case, it does not offer any flexibility in making the SPARQL and C-SPARQL queries context-aware in a fully automated and adaptive way.

Moreover, Jabbar et al. [39] and Ullah et al. [71] both presented an IoT-based Semantic Interoperability Model that provides interoperability among heterogeneous IoT devices in the healthcare domain. These models add semantic annotations to the IoT data, allowing SPARQL queries to easily extract concepts of interest. However, these illustrative SPARQL queries require manual configuration effort and are not automatically ensuring context-awareness in a dynamic environment. In addition, Ali et al. [3] present an ontology-aided recommendation system to efficiently monitor the patient's physiology based on wearable sensor data while recommending specific, personalized diets. Similarly, Subramaniaswamy et al. [67] present a personalized travel and food recommendation system based on real-time IoT data about the patient's physical conditions and activities. Again, these systems only work with static SPARQL queries to evaluate their system, not achieving context-awareness in an adaptive, dynamic environment.

In summary, many of the presented platforms are adopting a wide range of existing Semantic Web technologies to deal with the challenges associated with real-time IoT applications in complex IoT domains such as healthcare. These platforms typically combine different technologies that involve both stream processing and semantic reasoning components. They all have in common that the queries for the stream processing components are not yet configured and managed in a fully automated, adaptive and context-aware way.

Privacy by design is an approach that states that privacy must be incorporated into networked data systems and technologies, by default [21,45,60]. It approaches privacy from the design-thinking perspective, stating that the data controller of a system must implement technical measures for data regulation by default, within the applicable context. Privacy by design is a broad concept that is more concretely defined through seven principles that can be applied to the design of a system. One of these principles is that the privacy-preserving capabilities should be embedded into the design and architecture of IT systems. Another principle focuses on the importance of keeping privacy user-centric, ensuring that the design always considers the needs and interests of the users. Other principles focus on visibility and transparency, privacy as the default setting, proactive instead of reactive measures, avoiding

unnecessary privacy-related trade-offs, and end-to-end security through the lifecycle of the data. Privacy by design is a key principle of the General Data Protection Regulation (GDPR) of the European Union [45].

3. Use case description and set-up

To demonstrate how DIVIDE can be employed in a semantic IoT network to perform context-aware homecare monitoring, a detailed use case is presented in this section.

3.1. Use case description

The homecare monitoring use case scenario presented in this paper focuses on a rule-based service that recognizes the activities of elderly people in their homes.

Use case background More and more people live with chronic illnesses and are followed up at home by various healthcare actors such as their General Practitioner (GP), nursing organization, and volunteers. Patients in homecare are increasingly equipped with monitoring devices such as lifestyle monitoring devices, medical sensors, localization tags, etc. The shift to homecare makes it important to continuously assess whether an alarming situation occurs at the patient. If an alarm is generated, either automatically or initiated by the patient, a call operator at an alarm center should decide which intervention strategy is required. By reasoning on the measured parameters in combination with the medical domain knowledge, a system could help a human operator with choosing the most optimal intervention strategy.

A core building block of a homecare monitoring solution is an autonomous activity recognition (AR) service that detects and recognizes different in-home activities performed by the patient. Moreover, it should also monitor whether ongoing activities belong to a known regular routine of the patient, so that anomalies in the patient's daily activity pattern can be detected. Such a service could make use of the data collected by the different sensors and devices installed in the patient's home environment, as well as knowledge about AR rules and known routines of the patient. Given the heterogeneous nature of these different data sources, Semantic Web technologies are ideally suited to create this autonomous AR service.

Details of the activity recognition service The use case of routine and non-routine AR has been designed together with the home monitoring company Z-Plus. To properly perform knowledge-driven AR, AR rules should be known by the system. Z-Plus helped us with designing the rules.

An AR rule can be defined as a set of one or more value conditions defined on certain observable properties that are being analyzed for a certain entity type. An observable property is any property that can be measured by a sensor in the patient's environment, e.g., temperature, relative humidity, power consumption, door status (open vs. closed), indoor location, etc. Every sensor analyzes its property for a specific entity. Examples of analyzed entities are a room (e.g., for a humidity sensor), an electrical appliance such as a cooking stove (e.g., for a power consumption sensor), a cupboard (e.g., for a door contact sensor), or even the patient (e.g., for a wearable sensor).

In a realistic home environment with a wide range of sensors installed, many different AR rules will be defined. This makes it highly inefficient to continuously monitor all possible activities that can be recognized in the home, since this would require the continuous monitoring of all sensors that observe a certain property for an entity type associated with at least one rule. Hence, the AR service performs location-dependent activity monitoring: it only observes activities that are relevant to the room that the patient is currently located in. To enable this, an indoor location system should be installed that unambiguously knows the current room of the patient at every point in time. The activities relevant to the current room can be derived by considering all sensors that analyze this room or an entity in the room: all activity rules should be evaluated that have conditions (i) on observable properties that are measured by these sensors, and (ii) that are defined for the same entity type as analyzed by those sensors.

Activities recognized by the AR service should be labeled as belonging to the regular routine of this patient or not. If an ongoing activity in the patient's routine is recognized, the situation is normal and requires no more strict follow-up. Ideally, as long as an activity is going on, location changes in the home are less probable and should therefore be monitored less frequently. However, if an activity outside the routine of the patient is being

detected, more strict location monitoring is required since the situation is abnormal. If necessary, an alarm should automatically be generated by the system. To implement such a system, knowledge on the existing routines of the patient at different times of the day should exist.

Finally, an important requirement of the AR service is that it reduces the information that leaves the patient's home environment to a minimum, as a first step in preserving the patient's privacy. This implies that no actual raw sensor data should be sent over the network. To enable this, the AR service should largely run in-home, so that only the actual outputs such as detected activities are being sent. Obviously, data that is not contained in the HomeLab should always be sent over a secure, encrypted connection.

Running example To facilitate the methodological description of DIVIDE in Sections 4, 5 and 6, consider the following illustrative running example derived from the presented homecare monitoring use case.

Consider a smart home with an indoor location system detecting in which room the patient is present, and an environmental sensor system measuring the relative humidity in every room of the home. The smart home consists of multiple rooms including one bathroom. The patient living in the home has a morning routine that includes showering. To keep it simple, the AR service of the running example consists of a single rule. This rule detects when a person is showering, and is formulated as follows:

A person is showering if the person is present in a bathroom with a relative humidity of at least 57%.

This is a rule with a single condition, defined on a crossed lower threshold for the relative humidity observable property, for the bathroom entity type. Hence, given the presence of a humidity sensor in the patient's bathroom, the showering activity will be monitored by the AR service *if* the patient is located in the bathroom.

3.2. Activity recognition ontology

An Activity Recognition ontology has been designed to support the described use case scenario. This Activity Recognition ontology is linked to the DAHCC (Data Analytics for Healthcare and Connected Care) ontology [63], which is an in-house designed ontology that includes different modules connecting data analytics to healthcare knowledge. Specifically for the purpose of this semantic use case, it is extended with a module `KBActivityRecognition` supporting the knowledge-driven recognition of in-home activities.

The DAHCC ontology contains five main modules. The `SensorsAndActuators` and `SensorsAndWearables` modules describe the concepts that allow defining the observed properties, location, observations and/or actions of different sensors, wearables and actuators in a monitored environment such as a smart patient home. The `MonitoredPerson` and `CareGiver` modules contain concepts for the definition of a patient monitored inside a residence and the patient's caregivers. The `ActivityRecognition` module allows describing the activities performed by a monitored person that are predicted by an AR model.

The DAHCC ontology bridges the concepts of multiple existing ontologies in the data analytics and healthcare domains. These ontologies include SAREF (the Smart Applications REFerence ontology) [26] and its extensions SAREF4EHAW (SAREF extended with concepts of the eHealth Ageing Well domain) [37], SAREF4BLDG (an extension for buildings and building spaces) and SAREF4WEAR (an extension for wearables), as well as the Execution-Executor-Procedure (EEP) ontology [33].

Listing 1 shows how a knowledge-based AR model can be defined and configured. In the example, it is configured according to the use case's running example, i.e., with one activity rule for showering. Lines 13–17 of this listing contain the definition of the single condition of this rule.

In Section A.1 of Appendix A, additional listings detail multiple other definitions within the Activity Recognition ontology that support the knowledge-driven AR use case and its running example. This includes the ontological definitions that can be used by a semantic reasoner to define whether an activity prediction corresponds to a person's routine, as well as the semantic description of the example patient and home in the running use case example.

```

1 # define knowledge-based activity recognition model and its config with a specific rule
2 :KBActivityRecognitionModel rdf:type ActivityRecognition:ActivityRecognitionModel ;
3   eep:implements :KBActivityRecognitionModelConfig1 .
4 :KBActivityRecognitionModelConfig1 rdf:type ActivityRecognition:Configuration ;
5   :containsRule :showering_rule .
6
7 # define showering activity rule: detected by one specific condition
8 :showering_rule rdf:type :ActivityRule ;
9   ActivityRecognition:forActivity [ rdf:type ActivityRecognition:Showering ] ;
10  :hasCondition :showering_condition .
11
12 # define showering condition: relative humidity in the bathroom should be at least 57%
13 :showering_condition rdf:type :RegularThreshold ;
14   :forProperty [ rdf:type SensorsAndActuators:RelativeHumidity ] ;
15   Sensors:analyseStateOf [ rdf:type SensorsAndActuators:BathRoom ] ;
16   :isMinimumThreshold "true"^^xsd:boolean ;
17   saref-core:hasValue "57"^^xsd:float .
18
19 # define in system that conditions can be defined on relative humidity in a room
20 SensorsAndActuators:RelativeHumidity rdfs:subClassOf :ConditionableProperty .
21 SensorsAndActuators:Room rdfs:subClassOf :AnalyzableForCondition .

```

Listing 1. Example of how a knowledge-based AR model with an activity rule for showering can be described through triples in the KBActivityRecognition ontology module. All definitions are listed in RDF/Turtle syntax. To improve readability, the KBActivityRecognition: prefix is replaced by the : prefix.

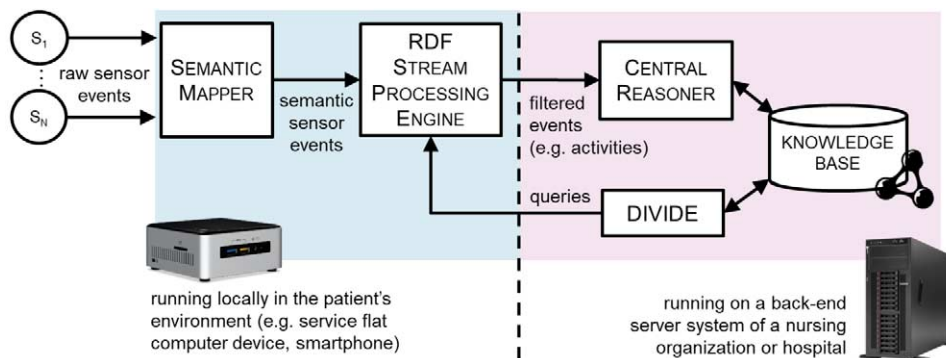


Fig. 1. Architectural set-up of the eHealth use case scenario.

3.3. Architectural use case set-up

To implement the use case scenario of a knowledge-driven routine and non-routine AR service, a cascading reasoning architecture is used [27]. An overview of the architectural cascading reasoning set-up for this use case is shown in Fig. 1. This architecture is generic and can be applied to different use case scenarios in the healthcare domain with similar requirements.

The architecture of the system is split up in a local and a central part. The local part consists of a set of components that are running on a local device in the patient's environment. This device could be any existing gateway that is already installed in the patient's home, such as the device for a deployed nurse call system. The local components are the Semantic Mapper and an RSP Engine. The components of the central part are deployed on a back-end server of an associated nursing home or hospital. They consist of a Central Reasoner, DIVIDE, and a Knowledge Base.

Knowledge Base The Knowledge Base contains the semantic representation of all domain knowledge and context data in the system, in an RDF-based knowledge graph. In the given use case scenario, this domain knowledge consists of the Activity Recognition ontology that is discussed in Section 3.2. It includes the AR model with its activity rules. The contextual information describes the different smart homes and their installed sensors, and patients.

Semantic Mapper The Semantic Mapper semantically annotates all raw observations generated by the sensors in the patient's environment. These semantic sensor observations are forwarded to the data streams of the RSP Engine.

RSP Engine The RSP engine continuously evaluates the registered queries on the RDF data streams, to filter relevant events. In this use case scenario, the filtered events are in-home locations and recognized activities both in and not in the patient's routine. Only these filtered events are encrypted and sent over the network to the Central Reasoner. By applying the cascading reasoning principles and installing the RSP Engine locally in the patient's environment, a first step in preserving the patient's privacy can be taken.

Central Reasoner The Central Reasoner is responsible for further processing the events received from the RSP Engine, and acting upon them. For example, it can aggregate the filtered events and save them to use for future call enrichment, or send an alarm to the patient's caregivers when necessary. In general, any action is possible, depending on what additional components are deployed and implemented on the central node.

Importantly, the Central Reasoner will also update relevant contextual information in the Knowledge Base, such as events occurring in the patients' environment. This information can then trigger a re-evaluation of the queries deployed on the local RSP engines. In the given use case scenario, relevant context changes that trigger a possible change in the deployed RSP queries are location updates and detected activities. When the in-home location of the patient changes, the set of activities that need to be monitored changes as well, since the AR service is location-dependent. Moreover, context information about recognized ongoing routine and non-routine activities directly defines the execution frequency of the location monitoring RSP query.

DIVIDE DIVIDE is the component that manages the queries executed by the local RSP Engine components. It updates the queries whenever triggered by context updates in the Knowledge Base. By aggregating contextual information with medical domain knowledge through semantic reasoning during the query derivation, the resulting RSP queries only involve filtering and do not require any more real-time reasoning. Moreover, it allows to dynamically manage the window parameters of the queries (i.e., the size of the data window and its sliding step) based on the current context. It is fully automated and adaptive, so that at all times, relevant queries are being executed given the context information about the patients in the Knowledge Base.

In the running example, DIVIDE will ensure that there is always a location monitoring query running on the RSP Engine component installed in the patient's home. The window parameters of this query will depend on whether or not an activity is currently going on, and whether or not this activity belongs to the current patient's routine. In addition, when the patient is located in the bathroom, an additional RSP query will be derived and installed that monitors when the patient is showering. When the query detects this activity, this would be considered a recognized routine activity as showering is included in the patient's morning routine.

4. Overview of the DIVIDE system

In Section 3, the general cascading reasoning architecture of the semantic system in the eHealth use case scenario is explained. This section zooms in on DIVIDE, the architectural component responsible for managing the queries running on the local RSP Engine components. It is the task of the DIVIDE system to ensure that these queries perform the relevant filtering given the current context, at any given time, for every RSP Engine known to DIVIDE.

The methodological design of DIVIDE contains of two main pillars: (i) the initialization of DIVIDE, involving the DIVIDE query parsing and ontology preprocessing steps, and (ii) the core of DIVIDE which is the query derivation. Figure 2 shows a schematic overview of the action steps, inputs and internal assets DIVIDE, in which the two main pillars can be distinguished. The following two sections, Section 5 and Section 6, provide more information on this initialization and query derivation, respectively. Throughout the descriptions of DIVIDE in these sections, the running eHealth use case example described in Section 3.1 is considered.

In terms of logic, DIVIDE works with the rule-based Notation3 Logic (N_3) [14]. The semantic reasoner used within DIVIDE should thus be a reasoner supporting N_3 . Such a reasoner can reason within the OWL 2 RL profile [53], which implies that a semantic system that uses DIVIDE in combination with an RSP engine is equivalent to a set-up involving a semantic OWL 2 RL reasoner. The reasoner should support the generation of all triples based on a set of input triples and rules, as well as generating a proof towards a certain goal rule. Such a proof should contain the chain of all rules used by the reasoner to infer new triples based on its inputs, described in N_3 logic.

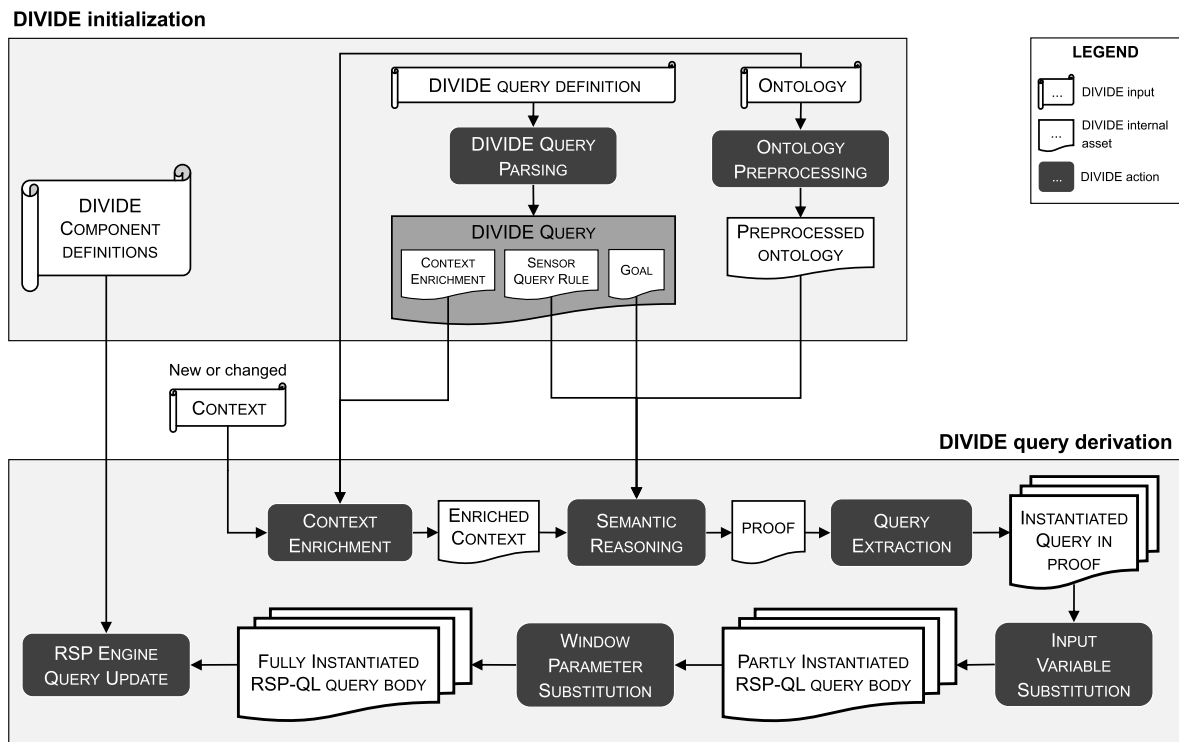


Fig. 2. Schematic overview of the DIVIDE system. It shows all actions that can be performed by DIVIDE with their inputs and outputs. A distinction is made between internal assets and external inputs to the system. The overview is split up in the two major parts: the inputs, steps and assets of the DIVIDE initialization, and those of the DIVIDE query derivation.

5. Initialization of the DIVIDE system

The core task of DIVIDE is the derivation and management of the queries running on the RSP engines of the semantic components in the system that are known to DIVIDE. To allow DIVIDE to effectively and efficiently perform the query derivation for one or more components upon context changes, different initialization steps are required. Three main steps can be distinguished from the upper part of the DIVIDE system overview in Fig. 2: (i) parsing and initializing the DIVIDE queries, (ii) preprocessing the system ontology, and (iii) initializing the DIVIDE components. This section zooms in on each of these three initialization tasks.

5.1. Initialization of the DIVIDE queries

A DIVIDE query is a generic definition of an RSP query that should perform a real-time processing task on the RDF data streams generated by the different local components in the system. The goal of DIVIDE is to instantiate this query in such a way that it can perform this task in a single query that simply filters the RDF data streams. To this end, the internal representation of a DIVIDE query contains a goal, a sensor query rule with a generic query pattern, and a context enrichment. These three items are essential for correctly deriving the relevant queries during the query derivation process. They will each be explained in detail in the first three subsections of this section.

In the running example, there is one RSP query that actively monitors the location of the patient in the home, and one query that detects a showering activity when the patient is located in the bathroom. This subsection will focus on the latter, which is an example of an actual AR query. Within DIVIDE, a generic DIVIDE query will be defined for each *type* of activity rule present in the system. This means that no dedicated DIVIDE query per activity should be defined, which would be too cumbersome and highly impractical in a real-world

```

{
  ?p rdf:type KBActivityRecognition:RoutineActivityPrediction ;
  ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
  ActivityRecognition:activityPredictionMadeFor ?patient ;
  ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?t .
  ?activityType rdfs:subClassOf KBActivityRecognition:DetectableActivity .
} => {
  _:p rdf:type KBActivityRecognition:RoutineActivityPrediction ;
  ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
  ActivityRecognition:activityPredictionMadeFor ?patient ;
  ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?t .
} .

```

Listing 2. Goal of the generic DIVIDE query detecting an ongoing activity in a patient's routine.

deployment. A rule type is a specific combination of conditions and the type of value they are defined on. For the showering rule, this means that the type is defined as follows: a rule with a single condition on a lower regular threshold that should be crossed. This means that the detailed specific RSP queries corresponding to activity rules of the same type will all be derived from the same generic DIVIDE query. The generic DIVIDE query corresponding to the type of the showering activity rule will be used as the running example DIVIDE query in this section. Note that the running example will only focus on the detection of this activity *in* the patient's routine.

5.1.1. Goal

The goal of a DIVIDE query defines the semantic output that should be filtered by the resulting RSP query. This required query output is translated to a valid N_3 rule. This rule is used in the DIVIDE query derivation to ensure that the resulting RSP query is filtering this required RSP query output.

For the generic query definition corresponding to the RSP query that detects the showering activity in the running example, the goal is specified in Listing 2. It is looking for any instance of a `RoutineActivityPrediction`.

5.1.2. Sensor query rule with generic query pattern

The sensor query rule is the core of the DIVIDE query definition. It is a complex N_3 rule that defines the generic pattern of the RSP query, together with semantic information on when and how to instantiate it. Its usage by the semantic rule reasoner during the DIVIDE query generation defines whether or not this generic query should be instantiated given the involved context.

The formalism of the sensor query rule builds further on SENSdesc, which is the result of previous research [6]. This theoretical work was the first step in designing a format that describes an RSP query in a generic way that can be combined with formal reasoning to obtain the relevant queries that filter patterns of interest. By generalizing this format and integrating it into DIVIDE, it has become practically usable.

Each sensor query rule consists of three main parts: the relevant context in the rule's antecedence, and the generic query and ontology consequences defined in the rule's consequence.

Relevant context In the antecedence of the sensor query, the context in which the generic RSP query might become relevant is generically described. For each set of query variables for which the antecedence is valid, there is a chance that the rule, instantiated with these query variables, will appear in the proof constructed by the semantic reasoner during the query derivation. If this is the case, the query will be instantiated for this set of variables.

To explain the different parts, consider the DIVIDE query corresponding to the running example detecting the showering activity. Listing 3 defines the sensor query rule for the corresponding type of activity rule. The rule's antecedence with the relevant context of the sensor query rule is described in lines 2–24. In short, it looks for AR rules relevant to the current room of the patient, following the definition of location-dependent activity monitoring in Section 3.1.

```

1  {
2    ?model rdf:type ActivityRecognition:ActivityRecognitionModel ;
3      <https://w3id.org/eep#implements> [ rdf:type ActivityRecognition:Configuration ;
4        ActivityRecognition:containsRule ?a ] .
5    ?a rdf:type KBActivityRecognition:ActivityRule ;
6      ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
7      KBActivityRecognition:hasCondition [
8        rdf:type KBActivityRecognition:RegularThreshold ;
9        KBActivityRecognition:isMinimumThreshold "true"^^xsd:boolean ;
10       saref-core:hasValue ?threshold ;
11       Sensors:analyseStateOf [ rdf:type ?analyzed ] ;
12       KBActivityRecognition:forProperty [ rdf:type ?prop ]
13     ] .
14
15    ?activityType rdfs:subClassOf KBActivityRecognition:DetectableActivity .
16
17    ?sensor rdf:type saref-core:Device ; saref-core:measuresProperty ?prop_o ;
18      Sensors:isRelevantTo ?room ; Sensors:analyseStateOf [ rdf:type ?analyzed ] .
19    ?prop_o rdf:type ?prop .
20
21    ?prop rdfs:subClassOf KBActivityRecognition:ConditionableProperty .
22    ?analyzed rdfs:subClassOf KBActivityRecognition:AnalyzableForCondition .
23
24    ?patient MonitoredPerson:hasIndoorLocation ?room .
25  }
26 =>
27 {
28   _:q rdf:type sd:Query ;
29     sd:pattern sd-query:pattern ;
30     sd:inputVariables ((" ?sensor" ?sensor) (" ?threshold" ?threshold)
31       (" ?activityType" ?activityType)
32       (" ?patient" ?patient) (" ?model" ?model) (" ?prop_o" ?prop_o)) ;
33     sd>windowParameters ((" ?range" 30 time:seconds) (" ?slide" 10 time:seconds)) .
34
35   _:p rdf:type ActivityRecognition:ActivityPrediction ;
36     ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
37     ActivityRecognition:activityPredictionMadeFor ?patient ;
38     ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp _:t .
39 } .
40
41 sd-query:pattern rdf:type sd:QueryPattern ;
42 sh:prefixes sd-query:prefixes-activity-showering ;
43 sh:construct """
44   CONSTRUCT {
45     _:p a KBActivityRecognition:RoutineActivityPrediction ;
46     ActivityRecognition:forActivity [ a ?activityType ] ;
47     ActivityRecognition:activityPredictionMadeFor ?patient ;
48     ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?now .
49   }
50   FROM NAMED WINDOW :win ON <http://protego.ilabt.imec.be/idlab.homelab>
51   [RANGE ?{range} STEP ?{slide}]
52   WHERE {
53     BIND (NOW() as ?now)
54     WINDOW :win { ?sensor saref-core:makesMeasurement [
55       saref-core:hasValue ?v ; saref-core:hasTimestamp ?t ;
56       saref-core:relatesToProperty ?prop_o ] .
57       FILTER (xsd:float(?v) > xsd:float(?threshold)) }
58   }
59   ORDER BY DESC(?t) LIMIT 1""" .
60
61 sd-query:prefixes-activity-showering rdf:type owl:Ontology ;
62 sh:declare [ sh:prefix "xsd" ;
63   sh:namespace "http://www.w3.org/2001/XMLSchema#"^^xsd:anyURI ] ;
64 sh:declare [ sh:prefix "saref-core" ;
65   sh:namespace "https://saref.etsi.org/core/"^^xsd:anyURI ] ;
66 sh:declare [ sh:prefix "ActivityRecognition" ; sh:namespace
67   "https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/"^^xsd:anyURI ] ;
68 sh:declare [ sh:prefix "KBActivityRecognition" ;
69   sh:namespace "https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/
70   KBActivityRecognition/"^^xsd:anyURI ] .

```

Listing 3. Sensor query rule of the generic DIVIDE query detecting an ongoing activity in a patient's routine, for an activity rule type with a single condition on a certain property of which a value should cross a lower threshold.

Generic query The generic query definition is contained inside the consequence of the sensor query rule. It consists of three main aspects: the generic query pattern, its input variables, and its static window parameters.

The generic query pattern is a string representation of the actual RSP-QL query that will be the result of the DIVIDE query derivation. This pattern is however still generic: some of its query variables still need to be substituted by actual values to obtain the correct and valid RSP-QL query. Similarly, the window parameters of the input stream windows of the RSP-QL query also need to be substituted.

The input variables that need to be substituted by the semantic reasoner in the generic query pattern are defined as a N_3 list. Every item in this list represents one input variable. This input variable is a list itself as well: the first item represents the string literal of the variable in the generic query pattern to be substituted, the second item is the query variable that should occur in the sensor query rule's antecedence so that it is instantiated by the semantic reasoner if the rule is applied in the proof during the query derivation.

Similarly, the definition of the static window parameters is also a list of lists. Static window parameters are variables that should also be substituted by the semantic reasoner during the query derivation, but in the stream window definition instead of the query body or output. They are static as their value is directly defined by the value of the corresponding variable. Every item of the outer list is an inner list of three items. The first item represents the string literal of the variable in a window definition of the generic query pattern. The second item can either be a query variable or literal defining the value of the window parameter. If this is a query variable, it will be substituted during the rule evaluation based on the matching value in the rule's antecedence, similarly to the input variables. The third item defines the unit of the value.

In Listing 3, the generic query definition is described in lines 28–33 and lines 41–69. More specifically, lines 28–29 and lines 41–69 define the generic query pattern, whereas lines 30–32 and line 33 define the input variables and static window parameters of the generic query, respectively.

Inspecting the example in Listing 3 in further detail, the generic RSP-QL query pattern string is defined in lines 44–59. The query filters observations on the defined stream data window `:win` of a certain sensor `?sensor` with a value for the observed property `?prop_o` that is higher than a certain threshold `?threshold` (WHERE clause in lines 52–58). For every match of this pattern, output triples are constructed that represent an ongoing activity of type `?activityType` in the routine of a patient `?patient`, predicted by the activity recognition model `?model` (CONSTRUCT clause in lines 44–49). These six variables are exactly the six input variables as defined in lines 30–32: their values will be instantiated during the query derivation. Note that the window parameter definitions specified in line 33 of Listing 3 define a window size of 30 seconds and a window sliding step of 10 seconds.

Ontology consequences The ontology consequences are the second main part of the sensor query rule's consequence. This part describes the *direct* effect of a query result in a real-time reasoning context. This effect is obtained when a stream window of the generic RSP query would fulfill the pattern of the WHERE clause but no additional reasoning has been done (yet) to know the *indirect* consequences of this matching pattern. This is an essential aspect to understand: the purpose of DIVIDE is to derive queries that can make conclusions that are valid with the given context, through a single RSP-QL query without any reasoning involved. In a context without DIVIDE, these same *indirect* conclusions could only be made by performing an additional semantic reasoning step, based on the *direct* conclusions that are directly known from the matching query pattern. In other words, the triples defining the ontology consequences *can* be the same as the output of the generic RSP-QL query and thus the consequence of the rule representing the DIVIDE query's goal. However, in practice, it will often require an additional semantic reasoning step to see whether the ontology consequences actually imply the output of the generic RSP-QL query.

In the running example, the *direct* consequences of a sensor observation matching the WHERE clause in lines 52–58 of Listing 3 would be the fact that an ongoing activity of the given type is detected for the given patient (lines 35–38). The *indirect* consequences represented by the definitions in the RSP-QL query output (lines 45–48) state that this is an activity *in the patient's routine*.

5.1.3. Context enrichment

Prior to the start of the query derivation with the semantic reasoner, the current context can still be enriched by executing one or more SPARQL queries on this context. The context enrichment of a DIVIDE query consists of this ordered set of valid SPARQL queries.

It is important to note that context-enriching queries are not only used to add general context to the model, but also for the dynamic window parameter substitution as will be explained in Section 6.5.

For the running example, no context-enriching queries are part of the DIVIDE query definition. However, Appendix A discusses the definition of a related DIVIDE query that does include a context enrichment.

5.1.4. DIVIDE query parser

As an end user of DIVIDE, it is not required to define a DIVIDE query according to its internal representation to properly initialize DIVIDE. Instead, the recommended way to define a DIVIDE query is by specifying an ordered collection of existing SPARQL queries that are applied in an existing rule-based stream reasoning system, or through an already existing RSP-QL query. Through DIVIDE, this set of ordered queries will be replaced in the semantic platform by a single RSP query that performs a semantically equivalent task. To enable this, DIVIDE contains a query parser, which converts such an external DIVIDE query definition into its internal representation. The goal of this approach is to make it easy for an end user to integrate DIVIDE into an existing semantic (stream) reasoning system, without having to know the details of how DIVIDE works.

DIVIDE is applied in a cascading system architecture. It considers its equivalent regular (stream) reasoning system as a semantic reasoning engine in which the set of SPARQL queries is executed sequentially on a data model containing the ontology (TBox) triples and rules, context (ABox) triples, and triples representing the sensor observations in the data stream. Each query in the ordered collection, except for the final one, should be a CONSTRUCT query, and its outputs are added to the data model on which (incremental) rule reasoning is applied before the next query in the chain is executed.

The definition of a DIVIDE query as an ordered set of SPARQL queries includes a context enrichment with zero or more context-enriching queries, exactly one stream query, zero or more intermediate queries, and either no or exactly one final query. Besides these queries, such a DIVIDE query definition also includes a set of stream windows (required), a solution modifier (optional), and a variable mapping from stream query to final query (optional). The remainder of this subsection will discuss these different inputs in this DIVIDE query definition.

Stream query and context enrichment In the ordered set of SPARQL queries, it is important that there is exactly one query that reads from the stream(s) of sensor observations. This query is called the *stream query*. In some cases, this query will be the first in the chain. If this is not the case, any preceding queries are defined as *context-enriching queries* in the DIVIDE query definition. Importantly, the WHERE clause conditions of the stream query should be part of named graphs defined as data inputs with a FROM clause, except for special SPARQL constructs such as a FILTER or BIND clause. The IRIs of the named graphs are used to distinguish which data is considered as part of the context, and which data will be put on the data stream. For the data streams, the named graph IRI should reflect the stream IRI. This stream IRI should also be defined as a stream window.

Final query The final query in the ordered set of SPARQL queries is called the *final query* in DIVIDE. A final query is optional: if it does not exist, the stream query is considered the final query.

Intermediate queries The intermediate queries are an ordered list of zero or more SPARQL queries. This list contains those queries in the original set of SPARQL queries that are executed between the stream and final query.

Stream windows Each data stream window that should be included as input in the resulting RSP-QL query should be explicitly defined. It consists of a stream IRI, a window definition, and a set of default window parameter values.

The stream IRI represents the IRI of the data stream. This IRI should exactly match the name of a named graph defined in the stream query. The window definition defines the specification of how the windows are created on the stream. If the user wants to define variable window parameters, named variables should be inserted into the places that will be instantiated during the query derivation. In DIVIDE, two types of variable window parameters exist: static and dynamic window parameters. Static window parameters *might* be substituted similarly to an input

variable during the DIVIDE query derivation. Hence, the variable name of this window parameter should appear in the WHERE clause of the stream query, in a named graph that is *not* corresponding to a stream window. This will ensure that the variable name can be substituted as a regular input variable. During the DIVIDE query derivation, dynamic window parameters are substituted before static parameters. A dynamic window parameter can be defined in the output of a context-enriching query. In case no context-enriching query yields a value for the dynamic window parameter variable, the value of the static window parameter with the same variable name will be substituted. If no such static window parameter is defined, a default value will be used. Hence, for each such variable in the window definition that is not defined as a static window parameter, this default value should be defined by the end user.

Solution modifier If the resulting RSP-QL query should have a SPARQL solution modifier, this can be included in the DIVIDE query definition. Any unbound variable names in the solution modifier should be defined in a named graph of the stream query's WHERE clause that represents a stream window.

Variable mapping of stream to final query If a final query is specified, it often occurs that certain query variables in both the stream and final query actually refer to the same individuals. To make sure that DIVIDE parses the DIVIDE query input correctly, the mapping of these variable names should be explicitly defined. This is a manual required step. Often, they will have the same variable names, making this mapping trivial.

Parsing the end user definition of a DIVIDE query to its internal representation The DIVIDE query parser can construct the goal, sensor query rule and context enrichment of a DIVIDE query from its end user definition. The context enrichment requires no parsing, while the goal and sensor query rule are composed from the different inputs.

The goal of the DIVIDE query is directly constructed from the final query. If it is a CONSTRUCT query, the content of the WHERE clause is put in the antecedence of the goal, while the content of the CONSTRUCT clause represents the goal's consequence. For any other query form, the WHERE clause of the final query is used for both the goal's antecedence and consequence. If no final query is available, the antecedence and consequence of the goal are copied from the result of the stream query. If the stream query is no CONSTRUCT query, the SELECT, ASK or DESCRIBE result clause is first converted to a triple pattern containing all its unbound variables.

The sensor query rule is the most complex part to construct. In the standard case, disregarding any exceptions, the antecedence of the rule is composed from all named graph patterns in the WHERE clause of the stream query that do *not* represent a stream graph. The ontology consequences in the consequence of the sensor query rule are copied from the stream query's output. The generic RSP-QL query pattern is constructed from different parts. Its resulting CONSTRUCT, SELECT, ASK or WHERE clause is directly copied from the result clause of the final query, or the stream query if no final query is present. Its input stream window definitions are constructed using the defined stream windows. The WHERE clause contains the content of the stream graphs in the stream query's WHERE clause, and the special SPARQL patterns that are not put inside a named graph pattern. If a solution modifier is specified, it is appended to the generic RSP-QL query pattern. The input variables and window parameters of the sensor query rule are derived by analyzing the stream query, final query and the variable mapping between both. Any intermediate queries are converted to additional semantic rules that are appended to the main sensor query rule.

Finally, it is worth noting that a DIVIDE query can alternatively also be defined through an existing RSP-QL query. Such a definition is quite similar to the definition described above, with a few differences. The main difference is that by definition, no intermediate and final queries will be present since the original system already uses RDF stream processing and individual RSP-QL queries. This means no variable mapping should be defined either. Hence, this definition is typically more simple than the definition of a DIVIDE query as a set of SPARQL queries.

For the running use case example, the DIVIDE query that performs the monitoring of the showering activity rule can be defined as a set of ordered SPARQL queries. The DIVIDE query parser will translate this definition into the internal representation of this DIVIDE query, exactly as discussed in the previous subsections. This end user definition is discussed in detail in Section A.2 of Appendix A.

5.2. Initialization of the DIVIDE ontology

To properly perform the query derivation, an ontology should be specified as input to DIVIDE by the end user. During initialization, this ontology will be loaded into the system. By definition, this ontology is considered not

to change often during the system's lifetime, in contrast with the context data. Therefore, the ontology should be preprocessed by the semantic reasoner wherever possible. This will speed up the actual query derivation process, since it avoids that the full ontology is loaded and processed every time the DIVIDE query derivation is triggered. To what extent the ontology can be preprocessed depends on the semantic reasoner used.

For the running example, the triples and axioms in the `KBActivityRecognition` module of the Activity Recognition ontology are preprocessed, including the definitions in all its imported ontologies.

5.3. Initialization of the DIVIDE components

To properly initialize DIVIDE, it should have knowledge about the components it is responsible for. A component is defined as an entity in the IoT network on which a single RSP engine runs. For each DIVIDE component, the following information should be specified by an end user for the correct initialization of DIVIDE:

- The name of the graph (ABox) pattern in the knowledge base that contains the context specific for the entity that this component's RSP engine is responsible for. A typical example in the eHealth scenario is a graph pattern of a specific patient, containing all patient information.
- A list of any additional graph patterns in the knowledge base that contain context relevant to the entity that this component's RSP engine is responsible for. An example is generic information on the layout of the environment in which the patient's smart home is situated. Such context information is relevant to multiple components, and is therefore stored in separate graphs in the knowledge base.
- The type of the RSP engine of this component (e.g., C-SPARQL).
- The base URL of the RSP engine's server API. This API should support registering and unregistering RSP queries, and pausing and restarting an RSP stream. It will be used during the DIVIDE query derivation.

Upon initialization, all component information is processed and saved by DIVIDE. For every graph pattern associated with at least one component, DIVIDE should actively monitor for any updates to this ABox in the knowledge base, to trigger the query derivation for the relevant components when updates occur.

6. DIVIDE query derivation

Whenever DIVIDE is alerted of a context change in the knowledge base, the DIVIDE query derivation is triggered for every DIVIDE query. Based on the name of the updated ABox graph and the components known by the system, DIVIDE knows for which components the query derivation process should be started. This process can be executed independently, i.e., in parallel, for each combination of component and DIVIDE query. Hence, this section will focus on the query derivation task for a single component and a single DIVIDE query.

The DIVIDE query of the running example, that performs the monitoring of the showering activity rule, will be further used in this section to illustrate the query derivation process. The query derivation is triggered if any relevant context for a given component is updated. For this example, this context consists of all information about the patient and the smart home. Moreover, it also contains the output of the RSP queries: the in-home patient location and the detected ongoing activities.

The DIVIDE query derivation task for one RSP engine and one DIVIDE query consists of several steps, which are executed sequentially: (i) enriching the context, (ii) semantic reasoning on the enriched context to construct a proof containing the details of derived queries and how to instantiate them, (iii) extracting these derived queries from the proof, (iv) substituting the instantiated input variables in the generic RSP-QL query pattern for every derived query, (v) substituting the window parameters in a similar way, and (vi) updating the active RSP queries on the corresponding RSP engine. The input of the query derivation is the updated context, which consists of the set of triples in the context graph(s) of the knowledge base that are associated with the given component's RSP engine. In the following subsections, the DIVIDE query derivation action steps are further detailed. Figure 2 shows a schematic overview of these steps on the bottom part. For every step, the inputs and outputs are detailed on the figure.

6.1. Context enrichment

Prior to actually deriving the RSP queries for the given DIVIDE query, the context data model can still be enriched. This is done by executing the ordered set of context-enriching queries corresponding to the DIVIDE query with a SPARQL query engine, if there are any, possibly after performing rule-based reasoning with the ontology axioms. The result of this step is a data model containing the original context triples and all triples in the output of any of the context-enriching queries, if there are any. Note that the output of the context-enriching queries can also contain dynamic window parameters to be used in the window parameter substitution step of the query derivation.

The generic DIVIDE query corresponding to the running example of detecting the showering activity does not contain any context-enriching query. Hence, the updated context will directly be sent to the input of the next step. In Section A.3 of Appendix A, two additional examples are discussed of DIVIDE queries related to the running example that do contain context-enriching queries.

6.2. Semantic reasoning to derive queries

Starting from the enriched context data model, the semantic reasoner used within DIVIDE is run to perform the actual query derivation. This way, the reasoner will define whether the DIVIDE query should be initialized for the given context. If so, it specifies with what values the input variables and static window parameters, as defined in the query's sensor query rule consequence, should be substituted in the generic query pattern of the DIVIDE query.

The inputs of the semantic reasoner in this step consist of the preprocessed ontology (i.e., all triples and rules extracted from its axioms), the enriched context triples, the sensor query rule and the goal of the DIVIDE query. Given these inputs, the reasoner performs semantic reasoning to construct and output a proof with all possible rule chains in which the goal of the DIVIDE query is the final rule applied. Every such rule chain will be (partially) different and correspond to a different set of instantiated query variables appearing in the goal's rule.

To allow the semantic reasoner to construct a rule chain that starts from the context and ontology triples and ends with the goal rule, the sensor query rule is crucial. If the inputs allow the reasoner to derive the set of triples in the antecedence of the sensor query rule for a certain set of query variables, the rule *can* be evaluated for this set of variables. However, the semantic reasoner *will* only actually evaluate the rule for this set and include it in the rule chain, *if* the triples in the consequence of the sensor query rule (and more specifically, the part with the ontology consequences) allow the semantic reasoner to derive the antecedence of the goal rule. This can be either directly (i.e., without semantic reasoning) or indirectly (i.e., after rule-based semantic reasoning). If this is not the case, the sensor query rule will not help the semantic reasoner in constructing a rule chain where the goal is the last rule applied, for the given set of sensor query rule variables. Hence, if the proof contains an instantiation of the sensor query rule for a given set of query variables, this implies that the generic RSP-QL query of this DIVIDE query should be instantiated for this set. This should be done with those query variables of this set that are present in the list of input variables or window parameters of the sensor query rule's consequence.

To reassure that this process works, consider the DIVIDE query parser's translation of the ordered set of SPARQL queries in the end user DIVIDE query definition into its internal representation. If the original stream query in the SPARQL input would yield a query result, the final query's WHERE clause *might* have a matching pattern, and thus an output. This is equivalent to the potential evaluation of the sensor query rule in the proof, depending on whether the sensor query rule's consequence directly or indirectly leads to a matching antecedence of the goal rule.

When the query derivation is executed for the DIVIDE query of the running example, the inputs will include the showering AR rule in Listing 1 that is defined in the preprocessed ontology. In the proof constructed by the semantic reasoner, the DIVIDE query's sensor query rule of Listing 3 would be instantiated once for the showering activity, *if* the current location of the patient is the bathroom. The step in the rule chain of the reasoner's proof in which this happens, is shown in Listing 4. This proof shows that the relative humidity sensor with the given ID can detect the showering activity for patient with ID 157 if its value is 57 or higher. If

```

@prefix r: <http://www.w3.org/2000/10/swap/reason#>.

<#lemma3> a r:Inference;
  r:gives {
    _:sk_0 a sd:Query.
    _:sk_0 sd:pattern sd-query:pattern.
    _:sk_0 sd:inputVariables (
      ("?sensor" <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/70:ee:50:67:3e:78>)
      ("?threshold" "57"^^xsd:float)
      ("?activityType" ActivityRecognition:Showering)
      ("?patient" patients:patient157)
      ("?model" :KBActivityRecognitionModel)
      ("?prop_o" <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/org.dyamand.types.common.
        RelativeHumidity>)
    ).
    _:sk_0 sd>windowParameters (("range" 30 time:seconds) ("slide" 10 time:seconds)).
    _:sk_1 a ActivityRecognition:ActivityPrediction.
    _:sk_1 ActivityRecognition:forActivity _:sk_2.
    _:sk_2 a ActivityRecognition:Showering.
    _:sk_1 ActivityRecognition:activityPredictionMadeFor patients:patient157.
    _:sk_1 ActivityRecognition:predictedBy :KBActivityRecognitionModel.
    _:sk_1 saref-core:hasTimestamp _:sk_3.
  };
r:evidence ( <#lemma8> [...] <#lemma31> );
[...]
r:rule <#lemma32>.

```

Listing 4. One step of the proof constructed by the semantic reasoner used in DIVIDE during the DIVIDE query derivation for the generic DIVIDE query of the running use case example. It shows how the sensor query rule in Listing 3 is instantiated in the proof’s rule chain. [...] is a placeholder for omitted parts that are not of interest.

the current context would describe another patient location than the bathroom, or would not define showering as part of the routine of the patient with ID 157, the proof would not contain this sensor query rule instantiation.

6.3. Query extraction

The proof in the output of the semantic reasoning step *can* contain instantiations of the sensor query rule. If not, the proof will be empty, since this means that the semantic reasoner has not found any rule chain that leads to an instantiation of the goal rule. Every sensor query rule instantiation in the proof contains the list of input variables and window parameters that need to be substituted into the generic RSP-QL query of the considered DIVIDE query. In the query extraction step, DIVIDE will extract these definitions from every sensor query rule instantiation in the proof. Hence, the output of this step is a set of zero, one or more extracted queries.

The query extraction happens through two forward reasoning steps with the semantic reasoner used in DIVIDE. The outputs of both steps are combined to construct the output of the query extraction. The first reasoning step extracts the relevant content from the sensor query rule instantiations in the proof. For each instantiation, this content includes the instantiated input variables and window parameters, as well as a reference to the query pattern in which they need to be substituted. The second forward reasoning step of the query extraction retrieves any defined window parameters from the enriched context that are associated with the instantiated RSP-QL query pattern. Such window parameters may have been added to the enriched context during the context enrichment step. They will be used as dynamic window parameters during the window parameter substitution, while the window parameters occurring in the sensor query rule instantiations are considered as static window parameters.

For the running example, the output of the extraction for the proof step in Listing 4 is presented in lines 4–12 of Listing 5. Line 15 of this listing presents the output of the second step. For this query example, there are no dynamic window parameters, which defaults the output of this second query extraction step to an empty list.

In Section A.3 of Appendix A, a related example is presented that does include dynamic window parameters.

```

1  @prefix : <file:///home/divide/.divide/query-derivation/10-10-129-31-8175-/activity-ongoing/20211220
   _194006/proof.n3#>.
2
3  # output of the first reasoning step of the query extraction
4  :lemma3 a sd:Query.
5  :lemma3 sd:inputVariables (
6    ("?sensor" <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/70:ee:50:67:3e:78>)
7    ("?threshold" "57"^^xsd:float) ("?activityType" ActivityRecognition:Showering)
8    ("?patient" patients:patient157) ("?model" :KBActivityRecognitionModel)
9    ("?prop_o" <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/org.dyamand.types.common.
   RelativeHumidity>)
10 ) .
11 :lemma3 sd:staticWindowParameters ("?range" 30 time:seconds) ("?slide" 10 time:seconds) .
12 :lemma3 sd:pattern sd-query:pattern.
13
14 # output of the second reasoning step of the query extraction
15 :lemma3 sd:dynamicWindowParameters () .

```

Listing 5. Output of the query extraction step of the DIVIDE query derivation, performed for the running example on the proof with a single sensor query rule instantiation presented in the proof step of Listing 4. The extraction of the dynamic window parameters (line 15) is done on the enriched context outputted by the context enrichment step.

6.4. Input variable substitution

In this step, DIVIDE substitutes the input variables of each query from the query extraction output into the associated RSP-QL query pattern. To achieve this, a collection of N_3 rules have been defined that allow to substitute the input variables into the query body in a deterministic way. Moreover, they ensure that the substitution is correct for IRIs and literals of any data type. To perform the substitution, the semantic reasoner used in DIVIDE performs a forward reasoning step. The input of this reasoning step consists of the substitution rules, the output of the query extraction step and the query pattern of the considered DIVIDE query. For each query in the query extraction output, the output of this step consists of a set of triples that define the partially substituted RSP-QL query body.

The output of the input variable substitution step for the running example is presented in Listing 6. This substitution is performed using the generic RSP-QL query body referenced in the output of the query extraction in Listing 5. This query body is shown in Listing 3. In the output, lines 1–13 redefine the prefixes, which will be required in a further step to construct the full RSP-QL query. Line 16 shows the current state of the instantiated RSP-QL query body: input variables have already been substituted, but the window parameters still need to be substituted. The static and dynamic window parameters that will be used for substitution in the following step, are propagated from the output of the query extraction step (lines 19–20).

6.5. Window parameter substitution

In this step, the window parameters are also substituted in the partially instantiated queries to obtain the resulting RSP-QL query bodies. This is the final step that is performed by the semantic reasoner used in DIVIDE.

In general, DIVIDE offers the possibility to define the window parameters of derived RSP queries using semantic definitions. Currently, context-aware window parameters can be defined by an end user via the definition of a DIVIDE query. By separating the window parameter substitution from the other query derivation steps, DIVIDE offers the flexibility to trigger this substitution for other reasons than a context change. An example of this could be a device monitor observing that the resources of the device cannot handle the current query execution frequency.

Currently, to enable the substitution of *use case dependent* window parameters, DIVIDE makes the distinction between static and dynamic window parameters. For a static window parameter, the variable behaves as a regular input variable. This means that it should be defined in the consequence of a DIVIDE query's sensor query rule with a triple similar to the following one:

```

1 sd-query:prefixes-activity-showering a owl:Ontology.
2 sd-query:prefixes-activity-showering sh:declare _:bn_1.
3 _:bn_1 sh:prefix "xsd".
4 _:bn_1 sh:namespace "http://www.w3.org/2001/XMLSchema#"^^xsd:anyURI.
5 sd-query:prefixes-activity-showering sh:declare _:bn_2.
6 _:bn_2 sh:prefix "saref-core".
7 _:bn_2 sh:namespace "https://saref.etsi.org/core/"^^xsd:anyURI.
8 sd-query:prefixes-activity-showering sh:declare _:bn_3.
9 _:bn_3 sh:prefix "ActivityRecognition".
10 _:bn_3 sh:namespace "https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/"^^xsd:anyURI.
11 sd-query:prefixes-activity-showering sh:declare _:bn_4.
12 _:bn_4 sh:prefix "KBActivityRecognition".
13 _:bn_4 sh:namespace "https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/KBActivityRecognition/"^^xsd:anyURI.
14
15 _:sk_20 a sd:Query.
16 _:sk_20 sd:queryBody " CONSTRUCT { _:p a KBActivityRecognition:RoutineActivityPrediction ;
    ActivityRecognition:forActivity [ a <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/
    Showering> ] ; ActivityRecognition:activityPredictionMadeFor <http://protego.ilabt.imec.be/idlab.
    homelab/patients/patient157> ; ActivityRecognition:predictedBy <https://dahcc.idlab.ugent.be/
    Ontology/ActivityRecognition/KBActivityRecognition/KBActivityRecognitionModel> ; saref-core:
    hasTimestamp ?now . } FROM NAMED WINDOW :win ON <http://protego.ilabt.imec.be/idlab.homelab> [
    RANGE ?{range} STEP ?{slide}] WHERE { BIND (NOW() as ?now) WINDOW :win { <https://dahcc.idlab.
    ugent.be/Homelab/SensorsAndActuators/70:ee:50:67:3e:78> saref-core:makesMeasurement [ saref-core:
    hasValue ?v ; saref-core:hasTimestamp ?t ; saref-core:relatesToProperty <https://dahcc.idlab.ugent
    .be/Homelab/SensorsAndActuators/org.dyamand.types.common.RelativeHumidity> ] . FILTER (xsd:float(?
    v) > xsd:float(\"57\"^^xsd:float)) } } ORDER BY DESC(?t) LIMIT 1 ".
17 _:sk_20 sh:prefixes sd-query:prefixes-activity-showering.
18 _:sk_20 sd:pattern sd-query:pattern.
19 _:sk_20 sd:staticWindowParameters ("?range" 30 time:seconds) ("?slide" 10 time:seconds).
20 _:sk_20 sd:dynamicWindowParameters ().

```

Listing 6. Output of the input variable substitution step of the DIVIDE query derivation, performed for the running example on the query extraction output presented in Listing 5. The substitution is done using the generic RSP-QL query body of the corresponding DIVIDE query presented in Listing 3.

```

_:q sd:windowParameters ("?range" ?var time:seconds) .

```

This requires the variable `?var` to occur in the sensor query rule's antecedence. When defining a DIVIDE query as an end user using an ordered set of existing SPARQL queries, this can be achieved by ensuring that the variable name of this window parameter appears in the WHERE clause of the stream query, in a named graph that is not corresponding to a stream window. By definition, static window parameter variables will always receive a value in the query extraction output that *can* be used for substitution. In addition, dynamic window parameters are dynamically defined as triples in the outputs of context-enriching queries, similar to the following ones:

```

sd-query:pattern sd:windowParameters (
    [ sd-window:variable "range" ; sd-window:value 30 ; sd-window:type time:seconds ] ) .

```

Importantly, dynamic window parameters will *always* overwrite static ones. This means that during the window parameter substitution, dynamic window parameters will be substituted first. Next, static window parameters are substituted for those window parameter variables in the RSP-QL query body that have not yet been substituted.

The substitution order of static and dynamic window parameters implies a few important things. Multiple dynamic window parameters can be defined in different context-enriching queries of the same DIVIDE query, to handle different situations. It is however the responsibility of the end user that no more than one definition occurs for each window parameter variable in the enriched context. If multiple values are defined for the same window parameter variable, the one that is substituted will be chosen arbitrarily. If no value is defined for a window parameter variable in the enriched context either, the value of the static window parameter with the same variable name will be substituted. However, if no static window parameter value is defined for this variable either, the default value in the end user definition of the DIVIDE query will be substituted. To make this work, DIVIDE will define a window parameter in the sensor query rule of the DIVIDE query with the given default value, for each such variable.

In the running example, the definition of the generic DIVIDE query associated with the detection of an ongoing showering activity does not contain any context-enriching query that defines a dynamic window parameter. However, Section A.3 of Appendix A discusses an example of a related DIVIDE query that does contain dynamic window parameter definitions in its context-enriching queries.

The actual substitution of window parameters is very similar to the input variable substitution. For both the static and dynamic window parameters, a forward reasoning step is performed with the semantic reasoner used in DIVIDE. The inputs of the reasoner are the output of the previous step and a collection of N_3 rules that ensure the correct substitution in a deterministic way. The unit of the window parameter, which is either a valid XML Schema duration string or a time unit, defines how the window parameter value is exactly substituted in the query body string.

6.6. RSP engine query update

The output of the window parameter substitution step is a set of instantiated, valid RSP-QL queries that are contextually relevant for the given component. These queries are however still presented as a series of semantic triples. This final step constructs the actual RSP-QL query string, translates the query to the correct query language and updates the registered queries at the component's RSP engine.

Query construction This step constructs an actual RSP-QL query from the set of prefixes and the instantiated query body triples in the output of the window parameter substitution step.

For the running example, the RSP-QL query resulting from the query construction step is presented in Listing 7. This query is the result of performing the window parameter substitution and query construction on the output of the input variable substitution step presented in Listing 6.

Query translation The definition of a DIVIDE component contains the query language of its RSP engine. If this language differs from RSP-QL, e.g., C-SPARQL, the RSP-QL query is translated in this step to this other language.

```

PREFIX ActivityRecognition: <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/>
PREFIX KBActivityRecognition: <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/
    KBActivityRecognition/>
PREFIX saref-core: <https://saref.etsi.org/core/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

CONSTRUCT {
  _:p a KBActivityRecognition:RoutineActivityPrediction ;
  ActivityRecognition:forActivity [
    a <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/Showering> ] ;
  ActivityRecognition:activityPredictionMadeFor
    <http://protego.ilabt.imec.be/idlab.homelab/patients/patient157> ;
  ActivityRecognition:predictedBy
    <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/KBActivityRecognition/
      KBActivityRecognitionModel> ;
  saref-core:hasTimestamp ?now .
}
FROM NAMED WINDOW :win ON <http://protego.ilabt.imec.be/idlab.homelab> [RANGE PT30S STEP PT10S]
WHERE {
  BIND (NOW() as ?now)
  WINDOW :win {
    <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/70:ee:50:67:3e:78>
      saref-core:makesMeasurement [
        saref-core:hasValue ?v ; saref-core:hasTimestamp ?t ;
        saref-core:relatesToProperty <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/org.
          dyamand.types.common.RelativeHumidity> ] .
    FILTER (xsd:float(?v) > xsd:float("57"^^xsd:float))
  }
}
ORDER BY DESC(?t) LIMIT 1

```

Listing 7. Final RSP-QL query that is the result of performing the window parameter substitution and query construction steps of the DIVIDE query derivation, performed for the running example on the input variable substitution output presented in Listing 6.

Query registration update The output of the previous step is a set of translated RSP queries for the given DIVIDE query. Since the DIVIDE query derivation is triggered because of a detected context change relevant to this component, the queries on the RSP engine of this component should be updated to reflect this new situation. To do so, DIVIDE keeps track of the queries that are currently registered on the RSP engine for the given DIVIDE query. In this step, DIVIDE semantically compares the new set of instantiated translated RSP queries with this existing set of registered queries. Based on this comparison, any registered queries that are no longer in the new set of contextually relevant RSP queries are unregistered. New queries that are not running yet on the RSP engine, are registered.

For completeness, it is important to mention that during the full DIVIDE query derivation, the query processing on the RSP engine of the corresponding component should ideally be temporarily paused. This is to avoid that incorrect filtering is done, since DIVIDE already knows that the active queries might no longer be contextually relevant as soon as DIVIDE is informed of a context change for this component. During the pause, incoming observations on the RSP engine's streams should be buffered temporarily. This way, the queries can be restarted as soon as the RSP query update step finishes, and the buffered stream data can be fed to the RSP engine with their original timestamps.

7. Implementation of the DIVIDE system

The previous sections have described DIVIDE from a methodological point of view, irregardless of implementation details. This section zooms in on our implementation of DIVIDE.

7.1. Technologies

DIVIDE is implemented in Java as a set of Java JAR modules. These modules include the DIVIDE engine, which is the core of DIVIDE, the DIVIDE reasoning module and the DIVIDE server.

The DIVIDE reasoning module implements the ontology preprocessing and the query derivation steps with the semantic reasoner used in DIVIDE. Our implementation uses the EYE reasoner [74], which fulfills the requirements of the semantic reasoner explained in Section 4. This N₃ reasoner runs in a Prolog virtual machine.

The DIVIDE server module is an executable that starts up DIVIDE. It exposes a REST API that allows to add, delete and request information about DIVIDE queries and DIVIDE components in the DIVIDE system.

7.2. Configuration of DIVIDE

The configuration of DIVIDE is provided through a main JSON file. It includes details about different aspects of DIVIDE. In addition, the DIVIDE components can be defined in a separate file. An example of the JSON configuration of the DIVIDE system is provided in Appendix B.

Knowledge base The type of the knowledge base (e.g., Apache Jena, RDFox) should be configured, *if* it is deployed by the DIVIDE server. This is the preferred option when deploying new systems. If DIVIDE is deployed in an existing system, an existing Knowledge Base can also be used. In that case, the system will be responsible for monitoring context updates relevant to components registered to DIVIDE, and triggering the query derivation in the DIVIDE engine for those components whenever such a context update occurs.

Ontology To configure the ontology used by DIVIDE, the relevant ontology files should be specified.

Reasoner and engine The configuration of the DIVIDE reasoner and engine consists of a series of flags that allow to change the default DIVIDE behavior. For example, DIVIDE can be configured to handle TBox definitions in context graphs during the query derivation. Moreover, the parser can be configured to automatically create a variable mapping between the stream and final query based on equal variable names.

DIVIDE queries For every DIVIDE query, a separate JSON file should be linked in the configuration. This file can include the items of the internal representation of a DIVIDE query, or the end user definition of a DIVIDE query. In the latter case, the implementation of the DIVIDE query parser ensures that the parsed DIVIDE queries result in valid RSP-QL queries after the query derivation. This is achieved by validating the inputs, renaming the query variables to avoid any mismatches, ordering the input variables and static window parameters to obtain a deterministic substitution, and handling query variables in special constructs such as GROUP BY clauses.

The JSON configuration of the DIVIDE query for the running use case example can be found in Appendix B.

Server For the DIVIDE server, the host and port of the exposed REST API is defined. If DIVIDE deploys the Knowledge Base as well, the port of the Knowledge Base REST API available for context updates is also specified.

DIVIDE components The components known by DIVIDE should be defined in an additional CSV file, which contains one entry per component. The properties of every component entry are separated by a semicolon.

7.3. Implementation of the ontology preprocessing

During the initialization of DIVIDE, the configured ontology is preprocessed with the EYE reasoner in three steps. First, an N_3 copy of the full ontology is created. Second, specialized ontology-specific rules are created from the original rules taken from the OWL 2 RL profile description [53]. Starting the EYE reasoning process from these rules will reduce the computational complexity of the reasoning [7]. Third, an image of the EYE reasoner, which has already loaded the ontology and specialized rules, is compiled within Prolog. This precompiled Prolog image is the result of the ontology preprocessing. By starting the semantic reasoning step of the query derivation process from this image, the triples and rules do not need to be loaded into the EYE reasoner each time it is called during the DIVIDE query derivation. This allows to make the semantic reasoning step significantly more efficient.

Although considered infrequent, ontology changes can be handled by DIVIDE. If DIVIDE is hosting the Knowledge Base, ontology changes can be made by using the Knowledge Base REST API. Any TBox change will result in DIVIDE reloading the ontology, redoing the ontology preprocessing, and triggering the query derivation for all DIVIDE queries and components. This is a computationally intensive operation.

7.4. Implementation of the DIVIDE query derivation

The DIVIDE query derivation is managed by the DIVIDE engine. To decouple the scheduler of query derivation tasks from their actual parallel execution, the DIVIDE engine manages a blocking task queue and a dedicated processing thread for every DIVIDE component in the system.

Different tasks can be scheduled by the DIVIDE engine in the blocking task queue of a DIVIDE component. The main task type is a query derivation for one or all DIVIDE queries. In case of a context change, the query derivation is scheduled for all DIVIDE queries. However, when a new DIVIDE query is added to the engine via the server API, the query derivation is only scheduled for the new DIVIDE query. In case the query execution should be performed for multiple DIVIDE queries, the query derivation steps are executed in parallel threads for every DIVIDE query. Another task type is the removal of a DIVIDE query from a component, which requires all related RSP queries to be unregistered from the component's RSP engine. This task is scheduled for all DIVIDE queries of a component when that component is removed, or for all components and one DIVIDE query when this DIVIDE query is deleted.

The following paragraphs present some further implementation details of some DIVIDE query derivation steps.

Context enrichment This step involves the execution of SPARQL queries prior to the actual query derivation. Hence, this is the only semantic step of the query derivation that is not necessarily performed by the EYE reasoner. This is the case if the queries contain SPARQL constructs that cannot be translated to a valid N_3 rule. In this case, the queries are executed in Java by using Apache Jena. In the other case, the queries are translated to N_3 rules which are then applied on the set consisting of triples and, if reasoning is enabled, also consisting of the ontology rules.

RSP engine query update This final step of the query derivation is not performed with the EYE reasoner. To update the query registrations at the RSP engines, the REST APIs of the RSP engine servers are used.

8. Evaluation set-ups

This section presents the set-up of two evaluations of the DIVIDE system. First, the performance of DIVIDE is evaluated by measuring the duration of the different key actions taken by DIVIDE during its initialization and query derivation. Second, the real-time execution of RSP-QL queries generated by the DIVIDE query derivation is evaluated. This is done by comparing the real-time DIVIDE set-up with other well-known real-time approaches.

General information about the collected data, the ontology and context, and activity rules used for these evaluations are presented in Section 8.1. The detailed set-ups of both individual evaluations are further described in Section 8.2 and Section 8.3, respectively. Supportive information relevant to the evaluation set-ups of this paper is publicly available at <https://github.com/IBCNServices/DIVIDE/tree/master/swj2022>.

8.1. Evaluation scenarios

All evaluations are performed on the eHealth use case described in Section 3.1. This section zooms in on the details of the evaluation scenarios of this use case.

8.1.1. Ontology

The ontology of the evaluation system is the Activity Recognition ontology as an extension of the existing DAHCC ontology [63], as presented in Section 3.2. This includes the `KBActivityRecognition` ontology and its imports. The imported ontologies include the `ActivityRecognition`, `MonitoredPerson`, `Sensors AndActuators` and `SensorsAndWearables` modules of the DAHCC ontology and its imported ontologies.

8.1.2. Realistic dataset for rule extraction and simulation

To properly perform the evaluations presented in this paper, a realistic data set is used that is the result of a large scale data collection process. This data collection took place in the imec-UGent HomeLab from June 2021 until October 2021. The HomeLab is an actual standalone house located on the UGent Campus Zwijnaarde, offering a unique residential test environment for IoT services, as it is equipped with all kinds of sensors and actuators. It contains different rooms that represent a typical home: an entry hallway, a ground floor toilet, a living room and kitchen, a staircase to the first floor, and a bathroom, master bedroom, hallway and toilet on the first floor. Prior to the data collection period, a literature study, observational studies and interviews with caregivers were performed to derive the activity types that are important to detect in a patient's home. Based on these activities, a list of properties was derived that could be of relevance to observe in order to detect these activities. These properties were then translated to the required sensors, which were all installed in the HomeLab. The data collected during the data collection period in the HomeLab is used for the evaluation in two ways: to extract realistic rules for activity recognition, and to create a realistic data set for simulation during the real-time evaluations.

Throughout the data collection, data was obtained from two sources relevant to this evaluation: a wearable device, and the in-home contextual sensors. For the former, the patient was equipped with an Empatica E4 wearable device [32]. It has a 3-axis accelerometer (32 Hz) as well as different sensors to measure a person's physiological data: blood volume pulse (64 Hz) and derived inter beat interval of heart rate, galvanic skin response (4 Hz) and skin temperature (4 Hz). For the latter, as explained, a wide range of sensors was installed in the HomeLab. These sensors measure localization, the number of people in a room, relative humidity, indoor temperature, motion, light intensity, sound, air quality, usage of water, electric power consumption of multiple devices, interaction with light switches and other buttons, the state of windows, doors, blinds, and others. During the data collection, participants labeled their activities, which were mapped by the researchers to the activities in the DAHCC ontology.

8.1.3. Context

The context for the evaluations, as considered by DIVIDE, consists of three main parts. The first part is the description of a patient living in a smart home, including the patient's wearables and a routine. For this part, the exact definitions in Listing 10 of Appendix A are used. The second part is a single triple representing the patient's location in the home, which is normally derived by a specific query. For the evaluation scenarios, the location of the patient in the home will always be the bathroom. The third part is the description of all sensors, actuators and wearables of the patient's smart home with the DAHCC ontology concepts. The smart home used in this evaluation scenario is

the HomeLab. The instantiated example modules `_Homelab` and `_HomelabWearables` of the DAHCC ontology contain an actual representation of all sensors, actuators and wearables used within the HomeLab. The ABox definitions in these ontology modules represent the second part of the context used for the evaluations. Note that for these evaluations, the small set of TBox definitions present in both modules are also considered part of the ontology.

8.1.4. Activity rules

From the data collected during the large scale data collection in the HomeLab, data-driven rule mining algorithms were created that have extracted some realistic rules that can recognize some of the DAHCC activities from the data. For the evaluations of DIVIDE in this paper, rules for three bathroom activities are considered: toileting, showering and brushing teeth. Based on the analysis, the following rules were extracted:

- Toileting: the person present in the HomeLab is going to the toilet if a sensor that analyzes the energy consumption of the water pump has a value higher than 0.
- Showering: the person present in the HomeLab bathroom is showering if the relative humidity in the bathroom is at least 57%.
- Brushing teeth: the person present in the HomeLab bathroom is performing the brushing teeth activity if in the same time window (a) the sensor that analyzes the water running in the bathroom sink measures water running, and (b) the activity index value of the person's acceleration (measured by a wearable) is higher than 30. The activity index based on acceleration is defined as the mean variance of the acceleration over the three axes [8].

These three activity rules have been semantically described using the Activity Recognition ontology. The resulting descriptions are part of the `KBActivityRecognition` ontology module. They are detailed in Appendix C.

The three given activity rules are the only rules present in the `KBActivityRecognition` ontology module during the evaluations. To represent each activity rule within DIVIDE, a DIVIDE query is created for each rule. Because of the completely similar definition, the generic DIVIDE query corresponding to the toileting and showering activity rules is the same. The DIVIDE queries are defined as an ordered collection of SPARQL queries.

8.2. Performance evaluation of DIVIDE

To derive the contextually relevant RSP queries, DIVIDE performs multiple steps, both during initialization and the query derivation. To evaluate the performance of DIVIDE, the duration of the main steps is measured for the given evaluation scenarios. In concrete, the duration of the following steps is measured:

1. the ontology preprocessing step with EYE of the Activity Recognition ontology;
2. the DIVIDE query parsing step with Java of the toileting DIVIDE query defined as SPARQL input;
3. the DIVIDE query derivation step for the toileting and brushing teeth DIVIDE queries separately, split up between the different EYE steps: semantic reasoning, query extraction, input variable substitution and window parameter substitution (note that the showering DIVIDE query is the same as the toileting DIVIDE query, and is therefore not evaluated twice).

The duration of these steps is measured from within the execution of the DIVIDE server Java JAR, which is configured with the scenario's ontology and DIVIDE queries to allow performing evaluation 1 and 2. To perform evaluation 3, the full context of the scenario is sent as new context to the DIVIDE Knowledge Base server.

While evaluating the performance of the DIVIDE query parser, the correctness of the parsing is also validated.

Technical specifications The evaluation is performed on a device with a 2800 MHz quad-core Intel Core i5-7440HQ CPU and 16 GB DDR4-2400 RAM.

8.3. Real-time evaluation of derived DIVIDE queries

The task of DIVIDE is to manage the RSP queries on the registered RSP engines. These RSP queries are characterized by the fact that they do not require any more reasoning during their continuous evaluation, since semantic reasoning during the query derivation ensures that they are contextually relevant at every point in time. This section compares the real-time performance of evaluating these derived RSP queries on the C-SPARQL RSP engine [9].

8.3.1. Evaluation of DIVIDE in comparison with real-time reasoning approaches

This evaluation compares the DIVIDE real-time approach with other traditional approaches that do require real-time reasoning. The goals of the evaluation are to understand how DIVIDE compares to these traditional set-ups in terms of processing performance, and to understand the differences, advantages and drawbacks of the approaches.

To have a fair comparison, the real-time reasoning approaches should all reason within the same reasoning profile as DIVIDE, i.e., OWL 2 RL. Most approaches use RDFox [55], as this is known as one of the fastest OWL 2 RL (Datalog) reasoning engines that exist in the current state-of-the-art. Others use the Apache Jena rule reasoner.

Set-ups Different set-ups are considered for this evaluation. Most of the set-ups are streaming set-ups, meaning that they operate on windows taken from data streams. For every streaming set-up, Esper is the technology used to manage the windowing and to generate the window triggers [34].

1. DIVIDE approach using C-SPARQL without reasoning: regular C-SPARQL engine [9]. No ontology or context data is loaded into the engine, and no reasoning is performed during the continuous query evaluation.
2. Streaming RDFox: streaming version of RDFox. Consists of one engine that pipes Esper for windowing with RDFox for reasoning, via a processing queue. Initially, the ontology and context data are loaded into the data store of the RDFox engine, and a reasoning step is performed. Upon every window trigger generated by Esper, the window content is added as one event to a processing queue. When available, RDFox takes an event from the queue, incrementally adds it to the RDFox data store (i.e., it performs incremental reasoning with the event scheduled for addition), and executes the registered queries in order. If there are multiple queries registered, query X incrementally adds its results to the data store, before query $X + 1$ is executed. Finally, RDFox performs incremental reasoning with the event and all previous query outputs scheduled for deletion (i.e., incremental deletion).
3. C-SPARQL piped with (non-streaming) RDFox: Initially, the RDFox data store contains the ontology and context data, and a reasoning step is performed. The queries registered on C-SPARQL listen to the observation stream, and run continuously on the stream window data and on the ontology and context triples. The axioms in the ontology are converted to a set of rules. Rule reasoning is performed during each query evaluation using these rules by C-SPARQL, which uses the Apache Jena rule reasoner with a hybrid forward and backward reasoning algorithm. C-SPARQL sends each query result to the event stream of a regular non-streaming RDFox engine, which adds it to a processing queue. Upon processing time, it incrementally adds the event to the data store, executes the registered queries in order, and incrementally deletes the event from the data store.
4. RDFox (non-streaming): RDFox engine wrapped into a server, that listens to the observation stream. Each incoming observation is added to a queue, which is processed by a separate thread. This thread takes an event from the queue, adds it to RDFox, performs incremental reasoning, and executes the registered queries in order. If there are multiple queries registered, query X incrementally adds its results to the data store, before query $X + 1$ is executed. Because this is a non-streaming version of RDFox, the event triples and triples constructed by the intermediate queries are not removed from the data store after processing.
5. Adapted Streaming RDFox: adapted streaming version of RDFox. This set-up only differs in one aspect from the original streaming RDFox set-up (2): before an event is added to RDFox, it checks the overlap between the event triples and existing triples in the data store. If overlapping triples are found, they are not added again to RDFox, and – most importantly – they are also not removed afterwards, so that no previously existing triples are removed from the data store after the event processing.
6. Semi-Streaming RDFox: mix between streaming RDFox set-up (2) and non-streaming RDFox set-up (4). This set-up only differs in one aspect from the original streaming RDFox set-up: the event triples and triples constructed by intermediate queries are not removed from the data store after processing. Hence, the only difference with the non-streaming RDFox set-up is that events are not added directly to the queue from the observation stream, but grouped together on Esper window triggers.
7. Streaming Jena: streaming version of the Apache Jena rule reasoner, similar to the streaming RDFox set-up (2). The only difference is the fact that during initialization, a set of rules is extracted from the ontology and loaded together with the ontology triples into the Apache Jena rule reasoner. Processing of events from the processing queue is done by this reasoner: it takes events, add them to the reasoner's data model, performs forward rule reasoning using the RETE algorithm, and executes the registered queries in order. Temporal query

results are also added to the reasoner's data model, which are removed after processing of the event together with the event triples, followed by a final reasoning step. This set-up uses Apache Jena v3.7.

Each set-up is deployed with an associated WebSocket server to which an external component can connect to send data to the registered data streams. Each set-up involving RDFox uses RDFox v5.2.1, via the JRDFox Java jar, which is the Java bridge to the native RDFox engine. The RDFox data store used is the default `par-complex-nn` store, indicating a parallel data store using a complex indexing scheme with 32-bit integers.

Simulated data To create a simulation dataset to use in the evaluations, an anonymous representative portion is extracted from the dataset obtained with the large-scale data collection in the HomeLab. It contains real sensor observations of all HomeLab sensors and an Empatica E4 wearable worn by a real person living in the HomeLab for a day. Hence, the frequencies and values of the different observations are representative for a real smart home.

The simulated data for the scenarios is changed in two aspects: (i) timestamps are shifted to real-time timestamps, and (ii) the values for the sensors relevant to the evaluated activity are modified to ensure that its conditions are fulfilled all the time. In other words, the simulation for the brushing teeth scenario described below will lead to a detected brushing teeth activity during the full course of the scenario, and similarly for the other activities.

One hour of data from the anonymous data set used in this evaluation contains data of 231 different sensors, together producing 670,118 observations in this hour. 605,090 of these observations are produced by the 4 sensors of the Empatica E4 wearable, the remaining 65,028 observations are produced by 227 sensors in the HomeLab.

Specific scenarios Three specific scenarios, one for each activity rule in the general scenario, are constructed for this evaluation:

- Toileting scenario: Simulated HomeLab data for a period of 30 minutes is replayed at real rate, in batches of 1 second. For the streaming set-ups, the streaming queries are evaluated on a sliding window of 60 seconds with a sliding step of 10 seconds.
- Showering scenario: Simulated HomeLab data for a period of 20 minutes is replayed at real rate, in batches of 1 second. For the streaming set-ups, the streaming queries are evaluated on a sliding window of 60 seconds with a sliding step of 10 seconds.
- Brushing teeth scenario: Simulated HomeLab data for a period of 30 minutes is replayed at real rate, in batches of 1 second. For the streaming set-ups, the streaming queries are evaluated on a sliding window of 30 seconds with a sliding step of 10 seconds.

The purpose of the evaluations is to measure and study the executions of the query evaluations and associated operations of the reasoner or engine, such as the semantic reasoning, both individually and progressively over time.

Replaying the data is performed by a data simulation component running on an external device in the same local network, to realistically represent the different sensor gateways. During simulation, this component connects as a client to the WebSocket server of the evaluated set-up, and sends the observations in each batch as a single message over the WebSocket connection to the appropriate data streams. This implies that one incoming event is received by the set-ups every second. Hence, the streaming set-ups will add such an event to Esper for windowing every second, while the non-streaming set-up will trigger the in-order evaluation of the registered set-ups every second.

Evaluation queries To properly compare the different set-ups for each specific scenario, different versions of the SPARQL and C-SPARQL queries are created. In concrete, the following adaptations are made:

- For set-up 1, the C-SPARQL query as outputted by DIVIDE is registered.
- For set-ups 2, 4, 5, 6 and 7, the SPARQL definition of the DIVIDE query is modified to obtain two queries registered to the single reasoning service. The first reasoning query is the stream query of the SPARQL definition, from which the graph specifications are removed. The second query is the final query of this definition. Hence, the evaluated queries that are executed with RDFox or Apache Jena are regular SPARQL queries that involve semantic reasoning and are not rewritten by DIVIDE.
- For set-up 3, the SPARQL definition of the DIVIDE query is modified to obtain two queries. The first reasoning query is derived from the stream query of the SPARQL definition: the graph specifications are removed, and the query is translated to C-SPARQL syntax by adding the relevant FROM clauses that specify the query input:

the static resources and the data stream window definition. This query is registered to the C-SPARQL engine. The second query is identical to set-up 2 and is registered to RDFox.

During an evaluation run, only the quer(y)(ies) related to the activity rule of the scenario are deployed on the engines. Queries related to other activity rules or aspects like location monitoring are not registered to the engines.

Measurements For each presented set-up, the *total execution time* metric is measured for each event. This metric is defined as the time starting from a *generated event* until the timestamp where an instance of the `RoutineActivityPrediction` is returned as output by the corresponding query. In a set-up with multiple queries that are executed in order, this is always the output of the final query in the chain. The definition of a *generated event* differs for each set-up: in the streaming set-ups, this is the time of an Esper window trigger; in the non-streaming set-up 4, this is the time of an incoming set of sensor observations.

Technical specifications All evaluations are run on a typical processing device in the IoT world: an Intel NUC, model D54250WYKH. It has a 1300 MHz dual-core Intel Core i5-4250U CPU (turbo frequency 2600 MHz) and 8 GB DDR3-1600 RAM.

8.3.2. Real-time evaluation of derived DIVIDE queries on a Raspberry Pi

DIVIDE is considered as a semantic component in a cascading reasoning set-up in an IoT network, which involves running RSP queries on local devices. These devices can be low-end devices with few resources in an IoT context. Hence, it is important to evaluate the real-time performance of continuously executing the RSP queries outputted by DIVIDE on a low-end device like a Raspberry Pi. This is the topic of the final evaluation.

For this evaluation, only the C-SPARQL baseline set-up (1) of the previous section is considered. All other properties of this evaluation are identical to those used for the evaluation in the previous section.

Technical specifications This evaluation is performed on a Raspberry Pi 3, Model B. This Raspberry Pi model has a Quad Core 1.2 GHz Broadcom BCM2837 64bit CPU, 1 GB RAM and MicroSD storage.

9. Evaluation results

This section presents the results of the three evaluations described in Section 8. All results contain data of multiple evaluation runs, always excluding 3 warm-up and 2 cool-down runs.

9.1. Performance evaluation of DIVIDE

Figure 3 shows the distribution of the duration of two initialization steps of the DIVIDE system: the preprocessing of the Activity Recognition ontology, and the parsing of the toileting query specified as SPARQL input. The preprocessing of the ontology on average takes 9,640 ms, with a standard deviation (SD) of only 42 ms. The average duration of the query parsing is only 64.87 ms (SD 2.76 ms). It was also validated that the parsing of the end user definition of the DIVIDE query to its internal representation was done correctly.

Figure 4 shows the performance results of the query derivation with DIVIDE, for the DIVIDE query corresponding to the toileting activity rule (also corresponds to the showering rule) and the DIVIDE query corresponding to the brushing teeth activity rule. Figure 4(a) shows the distribution of the duration of the query derivation for each individual query. The average durations of the query derivation are 3,578 ms (SD 38 ms) and 2,968 ms (SD 37 ms) for the toileting and brushing teeth DIVIDE queries, respectively.






Fig. 3. Performance results of the initialization of the DIVIDE system: boxplot distributions of total execution times per step.



(a) Boxplot distribution of query derivation time for the toileting and brushing teeth DIVIDE query



(b) Relative times for query derivation substeps with EYE reasoner, averaged over all runs of the three DIVIDE queries

Fig. 4. Performance results of the query derivation of the DIVIDE system.

Figure 4(b) shows the percentage of time taken up by the different substeps, averaged over all runs for the three DIVIDE queries. These substeps include all steps performed with the EYE reasoner: the reasoning (47.27% on average), the query extraction (27.32% on average), the input variable substitution (9.44% on average) and the window parameter substitution (10.93% on average). The remaining time (5.04% on average) is overhead of the DIVIDE implementation, including internal threading and memory operations.

9.2. Evaluation of DIVIDE in comparison with real-time reasoning approaches

Figure 5 shows the results of the comparison of the real-time evaluation with DIVIDE on a C-SPARQL engine with different real-time reasoning approaches, for the toileting query. The results show the evolution over time of the total execution time from the event generation until the routine activity prediction is generated by the engine. The measurements included in the graphs are averaged over the evaluation runs. For three setups, there are no measurements shown for the full time course of the evaluation, which takes 30 minutes. These set-ups are the pipe of C-SPARQL with RDFox set-up (3), the adapted streaming RDFox set-up (5) and the streaming Jena set-up (7). These missing measurements are caused by the systems running out of memory, causing them to stop evaluating the queries for the remainder of the scenario. The DIVIDE baseline set-up (1) has the lowest average total execution time from 960 seconds into the evaluation. Before this timestamp, the non-streaming RDFox set-up (4) is the quickest.

Figure 6 shows similar results for the real-time evaluation of the showering query. The same three set-ups run out of memory at a certain point, causing missing measurements for the remainder of the evaluation runs. Already after 550 seconds into the evaluation, the DIVIDE baseline set-up (1) has the lowest average total execution time.

Figure 7 shows similar results of the comparison of the real-time evaluation of DIVIDE with the real-time reasoning approaches, but for the brushing teeth query. The properties of the graph are similar to those of the graph presenting the results for the toileting query. In these results, only the non-streaming RDFox set-up (4) has no measurements for the full time course of the evaluation scenario due to the engine running out of memory.

In Appendix D, additional results of the evaluation runs over time are included. These results visualize the distribution of the total execution times for the different set-ups at different times during the evaluation runs.

9.3. Real-time evaluation of derived DIVIDE queries on a Raspberry Pi

Figure 8 shows the results of the evaluation of the DIVIDE set-up on the Raspberry Pi 3. These results visualize the distribution of the individual execution times of the RSP queries generated by DIVIDE with the C-SPARQL

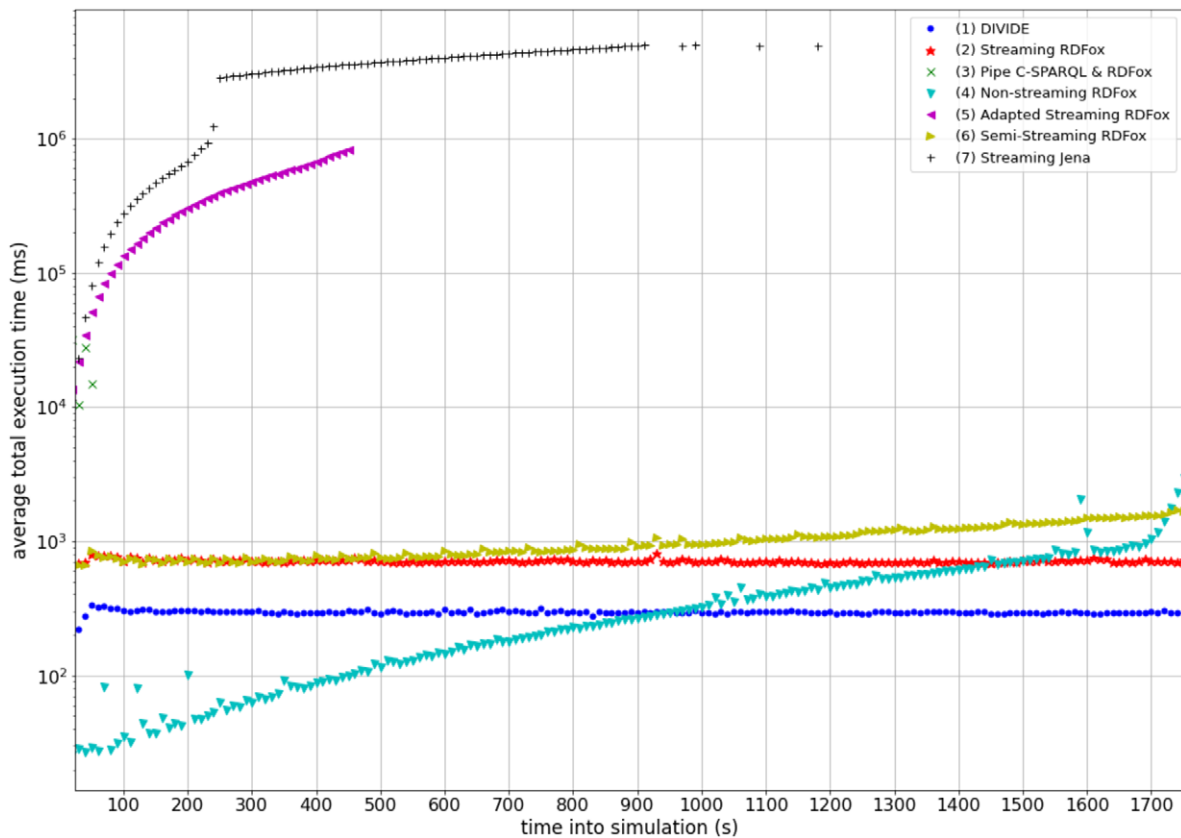


Fig. 5. Results of the comparison of the DIVIDE real-time query evaluation approach with real-time reasoning approaches, for the toileting query. For each evaluation set-up, the results show the evolution over time of the total execution time from the generated event (either a windowed event in a streaming set-up or an incoming event in a non-streaming set-up) until the routine activity prediction as output of the final query. For all set-ups, measurements are shown for the processed event, either incoming or windowed, at every 10 seconds. All plotted execution times are averaged over the evaluation runs.

baseline set-up, for the toileting, showering and brushing teeth scenarios. For the toileting query, the average total execution time is 3,666 ms (SD 318 ms). This average number is 3,699 ms (SD 286 ms) and 3,001 ms (SD 174 ms) for the showering and brushing teeth scenarios, respectively.

10. Discussion

Including DIVIDE as a component in a semantic IoT platform allows to perform context-aware monitoring of patients in homecare scenarios. This is possible because DIVIDE is designed to fit in a cascading architecture: it derives and manages contextually relevant RSP queries that require no additional reasoning while they are being executed, which makes them perfectly suitable to run on local low-end devices in the patient's home environment.

An end user of DIVIDE will design the IoT platform architecture for a specific use case within a certain application domain. By employing DIVIDE in a cascading reasoning architecture, DIVIDE enables privacy by design to a certain extent. As such, DIVIDE helps the end user to integrate privacy by design into the application, by following some of the privacy by design principles. More specifically, DIVIDE leaves its end users in full control to specifically define which privacy-sensitive data is exposed to the outside world. This data will typically consist of different levels of abstractions of the raw data observed by the IoT sensors. The end user control of exposed data directly follows from the definition of the DIVIDE queries: these queries exactly define which semantic concepts will be filtered by the local RSP engines, and sent over the IoT network to the central reasoner on a central server. Only the




Fig. 6. Results of the comparison of the DIVIDE real-time query evaluation approach with real-time reasoning approaches, for the showering query. For each evaluation set-up, the results show the evolution over time of the total execution time from the generated event (either a windowed event in a streaming set-up or an incoming event in a non-streaming set-up) until the routine activity prediction as output of the final query. For all set-ups, measurements are shown for the processed event, either incoming or windowed, at every 10 seconds. All plotted execution times are averaged over the evaluation runs.

outputs of those queries will ever leave the local environment of the patient; all other data will be kept locally. This way, DIVIDE helps its end users to adhere to the embedded, user-centric, and visibility and transparency principles of privacy by design. By embedding DIVIDE into a cascading architecture, the design can consider the interests of the patients and be transparent about the data being sent over the network (i.e., the outputs of the generic RSP queries in the DIVIDE query definitions). Nevertheless, the research, implementation and integration of additional privacy solutions into the application design is required to optimally achieve privacy preservation. For example, in the described use case scenario, the patient's in-home location and detected activities comprise the only information that is ever leaving the home. While this ensures all other privacy-sensitive data is kept locally, it does not guarantee the preservation of this small set of privacy-sensitive data that is leaving the patient's environment. Such additional privacy solutions that need to be built into the application design will often be use case specific, according to the use case requirements. Hence, this requires additional, use case specific privacy research that is considered out of scope of the presented research.

With respect to security, note that the integration of DIVIDE into a semantic IoT platform does not guarantee any additional security to the system. Currently, the communication within DIVIDE is only protected via the standard SSL/TLS encryption associated with the HTTPS protocol, which is not sufficient to ensure maximum security. Hence, an additional security system or framework should be integrated into a semantic IoT platform architecture that involves DIVIDE. Existing security systems and frameworks should be researched to achieve this. However, this is considered to be out of scope of the presented research.




Fig. 7. Results of the comparison of the DIVIDE real-time query evaluation approach with real-time reasoning approaches, for the brushing teeth query. For each evaluation set-up, the results show the evolution over time of the total execution time from the generated event (either a windowed event in a streaming set-up or an incoming event in a non-streaming set-up) until the routine activity prediction as output of the final query. For all set-ups, measurements are shown for the processed event, either incoming or windowed, at every 10 seconds. All plotted execution times are averaged over the evaluation runs.




Fig. 8. Results of evaluating the DIVIDE real-time query evaluation approach with the C-SPARQL baseline set-up (1), on a Raspberry Pi 3, Model B. The results show the total execution time distribution over the engine's runtime and multiple runs, for the toileting, showering and brushing teeth DIVIDE queries.

A DIVIDE query is generic by nature, which ensures that you should not define one DIVIDE query for every individual reasoning or filtering task that should be performed in the use case. In the activity recognition use case scenario discussed in this paper, one should only define a generic DIVIDE query per *type* of activity rule, instead of per activity rule individually. The generic nature of a DIVIDE query ensures that DIVIDE can derive the instantiated queries from it that are contextually relevant at any given point in time. This is achieved by listening to context updates in the knowledge base, and automatically triggering the query derivation upon a context change for all

components that are affected by this context change. This is an improvement compared to systems where the management of the queries on the stream processing components of the IoT platform is still a manual, labor-intensive and thus highly impractical task. On the other hand, generic semantic queries can also be processed by reasoning engines, but while this is certainly feasible with current existing semantic reasoners for a single home environment, it might become more complex if this needs to be managed for a full network with for example many smart homes.

By deploying DIVIDE in a cascading architecture, more benefits are obtained than solely the privacy control for the end user, generic query definition and context-awareness. Since the high frequency and high volume data streams are processed locally, this data should not be transferred over the network. This significantly reduces network bandwidth usage and network delay impacting the system's performance. In addition, the data does not need to be processed by the central reasoner, which now only receives the outputs of the local RSP queries to do further processing. As such, the main resources of the server can be saved for the high-priority situations. In the presented use case scenario, an example is when an activity is detected that is not in the patient's routine: when this prediction is received by the central reasoner, it can investigate the cause of the issue and trigger further actions such as generating an alarm when needed. Meanwhile, the server resources can also be used by DIVIDE to derive the updated location monitoring query to ensure that the patient's location is followed up more closely.

When DIVIDE is used as a component in a semantic IoT platform to derive and manage the local RSP queries, it is of course important that the queries derived by DIVIDE have a good performance that is comparable to existing state-of-the-art stream reasoning systems. The results of this comparison with the C-SPARQL RSP engine running on an Intel NUC device demonstrate that the filtering RSP queries perform very well for the different activity detection queries that each correspond to a generic real-time reasoning set-up. The results show how the C-SPARQL queries are only outperformed by the classic non-streaming RDFox reasoning engine if you only look at the processing of single events. This can easily be explained by the fact that the events processed by this RDFox set-up contain fewer observations, and thus triples, than the events processed by C-SPARQL, which are larger batches of data grouped in data windows. Hence, due to the incremental reasoning in RDFox, this set-up initially performs best. However, looking at the evolution of the total execution times over time, the DIVIDE baseline set-up starts to perform better after a while. This is because the performance of the DIVIDE set-up stays constant over time, while the total execution time of the queries on the RDFox set-up increases over time because events are not removed from the data store, increasing the size of the data store on every execution. Therefore, we have also included a comparison with a streaming version of RDFox. This set-up also performs constant over time, and is outperformed by a slight margin only by DIVIDE. This is mainly because RDFox still has to do some reasoning, which, even though this happens very efficiently with RDFox, is not required for the evaluation of the RSP queries with C-SPARQL in the baseline set-up. The streaming set-up of RDFox used in the evaluations makes a few assumptions that can still be optimized by looking at overlapping events and ensuring they are not removed after the processing of an event. However, in this optimized, adapted streaming RDFox setup, the processing of incoming events cannot keep up with the rate of the windowed events, causing the processing delay to build up. This leads to very long query execution times and memory issues in some cases. Moreover, looking at the results that involve reasoning with Apache Jena, it is clear that the set-ups using this semantic reasoner perform way worse than the DIVIDE and optimal RDFox set-ups. This is also true for the pipe of C-SPARQL with RDFox, in which C-SPARQL is performing rule-based reasoning with Apache Jena in the first query. This reasoning step causes the bad performance entirely on its own. This learns that using the built-in rule reasoning support of C-SPARQL is not efficient compared to alternative set-ups. As a conclusion, over time, DIVIDE performs comparable or even slightly better than the best RDFox set-ups, making it an ideal solution to integrate in a semantic platform to manage the local RSP queries, given the other main advantages. Ideally, this is combined in the cascading architecture with a central reasoner that does use a performant semantic reasoner such as RDFox.

In IoT networks, devices with resources comparable to those of an Intel NUC are often unavailable locally. Therefore, it is important that the RSP queries can also be continuously executed on low-end devices with fewer resources. Otherwise, the data would still have to be sent to other devices with more resources running more centrally in the network that would then host the RSP engines. This would imply that all other advantages related to privacy, network usage and server resources do no longer apply. Therefore, the evaluation of the C-SPARQL baseline set-up was also performed on a Raspberry Pi. The results demonstrate that the queries can still be efficiently and consistently executed on such devices with way fewer resources than an Intel NUC. Specifically for this evaluation,

the queries take approximately 10 times longer than on the Intel NUC, but take still well below the query execution frequency of 10 seconds. This is an additional advantage when deploying a system involving DIVIDE, as no large scale investment in expensive high-end hardware is required. In real set-ups, actually deploying a Raspberry Pi may however not be very practical or realistic. However, the resources of a Raspberry Pi are very comparable to other local devices such as wearable devices like the smartwatches in Samsung Galaxy Watch or Apple Watch series. Note that RDFox can also be used instead of C-SPARQL to run the queries derived by DIVIDE on a local low-end device, since RDFox can successfully run on an ARM based edge device like a Raspberry Pi or a smartphone as well [48]. This implies that the use of a RDFox set-up would also ensure that data can be processed locally instead of being sent to a server. Also note that RDFox is able to handle any arbitrary OWL 2 RL ontology, including recursive ones.

Up to now, we have only looked at the real-time evaluation of RSP queries derived by DIVIDE. They perform well in realistic homecare monitoring environments, but another important aspect is the performance of DIVIDE itself. The results of the DIVIDE performance evaluation show that the main portion of time during the initialization of DIVIDE is taken by the preprocessing of the ontology. Of course, the duration of the preprocessing depends on the number of triples and axioms defined in the ontology, which is use case specific. In any case, this is a task that should only happen once, given the assumption in DIVIDE that ontology updates do not happen. Nevertheless, DIVIDE does support ontology updates, but they require the ontology preprocessing to be redone. Besides the initialization, it is important to inspect the duration the query derivation process when a context change is observed. For this step, the performance results show that for the given evaluation use case scenario, the query derivation typically takes around 3 to 4 seconds. This is an order of magnitude higher than the time needed to perform real-time reasoning with RDFox during the query evaluation on an incoming event. However, the execution frequency of the query derivation is a few orders of magnitude smaller than the frequency of the event processing: events are processed on every window trigger or incoming observation, which is every 10 seconds or every second in the evaluation use case scenario. As you are not expecting a context change every 10 seconds, this shows that the performance results of the query derivation step are perfectly acceptable. In addition, the results show that the largest portion of the time is taken up by the different steps performed with the EYE reasoner. The biggest portion of the time, almost 50%, is spent on generating the proof with the EYE reasoner. The results show that only less than 5% of the query derivation step is overhead induced by the DIVIDE implementation.

When integrating DIVIDE into a semantic IoT platform, it is important to note that DIVIDE considers all semantic specifications to be accurate. Hence, DIVIDE considers it the responsibility of its end users to ensure that the semantic definitions in the knowledge base and the DIVIDE queries are correctly defined. For example, in the use case scenario described in this paper, DIVIDE assumes that all activity rules defined in the Activity Recognition ontology correctly detect the corresponding activity types. Thus, DIVIDE will not take any measures itself to avoid any misleading of the system: if the end user wants to abuse DIVIDE to generate incorrect outputs, such as incorrectly detected activities in the described use case scenario, this is possible. Hence, it is important that all semantic definitions and DIVIDE queries of a use case are validated before they are integrated into DIVIDE.

To be able to use DIVIDE in a real IoT platform set-up, it is important that DIVIDE is practically usable. Therefore, we have implemented DIVIDE in a way that tries to maximize its practical usability. First, DIVIDE is available as an executable Java JAR component that can easily be run in a server environment, allowing for easy integration into an existing IoT platform. The main configuration of the server, engine, DIVIDE queries and components can be easily created and modified with straightforward JSON and CSV files. Importantly, DIVIDE also does not hinder RSP engines to have active queries managed manually or by other system components, ensuring that the inclusion of DIVIDE into a semantic platform is not an all-or-nothing choice. In addition, the REST API exposed by the DIVIDE server implies that the configuration of DIVIDE is not fixed: components and DIVIDE queries can be easily added or removed, increasing the flexibility of the system. The internal implementation ensures that such changes are correctly handled and reflected on the RSP engines as well. Moreover, the implementation of the query parser allows the flexible and straightforward end user definition of a DIVIDE query. This allows existing sets of queries to be used with DIVIDE to perform semantically equivalent tasks after only a small configuration effort. This way, no inner details of DIVIDE need to be known by end users who want to integrate it into their system. Our implementation of the parser also validates the DIVIDE query definitions given by the end user, and provides a human-friendly explanation about what is wrong in case the input is invalid. As a result, we believe that DIVIDE is perfectly suited in an IoT set-up where it is deployed in a cascading architecture.

11. Conclusion

This paper has presented the DIVIDE system. DIVIDE is designed as a semantic component that can automatically and adaptively derive and manage the queries of the stream processing components in a semantic IoT platform, in a context-aware manner. Through a specific homecare monitoring use case, this paper has shown how DIVIDE can divide the active queries across a cascading IoT set-up, and conquer the issues of existing systems by fulfilling important requirements related to data privacy preservation, performance, and usability.

Reaching back to the research objectives outlined in Section 1, we have achieved these in this paper with DIVIDE in the following ways:

1. DIVIDE automatically triggers the derivation of the semantic queries of a stream processing component when changes are observed to context information that is relevant to that specific component. This way, DIVIDE automatically ensures that the active queries on each component are contextually relevant at all times. This process is context-aware and adaptive by design, minimizing the manual configuration effort for the end user to the initial query definition only. Once the system is deployed, no configuration changes are required anymore.
2. By performing semantic reasoning on the current context during the query derivation, DIVIDE ensures that the resulting stream processing queries can perform all relevant monitoring tasks without doing real-time reasoning. The evaluations on the use case scenario demonstrate how this ensures that DIVIDE performs comparable or even slightly better than state-of-the-art stream reasoning set-ups involving RDFS in terms of query execution times. This implies that the queries can also be executed in real-time on low-end devices with few resources, as demonstrated by the evaluations. The cascading architecture in which DIVIDE is adopted ensures minimal network congestion and optimal usage of the central resources of the network.
3. Through the definition of a DIVIDE query, an end user can make the window parameters of the stream processing queries context-dependent with DIVIDE.
4. By adopting a cascading reasoning architecture, DIVIDE manages the queries for the stream processing components that are running on local IoT devices. Integrating DIVIDE into a semantic IoT platform enables privacy by design to a certain extent: it leaves the end users, who design the platform architecture for a specific use case, in full control to specify in the DIVIDE query definitions which privacy-sensitive data is kept locally by the local stream processing queries, and which data (abstractions) in the query outputs are sent over the IoT network to the central services.
5. Generic queries in DIVIDE can be easily defined by only slightly adapting existing SPARQL or RSP-QL queries, ensuring DIVIDE is practically usable.

There are multiple interesting future pathways related to DIVIDE that are worth investigating. First, the cascading architectural set-up in which DIVIDE is ideally deployed can be further exploited. By including the monitoring of device, network and/or stream characteristics into DIVIDE, the distribution of semantic stream processing queries across the IoT network could be dynamically adapted to optimize both local and global system performance. Such a monitor could also exploit the dynamic window parameter substitution functionality of DIVIDE to adapt these parameters to the monitored conditions. Second, the current implementation of DIVIDE only supports use cases that reason in the OWL 2 RL profile. However, the EYE reasoner used supports extending the rule set to obtain higher expressivity. Doing so would introduce support for higher expressivity use cases in DIVIDE.

Acknowledgements

This research is part of the imec.ICON project PROTEGO (HBC.2019.2812), co-funded by imec, VLAIO, Televic, Amaron, Z-Plus and ML2Grow. Bram Steenwinckel (1SA0219N) is funded by a strategic base research grant of Fund for Scientific Research Flanders (FWO), Belgium. Pieter Bonte (1266521N) is funded by a postdoctoral fellowship of FWO.

Availability of data and materials

Supportive information relevant to the evaluation set-ups of this paper is publicly available at <https://github.com/IBCNServices/DIVIDE/tree/master/swj2022>. This page also refers to the source code of DIVIDE at <https://github.com/IBCNServices/DIVIDE/tree/master/src/divide-central>, additional details of the DAHCC ontology at <https://dahcc.idlab.ugent.be>, and the described dataset used to construct the evaluation rules and simulation data at <https://dahcc.idlab.ugent.be/dataset.html>.

Appendix A. Additional details of homecare monitoring use case and running example

This appendix includes additional details about the homecare monitoring use case, which is described in Section 3, and its running example that is used in the discussion of the DIVIDE methodology in Section 4, 5 and 6.

A.1. Semantic representation of use case and running example

This part of the appendix provides additional details of how the homecare monitoring use case and its running example are semantically represented with the Activity Recognition ontology.

- Listing 8 gives an overview of all prefixes used in the listings with semantic content in this paper.
- Listing 9 lists some ontology definitions that specify when a showering activity prediction corresponds to the routine of a patient and when it does not, based on the activities defined in this patient's routine. Based on these definitions, a semantic reasoner can define a recognized activity as an instance of either `RoutineActivityPrediction` or `NonRoutineActivityPrediction`. The desired output of the semantic AR service consists of instances of these classes and their relations.

```
# Activity Recognition ontology including DAHCC ontology modules
@prefix KBActivityRecognition: <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/
  KBActivityRecognition/> .
@prefix ActivityRecognition: <https://dahcc.idlab.ugent.be/Ontology/ActivityRecognition/> .
@prefix MonitoredPerson: <https://dahcc.idlab.ugent.be/Ontology/MonitoredPerson/> .
@prefix SensorsAndActuators: <https://dahcc.idlab.ugent.be/Ontology/SensorsAndActuators/> .
@prefix SensorsAndWearables: <https://dahcc.idlab.ugent.be/Ontology/SensorsAndWearables/> .
@prefix Sensors: <https://dahcc.idlab.ugent.be/Ontology/Sensors/> .

# instances in use case scenario
@prefix : <http://divide.ilabt.imec.be/idlab.homelab/> .
@prefix patients: <http://divide.ilabt.imec.be/idlab.homelab/patients/> .
@prefix Homelab: <https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/> .
@prefix HomelabWearable: <https://dahcc.idlab.ugent.be/Homelab/SensorsAndWearables/> .

# SAREF and extensions
@prefix saref-core: <https://saref.etsi.org/core/> .
@prefix saref4ehaw: <https://saref.etsi.org/saref4ehaw/> .
@prefix saref4bldg: <https://saref.etsi.org/saref4bldg/> .
@prefix saref4wear: <https://saref.etsi.org/saref4wear/> .

# other imports
@prefix time: <http://www.w3.org/2006/time#> .
@prefix eep: <https://w3id.org/eep#> .

# generic prefixes
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix xml: <http://www.w3.org/XML/1998/namespace> .

# definitions within DIVIDE
@prefix sd: <http://idlab.ugent.be/sensdesc#> .
@prefix sd-query: <http://idlab.ugent.be/sensdesc/query#> .
@prefix sh: <http://www.w3.org/ns/shacl#> .
```

Listing 8. Overview of all prefixes used in the listings with semantic content in this document.

```

:RoutineActivityPrediction SubClassOf: ActivityRecognition:ActivityPrediction
:NonRoutineActivityPrediction SubClassOf: ActivityRecognition:ActivityPrediction

:RoutineShoweringActivityPrediction SubClassOf: :RoutineActivityPrediction
:RoutineShoweringActivityPrediction EquivalentTo:
  :RoutineActivityPrediction and :ShoweringActivityPrediction
:RoutineShoweringActivityPrediction SubClassOf: :ShoweringActivityPrediction
:RoutineShoweringActivityPrediction EquivalentTo:
  :ShoweringActivityPrediction and
  (ActivityRecognition:activityPredictionMadeFor some :UserWithShoweringRoutine)

:NonRoutineShoweringActivityPrediction SubClassOf: :NonRoutineActivityPrediction
:NonRoutineShoweringActivityPrediction EquivalentTo:
  :NonRoutineActivityPrediction and :ShoweringActivityPrediction
:NonRoutineShoweringActivityPrediction SubClassOf: :ShoweringActivityPrediction
:NonRoutineShoweringActivityPrediction EquivalentTo:
  :ShoweringActivityPrediction and
  (ActivityRecognition:activityPredictionMadeFor some :UserWithoutShoweringRoutine)

:ShoweringActivityPrediction SubClassOf: ActivityRecognition:ActivityPrediction
:ShoweringActivityPrediction EquivalentTo:
  ActivityRecognition:ActivityPrediction and
  (ActivityRecognition:forActivity some ActivityRecognition:Showering)

:UserWithShoweringRoutine EquivalentTo:
  saref4ehaw:User and
  (MonitoredPerson:hasRoutine some (
    ActivityRecognition:Routine and
    (ActivityRecognition:consistsOf some ActivityRecognition:Showering)))
:UserWithoutShoweringRoutine EquivalentTo:
  saref4ehaw:User and
  (:doesNotHaveActivityInRoutine some ActivityRecognition:Showering)

```

Listing 9. Example of how different subclass and equivalence relations between concepts are defined in the `KBActivityRecognition` ontology module of the Activity Recognition ontology, allowing a semantic reasoner to derive whether an activity prediction corresponds to a person’s routine or not. To improve readability, all definitions are listed in Manchester syntax and the `KBActivityRecognition:` prefix is replaced by `:`.

- Listing 10 gives an example context description of a patient in the described use case scenario. The current location of this patient in the service flat is the bathroom.
- The description of the HomeLab service flat is given in the instantiated example modules `_HomeLab` and `_HomeLabWearable` of the DAHCC ontologies. The most relevant descriptions of these modules with respect to the running example are presented in Listing 11. As can be observed, for each sensor in the home, the observed properties are defined through the `measuresProperty` object property, and the analyzed entity is specified with the `analyseStateOf` property.

A.2. End user definition of running example’s DIVIDE query as an ordered collection of SPARQL queries

The DIVIDE query corresponding to the running example is detailed in Section 5.1. This query can be defined by an end user as an ordered collection of existing SPARQL queries. This definition can then be translated by the DIVIDE query parser to its internal representation. This appendix section details this end user definition.

The stream and final queries of the definition are shown in Listing 12 and 13, respectively. There are no intermediate queries. The context enrichment also consists of an empty set of queries, since the stream query is the first query in the ordered set of SPARQL queries used in the stream reasoning system. However, Section A.3 of this appendix discusses a related DIVIDE query that does include a context enrichment and intermediate queries.

Moreover, the DIVIDE query definition contains one stream window with the following properties:

- Stream IRI: `http://protego.ilabt.imec.be/idlab.homelab`
- Window definition: `RANGE PT?rangeS STEP PT?slideS`
- Default window parameter values: `?range` has a default value of 30, `?slide` has default value 10

```

# patient with ID 157 lives in a smart home called the HomeLab
patients:patient157 rdf:type saref4ehaw:Patient ;
    MonitoredPerson:livesIn Homelab:homelab .

# patient has a location tag
patients:patient157 rdf:type saref4wear:Wearer .
Homelab:AQURA_10_10_145_9 saref4wear:isLocatedOn patients:patient157 ;
    MonitoredPerson:hasLocation Homelab:homelab .

# patient has a morning routine consisting of a series of activities
patients:patient157 MonitoredPerson:hasRoutine :MorningRoutine_patient157 .
:MorningRoutine_patient157 rdf:type ActivityRecognition:MorningRoutine ;
    ActivityRecognition:consistsOf _:A1, _:A2, _:A3, _:A4, _:A5, _:A6 ;
    ActivityRecognition:nextActivity _:A1 .
_:A1 rdf:type ActivityRecognition:WakingUp ;
_:A2 rdf:type ActivityRecognition:Toileting .
_:A3 rdf:type ActivityRecognition:Showering .
_:A4 rdf:type ActivityRecognition:BrushingTeeth .
_:A5 rdf:type ActivityRecognition:EatingMeal .
_:A6 rdf:type ActivityRecognition:WatchingTVActively .
_:A1 ActivityRecognition:nextActivity _:A2 .
_:A2 ActivityRecognition:nextActivity _:A3 .
_:A3 ActivityRecognition:nextActivity _:A4 .
_:A4 ActivityRecognition:nextActivity _:A5 .
_:A5 ActivityRecognition:nextActivity _:A6 .

# patient is currently located in the bathroom
patients:patient157 MonitoredPerson:hasIndoorLocation Homelab:bathroom .

```

Listing 10. Context description of the example patient in the use case scenario and corresponding running example, in RDF/Turtle syntax. Only a selected set of context definitions are presented, some are omitted.

```

# the HomeLab building consists of a bathroom on the first floor
Homelab:homelab rdf:type saref4bldg:Building .
Homelab:firstfloor rdf:type SensorsAndActuators:Floor ;
    saref4bldg:isSpaceOf Homelab:homelab .
Homelab:bathroom rdf:type SensorsAndActuators:BathRoom ;
    saref4bldg:isSpaceOf Homelab:firstfloor .

# the bathroom contains a Netatmo sensor that measures, among others, relative humidity
<https://dahcc.idlab.ugent.be/Homelab/SensorsAndActuators/70:ee:50:67:3e:78>
    rdf:type Homelab:Netatmo ;
    core:measuresProperty Homelab:org.dyamand.types.airquality.CO2 ,
        Homelab:org.dyamand.types.common.AtmosphericPressure ,
        Homelab:org.dyamand.types.common.Loudness ,
        Homelab:org.dyamand.types.common.RelativeHumidity ,
        Homelab:org.dyamand.types.common.Temperature ;
    Sensors:analyseStateOf Homelab:bathroom ;
    saref4bldg:isContainedIn Homelab:bathroom .
Homelab:Netatmo rdf:type owl:Class ;
    rdfs:subClassOf saref-core:Sensor .
Homelab:org.dyamand.types.common.RelativeHumidity
    rdf:type SensorsAndActuators:RelativeHumidity .

# the HomeLab consists of a location system that can detect the room in which
# the patient is located based on a tag system
Homelab:AQURA_10_10_145_9 core:consistsOf Homelab:AQURA_10_10_145_9.Tag .
Homelab:AQURA_10_10_145_9.Tag rdf:type saref4bldg:Sensor ;
    Sensors:analyseStateOf Homelab:AQURA_10_10_145_9 ;
    core:measuresProperty Homelab:org.dyamand.agura.AquraLocationState_Protego_User .
Homelab:org.dyamand.agura.AquraLocationState_Protego_User
    rdf:type SensorsAndActuators:Localisation .

```

Listing 11. Context description of the service flat of the example patient in the use case scenario and corresponding running example, in RDF/Turtle syntax. Only a selected set of context definitions are presented, some are omitted.

This window definition contains the two variable window parameters `?range` and `slide`. The definition of a default value for both window parameters implies that they are not used as static window parameters. This can be confirmed by observing their absence in the WHERE clause of the stream query in Listing 12.

```

1  CONSTRUCT {
2    _:p rdf:type ActivityRecognition:ActivityPrediction ;
3      ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
4      ActivityRecognition:activityPredictionMadeFor ?patient ;
5      ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?now .
6  }
7  FROM NAMED <http://protego.ilabt.imec.be/idlab.homelab>
8  FROM NAMED <http://protego.ilabt.imec.be/context>
9  WHERE {
10   BIND (NOW() as ?now)
11
12   GRAPH <http://protego.ilabt.imec.be/idlab.homelab> {
13     ?sensor saref-core:makesMeasurement [
14       saref-core:hasValue ?v ; saref-core:hasTimestamp ?t ;
15       saref-core:relatesToProperty ?prop_o ] .
16   }
17
18   GRAPH <http://protego.ilabt.imec.be/context> {
19     ?model rdf:type ActivityRecognition:ActivityRecognitionModel ;
20       <https://w3id.org/eep#implements> [
21         rdf:type ActivityRecognition:Configuration ;
22         KBActivityRecognition:containsRule ?a ] .
23     ?a rdf:type KBActivityRecognition:ActivityRule ;
24       ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
25       KBActivityRecognition:hasCondition [
26         rdf:type KBActivityRecognition:RegularThreshold ;
27         KBActivityRecognition:isMinimumThreshold "true"^^xsd:boolean ;
28         saref-core:hasValue ?threshold ;
29         Sensors:analyseStateOf [ rdf:type ?analyzed ] ;
30         KBActivityRecognition:forProperty [ rdf:type ?prop ]
31       ] .
32
33     ?activityType rdfs:subClassOf KBActivityRecognition:DetectableActivity .
34   }
35
36   FILTER (xsd:float(?v) > xsd:float(?threshold))
37
38   GRAPH <http://protego.ilabt.imec.be/context> {
39     ?sensor rdf:type saref-core:Device ; saref-core:measuresProperty ?prop_o ;
40       Sensors:isRelevantTo ?room ; Sensors:analyseStateOf [ rdf:type ?analyzed ] .
41     ?prop_o rdf:type ?prop .
42
43     ?prop rdfs:subClassOf KBActivityRecognition:ConditionableProperty .
44     ?analyzed rdfs:subClassOf KBActivityRecognition:AnalyzableForCondition .
45
46     ?patient MonitoredPerson:hasIndoorLocation ?room .
47   }
48 }

```

Listing 12. Stream query of the end user definition of the DIVIDE query of the running example that performs the monitoring of the showering activity rule.

```

CONSTRUCT {
  _:p rdf:type KBActivityRecognition:RoutineActivityPrediction ;
    ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
    ActivityRecognition:activityPredictionMadeFor ?patient ;
    ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?t .
}
WHERE {
  ?p rdf:type KBActivityRecognition:RoutineActivityPrediction ;
    ActivityRecognition:forActivity [ rdf:type ?activityType ] ;
    ActivityRecognition:activityPredictionMadeFor ?patient ;
    ActivityRecognition:predictedBy ?model ; saref-core:hasTimestamp ?t .
  ?activityType rdfs:subClassOf KBActivityRecognition:DetectableActivity .
}

```

Listing 13. Final query of the end user definition of the DIVIDE query of the running example that performs the monitoring of the showering activity rule.

In addition, the DIVIDE query definition contains the following solution modifier: `ORDER BY DESC(?t) LIMIT 1`. This solution modifier contains the unbound variable name `?t`, which is allowed since it is present in a stream graph of the stream query (Listing 12, line 14).

The variable mapping of stream to final query consists of the stream query variables `?activityType`, `?patient`, and `?model`, which are all mapped to the same variable name in the final query. In addition, it contains the mapping of the variable `?now` in the stream query to the variable `?t` in the final query. The reason for this final mapping becomes clear when inspecting the corresponding generic RSP-QL query pattern in the internal representation of this DIVIDE query (Listing 3, lines 44–59): the literal object of the `hasTimestamp` property in the resulting query output indeed corresponds to the `?now` variable.

A.3. Additional use case examples associated with running example

The running example of the homecare monitoring use case focuses on the detection of in-home activities that are part of the patient’s routine. However, this example does not cover three aspects of the DIVIDE query derivation: the associated DIVIDE query does not have context-enriching queries, no intermediate queries, and no definitions of variable window parameters. Therefore, this appendix zooms in on those aspects for two DIVIDE queries that relate to the DIVIDE query of the running example.

A.3.1. Additional example 1: Query detecting activities not in the patient’s routine

The first additional example focuses on the DIVIDE query that performs the monitoring of the showering activity rule in case the activity is *not* part of the patient’s routine. This DIVIDE query is very similar to the DIVIDE query of the running example. However, the output of this DIVIDE query should contain instances of the class `NonRoutineActivityPrediction`. From the ontology definitions in Listing 9, it follows that the derivation of such instances requires the association between patient and activity with the `doesNotHaveActivityInRoutine` property for every activity type that is not in the patient’s routine. However, such definitions are not present in the regular patient context described in Listing 10. Hence, in an existing stream reasoning system applying the DIVIDE query’s equivalent as a set of ordered SPARQL queries, the evaluation of the stream query would be preceded by an additional query that is enriching the context with this information. In the DIVIDE query definition, this first SPARQL query would be defined as a context-enriching query. It is presented in Listing 14 for illustration purposes.

A.3.2. Additional example 2: Indoor location monitoring query

The second additional example focuses on the DIVIDE query that corresponds to the monitoring of the patient’s location in the home. This DIVIDE query includes variable dynamic window parameters and an intermediate query.

Dynamic window parameters The DIVIDE query contains two context-enriching queries that define multiple dynamic window parameters. These queries are shown in Listing 15. The dynamic window parameters defined in the output of these queries are constructed based on the current context concerning any ongoing activity for this patient. It makes a distinction between two scenarios: when an activity *not in* the patient’s routine is ongoing (first query),

```

CONSTRUCT {
  ?p KBActivityRecognition:doesNotHaveActivityInRoutine [ rdf:type ?activityType ] .
}
WHERE {
  ?p rdf:type saref4ehaw:Patient .

  ?activityType rdf:type owl:Class ;
    rdfs:subClassOf KBActivityRecognition:DetectableActivity .

  FILTER NOT EXISTS {
    ?p MonitoredPerson:hasRoutine ?routine .
    ?routine ActivityRecognition:consistsOf ?routineActivity .
    ?routineActivity rdf:type ?activityType .
  }
}

```

Listing 14. Context-enriching query in the definition of the DIVIDE query that detects an ongoing activity that is *not* in a patient’s routine. It enriches the context with all activity types that are not part of the patient’s routines.

```

# first context-enriching query
CONSTRUCT {
  sd-query:pattern sd:windowParameters (
    [ sd-window:variable "range" ; sd-window:value 30 ; sd-window:type time:seconds ]
    [ sd-window:variable "slide" ; sd-window:value 30 ; sd-window:type time:seconds ] )
} WHERE {
  ?patient rdf:type saref4ehaw:Patient ; MonitoredPerson:livesIn ?home .
  ?prediction1 rdf:type KBActivityRecognition:RoutineActivityPrediction ;
    ActivityRecognition:activityPredictionMadeFor ?patient .
  FILTER NOT EXISTS {
    ?prediction2 rdf:type KBActivityRecognition:NonRoutineActivityPrediction ;
      ActivityRecognition:activityPredictionMadeFor ?patient . }
}

# second context-enriching query
CONSTRUCT {
  sd-query:pattern sd:windowParameters (
    [ sd-window:variable "range" ; sd-window:value 5 ; sd-window:type time:seconds ]
    [ sd-window:variable "slide" ; sd-window:value 5 ; sd-window:type time:seconds ] )
} WHERE {
  ?patient rdf:type saref4ehaw:Patient ; MonitoredPerson:livesIn ?home .
  ?prediction1 rdf:type KBActivityRecognition:NonRoutineActivityPrediction ;
    ActivityRecognition:activityPredictionMadeFor ?patient .
  FILTER NOT EXISTS {
    ?prediction2 rdf:type KBActivityRecognition:RoutineActivityPrediction ;
      ActivityRecognition:activityPredictionMadeFor ?patient . }
}

```

Listing 15. Context-enriching queries that define dynamic window parameters for the DIVIDE query that performs the monitoring of the patient's location in the home. They define the window parameters of this location query based on the current context about any ongoing activity for this patient that *is* or *is not* part of the patient's known routine.

and when an activity *in* the patient's routine is ongoing (second query). Note that for the default case when no activity is currently ongoing, no dynamic window parameters are defined: in those cases, default values for the window parameter variables will be substituted as static window parameters. Moreover, note that the two graph patterns in the WHERE clauses of the queries are semantically distinct: there will never be more than one query for which the graph pattern in the WHERE clause has a matching set of variables. This ensures that there is at most one value defined for the two window parameter variables in the enriched context.

Intermediate query The output constructed by the stream query in this DIVIDE query's definition, is the following:

```

?patient MonitoredPerson:hasIndoorLocationString ?v ;
  saref-core:hasTimestamp ?t .

```

The value of `?v` contains the string representation of the indoor location, as measured by the localization system. However, this does not yet define the location with its actual ontology entity IRI. Therefore, an intermediate query could be used to make this translation. This way, the final query can look for the most recent location IRI. The example of such a combination of intermediate and final query is presented in Listing 16. Note that it would also be possible and semantically equivalent to integrate the translation done in the intermediate query into the final query. However, for readability purposes, it is often better to have multiple, simpler SPARQL queries like in this example.

Appendix B. Configuration of the DIVIDE implementation

This appendix gives some examples of how our implementation of DIVIDE, which is presented in Section 7, should be concretely configured.

- Listing 17 shows an example of the JSON configuration of the DIVIDE system.
- Listing 18 contains the JSON configuration of the DIVIDE query for the running use case example discussed in Section 5.1. In other words, parsing the configured DIVIDE query with the DIVIDE query parser leads to the DIVIDE query goal in Listing 2 and the sensor query rule in Listing 3.

```

# intermediate query
CONSTRUCT {
  ?patient MonitoredPerson:hasIndoorLocationOfInterest [
    saref-core:hasValue ?room; saref-core:hasTimestamp ?t ] .
} WHERE {
  ?patient MonitoredPerson:hasIndoorLocationString [
    saref-core:hasValue ?l ; saref-core:hasTimestamp ?t ] .

  ?room rdf:type saref4bldg:BuildingSpace ; rdfs:label ?roomLabel .
  FILTER (xsd:string(?roomLabel) = xsd:string(?l))
}

# final query
CONSTRUCT {
  ?patient MonitoredPerson:hasIndoorLocation ?room .
} WHERE {
  ?patient MonitoredPerson:hasIndoorLocationOfInterest [
    saref-core:hasValue ?room; saref-core:hasTimestamp ?t ] .
}

```

Listing 16. Example of intermediate query and final query in the end user definition of the DIVIDE query that performs the monitoring of the patient's location in the home. The solution modifier of the final query would be `ORDER BY DESC(?t) LIMIT 1` to retrieve the most recent location only.

```

{
  "divide": {
    "kb": { "type": "Jena", "baseIri": "http://protego.ilabt.imec.be/idlab.homelab/" },
    "ontology": {
      "dir": "definitions/ontology/",
      "files": [ "KBActivityRecognition.ttl", "ActivityRecognition.ttl", "MonitoredPerson.ttl",
        "Sensors.ttl", "SensorsAndActuators.ttl", "SensorsAndWearables.ttl",
        "_Homelab_tbox.ttl", "_HomelabWearable_tbox.ttl",
        "imports/eep.ttl", "imports/affectedBy.ttl", "imports/cpannotationschema.ttl",
        "imports/saref.ttl", "imports/saref4bldg.ttl",
        "imports/saref4ehaw.ttl", "imports/saref4wear.ttl" ]
    },
    "queries": { "sparql": [ "divide-queries/activity-showering.json" ] },
    "reasoner": { "handleTboxDefinitionsInContext": false },
    "engine": {
      "parser": {
        "processUnmappedVariableMatches": false, "validateUnboundVariablesInRspQlQueryBody": true
      },
      "stopRspEngineStreamsOnContextChanges": true
    }
  },
  "server": { "host": "localhost", "port": { "divide": 8342, "kb": 8343 } }
}

```

Listing 17. Example JSON configuration of the DIVIDE system.

```

{
  "streamWindows": [{
    "streamIri": "http://protego.ilabt.imec.be/idlab.homelab",
    "windowDefinition": "RANGE PT?{range}S STEP PT?{slide}S",
    "defaultWindowParameterValues": { "?range": "30", "?slide": "10" }
  }],
  "streamQuery": "stream-query.sparql",
  "finalQuery": "final-query.sparql",
  "solutionModifier": "ORDER BY DESC(?t) LIMIT 1",
  "streamToFinalQueryVariableMapping": {
    "?activityType": "?activityType", "?patient": "?patient", "?model": "?model", "?now": "?t"
  },
  "contextEnrichment": {
    "queries": [], "doReasoning": true, "executeOnOntologyTriples": true
  }
}

```

Listing 18. End user definition of the DIVIDE query of the running example that performs the monitoring of the showering activity rule. The content of the file named `stream-query.sparql` is presented in Listing 12, the content of the file named `final-query.sparql` is presented in Listing 13.

Appendix C. Semantic activity rules of the DIVIDE evaluation scenarios

This appendix contains the semantic description of the activity rules used in the evaluation of the DIVIDE system, as presented in Section 8.1.4. These rules include a rule for the toileting, showering and brushing teeth activity. They are semantically defined using the Activity Recognition ontology presented in Section 3.2, in the `KBActivityRecognition` ontology module. To improve readability, the `KBActivityRecognition:` prefix is replaced by the `:` prefix in all semantic listings of this appendix.

- Toileting: the person present in the HomeLab is going to the toilet if a sensor that analyzes the energy consumption of the water pump has a value higher than 0. This translates into the following activity rule definition:

```
:toileting_rule rdf:type :ActivityRule ;
  ActivityRecognition:forActivity :_Toileting ;
  hasCondition :toileting_condition01 .

:toileting_condition01 rdf:type :RegularThreshold ;
  :forProperty :_EnergyConsumption ;
  Sensors:analyseStateOf :_Pump ;
  isMinimumThreshold "true"^^xsd:boolean ;
  saref-core:hasValue "1.0E-5"^^xsd:float .
```

- Showering: the person present in the HomeLab bathroom is showering if the relative humidity in the bathroom is at least 57%. This translates into the following activity rule definition:

```
:showering_rule rdf:type :ActivityRule ;
  ActivityRecognition:forActivity :_Showering ;
  hasCondition :showering_condition01 .

:showering_condition01 rdf:type :RegularThreshold ;
  :forProperty :_RelativeHumidity ;
  Sensors:analyseStateOf :_BathRoom ;
  isMinimumThreshold "true"^^xsd:boolean ;
  saref-core:hasValue "57.0"^^xsd:float .
```

- Brushing teeth: the person present in the HomeLab bathroom is performing the brushing teeth activity if in the same time window (a) the sensor that analyzes the water running in the bathroom sink measures water running, and (b) the activity index value of the person's acceleration (measured by a wearable) is higher than 30. The activity index based on acceleration is defined as the mean variance of the acceleration over the three axes. This translates into the following activity rule definition:

```
:brushing_teeth_rule rdf:type :ActivityRule ;
  ActivityRecognition:forActivity :_BrushingTeeth ;
  hasCondition :brushing_teeth_condition01 .

:brushing_teeth_condition01 rdf:type :AndCondition ;
  firstCondition :brushing_teeth_condition02 ;
  secondCondition :brushing_teeth_condition03 .

:brushing_teeth_condition02 rdf:type :RegularThreshold ;
  :forProperty :_WaterRunning ;
  Sensors:analyseStateOf :_Room ;
  isMinimumThreshold "true"^^xsd:boolean ;
  saref-core:hasValue "1.0E-5"^^xsd:float .

:brushing_teeth_condition03 rdf:type :MeanVarianceThreshold ;
  :forProperty :_WearableAcceleration ;
  Sensors:analyseStateOf :_Patient ;
  isMinimumThreshold "true"^^xsd:boolean ;
  saref-core:hasValue "30.0"^^xsd:float .
```




Fig. 9. Results of the comparison of the DIVIDE real-time query evaluation approach with real-time reasoning approaches, for the toileting query. For each evaluation set-up, the results show a boxplot distribution of the total execution time from the generation event (either a windowed event in a streaming set-up or an incoming event in a non-streaming set-up) until the routine activity prediction as output of the (final) query. The distribution is shown for two timestamps corresponding to the mean values for this timestamp plotted in Fig. 5.




Fig. 10. Results of the comparison of the DIVIDE real-time query evaluation approach with real-time reasoning approaches, for the brushing teeth query. For each evaluation set-up, the results show a boxplot distribution of the total execution time from the generation event (either a windowed event in a streaming set-up or an incoming event in a non-streaming set-up) until the routine activity prediction as output of the (final) query. The distribution is shown for two timestamps corresponding to the mean values for this timestamp plotted in Fig. 7.

Appendix D. Additional results of the evaluation of DIVIDE in comparison with real-time reasoning approaches

This appendix contains additional results of comparing the real-time evaluation of RSP queries derived by DIVIDE on a C-SPARQL engine, with the other evaluation set-ups that do involve real-time reasoning. These results are complementary to the results shown in Section 9.2, for the evaluation set-up as discussed in Section 8.3.1.

Figure 9 includes two boxplots that show the distribution of the total query execution times for the evaluation of the toileting DIVIDE query, for each set-up over the multiple evaluation runs. The distribution is shown for two timestamps corresponding to the mean values that are visualized in the timeline of Fig. 5. Hence, the distributions correspond to the total execution times measured during the same corresponding evaluation runs. Figure 9(a) shows the distribution for the total execution times for the event, either streaming or incoming, generated 60 seconds after starting the data simulation. Figure 9(b) visualizes this distribution for the event generated 1300 seconds after the start of the data simulation. The results show how the non-streaming RDFox set-up has the smallest total execution times in the beginning of the simulation after only 60 seconds, while DIVIDE has smaller total execution times after 1300 seconds. Note that the boxplot distributions after 1300 seconds do not include results for the pipe of C-SPARQL with RDFox set-up (3), the adapted streaming RDFox set-up (5) and the streaming Jena set-up (7) due to those systems running out of memory before reaching this timestamp in the evaluation.

Figure 10 shows results completely similar to the results in Fig. 9, but for the brushing teeth query. The distributions that are visualized correspond to the mean values that are visualized in the timeline of Fig. 7. Additional results for the showering query are omitted due to their high similarity with the results of the other queries.

References

- [1] K. Abouelmehdi, A. Beni-Hssane, H. Khaloufi and M. Saadi, Big data security and privacy in healthcare: A review, *Procedia Computer Science* **113** (2017), 73–80. doi:[10.1016/j.procs.2017.08.292](https://doi.org/10.1016/j.procs.2017.08.292).
- [2] C.C. Aggarwal, N. Ashish and A. Sheth, The Internet of things: A survey from the data-centric perspective, in: *Managing and Mining Sensor Data*, C.C. Aggarwal, ed., Springer, US, 2013, pp. 383–428. doi:[10.1007/978-1-4614-6309-2_12](https://doi.org/10.1007/978-1-4614-6309-2_12).
- [3] F. Ali, S.R. Islam, D. Kwak, P. Khan, N. Ullah, S. Yoo and K.S. Kwak, Type-2 fuzzy ontology-aided recommendation systems for IoT-based healthcare, *Computer Communications* **119** (2018), 138–155. doi:[10.1016/j.comcom.2017.10.005](https://doi.org/10.1016/j.comcom.2017.10.005).
- [4] D. Anicic, P. Fodor, S. Rudolph and N. Stojanovic, EP-SPARQL: A unified language for event processing and stream reasoning, in: *Proceedings of the 20th International Conference on World Wide Web (WWW 2011)*, Association for Computing Machinery (ACM), New York, NY, USA, 2011, pp. 635–644. doi:[10.1145/1963405.1963495](https://doi.org/10.1145/1963405.1963495).
- [5] D. Anicic, S. Rudolph, P. Fodor and N. Stojanovic, Stream reasoning and complex event processing in ETALIS, *Semantic web* **3**(4) (2012), 397–407. doi:[10.3233/SW-2011-0053](https://doi.org/10.3233/SW-2011-0053).
- [6] D. Arndt, P. Bonte, A. Dejonghe, R. Verborgh, F. De Turck and F. Ongenaes, SENSdesc: Connect sensor queries and context, in: *11th International Joint Conference on Biomedical Engineering Systems and Technologies*, 2018, pp. 1–8. doi:[10.5220/0006733106710679](https://doi.org/10.5220/0006733106710679).
- [7] D. Arndt, B. De Meester, P. Bonte, J. Schaballie, J. Bhatti, W. Dereuddre, R. Verborgh, F. Ongenaes, F. De Turck, R. Van de Walle et al., Improving OWL RL reasoning in N3 by using specialized rules, in: *International Experiences and Directions Workshop on OWL (OWLED) 2015*, Springer, 2015, pp. 93–104. doi:[10.1007/978-3-319-33245-1_10](https://doi.org/10.1007/978-3-319-33245-1_10).
- [8] J. Bai, C. Di, L. Xiao, K.R. Evenson, A.Z. LaCroix, C.M. Crainiceanu and D.M. Buchner, An activity index for raw accelerometry data and its comparison with other activity metrics, *PLoS one* **11**(8) (2016), e0160644. doi:[10.1371/journal.pone.0160644](https://doi.org/10.1371/journal.pone.0160644).
- [9] D.F. Barbieri, D. Braga, S. Ceri, E. Della Valle and M. Grossniklaus, C-SPARQL: A continuous query language for RDF data streams, *International Journal of Semantic Computing* **4**(1) (2010), 3–25. doi:[10.1142/S1793351X10000936](https://doi.org/10.1142/S1793351X10000936).
- [10] D.F. Barbieri, D. Braga, S. Ceri, E.D. Valle and M. Grossniklaus, Incremental reasoning on streams and rich background knowledge, in: *The Semantic Web: Research and Applications: Proceedings, Part I of the 7th Extended Semantic Web Conference, ESWC 2010*, Springer, Cham, Switzerland, 2010, pp. 1–15. doi:[10.1007/978-3-642-13486-9_1](https://doi.org/10.1007/978-3-642-13486-9_1).
- [11] P. Barnaghi, W. Wang, C. Henson and K. Taylor, Semantics for the Internet of Things: Early progress and back to the future, *International Journal on Semantic Web and Information Systems (IJSWIS)* **8**(1) (2012), 1–21. doi:[10.4018/jswis.2012010101](https://doi.org/10.4018/jswis.2012010101).
- [12] H.R. Bazoobandi, H. Beck and J. Urbani, Expressive stream reasoning with laser, in: *The Semantic Web – ISWC 2017: Proceedings, Part I of the 16th International Semantic Web Conference*, Springer, Cham, Switzerland, 2017, pp. 87–103. doi:[10.1007/978-3-319-68288-4_6](https://doi.org/10.1007/978-3-319-68288-4_6).
- [13] H. Beck, M. Dao-Tran and T. Eiter, LARS: A logic-based framework for analytic reasoning over streams, *Artificial Intelligence* **261** (2018), 16–70. doi:[10.1007/978-3-319-73117-9_6](https://doi.org/10.1007/978-3-319-73117-9_6).
- [14] T. Berners-Lee, D. Connolly, L. Kagal, Y. Scharf and J. Hendler, N3Logic: A logical framework for the World Wide Web, *Theory and Practice of Logic Programming* **8**(3) (2008), 249–269. doi:[10.1017/S1471068407003213](https://doi.org/10.1017/S1471068407003213).

- [15] H.K. Bharadwaj, A. Agarwal, V. Chamola, N.R. Lakkaniga, V. Hassija, M. Guizani and B. Sikdar, A review on the role of machine learning in enabling IoT based healthcare applications, *IEEE Access* **9** (2021), 38859–38890. doi:10.1109/ACCESS.2021.3059858.
- [16] P. Bonte, F. Ongenae and F. De Turck, Subset reasoning for event-based systems, *IEEE Access* **7** (2019), 107533–107549. doi:10.1109/ACCESS.2019.2932937.
- [17] P. Bonte, R. Tommasini, F. De Turck, F. Ongenae and E.D. Valle, C-sprite: Efficient hierarchical reasoning for rapid RDF stream processing, in: *Proceedings of the 13th ACM International Conference on Distributed and Event-Based Systems*, 2019, pp. 103–114. doi:10.1145/3328905.3329502.
- [18] P. Bonte, R. Tommasini, E. Della Valle, F. De Turck and F. Ongenae, Streaming MASSIF: Cascading reasoning for efficient processing of IoT data streams, *Sensors* **18**(11) (2018), 3832. doi:10.3390/s18113832.
- [19] A. Bröring, S. Schmid, C.-K. Schindhelm, A. Khelil, S. Käbisch, D. Kramer, D. Le Phuoc, J. Mitic, D. Anicic and E. Teniente, Enabling IoT ecosystems through platform interoperability, *IEEE software* **34**(1) (2017), 54–61. doi:10.1109/MS.2017.2.
- [20] J.-P. Calbimonte, J. Mora and O. Corcho, Query rewriting in RDF stream processing, in: *The Semantic Web: Latest Advances and New Domains: Proceedings of the 13th International Conference, ESWC 2016*, Springer, Cham, Switzerland, 2016, pp. 486–502. doi:10.1007/978-3-319-34129-3_30.
- [21] A. Cavoukian, Privacy by design, Office of the Information and Privacy Commissioner, 2009. <https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf>.
- [22] P. Chamoso, A. González-Briones, F. De La Prieta, G.K. Venyagamoorthy and J.M. Corchado, Smart city as a distributed platform: Toward a system for citizen-oriented management, *Computer Communications* **152** (2020), 323–332. doi:10.1016/j.comcom.2020.01.059.
- [23] A. Cimmino, V. Oravec, F. Serena, P. Kostelnik, M. Poveda-Villalón, A. Tryferidis, R. García-Castro, S. Vanya, D. Tzovaras and C. Grimm, VICINITY: IoT semantic interoperability based on the web of things, in: *15th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, IEEE, New York, NY, USA, 2019, pp. 241–247. doi:10.1109/DCOSS.2019.00061.
- [24] F. Cirillo, G. Solmaz, E.L. Berz, M. Bauer, B. Cheng and E. Kovacs, A standard-based open source IoT platform: FIWARE, *IEEE Internet of Things Magazine* **2**(3) (2019), 12–18. doi:10.48550/arXiv.2005.02788.
- [25] R. Cyganiak, D. Wood, M. Lanthaler, G. Klyne, J.J. Carroll and B. McBride, RDF 1.1 concepts and abstract syntax, W3C Recommendation, World Wide Web Consortium (W3C), 2014. <https://www.w3.org/TR/rdf11-concepts/>.
- [26] L. Daniele, F. den Hartog and J. Roes, Created in close interaction with the industry: The Smart Appliances REference (SAREF) ontology, in: *Formal Ontologies Meet Industry*, Springer, Cham, Switzerland, 2015, pp. 100–112. ISBN 978-3-319-21545-7. doi:10.1007/978-3-319-21545-7_9.
- [27] M. De Brouwer, F. Ongenae, P. Bonte and F. De Turck, Towards a cascading reasoning framework to support responsive ambient-intelligent healthcare interventions, *Sensors* **18**(10) (2018), 3514. doi:10.3390/s18103514.
- [28] D. Dell’Aglío, M. Dao-Tran, J.-P. Calbimonte, D. Le Phuoc and E. Della Valle, A query model to capture event pattern matching in RDF stream processing query languages, in: *Knowledge Engineering and Knowledge Management: Proceedings of the 20th International Conference, EKAW 2016*, Springer, Cham, Switzerland, 2016, pp. 145–162. doi:10.1007/978-3-319-49004-5_10.
- [29] D. Dell’Aglío, E. Della Valle, J.-P. Calbimonte and O. Corcho, RSP-QL semantics: A unifying query model to explain heterogeneity of RDF stream processing systems, *International Journal on Semantic Web and Information Systems (IJSWIS)* **10**(4) (2014), 17–44. <https://dl.acm.org/doi/10.5555/2795081.2795083>. doi:10.4018/ijswis.2014100102.
- [30] D. Dell’Aglío, E. Della Valle, F. van Harmelen and A. Bernstein, Stream reasoning: A survey and outlook, *Data Science* **1**(1–2) (2017), 59–83. doi:10.3233/DS-170006.
- [31] M. Dürst and M. Suignard, Internationalized Resource Identifiers (IRIs), RFC – Proposed Standard, Internet Engineering Task Force (IETF), 2005. <https://datatracker.ietf.org/doc/rfc3987/>.
- [32] Empatica, E4 wristband, 2020, Accessed: 2020-10-23. <https://www.empatica.com/research/e4>.
- [33] I. Esnaola-Gonzalez, J. Bermúdez, I. Fernández and A. Arnaiz, Two ontology design patterns toward energy efficiency in buildings, in: *Proceedings of the 9th Workshop on Ontology Design and Patterns (WOP 2018), Co-Located with 17th International Semantic Web Conference (ISWC 2018)*, CEUR Workshop Proceedings, 2018, pp. 14–28. https://ceur-ws.org/Vol-2195/pattern_paper_2.pdf.
- [34] EsperTech, Esper, Accessed: 2022-03-29. <https://www.espertech.com/esper>.
- [35] A. Felernig, S.P. Erdeniz, P. Azzoni, M. Jeran, A. Akcay and C. Doukas, Towards configuration technologies for IoT gateways, in: *Proceedings of the 18th International Configuration Workshop*, 2016, pp. 73–76. <https://ase.ist.tugraz.at/wp-content/uploads/sites/34/2016/07/configuration-technologies-iot-16.pdf>.
- [36] M. Ganzha, M. Paprzycki, W. Pawlowski, P. Szmaja and K. Wasielewska, Semantic interoperability in the Internet of Things: An overview from the INTER-IoT perspective, *Journal of Network and Computer Applications* **81** (2017), 111–124. doi:10.1016/j.jnca.2016.08.007.
- [37] M. Girod-Genet, L.N. Ismail, M. Lefrançois and J. Moreira, ETSI TS 103 410-8 V1.1.1 (2020-07): SmartM2M; Extension to SAREF; Part 8: eHealth/Ageing-well Domain, Technical report, ETSI Technical Committee Smart Machine-to-Machine communications (SmartM2M), 2020. https://www.etsi.org/deliver/etsi_ts/103400_103499/10341008/01.01.01_60/ts_10341008v010101p.pdf.
- [38] F. Heintz, J. Kvarnström and P. Doherty, Bridging the sense-reasoning gap: DyKnow – stream-based middleware for knowledge processing, *Advanced Engineering Informatics* **24**(1) (2010), 14–26. doi:10.1016/j.aei.2009.08.007.
- [39] S. Jabbar, F. Ullah, S. Khalid, M. Khan and K. Han, Semantic interoperability in heterogeneous IoT infrastructure for healthcare, *Wireless Communications and Mobile Computing* **2017** (2017). doi:10.1155/2017/9731806.
- [40] K. Jaiswal and V. Anand, A survey on IoT-based healthcare system: Potential applications, issues, and challenges, in: *Advances in Biomedical Engineering and Technology*, A.A. Rizvanov, B.K. Singh and P. Ganasala, eds, Springer, Singapore, 2021, pp. 459–471. doi:10.1007/978-981-15-6329-4_38.

- [41] M. Javaid and I.H. Khan, Internet of Things (IoT) enabled healthcare helps to take the challenges of COVID-19 pandemic, *Journal of Oral Biology and Craniofacial Research* **11**(2) (2021), 209–214. doi:10.1016/j.jobcr.2021.01.015.
- [42] A. Javed, S. Kubler, A. Malhi, A. Nurminen, J. Robert and K. Främling, bIoTope: Building an IoT open innovation ecosystem for smart cities, *IEEE Access* **8** (2020), 224318–224342. doi:10.1109/ACCESS.2020.3041326.
- [43] I. Kalamaras, N. Kaklanis, K. Votis and D. Tzovaras, Towards big data analytics in large-scale federations of semantically heterogeneous IoT platforms, in: *Artificial Intelligence Applications and Innovations: Proceedings of AIAI 2018 IFIP 12.5 International Workshops*, L. Iliadis, I. Maglogiannis and V. Plagianakos, eds, Springer, Cham, Switzerland, 2018, pp. 13–23. doi:10.1007/978-3-319-92016-0_2.
- [44] S. Komazec, D. Cerri and D. Fensel, Sparkwave: Continuous schema-enhanced pattern matching over RDF data streams, in: *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems (DEBS 2012)*, Association for Computing Machinery (ACM), New York, NY, USA, 2012, pp. 58–68. doi:10.1145/2335484.2335491.
- [45] C. Kurtz, M. Semmann and T. Böhm, Privacy by design to comply with GDPR: A review on third-party data processors, in: *Proceedings of the 24th Americas Conference on Information Systems (AMCIS) 2018*, 2018. <https://aisel.aisnet.org/amcis2018/Security/Presentations/36/>.
- [46] 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: *The Semantic Web – ISWC 2011: Proceedings, Part I of the 10th International Semantic Web Conference*, Springer, Berlin, Heidelberg, 2011, pp. 370–388. doi:10.1007/978-3-642-25073-6_24.
- [47] F. Lécué, Diagnosing changes in an ontology stream: A DL reasoning approach, in: *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, OJS/PPK, 2012. doi:10.1609/aaai.v26i1.8113.
- [48] J. Lee, T. Hwang, J. Park, Y. Lee, B. Motik and I. Horrocks, A context-aware recommendation system for mobile devices, in: *Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice, Co-Located with 19th International Semantic Web Conference (ISWC 2020)*, K. Taylor, R. Goncalves, F. Lecue and J. Yan, eds, CEUR Workshop Proceedings, 2020. <https://ceur-ws.org/Vol-2721/paper489.pdf>.
- [49] D. Luckham, *The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems*, Addison-Wesley Professional, 2002.
- [50] D.C. Luckham and B. Frasca, Complex event processing in distributed systems, Computer systems laboratory technical report CSL-TR-98-754, Stanford University, 1998. https://www.unix.com/pdf/CEP_in_distributed_systems.pdf.
- [51] G. Marques, A.K. Bhoi and K.S. Hareesha (eds), *IoT in Healthcare and Ambient Assisted Living*, Springer, Singapore, 2021. doi:10.1007/978-981-15-9897-5.
- [52] A. Mileo, A. Abdelrahman, S. Policarpio and M. Hauswirth, StreamRule: A nonmonotonic stream reasoning system for the semantic web, in: *Web Reasoning and Rule Systems: Proceedings of the 7th International Conference, RR 2013*, Springer, Berlin, Heidelberg, 2013, pp. 247–252. doi:10.1007/978-3-642-39666-3_23.
- [53] B. Motik, B.C. Grau, I. Horrocks, Z. Wu, A. Fokoue, C. Lutz et al., OWL 2 web ontology language profiles (second edition), W3C Recommendation, World Wide Web Consortium (W3C), 2012. <https://www.w3.org/TR/owl2-profiles/>.
- [54] B. Motik, Y. Nenov, R.E.F. Piro and I. Horrocks, Incremental update of datalog materialisation: The backward/forward algorithm, in: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, OJS/PPK, 2015. doi:10.1609/aaai.v29i1.9409.
- [55] 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: Proceedings, Part II of the 14th International Semantic Web Conference*, Springer, Cham, Switzerland, 2015, pp. 3–20. doi:10.1007/978-3-319-25010-6_1.
- [56] 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: Proceedings of the 37th Annual German Conference on AI*, Springer, Cham, Switzerland, 2014, pp. 183–194. doi:10.1007/978-3-319-11206-0_18.
- [57] T.-L. Pham, M.I. Ali and A. Mileo, Enhancing the scalability of expressive stream reasoning via input-driven parallelization, *Semantic Web* **10**(3) (2019), 457–474. doi:10.3233/SW-180330.
- [58] D. Puiu, P. Barnaghi, R. Tönjes, D. Kümper, M.I. Ali, A. Mileo, J.X. Parreira, M. Fischer, S. Kolozali, N. Farajidavar et al., CityPulse: Large scale data analytics framework for smart cities, *IEEE Access* **4** (2016), 1086–1108. doi:10.1109/ACCESS.2016.2541999.
- [59] X. Ren, O. Curé, H. Naacke and G. Xiao, BigSR: Real-time expressive RDF stream reasoning on modern big data platforms, in: *2018 IEEE International Conference on Big Data (Big Data)*, IEEE, New York, NY, USA, 2018, pp. 811–820. doi:10.1109/BigData.2018.8621947.
- [60] P. Schaar, Privacy by design, *Identity in the Information Society* **3**(2) (2010), 267–274. doi:10.1007/s12394-010-0055-x.
- [61] Sofia2, Sofia2 – Technology for innovators, 2020, Accessed: 2022-03-10. <https://sofia2.com>.
- [62] S. Soursos, I.P. Žarko, P. Zwickl, I. Gojmerac, G. Bianchi and G. Carrozzo, Towards the cross-domain interoperability of IoT platforms, in: *2016 European Conference on Networks and Communications (EuCNC)*, IEEE, New York, NY, USA, 2016, pp. 398–402. doi:10.1109/EuCNC.2016.7561070.
- [63] B. Steenwinckel, M. De Brouwer, M. Stojchevska, J. Van Der Donckt, J. Nelis, J. Ruysinck, J. van der Hertten, K. Casier, J. Van Ooteghem, P. Crombez, F. De Turck, S. Van Hoecke and F. Ongenaë, Data analytics for health and connected care: Ontology, knowledge graph and applications, in: *Proceedings of the 16th EAI Pervasive Healthcare Conference*, 2022. <https://dahcc.idlab.ugent.be>.
- [64] H. Stuckenschmidt, S. Ceri, E. Della Valle and F. Van Harmelen, Towards expressive stream reasoning, in: *Semantic Challenges in Sensor Networks, Dagstuhl Seminar Proceedings*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2010. doi:10.4230/DagSemProc.10042.4.
- [65] X. Su, E. Gilman, P. Wetz, J. Riecki, Y. Zuo and T. Leppänen, Stream reasoning for the Internet of Things: Challenges and gap analysis, in: *Proceedings of the 6th International Conference on Web Intelligence, Mining and Semantics (WIMS 2016)*, Association for Computing Machinery (ACM), New York, NY, USA, 2016, pp. 1–10. doi:10.1145/2912845.2912853.

- [66] X. Su, J. Riekk, J.K. Nurminen, J. Nieminen and M. Koskimies, Adding semantics to Internet of Things, *Concurrency and Computation: Practice and Experience* **27**(8) (2015), 1844–1860. doi:[10.1002/cpe.3203](https://doi.org/10.1002/cpe.3203).
- [67] V. Subramaniaswamy, G. Manogaran, R. Logesh, V. Vijayakumar, N. Chilamkurti, D. Malathi and N. Senthilselvan, An ontology-driven personalized food recommendation in IoT-based healthcare system, *The Journal of Supercomputing* **75**(6) (2019), 3184–3216. doi:[10.1007/s11227-018-2331-8](https://doi.org/10.1007/s11227-018-2331-8).
- [68] E. Thomas, J.Z. Pan and Y. Ren, TrOWL: Tractable OWL 2 reasoning infrastructure, in: *The Semantic Web: Research and Applications: Proceedings, Part II of the 7th Extended Semantic Web Conference, ESWC 2010*, Springer, Berlin, Heidelberg, 2010, pp. 431–435. doi:[10.1007/978-3-642-13489-0_38](https://doi.org/10.1007/978-3-642-13489-0_38).
- [69] R. Tommasini, P. Bonte, F. Ongenaes and E. Della Valle, RSP4J: An API for RDF stream processing, in: *The Semantic Web: Proceedings of the 18th International Conference, ESWC 2021*, R. Verborgh, K. Hose, H. Paulheim, P.-A. Champin, M. Maleshkova, O. Corcho, P. Ristoski and M. Alam, eds, Springer, Cham, Switzerland, 2021, pp. 565–581. doi:[10.1007/978-3-030-77385-4_34](https://doi.org/10.1007/978-3-030-77385-4_34).
- [70] R. Tommasini and E. Della Valle, Yasper 1.0: Towards an RSP-QL engine, in: *Proceedings of the ISWC 2017 Posters & Demonstrations and Industry Tracks, Co-Located with 16th International Semantic Web Conference (ISWC 2017)*, CEUR Workshop Proceedings, 2017. <https://ceur-ws.org/Vol-1963/paper487.pdf>.
- [71] F. Ullah, M.A. Habib, M. Farhan, S. Khalid, M.Y. Durrani and S. Jabbar, Semantic interoperability for big-data in heterogeneous IoT infrastructure for healthcare, *Sustainable cities and society* **34** (2017), 90–96. doi:[10.1016/j.scs.2017.06.010](https://doi.org/10.1016/j.scs.2017.06.010).
- [72] J. Urbani, C. Jacobs and M. Krötzsch, Column-oriented datalog materialization for large knowledge graphs, in: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, OJS/PKP, 2016. doi:[10.1609/aaai.v30i1.9993](https://doi.org/10.1609/aaai.v30i1.9993).
- [73] J. Urbani, A. Margara, C. Jacobs, F.v. Harmelen and H. Bal, Dynamite: Parallel materialization of dynamic RDF data, in: *The Semantic Web – ISWC 2013: Proceedings, Part I of the 12th International Semantic Web Conference*, Springer, Berlin, Heidelberg, 2013, pp. 657–672. doi:[10.1007/978-3-642-41335-3_41](https://doi.org/10.1007/978-3-642-41335-3_41).
- [74] R. Verborgh and J. De Roo, Drawing conclusions from linked data on the web: The EYE reasoner, *IEEE Software* **32**(3) (2015), 23–27. doi:[10.1109/MS.2015.63](https://doi.org/10.1109/MS.2015.63).
- [75] G. Xiao, L. Ding, B. Cogrel and D. Calvanese, Virtual knowledge graphs: An overview of systems and use cases, *Data Intelligence* **1**(3) (2019), 201–223. doi:[10.1162/dint_a_00011](https://doi.org/10.1162/dint_a_00011).
- [76] R. Zgheib, S. Kristiansen, E. Conchon, T. Plageman, V. Goebel and R. Bastide, A scalable semantic framework for IoT healthcare applications, *Journal of Ambient Intelligence and Humanized Computing* (2020), 1–19. doi:[10.1007/s12652-020-02136-2](https://doi.org/10.1007/s12652-020-02136-2).