eSPARQL: Representing and Reconciling Agnostic and Atheistic Beliefs in RDF-star Knowledge Graphs

Xinyi Pan[2], Daniel Hernandez[2], Philipp Seifer[1], Ralf Lämmel[1], and Steffen Staab[2,3]

[1] Institute for Artifical Intelligence, University of Stuttgart, Germany
`xinyi.pan@ms.informatik.uni-stuttgart.de`
`{daniel.hernandez, steffen.staab}@ki.uni-stuttgart.de`
[2] The Software Languages Team, University of Koblenz, Germany
`{pseifer, laemmel}@uni-koblenz.de`
[3] Web and Internet Science Research Group, University of Southampton, United Kingdom

Abstract. Over the past few years, we have seen the emergence of large knowledge graphs combining information from multiple sources. Sometimes, this information is provided in the form of assertions about other assertions, defining contexts where assertions are valid. A recent extension to RDF which admits statements over statements, called RDF-star, is in revision to become a W3C standard. However, there is no proposal for a semantics of these RDF-star statements nor a built-in facility to operate over them. In this paper, we propose a query language for epistemic RDF-star metadata based on a four-valued logic, called eSPARQL. Our proposed query language extends SPARQL-star, the query language for RDF-star, with a new type of `FROM` clause to facilitate operating with multiple and sometimes conflicting beliefs. We show that the proposed query language can express four use case queries, including the following features: (i) querying the belief of an individual, (ii) the aggregating of beliefs, (iii) querying who is conflicting with somebody, and (iv) beliefs about beliefs (i.e., nesting of beliefs).

Keywords: SPARQL, RDF-star, knowledge graphs, epistemic querying

1 Introduction

Over the past few years, we have seen the emergence of large knowledge graphs combining information from multiple sources, such as Wikidata [21], YAGO [20], and DBpedia [3]. These knowledge graphs provide information about a great variety of entities, such as people, countries, universities, as well as facts over these entities, which are codified as triples. For example, the fact that Albert Einstein was born in Germany can be encoded as a triple (Einstein, wasBorn, Germany). Triples like this are the information units of the Resource Description Framework (RDF) [12], a data model for representing information about World Wide

Web resources, and SPARQL [10], the query language proposed by the W3C to query RDF data.

RDF and SPARQL have been built on the simplifying assumption that statements (triples) are either *true* or *unknown*, leading to a representation of knowledge that can never conflict with each other, and thus can never lead to inconsistencies. To present competing statements, knowledge bases include reified statements without introducing them as a stated statement. For example, the statement "Jesus is a deity" is provided in a reified form in Wikidata, and annotated as "supported by Christianity, Messianic Judaism, and Manichaeism" but "disputed by Islam, Atheism, and Judaism". Acknowledging the need to make statements about statements, two recent extensions, namely RDF-star and SPARQL-star [11], were proposed. However, these extensions do not deviate from the assumption that statements are either true or unknown.

To illustrate the need to work with statements that are *true*, *unknown*, *false*, or *conflicted*, let us consider some disagreements on beliefs in the Christian religion. Several councils were convened to establish consensus about different aspects of the nature of Jesus, his birth, his mother's birth, his father, and his existence before his birth. By stating a dogma, a council established what should be considered true and what is a heresy. For example, one of these beliefs can be encoded in RDF-star as a triple where the last element is also a triple:

$$(\mathsf{PopeDI}, \mathsf{believesToBeTrue}, (\mathsf{Jesus}, \mathsf{a}, \mathsf{FullDeity})).$$

RDF-star does not provide a semantics for this statement, but we can assume that it means "Pope Damasus the First believes that Jesus is a full Deity." Indeed, the Pope presided over the Council of Nice, where this statement became a dogma. This belief was not shared by all participants of the Council. Arius believed that Jesus was created by decision of God, and thus his deity is not as full as the deity of his father. Arius' belief can be encoded as follows:

$$(\mathsf{Arius}, \mathsf{believesToBeFalse}, (\mathsf{Jesus}, \mathsf{a}, \mathsf{FullDeity})).$$

The two statements above do not entail that "Jesus is a full deity" nor the contrary, but ascribe some different people's beliefs. Since we have contradictory opinions between Christians, we can say that this statement is conflicted according to Christianity:

$$(\mathsf{Christianity}, \mathsf{believesToBeConflicted}, (\mathsf{Jesus}, \mathsf{a}, \mathsf{FullDeity})).$$

On the other hand, the mathematician Bertrand Russell described his religious beliefs as follows: *"In regard to the Olympic gods, speaking to a purely philosophical audience, I would say that I am an Agnostic. But speaking popularly, I think that all of us would say in regard to those gods that we were Atheists. In regard to the Christian God, I should, I think, take exactly the same line."* Regarding the aforementioned statement, Russell's beliefs can be encoded as follows:

$$(\mathsf{Russell}, \mathsf{believesToBeUnknown}, (\mathsf{Jesus}, \mathsf{a}, \mathsf{FullDeity})),$$
$$(\mathsf{Russell}, \mathsf{believesToBeUnknown}, (\mathsf{Zeus}, \mathsf{a}, \mathsf{FullDeity})).$$

Russell defines himself as agnostic, since he cannot decide whether the statement is false or true, and he also describes his beliefs about the beliefs of Christians who may be believers regarding Jesus, but atheists regarding an Olympic deity. This Russell's argument applied to PopeDI is encoded as follows:

(Russell, believesToBeTrue, (PopeDI, believesToBeTrue, (Jesus, a, FullDeity))),

(Russell, believesToBeTrue, (PopeDI, believesToBeFalse, (Zeus, a, FullDeity))).

In order to deal with multiple views of the world, it becomes necessary to operate on sets of facts. For example, we would like to execute a SPARQL query against the set of all facts that are valid according to Christianity, or a given person. However, SPARQL does hardly support working with sets of facts. Indeed, SPARQL has an ambivalent attitude towards sets of facts, as it allows constructing such using `FROM`, but its main operators all work on graph pattern matching. Aside from the SPARQL `UPDATE` functionality, the only operators to do so are `FROM` and `FROM NAMED`. We propose to extend the `FROM` functionality of SPARQL and call this extension eSPARQL to establish actual sets of facts based on which triple pattern matching and downstream operations (`FILTER`, `AGGREGATE`, etc.) may be applied.

eSPARQL allows constructing graphs in `FROM` clauses filtering triples with patterns that can include epistemic conditions. For example, the following eSPARQL query asks for all full deities according to Christianity.

```
1  SELECT ?deity
2  FROM BELIEF <Christianity>
3  WHERE { ?deity a <FullDeity> }
```

Intuitively, the clause `FROM BELIEF <Christianity>` generates a graph whose triples are marked as either *true*, *false*, *unknown*, or *conflicting*. These truth values represent the `<Christianity>` beliefs. The graph that results from extracting the Pope beliefs is then used for matching the triple pattern `?deity a <FullDeity>`, and returning the bindings for variable `?deity`. The answers for Jesus, his father God, and his mother Maria will be annotated as conflicted, true, and false, respectively, and individuals that are not mentioned will be annotated as unknown.

This paper makes the following contributions:

1. We describe use case requirements for an epistemic query language (Section 2).
2. We provide a semantics for RDF-star and SPARQL-star as a $\mathcal{K}$-annotated algebra, which operates over functions from SPARQL-star solution mappings to $\mathcal{K}$ elements (Section 4). This algebra extends the Geerts et al. $\mathcal{K}$-annotated algebra [6], which is used as the foundation for how-provenance in SPARQL.
3. We propose a query language, eSPARQL, an extension to the $\mathcal{K}$-annotated SPARQL-star where the abstract semiring structure $\mathcal{K}$ is substituted by the concrete semirings in the $\mathcal{FOUR}$ structure (Section 5).
4. We show that eSPARQL can express the use case queries, and present an user interface for the eSPARQL algebra (Section 6).

5. We show that SPARQL queries can lead to query results that cannot be encoded with finite expressions using the support of the solution (Section 7). We propose an approach to deal with these infinite results.

2   Use Case Requirements

We have explored a broad range of queries that an epistemic query language for RDF-star should be able to answer. We distil four query examples that this language must allow formulating easily. They require (i) to query for people's beliefs, (ii) to query for aggregated or integrated belief states, (iii) to query for subjects whose beliefs coincide, conflict, are compatible or ignorant of other belief sets, (iv) allow for nested belief states.

Use Case U1 (Query for people's beliefs) We should be able to query people's beliefs regarding a given statement, and the query should return the bindings to the possible variables in the statement and the respective four-valued logic truth values for the variable bindings. The expression for these queries should be simple (i.e., do not include any algebraic operators like AND, UNION, FILTER, or SELECT).
   Example: Return the full deities ?x according to Pope Damasus the First. All returned individuals should be annotated with a belief value that indicates the belief of the Pope regarding the individual. For example, Jesus must be annotated as *true* if the Pope believes Jesus is a full deity, *false* if the Pope believes Jesus is not a full deity, *unknown* if there is no information about the belief of the Pope regarding the nature of Jesus, and *conflicted* if either the Pope believes this statement is conflicted, or the belief is both *true* and *false.*

Use Case U2 (Query for combined belief values) We should be able to combine people's beliefs to represent the belief of a group, or operations between different people's beliefs.
   Example: Return the beliefs of Christian people regarding the full deities. For example, Jesus must be annotated with true if all Christians think Jesus is a *full* deity, *false* if all Christians think Jesus is not a full deity, *conflicted* if they have contradictory opinions, and unknown if no Christian has an opinion regarding the nature of Jesus.

Use Case U3 (Query for people whose beliefs coincide, conflict, or are compatible) Given a set of statements, we should be able to query people who coincide (they all have the same belief for all the statement), conflict (some people believe a statement is true, whereas others believe it is false), or they are compatible (if a person believes a fact is true, the others can believe it is true or unknown).
   Example: Return all people whose opinion is conflicted with Pope Damasus the First in at least one statement. For example, since Pope Damasus the First believes that Jesus is a full deity, and Arius believes it is not, then Arius should

be returned (i.e., annotated with true). Since Russell is agnostic regarding an opinion of Pope Damasus the First, Russell is not in conflict with the Pope, so Russell must be annotated with false.

Use Case U4 (Query nested belief states) We should be able to query about what people believe other people believe.

Example: Return all the people x such that there exists a person y that believes that x believes that Zeus is not a full deity. In our example in the introduction, we must annotate Pope Damasus the First with true, because Russell believes the pope believes Zeus is not a full deity. Other individuals should be annotated with unknown.
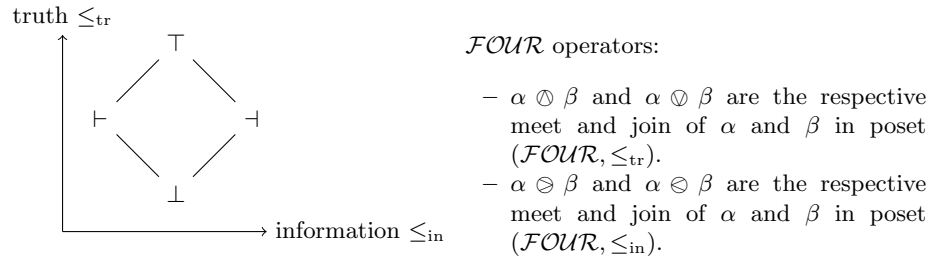
3   Preliminaries

The $\mathcal{FOUR}$ structure. A *partially ordered set* (or *poset*) is a set taken together with a partial order on it. Formally, a partially ordered set is defined as a pair $(A, \leq)$, where A is called the *ground set* of the poset and $\leq$ is the *partial order* of the poset.

A poset $(A, \leq)$ defines a lattice if for every two elements $a, b \in A$, there is a unique maximum element $c \in A$ such that $c \leq a$ and $c \leq b$, called the *meet* of a and b, and there is a unique minimum element $c \in A$ such that $a \leq c$ and $b \leq c$, called the *join* of a and b.

We write $\mathcal{FOUR}$ to denote the set $\{\bot, \top, \vdash, \dashv\}$, and we use $\mathcal{FOUR}$ *elements* to refer to the elements in this set. We define the posets $(\mathcal{FOUR}, \leq_{\mathrm{tr}})$ and $(\mathcal{FOUR}, \leq_{\mathrm{in}})$ as the minimum posets satisfying the inequalities $\bot \leq_{\mathrm{tr}} \vdash$, $\bot \leq_{\mathrm{tr}} \dashv$, $\vdash \leq_{\mathrm{tr}} \top$, $\dashv \leq_{\mathrm{tr}} \top$, $\vdash \leq_{\mathrm{in}} \bot$, $\vdash \leq_{\mathrm{in}} \top$, $\bot \leq_{\mathrm{in}} \dashv$, and $\top \leq_{\mathrm{in}} \dashv$. These two posets with ground set $\mathcal{FOUR}$, define the bilattice depicted in Figure 1. The axis can be interpreted as the orderings regarding truth and information, and the $\mathcal{FOUR}$ elements are the respective elements of the 4-valued logic *false* ($\bot$), *true* ($\top$), *unknown* ($\vdash$), and *conflicting* ($\dashv$). We use $\mathcal{FOUR}$ *operators* to refer to the operators $\oslash$, $\varobar$, $\ominus$, and $\oslash$ (see Figure 1), and $\mathcal{FOUR}$ *structure* to refer to the algebraic structure that consists of the $\mathcal{FOUR}$ operators over the $\mathcal{FOUR}$ set. Given a $\mathcal{FOUR}$ operator $\circ$ and a set $A = \{\alpha_1, \ldots, \alpha_n\} \subseteq \mathcal{FOUR}$, we write $\mathrm{Id}(\circ)$ to denote the identity element for an operator $\circ$ in the $\mathcal{FOUR}$ structure (i.e., $\mathrm{Id}(\oslash) = \top$, $\mathrm{Id}(\varobar) = \bot$, $\mathrm{Id}(\ominus) = \dashv$, and $\mathrm{Id}(\oslash) = \vdash$), we write $\mathrm{Abb}(\circ)$ to denote the absorbing element for a $\mathcal{FOUR}$ operator $\circ$, (i.e., $\mathrm{Abb}(\oslash) = \bot$, $\mathrm{Abb}(\varobar) = \top$, $\mathrm{Abb}(\ominus) = \vdash$, and $\mathrm{Abb}(\oslash) = \dashv$), and we write $\circ(A)$ for the $\mathcal{FOUR}$ element $\mathrm{Id}(\circ) \circ \alpha_1 \circ \cdots \circ \alpha_n$.

A *monoid* $\mathcal{M}$ is an algebraic structure $(M, +_{\mathcal{M}}, 0_{\mathcal{M}})$ where M is a non-empty set, $+_{\mathcal{M}}$ is a closed associative operation on set M, and $0_{\mathcal{M}} \in M$ is the identity element of operation $+_{\mathcal{M}}$ (i.e., $0_{\mathcal{M}} +_{\mathcal{M}} a = a$ and $a +_{\mathcal{M}} 0_{\mathcal{M}} = a$, for every $a \in M$). The monoid $\mathcal{M}$ is said to be *commutative* if operation $+_{\mathcal{M}}$ is commutative.

A *semiring* $\mathcal{K}$ is an algebraic structure $(K, +_{\mathcal{K}}, \times_{\mathcal{K}}, 0_{\mathcal{K}}, 1_{\mathcal{K}})$ where $(K, +_{\mathcal{K}}, 0_{\mathcal{K}})$ and $(K, \times_{\mathcal{K}}, 1_{\mathcal{K}})$ are monoids, and $\times_{\mathcal{K}}$ distributes over $+_{\mathcal{K}}$ (i.e., $a \times_{\mathcal{K}} (b +_{\mathcal{K}} c) =$

truth $\leq_{\mathrm{tr}}$

$\mathcal{FOUR}$ operators:

- $\alpha \oslash \beta$ and $\alpha \oslash \beta$ are the respective meet and join of α and β in poset $(\mathcal{FOUR}, \leq_{\mathrm{tr}})$.
- $\alpha \oslash \beta$ and $\alpha \oslash \beta$ are the respective meet and join of α and β in poset $(\mathcal{FOUR}, \leq_{\mathrm{in}})$.

information $\leq_{\mathrm{in}}$

Fig. 1. Bilattice of the $\mathcal{FOUR}$-structure.

$(a \times_{\mathcal{K}} b) +_{\mathcal{K}} (a \times_{\mathcal{K}} c)$ and $(b +_{\mathcal{K}} c) \times_{\mathcal{K}} a = (b \times_{\mathcal{K}} a) +_{\mathcal{K}} (c \times_{\mathcal{K}} a))$. The semiring $\mathcal{K}$ is said to be *commutative* if operations $+_{\mathcal{K}}$ and $\times_{\mathcal{K}}$ are commutative. For example, the algebraic structures $(\mathcal{FOUR}, \oslash, \oslash, \bot, \top)$, $(\mathcal{FOUR}, \oslash, \oslash, \vdash, \dashv)$, $(\{\bot, \top\}, \oslash, \oslash, \bot, \top)$, and $(\mathbb{N}, +, \cdot, 0, 1)$ are commutative semirings.

We will write $\mathcal{FOUR}_{\mathrm{th}}$ and $\mathcal{FOUR}_{\mathrm{in}}$ for the semirings $(\mathcal{FOUR}, \oslash, \oslash, \bot, \top)$ and $(\mathcal{FOUR}, \oslash, \oslash, \vdash, \dashv)$, respectively.

4 $\mathcal{K}$-annotated SPARQL-star

In this section, we extend a fragment of the Geerts et al. [6] $\mathcal{K}$-annotated SPARQL algebra to support RDF-star and SPARQL-star (following the notions from the working draft [11]). We consider a fragment without the non-monotonic operator MINUS because the non-monotonic operators go beyond the semiring structure of each of the two $\mathcal{FOUR}$ lattices.

We assume two countable pairwise disjoint sets $\mathbf{I}$ and $\mathbf{V}$, called the set of *IRIs* and the set of *variables*. We call the elements of $\mathbf{I}^3$ *RDF triples*[4]. Given an RDF triple $(s, p, o) \in \mathbf{I}^3$, we say that s is the *subject*, p is the *predicate*, and o is the *object* of the triple. According to the working draft [11], *SPARQL-star triples* and *SPARQL-star triple patterns* are defined recursively as follows. Given two SPARQL-star triple patterns T_1 and T_2, and an RDF triple pattern (S, P, O), the triples (S, P, O), (T_1, P, O), (S, P, T_2), and (T_1, P, T_2) are SPARQL-star triple patterns. A SPARQL-star triple pattern without variables is an *RDF-star triple* (and $\mathcal{T}$ is the set of all RDF-star triples), and an *RDF-star graph* is a set of RDF-star triples.

A *SPARQL-star solution mapping* (or a *mapping*) is a partial function $\mu :$ $\mathbf{V} \to \mathbf{I} \cup \mathcal{T}$ with finite domain $\mathrm{dom}(\mu)$. Two mappings μ and μ' are *compatible* if $\mu(?\mathtt{x}) = \mu'(?\mathtt{x})$ for every variable $?\mathtt{x} \in \mathrm{dom}(\mu) \cap \mathrm{dom}(\mu')$. If mappings μ and μ' are compatible, $\mu \cup \mu'$ is the mapping with domain $\mathrm{dom}(\mu) \cup \mathrm{dom}(\mu')$ that is compatible with μ and μ'. Given a set of variables $W \subset \mathbf{V}$, we write $\Omega|_W$ to denote the set of mappings $\{\mu : \mathrm{dom}(\mu) = W\}$.

---

[4] In RDF, triples can also have other element types, namely blank nodes and literals. For simplicity, we do not consider them, since they do not change our results.

We next extend RDF-star and SPARQL-star to support annotations over a semiring $\mathcal{K}$. Given a commutative semiring $\mathcal{K} = (K, +_{\mathcal{K}}, \times_{\mathcal{K}}, 0, 1)$, an *RDF-star $\mathcal{K}$-graph* is a function $G : \mathcal{T} \to K$ that maps every RDF-star triple to an element of the semiring $\mathcal{K}$.

Remark 1. Notice that $\mathcal{K}$ denotes an arbitrary semiring, whereas $\mathcal{FOUR}_{\text{th}}$ and $\mathcal{FOUR}_{\text{in}}$ are two concrete semirings. Thus, the $\mathcal{K}$-annotated SPARQL algebra by Geerts et al. [6] is defined in abstract, whereas the $\mathcal{FOUR}_{\text{th}}$-annotated and the $\mathcal{FOUR}_{\text{in}}$-annotated SPARQL-star algebras are concrete instances of the abstract algebra proposed by Geerts et al.

Definition 1 (Syntax of $\mathcal{K}$-annotated SPARQL-star). *We next present a recursive definition for the syntax of $\mathcal{K}$-annotated SPARQL-star queries Q and a finite set of variables, called the* in-scope variables, *denoted* inScope(Q).

1. *A SPARQL-star triple pattern T is a $\mathcal{K}$-annotated SPARQL-star query, whose in-scope variables are the variables occurring in triple pattern T.*
2. *Given a $\mathcal{K}$-annotated SPARQL-star query Q and a set of variables $W \subseteq$ inScope(Q), the expression (SELECT W Q) is a $\mathcal{K}$-annotated SPARQL-star query with in-scope variables W.*
3. *Given a $\mathcal{K}$-annotated SPARQL-star query Q and a formula φ, the expression (Q FILTER φ) is a $\mathcal{K}$-annotated SPARQL-star query with in-scope variables inScope(Q).*
4. *Given two $\mathcal{K}$-annotated SPARQL-star queries Q_1 and Q_2, the expression (Q_1 AND Q_2) is a $\mathcal{K}$-annotated SPARQL-star query with in-scope variables inScope(Q_1) $\cup$ inScope(Q_2).*
5. *Given two $\mathcal{K}$-annotated SPARQL-star queries Q_1 and Q_2 with the same set W of in-scope variables, the expression (Q_1 UNION Q_2) is a $\mathcal{K}$-annotated SPARQL-star query with in-scope variables W.*

Definition 2 (Semantics of $\mathcal{K}$-annotated SPARQL-star). *Given a semiring $\mathcal{K} = (K, +_{\mathcal{K}}, \times_{\mathcal{K}}, 0_{\mathcal{K}}, 1_{\mathcal{K}})$, the result of evaluating a $\mathcal{K}$-annotated SPARQL-star query Q over a $\mathcal{K}$-graph G, is a function $[\![Q]\!]_G : \Omega|_{\text{inScope}(Q)} \to K$, called $\mathcal{K}$-relation, defined recursively as follows:*

1. *Given a SPARQL-star triple pattern T, if $\text{dom}(\mu) = \text{inScope}(T)$ then $[\![Q]\!]_G(\mu)$ is $G(t)$, where t is the RDF-star triple resulting from replacing every variable $?x$ occurring in T with $\mu(?x)$; otherwise, $[\![Q]\!]_G(\mu) = 0_{\mathcal{K}}$.*
2. *For the recursive cases, the semantics is the following:*
   (a) $[\![(\text{SELECT } W\, Q)]\!]_G(\mu) = \sum_{\mu' : \mu'|_W = \mu} [\![Q]\!]_G(\mu')$,
   (b) $[\![(Q \text{ FILTER } \varphi)]\!]_G(\mu) = [\![Q]\!]_G(\mu) \times_{\mathcal{K}} k_{\mu \models \varphi}$,
   (c) $[\![(Q_1 \text{ UNION } Q_2)]\!]_G(\mu) = [\![Q_1]\!]_G(\mu) +_{\mathcal{K}} [\![Q_2]\!]_G(\mu)$,
   (d) $[\![(Q_1 \text{ AND } Q_2)]\!]_G(\mu) = [\![Q_1]\!]_G(\mu|_{\text{inScope}(Q_1)}) \times_{\mathcal{K}} [\![Q_2]\!]_G(\mu|_{\text{inScope}(Q_2)})$,
   where $\sum$ denotes sums using $+_{\mathcal{K}}$, $\mu'|_W$ is the projection of mapping μ' on the variables in W, and $k_{\mu \models \varphi}$ is $1_{\mathcal{K}}$ if replacing the variables in φ according to μ leads to a true formula, and $k_{\mu \models \varphi}$ is $0_{\mathcal{K}}$, otherwise.

Table 1. RDF-star triples in the running example graph $G = \{t_1, \ldots, t_8\}$. The RDF-star triples t_9, t_{10}, and t_{11} are not stated in G but occur nested in the triples in G.

Id	Triple
t_1	(PopeDI, believesToBeTrue, (Jesus, a, FullDeity)),
t_2	(Arius, believesToBeFalse, (Jesus, a, FullDeity)),
t_3	(Christianity, believesToBeConflicted, (Jesus, a, FullDeity)),
t_4	(Russell, believesToBeUnknown, (Jesus, a, FullDeity))
t_5	(Russell, believesToBeTrue, (PopeDI, believesToBeTrue, (Jesus, a, FullDeity)))
t_6	(Russell, believesToBeTrue, (PopeDI, believesToBeFalse, (Zeus, a, FullDeity)))
t_7	(PopeDI, a, Christian)
t_8	(Arius, a, Christian)
t_9	(Jesus, a, FullDeity),
t_{10}	(PopeDI, believesToBeFalse, (Zeus, a, FullDeity))
t_{11}	(Zeus, a, FullDeity)

Example 1 (Running Example). Let $G = \{t_1 \mapsto \top, \ldots, t_8 \mapsto \top, * \mapsto \vdash\}$ be the RDF-star $\mathcal{FOUR}_{\text{in}}$-graph such that $G(t_i) = \top$ for every RDF-star triple t_i, with $1 \leq i \leq 8$, listed in Table 1, and $G(t') = \vdash$ for the rest of all possible RDF-star triples t'. Then, $[\![(?\mathsf{x}, \mathsf{believesToBeFalse}, (?\mathsf{y}, \mathsf{a}, \mathsf{FullDeity}))]\!]_G$ returns the $\mathcal{K}$-relation $R = \{\{?\mathsf{x} \mapsto \mathsf{Arius}, ?\mathsf{y} \mapsto \mathsf{Jesus}\} \mapsto \top, * \mapsto \vdash\}$. Intuitively, this query returns all people who believe it to be false that somebody is a full deity.

Support of functions. So far, we have described the $\mathcal{FOUR}$ structure, and the extended Geerts et al. [6] $\mathcal{K}$-annotated algebra. In principle, the elements of this algebra, $\mathcal{K}$-graphs and $\mathcal{K}$-relations, can have no finite representation. One of the concepts used to characterize $\mathcal{K}$-graphs and $\mathcal{K}$-relations with a finite representation is the support of a function, which we describe next.

Given two functions $f : A \to B$ and $g : A \to B$, and a binary operation $\diamond$ closed in B, $f \diamond g$ is the function such that $(f \diamond g)(x) = f(x) \diamond g(x)$. If $\diamond$ has an identity element $\mathrm{Id}(\diamond)$, the *support* of f regarding operation $\diamond$ is the set $\mathrm{supp}_\diamond(f) = \{a \in A : f(a) \neq \mathrm{Id}(\diamond)\}$. If the support of f is a finite set $\{a_1, \ldots, a_n\}$, we encode f with the finite mapping $\{a_1 \mapsto f(a_1), \ldots, a_n \mapsto f(a_n), * \mapsto \mathrm{Id}(\diamond)\}$, where $*$ is a wildcard for the (possibly infinite) remaining mappings that are mapped to the identity of the operation.

Remark 2. Intuitively, the support represents all values that are assigned to *non-zero* values. In the case of a semiring $\mathcal{K} = (K, +_\mathcal{K}, \times_\mathcal{K}, 0_\mathcal{K}, 1_\mathcal{K})$ the zero value is $0_\mathcal{K}$, whereas in the $\mathcal{FOUR}$ structure we can consider two zero values, namely $\perp$ and $\vdash$, depending on which semiring we are considering. These two values are the corresponding bottoms of the two lattices in the $\mathcal{FOUR}$ structure. Intuitively, the zero values in $\mathcal{FOUR}$ represent the default value for statements that are not included in a knowledge base, according to either the *closed-world assumption* or the *open-world assumption*.

Remark 3. According to the specification, an RDF-star graph is a finite set of RDF-star triples $\{T_1, \ldots, T_n\}$ and under the open-world assumption (which is the standard for RDF-star) a triple T is true if $T \in G$ and unknown if $T \notin G$. Hence, the RDF-star graph G can be interpreted as the $\mathcal{FOUR}_{\text{in}}$-graph that is represented by the finite mapping $\{T_1 \mapsto \top, \ldots, T_n \mapsto \top, * \mapsto \vdash\}$. Furthermore, the evaluation of a $\mathcal{K}$-annotated SPARQL-star query over a $\mathcal{K}$-graph with finite support always returns a $\mathcal{K}$-relation with finite support. This finite property, guarantees that the algebra can be implemented.

Proposition 1. *Given a semiring $\mathcal{K}$, for every RDF-star $\mathcal{K}$-graph G and every SPARQL-star query Q, the $\mathcal{K}$-relation $[\![Q]\!]_G$ has finite support.*

Proof. This can be shown by induction on the structure of the query.

Remark 4. The $\mathcal{FOUR}_{\text{in}}$-annotated SPARQL-star algebra allows us to query $\mathcal{FOUR}_{\text{in}}$-graphs, but does not provide the means to easily formulate the epistemic queries described in Section 2. This motivates the proposal of the query language we describe in the next section.

5   Epistemic SPARQL

This section presents the syntax and semantics of Epistemic SPARQL (eS-PARQL), the language designed for the use cases described in Section 2.

So far, we have defined the notion of SPARQL-star $\mathcal{K}$-graphs, where $\mathcal{K}$ is a semiring. However, a $\mathcal{K}$-graph is an abstract notion, since no concrete semiring is provided. If concrete semirings like $\mathcal{FOUR}_{\text{th}}$ and $\mathcal{FOUR}_{\text{in}}$ are considered, then we have concrete notions as $\mathcal{FOUR}_{\text{th}}$-graphs and $\mathcal{FOUR}_{\text{in}}$-graphs (see Example 1). In what follows, a $\mathcal{FOUR}$-graph will be a function that associates RDF-star triples to elements in set $\mathcal{FOUR}$, without choosing one of the two semirings ($\mathcal{FOUR}_{\text{th}}$ or $\mathcal{FOUR}_{\text{in}}$) defined by the $\mathcal{FOUR}$ structure.

A key characteristic of $\mathcal{FOUR}$-graphs is that they can contain epistemic metadata encoded using four predicates: believesToBeTrue, believesToBeFalse, believesToBeUnknown, and believesToBeConflicted. Each of these predicates encodes a belief. For example, triple $t_1 = (\text{PopeDI}, \text{believesToBeTrue}, t_9)$ from Table 1 encodes that PopeDI believes statement t_9 to be true.

A basic operation over a $\mathcal{FOUR}$-graph G is thus extracting beliefs for each of these predicates. We next present *belief queries*, which are expressions to extract a $\mathcal{FOUR}$-graph G' representing the belief of one or more people according to the information of an input $\mathcal{FOUR}$-graph G. An atomic belief query $[\text{PopeDI}, \top, \vdash]$ represents all statements PopeDI believe to be $\top$, assuming that the rest are annotated with the state $\vdash$. Compound belief queries are generated by combining atomic statements with the $\mathcal{FOUR}$ operators.

Definition 3 (Syntax of a Belief Query). *A* belief query E *is recursively defined as follows:*

1. *Given an element $u \in \mathbf{I} \cup \mathbf{V}$, and two $\mathcal{FOUR}$ elements α and β, the triple $[u, \alpha, \beta]$ is an* atomic *belief query.*

2. *Given two belief queries E_1 and E_2, and a $\mathcal{FOUR}$-operator $\circ$, the expression $(E_1 \circ E_2)$ is a* compound *belief query.*

We write $\text{var}(E)$ *to denote the set of variables occurring in a belief query E. A belief query with no variables is said to be* ground. *We write $(u, \circ)$ to denote the belief query $[u, \top, \text{Id}(\circ)] \circ [u, \bot, \text{Id}(\circ)] \circ [u, \vdash, \text{Id}(\circ)] \circ [u, \dashv, \text{Id}(\circ)]$.*

Definition 4 (Semantics of a Ground Belief Query). *Let $\circ$ be a $\mathcal{FOUR}$-operator. The semantics of a belief query E over a $\mathcal{FOUR}$-graph G is given by the $\mathcal{FOUR}$-graph, denoted $G|_E$, defined recursively as follows:*

1. *Given an element $a \in \mathbf{I}$ and a $\mathcal{FOUR}$ operation $\circ$,*

$$G|_{[a,\top,\beta]}(t) = \begin{cases} \top & \text{if } G((a, \mathsf{believesToBeTrue}, t)) \in \{\top, \dashv\}, \\ \beta & \text{otherwise.} \end{cases}$$

$$G|_{[a,\bot,\beta]}(t) = \begin{cases} \bot & \text{if } G((a, \mathsf{believesToBeFalse}, t)) \in \{\top, \dashv\}, \\ \beta & \text{otherwise.} \end{cases}$$

$$G|_{[a,\vdash,\beta]}(t) = \begin{cases} \vdash & \text{if } G((a, \mathsf{believesToBeUnknown}, t)) \in \{\top, \dashv\}, \\ \beta & \text{otherwise.} \end{cases}$$

$$G|_{[a,\dashv,\beta]}(t) = \begin{cases} \dashv & \text{if } G((a, \mathsf{believesToBeConflicted}, t)) \in \{\top, \dashv\}, \\ \beta & \text{otherwise.} \end{cases}$$

2. $G|_{(E_1 \circ E_2)}(t) = G|_{E_1}(t) \circ G|_{E_2}(t)$.

The reader may wonder why we find believes on statements that are annotated as $\top$ or $\dashv$ instead of only $\top$. The reason is that, regarding the information dimension, a statement annotated with $\dashv$ is equivalent to a statement annotated twice, as $\top$ and $\bot$. Hence, the $\dashv$ annotated statement includes the $\top$ annotated statement.

Example 2. Let G be the $\mathcal{FOUR}$-graph described in Example 1. Then,

1. $G|_{[\mathsf{PopeDI}, \top, \vdash]} = \{t_9 \mapsto \top, * \mapsto \vdash\}$. Intuitively, the resulting graph has all the statements that are true according to PopeDI and all unstated statements to be unknown.
2. $G|_{(\mathsf{PopeDI}, \oslash)} = \{t_9 \mapsto \top, * \mapsto \vdash\}$, $G|_{(\mathsf{Arius}, \oslash)} = \{t_9 \mapsto \bot, * \mapsto \vdash\}$, and $G|_{(\mathsf{Russell}, \oslash)} = \{t_1 \mapsto \top, t_{10} \mapsto \bot, * \mapsto \vdash\}$. Intuitively, these are the respective beliefs of PopeDI, Arius, and $\mathsf{Russell}$.
3. $G|_{((\mathsf{PopeDI}, \oslash) \oslash (\mathsf{Arius}, \oslash))} = \{t_9 \mapsto \dashv, * \mapsto \vdash\}$. Intuitively, PopeDI and Arius are conflicted regarding the nature of Jesus deity.
4. $G|_{((\mathsf{PopeDI}, \oslash) \oslash (\mathsf{Russell}, \oslash))} = \{t_1 \mapsto \top, t_9 \mapsto \top, t_{10} \mapsto \top, * \mapsto \vdash\}$. Intuitively, PopeDI and $\mathsf{Russell}$ are not conflicted regarding the nature of Jesus deity, so the most informative beliefs is assigned to triple t_9.

Definition 5 (Syntax of eSPARQL Filter Formulas). eSPARQL *filter formulas are defined recursively as follows:*

1. *Given two variables $\textbf{?x}, \textbf{?y} \in \mathbf{V}$ and two IRIs $a, b \in \mathbf{I}$, the expressions $\textbf{?x} = \textbf{?y}$, $\textbf{?x} = a$, $a = b$, and* $\text{bound}(\textbf{?x})$ *are atomic eSPARQL filter formulas.*

2. *Given a $\mathcal{FOUR}$ value α, $\mathrm{state}(\alpha)$ is an atomic eSPARQL filter formula.*
3. *Given two eSPARQL filter formulas φ and ψ, the expressions $\neg\varphi$, $(\varphi \wedge \psi)$ and $(\varphi \vee \psi)$ are compound eSPARQL filter formulas.*

Definition 6 (Semantics of eSPARQL Filter Formulas). *Given a $\mathcal{FOUR}$-relation R, a mapping μ in the domain of R, and a eSPARQL filter formula φ, the truth value of φ on μ is defined recursively as follows:*

1. *If φ has the form $u = v$ then $\mu(\varphi) = \mathsf{error}$ if there is a variable $?\boldsymbol{x}$ in φ such that $?\boldsymbol{x} \notin \mathrm{dom}(\mu)$, $\mu(\varphi) = \mathsf{true}$ if replacing every variable $?\boldsymbol{x}$ in φ with $\mu(\varphi)$ leads to a true entity, and $\mathrm{dom}(\mu)$, otherwise, $\mu(\varphi) = \mathsf{false}$.*
2. *If φ has the form $\mathrm{bound}(?\boldsymbol{x})$ and $?\boldsymbol{x} \in \mathrm{dom}(\mu)$, $\mu(\varphi) = \mathsf{true}$; otherwise $\mu(\varphi) = \mathsf{false}$.*
3. *If φ has the form $\mathrm{state}(\alpha)$, $\mu(\varphi) = \mathsf{true}$ if $R(\mu) = \alpha$; otherwise, $\mu(\varphi) = \mathsf{false}$.*
4. *If φ is a compound eSPARQL filter formula then φ is evaluated with the standard SPARQL three-valued logic semantics of the logical connectives $\neg$, $\wedge$, and $\vee$ for the values true, false, and error (see [16]).*

We define the relation $\mu \models_R \varphi$ to be true if $\mu(\varphi) = \mathsf{true}$ for R.

Intuitively, eSPARQL filter formulas extend SPARQL filter formulas with the ability to check the state of the current mapping.

Definition 7 (Syntax of eSPARQL). *The* syntax of eSPARQL queries *and their* in-scope *variables is defined recursively as follows:*

1. *An RDF-star triple pattern T is an eSPARQL query whose in-scope variables are the variables occurring in T.*
2. *Given two eSPARQL queries Q_1 and Q_2 and a $\mathcal{FOUR}$ operation $\circ \in \{\oslash, \oslash\}$, the expression $(Q_1 \circ Q_2)$ is an eSPARQL query whose in-scope variables are $\mathrm{inScope}(Q_1) \cup \mathrm{inScope}(Q_2)$.*
3. *Given two eSPARQL queries Q_1 and Q_2 with the same set W of in-scope variables and a $\mathcal{FOUR}$ operation $\circ \in \{\oslash, \oslash\}$, the expression $(Q_1 \circ Q_2)$ is an eSPARQL query whose in-scope variables are W.*
4. *Given a eSPARQL query Q, a $\mathcal{FOUR}$ operation $\circ$, and a filter formula φ, the expression $(Q\,\mathsf{FILTER}_\circ\,\varphi)$ is an eSPARQL query with in-scope variables $\mathrm{inScope}(Q)$.*
5. *Given an eSPARQL query Q, a $\mathcal{FOUR}$ operator $\circ$, and a finite set of variables $W \subset \mathrm{inScope}(Q)$, the expression $(\mathsf{SELECT}_\circ\,W\,Q)$ is an eSPARQL query with in-scope variables W.*
6. *Given an eSPARQL query Q, a filter formula φ, and two $\mathcal{FOUR}$ elements α and β, the expression $(Q\,\mathsf{MAP}\,(\varphi, \alpha, \beta))$ is an eSPARQL query with in-scope variables $\mathrm{inScope}(Q)$.*
7. *Given a belief query E, and an eSPARQL query Q such that the in-scope variables of Q do not appear in E, the expression $(\mathsf{BELIEF}\ E\ Q)$ is an eSPARQL query with in-scope variables $\mathrm{var}(E) \cup \mathrm{inScope}(Q)$.*

Definition 8 (Semantics of eSPARQL). *The result of evaluating an eSPARQL query* query Q on a $\mathcal{FOUR}$-graph G is a function $\langle\!|Q|\!\rangle_G : \Omega_{\mathrm{inScope}(Q)} \to \mathcal{FOUR}$, called $\mathcal{FOUR}$-relation, defined recursively as follows:

1. *Given an RDF-star triple pattern T and a mapping $\mu \in \Omega_{\text{inScope}(Q)}$, $(\!(T)\!)_G(\mu)$ is $G(t)$, where t is the RDF-star triple pattern resulting from substituting every variable $\textbf{?x} \in \text{inScope}(Q)$ with $\mu(\textbf{?x})$.*

2. $(\!(Q_1 \oslash Q_2)\!)_G(\mu) = (\!(Q_1)\!)_G(\mu|_{\text{inScope}(Q_1)}) \oslash (\!(Q_2)\!)_G(\mu|_{\text{inScope}(Q_1)}).$

3. $(\!(Q_1 \ominus Q_2)\!)_G(\mu) = (\!(Q_1)\!)_G(\mu|_{\text{inScope}(Q_1)}) \ominus (\!(Q_2)\!)_G(\mu|_{\text{inScope}(Q_1)}).$

4. $(\!(Q_1 \oslash Q_2)\!)_G(\mu) = (\!(Q_1)\!)_G(\mu) \oslash (\!(Q_2)\!)_G(\mu).$

5. $(\!(Q_1 \oslash Q_2)\!)_G(\mu) = (\!(Q_1)\!)_G(\mu) \oslash (\!(Q_2)\!)_G(\mu).$

6. $(\!(Q_1 \text{ FILTER}_\circ \varphi)\!)_G(\mu) = (\!(Q_1)\!)_G(\mu) \circ \alpha_{\mu\models\varphi}$, *where* $\alpha_{\mu\models\varphi} = \text{Id}(\circ)$ *if formula* $\mu \models_R \varphi$ *in the $\mathcal{FOUR}$-relation* $R = (\!(Q_1)\!)_G$, *and* $\alpha_{\mu\models\varphi} = \text{Abb}(\circ)$, *otherwise.*

7. $(\!(\text{SELECT}_\circ\ W\ Q_1)\!)_G(\mu) = \circ(\{(\!(Q_1)\!)_G(\mu_1)\ :\ \mu_1|_W = \mu\}).$

8. $(\!(Q_1 \text{ MAP } (\varphi, \alpha, \beta))\!)_G(\mu) = \gamma$, *where* $\gamma = \alpha$ *if* $\mu \models_{(\!(Q_1)\!)_G} \varphi$, *and* $\gamma = \beta$, *otherwise.*

9. $(\!(\text{BELIEF } E\ Q)\!)_G(\mu) = (\!(Q)\!)_{G|_{E'}}(\mu|_{\text{inScope}(Q)})$, *where* E' *is the belief query resulting from replacing every variable* $\textbf{?x}$ *in* E *with* $\mu(\textbf{?x})$.

Intuitively, each of the SPARQL operators corresponds to two eSPARQL operators (e.g., AND corresponds to $\oslash$ and $\ominus$, and UNION corresponds to $\oslash$ and $\oslash$), and each of these operators are evaluated according to the Geerts et al. [6] $\mathcal{K}$-annotated SPARQL algebra by choosing $\mathcal{K}$ to be one of the two $\mathcal{FOUR}$ semirings, namely $\mathcal{FOUR}_{\text{th}}$ and $\mathcal{FOUR}_{\text{in}}$. Except for the operator BELIEF, which changes the context graph where the query is evaluated.

Example 3. Let $G = \{t_1, \ldots, t_8\}$ be the annotated graph mentioned in Table 1. Consider the query $Q = (\text{BELIEF } (\textbf{?x}, \oslash)\ (\textbf{?y}, \textsf{a}, \textsf{FullDeity}))$. According to Definition 8, the result of query Q is a $\mathcal{FOUR}$-relation with mappings μ whose domain includes the variables $\textbf{?x}$ and $\textbf{?y}$. Given such a mapping μ, $(\!(Q)\!)_G(\mu) = (\!((\textbf{?y}, \textsf{a}, \textsf{FullDeity}))\!)_{G|_{(\mu(\textbf{?x}), \oslash)}}(\mu|_{\{\textbf{?y}\}})$. For example, if $\mu(\textbf{?x}) = \textsf{PopeDI}$ and $\mu(\textbf{?y}) = \textsf{Jesus}$, then

$$(\!(Q)\!)_G(\mu) = (\!((\textbf{?y}, \textsf{a}, \textsf{FullDeity}))\!)_{G|_{(\textsf{PopeDI}, \oslash)}}(\mu|_{\{\textbf{?y}\}})$$
$$= G|_{(\textsf{PopeDI}, \oslash)}((\textsf{Jesus}, \textsf{a}, \textsf{FullDeity}))$$
$$= \top.$$

It is not difficult to see that if variable $\textbf{?x}$ had been bound to an individual whose beliefs are not encoded in the graph, then $(\!(Q)\!)_G(\mu) = \bot$ because $G|_{(\textsf{Jesus}, \oslash)}$ would have hold no information (i.e., it would have annotated all triples with $\bot$).

6   Use Case Requirements Discussion

In this section, we show a eSPARQL query for each of the use cases U1–U4. Since the notation of eSPARQL queries in an algebraic format is not suitable for end-users, we additionally present how this query can be written as an extension of the user SPARQL syntax.

Use Case U1 Query

$$(\mathsf{BELIEF}\ (\mathsf{PopeDI}, \oslash)\ (\texttt{?deity}, \mathsf{a}, \mathsf{FullDeity}))$$

The clause $(\mathsf{BELIEF}\ (\mathsf{PopeDI}, \oslash)\ \cdot)$ generates the graph G' consisting of all beliefs of PopeDI. The states of duplicate statements are aggregated with operation $\oslash$, and statements that are not mentioned are defined to have state $\vdash$. Then, the nested query $(\texttt{?deity}, \mathsf{a}, \mathsf{FullDeity})$ returns a $\mathcal{FOUR}$-relation R whose domain consists of all mappings $\Omega|_{\texttt{?deity}}$. This domain includes all mappings of the form $\mu_u = \{\texttt{?deity} \mapsto u\}$ where $u \in \mathbf{I}$. For each $u \in \mathbf{I}$, $R(\mu) = G'((u, \mathsf{a}, \mathsf{FullDeity}))$. Since the only belief of PopeDI is indicated in triple t_1, we conclude that $R(\mu)$ is $\top$ if u is Jesus, and $R(\mu)$ is $\vdash$, otherwise (see Example 8).

Observe that the query includes the operator $\oslash$. This means that the operations are done over the information lattice. In the user eSPARQL syntax, this is indicated in a simple way by introducing the modifier `INFO` to the `SELECT` clause. The next listing shows how the graph G' is specified using the `FROM BELIEF` clause.

```
1  SELECT INFO ?deity
2  FROM BELIEF <PopeDI>
3  WHERE { ?deity a <FullDeity> }
```

Use Case U2 Query

$$(\mathsf{SELECT}_\oslash\ \texttt{?deity}\ (((\texttt{?x}, \mathsf{a}, \mathsf{Christian})\ \mathsf{MAP}\ \mathrm{state}(\top)\ \dashv \vdash) \oslash$$
$$(\mathsf{BELIEF}\ (\texttt{?x}, \oslash)\ (\texttt{?deity}, \mathsf{a}, \mathsf{FullDeity}))))$$

The internal clause $Q_2 = (\mathsf{BELIEF}\ (\texttt{?x}, \oslash)\ \cdot)$ works similarly to the previous example, but the ground belief queries are constructed over variable $\texttt{?x}$. Thus, we need to bind this variable to entities whose belief are encoded in the graph (see Example 8). There are then four possible bindings for variable $\texttt{?x}$, namely PopeDI, Arius, $\mathsf{Christianity}$, and $\mathsf{Russell}$. The answers μ_2 of the query Q_2 will be joined with the answers μ_1 of the nested query $Q_1 = ((\texttt{?x}, \mathsf{a}, \mathsf{Christian})\ \mathsf{MAP}\ \mathrm{state}(\top)\ \dashv \vdash)$. This imposes a further restriction of the binding for variable $\texttt{?x}$ (remember that $\oslash$ acts as `AND`). Only for PopeDI and Arius, the state $(\!|Q_1 \oslash Q_2|\!)_G(\mu_1 \cup \mu_2)$ is not $\bot$. Thus, the $(\mathsf{SELECT}_\oslash\ \texttt{?deity}\ \cdot)$ aggregates the beliefs of these two instances of variable $\texttt{?x}$ with operation $\oslash$. If $\texttt{?y}$ is bound to Jesus, then for PopeDI the mapping μ_2 is annotated as $\top$, whereas for Arius it is annotated as $\bot$. Hence, the belief of Christians regarding the nature of Jesus is $\dashv$.

In user eSPARQL syntax, this query can be expressed using `FROM BELIEF` with a variable from within a nested `SELECT INFO` query.

```
1  SELECT INFO ?deity
2  WHERE {
3    ?x a <Christian> .
4    MAP IF (STATE IS TRUE) TO CONFLICTED ELSE UNKNOWN .
5    {
6      SELECT INFO ?deity
7      FROM BELIEF ?x
8      WHERE { ?deity a <FullDeity> }
9    }
10 }
```

Use Case U3 Query

$$((\mathsf{BELIEF}\ ((\mathsf{PopeDI}, \oslash)\ \oslash\ (\mathtt{?x}, \oslash))\ (\mathtt{?s}, \mathtt{?p}, \mathtt{?o}))\ \mathsf{MAP}\ \mathrm{state}(\dashv)\ \top\ \bot)$$

We first consider the believe query $((\mathsf{PopeDI}, \oslash)\ \oslash\ (\mathtt{?x}, \oslash))$, which joins the beliefs of PopeDI with all other ?x using $\oslash$. The $\mathcal{FOUR}$-relation constructed by (?s, ?p, ?o) is unconstrained; thus, we obtain the join of the information lattice over all triples. Before projecting to ?x via $\mathsf{SELECT}_\oslash$ we use MAP to mark all mappings that include any conflict, which are aggregated by the $\mathsf{SELECT}_\oslash$. Finally, we map $\dashv$ to $\top$ and everything else to $\bot$.

In the user eSPARQL syntax, we use SELECT without any modifiers, which defaults to the truth lattice.

```
1  SELECT ?x
2  WHERE {
3    {
4      SELECT INFO ?x
5      FROM BELIEF <PopeDI> ?x
6      WHERE { ?s ?p ?o }
7    }
8    MAP IF (STATE IS CONFLICTED) TO TRUE ELSE FALSE
9  }
```

Use Case U4 Query

$$(\mathsf{SELECT}_\oslash\ \mathtt{?x}\ ((\mathsf{BELIEF}\ (\mathtt{?y}, \oslash)\ (\mathsf{BELIEF}\ (\mathtt{?x}, \oslash)\ (\mathsf{Zeus}, \mathsf{a}, \mathsf{FullDeity}))))$$

The nested query $(\mathsf{BELIEF}\ (\mathtt{?y}, \oslash)\ \cdot)$ generates the $\mathcal{FOUR}$-annotated graph of all beliefs by ?y aggregating them on the information lattice. Similarly, the nested query $(\mathsf{BELIEF}\ (\mathtt{?x}, \oslash)\ \cdot)$ generates the $\mathcal{FOUR}$-annotated graph of all beliefs by ?x according the beliefs of ?y. Then, the innermost clause (Zeus, a, FullDeity) obtains a $\mathcal{FOUR}$-relation R with mappings μ whose in-scope variables are ?y and ?x. For example, $R(\{\mathtt{?y} \mapsto \mathsf{Russell}, \mathtt{?x} \mapsto \mathsf{PopeDI}\}) = \top$ because according to the beliefs of Russell, PopeDI beliefs that (Zeus, a, FullDeity). On the other hand, $R(\{\mathtt{?y} \mapsto \mathsf{Arius}, \mathtt{?x} \mapsto \mathsf{PopeDI}\}) = \vdash$ because the graph does not contain information about the beliefs of Arius regarding the beliefs of PopeDI. Finally, the clause $(\mathsf{SELECT}_\oslash\ \mathtt{?x}\ \cdot)$ aggregates the mappings on such a relation R for each instance of the variable ?x.

In the user eSPARQL syntax, we use a nested SELECT query for each level of nesting in the beliefs.

```
1   SELECT INFO ?x
2   WHERE {
3     {
4       SELECT INFO *
5       FROM BELIEF ?y
6       WHERE {
7         SELECT INFO *
8         FROM BELIEF ?x
9         WHERE { <Zeus> a <FullDeity> }
10      }
11    }
12  }
```

7 Finitely Supported eSPARQL

To implement eSPARQL, the query results must be finitely encoded. As we have already shown, we can encode an infinite function with a finite mapping that has a finite support by encoding the non-zero states only.

Definition 9. *An eSPARQL query Q is said to be finitely supported if, for every $\mathcal{FOUR}$-graph with finite support, there exists an element $\alpha \in \mathcal{FOUR}$, called a zero for Q and G, such that the set $\{\mu \mid (\!|Q|\!)_G(\mu) \neq \alpha\}$ is finite.*

To know if eSPARQL queries are finitely supported, consider two $\mathcal{FOUR}$-relations R_1 and R_2 with finite support that result of evaluating the respective queries Q_1 and Q_2. That is, there are two $\mathcal{FOUR}$ operations $\circ$ and $\diamond$, such that the sets $\mathrm{supp}_\circ(R_1)$ and $\mathrm{supp}_\diamond(R_2)$ are finite. Given a $\mathcal{FOUR}$ operation $\bullet$, does $R_3 = R_1 \bullet R_2$ have a finite support? To answer this question, lets start figuring what is needed for R_3 to have a finite support for each operator combination of the operators $\circ$, $\diamond$, and $\bullet$.

The first observation is that R_1 and R_2 have infinitely many mappings whose states are respectively $\mathrm{Id}(\circ)$ and $\mathrm{Id}(\diamond)$. If $\bullet$ is a join operator (i.e., $\mathbb{\vee}$ or $\otimes$), then we will have infinitely many mappings μ such that $R_3(\mu) = \mathrm{Id}(\circ) \bullet \mathrm{Id}(\diamond)$. Since there is exactly one $\mathcal{FOUR}$-operator, namely $*$ such that $\mathrm{Id}(*) = \mathrm{Id}(\circ)\bullet\mathrm{Id}(\diamond)$, we want to know if $\mathrm{supp}_*(R_3)$ is finite. In this case, the answer is straightforward. If we have a mapping μ' such that $R_3(\mu) \neq \mathrm{Id}(*)$ then it must happen that $R_1(\mu) \neq \mathrm{Id}(\circ)$ or $R_2(\mu) \neq \mathrm{Id}(\diamond)$. Since R_1 and R_2 have finite support, there are finitely many mappings μ satisfying that condition. Thus, $\mathrm{supp}_*(R_3)$ is finite. Hence, for the two operations corresponding to the two lattice join (i.e., the generalizations of the SPARQL `UNION`), two $\mathcal{FOUR}$-relations with finite support result in a $\mathcal{FOUR}$-relation with finite support.

We next show that if $\bullet$ is a meet operator (i.e., $\otimes$ or $\ominus$), then the resulting $\mathcal{FOUR}$-relation can have no finite support. To this end, consider the query $Q_1 \ominus Q_2$ where Q_1 and Q_2 are two queries with respective in-scope variables `?x` and `?y`. Given a $\mathcal{FOUR}$-annotated graph G, let R_1 be $(\!|Q_1|\!)_G$ and R_2 be $(\!|Q_2|\!)_G$, $\mathrm{supp}_\otimes(R_2)$ be finite, $R_1(\{\texttt{?x} \mapsto a\}) = \top$, $R_1(\{\texttt{?x} \mapsto b\}) = \bot$, and assume that there are infinitely many values c such that $R_2(\{\texttt{?y} \mapsto c\}) = \top$. Then,

$$(\!|Q_1 \ominus Q_2|\!)_G(\{\texttt{?x} \mapsto a, \texttt{?y} \mapsto c\}) = \top,$$
$$(\!|Q_1 \ominus Q_2|\!)_G(\{\texttt{?x} \mapsto b, \texttt{?y} \mapsto c\}) = \vdash.$$

Since we can take infinitely many values of c to produce these two different states, we conclude that the answer of query $Q_1 \ominus Q_2$ has no finite support.

Proposition 2. *Every fragment of eSPARQL that includes the operations $\otimes$ or $\ominus$ can include queries that are not finitely supported.*

This negative result is not necessarily an impediment for implementing eSPARQL. In general, since we are interested in statements about the individuals who appear in a knowledge graph, it suffices to consider mappings that range to the active domain of the graph (i.e., individuals that occur in triples in the support of the graph). Since this subset of $\mathbf{I}$ is finite, using the active domain will lead to a query language whose results can be finitely encoded.

8   Related Work

There are many works on four-valued logics [8, 17, 15], but no one of them considers them for the semantics and query evaluation in SPARQL. Other works annotate SPARQL answers with lattice elements [13, 2], but they do not provide a mean to operate with sets of statements encoding beliefs. Arnout et al. [1] consider knowledge graphs with negative facts. However, they do not consider conflicted statements as we did. Works on distributed knowledge contexts are all assuming that knowledge was represented using different ontologies but with the same epistemological status [4, 7, 9]. Some of them used SPARQL queries as mappings [19], but did not address epistemological status as a key concern. Schenk et al. [18] studies the semantics of trust and caching in the Semantic Web considering the $\mathcal{FOUR}$ structure, however did not consider the problem of querying.

9   Conclusions and Future Work

We presented eSPARQL, a novel approach that allows for the description of epistemic information using RDF-star and formulating epistemic queries using a query language which extends SPARQL-star. This query language is based on the concrete $\mathcal{FOUR}$ bilattice, but we expect to generalize it to include more general bilattices.

Future work could include the study of different ways to implement eSPARQL. The most direct way is to build on top of a standard SPARQL-star engine. Indeed, the functionality of `FROM BELIEF` clauses to generate a new graph representing people's beliefs can be implemented with SPARQL-star CONSTRUCT queries, which use aggregate operations to compute states of duplicated statements on a set of beliefs. Then, SPARQL-star SELECT queries can be executed on top of the results from these CONSTRUCT queries. Other implementations include: (i) the whole rewriting of a eSPARQL query as a single SPARQL-star SELECT query, and (ii) using specialized indexes and algorithms to implement this query language.

Implementation.

We implemented eSPARQL on top of Apache Jena [5], an open source SPARQL-star engine. Our implementation [14] is available under a free software license.

Acknowledgements

References

[1]   Hiba Arnaout et al. "Wikinegata: a Knowledge Base with Interesting Negative Statements". In: *Proc. VLDB Endow.* 14.12 (2021), pp. 2807–2810. DOI: 10.14778/3476311.3476350.

[2]   Zubaria Asma et al. "NPCS: Native Provenance Computation for SPARQL". In: *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024.* ACM, 2024, pp. 2085–2093. DOI: 10.1145/3589334.3645557.

[3]   Sören Auer et al. "DBpedia: A Nucleus for a Web of Open Data". In: *ISWC/ASWC.* Vol. 4825. LNCS. Springer, 2007, pp. 722–735. DOI: 10.1007/978-3-540-76298-0_52.

[4]   Paolo Bouquet et al. "C-OWL: Contextualizing Ontologies". In: *The Semantic Web - ISWC 2003, Sanibel Island, FL, USA, October 20-23, 2003, Proceedings.* Vol. 2870. LNCS. Springer, 2003, pp. 164–179. DOI: 10.1007/978-3-540-39718-2_11.

[5]   Apache Software Foundation. *Apache Jena.* Available at: https://jena.apache.org/.

[6]   Floris Geerts et al. "Algebraic Structures for Capturing the Provenance of SPARQL Queries". In: *J. ACM* 63.1 (2016), 7:1–7:63. DOI: 10.1145/2810037.

[7]   Chiara Ghidini and Luciano Serafini. "Distributed First Order Logic". In: *Artif. Intell.* 253 (2017), pp. 1–39. DOI: 10.1016/j.artint.2017.08.008.

[8]   Gösta Grahne and Ali Moallemi. "A useful four-valued database logic". In: *Proceedings of the 22nd International Database Engineering & Applications Symposium, IDEAS 2018, Villa San Giovanni, Italy, June 18-20, 2018.* ACM, 2018, pp. 22–30. DOI: 10.1145/3216122.3216157.

[9]   Bernardo Cuenca Grau, Bijan Parsia, and Evren Sirin. "Combining OWL ontologies using epsilon-Connections". In: *J. Web Semant.* 4.1 (2006), pp. 40–59. DOI: 10.1016/j.websem.2005.09.010.

[10] Steve Harris and Andy Seaborne. *SPARQL 1.1 Query Language.* Tech. rep. W3C Recommendation, Mar. 2013.

[11] Olaf Hartig et al. *RDF-star and SPARQL-star.* Tech. rep. W3C Final Community Group Report, 2021.

[12] Patrick J. Hayes and Peter F. Patel-Schneider. *RDF 1.1 Semantics.* Tech. rep. W3C Recommendation, 2014.

[13] Daniel Hernández, Luis Galárraga, and Katja Hose. "Computing How-Provenance for SPARQL Queries via Query Rewriting". In: *Proc. VLDB Endow.* 14.13 (2021), pp. 3389–3401. DOI: 10.14778/3484224.3484235.

[14] Xiny Pan et al. *eSPARQL: SPARQL for Epistemic Queries.* Version v1. 2024. DOI: 10.18419/darus-4344.

[15] Peter F. Patel-Schneider. "A Four-Valued Semantics for Terminological Logics". In: *Artif. Intell.* 38.3 (1989), pp. 319–351. DOI: 10.1016/0004-3702(89)90036-2.

[16]   Jorge Pérez, Marcelo Arenas, and Claudio Gutiérrez. "Semantics and complexity of SPARQL". In: *ACM Trans. Database Syst.* 34.3 (2009), 16:1–16:45. DOI: 10.1145/1567274.1567278.

[17]   Greg Restall. "Four-valued semantics for relevant logics (and some of their rivals)". In: *J. Philos. Log.* 24.2 (1995), pp. 139–160. DOI: 10.1007/BF01048529. URL: https://doi.org/10.1007/BF01048529.

[18]   Simon Schenk. "On the Semantics of Trust and Caching in the Semantic Web". In: *The Semantic Web - ISWC 2008, Karlsruhe, Germany, October 26-30, 2008. Proceedings.* Vol. 5318. LNCS. Springer, 2008, pp. 533–549. DOI: 10.1007/978-3-540-88564-1_34.

[19]   Simon Schenk and Steffen Staab. "Networked graphs: a declarative mechanism for SPARQL rules, SPARQL views and RDF data integration on the web". In: *Proceedings of the 17th International Conference on World Wide Web, WWW 2008, Beijing, China, April 21-25, 2008.* ACM, 2008, pp. 585–594. DOI: 10.1145/1367497.1367577.

[20]   Thomas P. Tanon, Gerhard Weikum, and Fabian M. Suchanek. "YAGO 4: A Reason-able Knowledge Base". In: *ESWC.* Vol. 12123. LNCS. Springer, 2020, pp. 583–596. DOI: 10.1007/978-3-030-49461-2_34.

[21]   Denny Vrandecic. "Wikidata: a new platform for collaborative data collection". In: *WWW (Companion Volume).* ACM, 2012, pp. 1063–1064. DOI: 10.1145/2187980.2188242.