

Curating the Specificity of Ontological Descriptions under Ontology Evolution

Yannis Tzitzikas^{1,2}, Mary Kampouraki¹ and Anastasia Analyti¹

¹ Institute of Computer Science, FORTH-ICS, Greece

² Department of Computer Science, University of Crete, Greece

{tzitzik | analyti}@ics.forth.gr, mary.kampouraki@gmail.com

Abstract. Semantic Web Ontologies are not static, but evolve, however this evolution usually happens independently of the ontological instance descriptions (usually referred as “metadata”) which are stored in the various Metadata Repositories (MRs) or Knowledge Bases (KBs). Nevertheless, it is a common practice for a MR/KB to periodically update its ontologies to their latest versions by “migrating” the available instance descriptions to the latest ontology versions. Such migrations incur gaps regarding the specificity of the migrated metadata, i.e. inability to distinguish those descriptions that should be reexamined (for possible specialization as consequence of the migration) from those for which no reexamination is justified. Consequently, there is a need for principles, techniques, and tools for managing the uncertainty incurred by such migrations, specifically techniques for (a) identifying automatically the descriptions that are candidates for specialization, (b) computing, ranking and recommending possible specializations, and (c) flexible interactive techniques for updating the available descriptions (and their candidate specializations), after the user (curator of the repository) accepts/rejects such recommendations. This problem is especially important for curated knowledge bases which have increased quality requirements (as in E-science). In this paper, we formalize the problem and propose principles, rules, and algorithms for realizing the aforementioned scenario, for the RDF/S framework. Finally, we report experimental results demonstrating the feasibility of the approach.

1 Introduction

Ontologies provide a shared conceptualization of a domain by defining the concepts in the domain and describing how those concepts are related to each other. Several reasons for changing an ontology have been identified in the literature [9]: an ontology may need to change because it offers a richer conceptualization of the problem domain, the domain of interest has been changed, the perspective under which the domain is viewed has changed, or the user/application needs have changed.

An important observation is that this evolution happens *independently* of the ontological instance descriptions which are stored in the various Metadata Repositories (MRs) or Knowledge Bases (KBs). With the term *ontological instance description*, we refer to RDF/S [3] descriptions that classify an instance o to a class c or relate two instances o, o' with a property pr . With the term MR or KB, we refer to a stored corpus of ontological instance descriptions, which can be stored in files, in RDF/S databases (i.e. RDF triple-stores [18]), or in the rapidly growing *Linked Open Data* (LOD) cloud [2]. Due to the distributed nature of the Web and the Semantic Web, the

evolution of ontologies³ happens independently of the ontological instance descriptions, e.g. this is the case with ontologies maintained by standardization authorities. However, it is a common practice (mainly for interoperability purposes) for a KB to periodically update its ontologies to their latest versions by “migrating” the stored instance descriptions to the latest ontology versions. Such migrations are in many cases not difficult, because newer versions are (or constructed to be) compatible with past ones. Nevertheless, they incur gaps regarding the specificity of the migrated instance descriptions, i.e. inability to distinguish those that should be reexamined (for possible specialization as consequence of the migration) from those for which no reexamination is justified. It follows that quality control is very laborious and error-prone.

To start with, consider a corpus of instance descriptions and suppose that at certain points in time we can make the assumption that the available instance descriptions are the *most specific and detailed descriptions* that are possible with respect to the employed ontology. For instance, we can make such an assumption after explicit human (e.g. by the curator of the KB) inspection and verification [4], or in cases where the descriptions have been produced automatically by a method that is guaranteed to produce specific descriptions (e.g. by transforming curated relational data to RDF/S descriptions [25], or by automatic classification to categories each defined by sufficient and necessary conditions, etc.). We will hereafter refer to this assumption by the name *maximum specificity assumption* (for short *MSA*). It is not hard to see that if the new version of the ontology is *richer* than the past one, then the corpus of the migrated instance descriptions *may no longer satisfy the MSA with respect to the new ontology*. This paper elaborates on this problem.

The ability to identify the instance descriptions that satisfy the *MSA* and those that do not, is useful in order to address questions of the form: (a) for what descriptions can we make the *MSA*? (b) what (class or property) instances should probably be reclassified (to more refined classes or properties), and (c) which are the candidate new classes or properties (refinements) of such instances? The above questions are very useful for curating a corpus of instance descriptions, i.e. for managing its specificity as the corpus (and its ontologies) evolves over time. Without special support, such tasks would be unacceptably expensive and vulnerable to omissions, for large datasets. The problem occurs in various domains, including Digital Libraries (e.g. as the *Library of Congress Subject Headings LCSH* evolves), in Biomedicine/Bioinformatics (*Gene Ontology*), in e-Government (*oeGOV Ontologies*), etc.

Introductory (toy) example. We will explain the main idea of our approach using the toy example depicted at Figure 1. Consider an e-commerce portal that sells various kinds of products. Suppose a car *c1* that has been classified under the class *Car*, and a person *p1* that has been classified under the class *Person*, defined in an ontology *Ont1*, and suppose that both classes have no subclasses. Assume that for the current set of instance descriptions according to *Ont1* the *MSA* holds (i.e. they are complete with respect to specificity). We can infer, from this knowledge, that *c1 is not a Person* and *p1 is not a Car*. Let *Ont2* be a new version of that ontology which, among other, defines the subclasses of the classes *Car* and *Person*, shown at Figure 1 (right). All subclasses of *Car* are *possible classes* for *c1*. *Adult is not a possible class* for *c1*, since *c1* was not a person according to *Ont1*. Analogously, none of the

³ In this paper, by the term of ontology we refer only to schema information.

subclasses of `Car` is a possible class for `p1`, since `p1` was not a car according to `Ont1`. Moreover, notice that `Ont1` defines a property `owns` and suppose that `(p1 owns c1)` is an instance description. Also notice that `Ont2` defines a subproperty `sells` of `owns` between `Person` and `Car`. This property will be prompted as a *possible specialization of the association* between `p1` and `c1`.

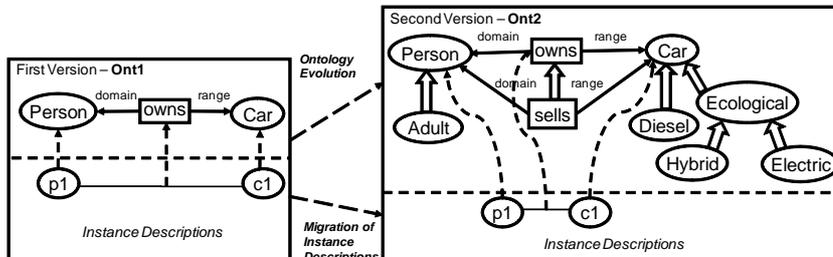


Fig. 1. Introductory example

The Difficulties. In real cases the computation of possibilities is more complex than the case of the toy example, as we can have conflicts among (a) new positive knowledge inferable from the instance descriptions and the new schema, (b) new “negative” information inferable from the past negative instance descriptions and the new schema, and (c) the previously computed possible instance descriptions (possible refinements). In fact, our approach resolves such conflicts by considering that (a) has higher priority than (b), and (b) has higher priority than (c). Our resolution of conflicts is illustrated in Figures 2 and 3.

In addition, it should be possible to update correctly the set of possibilities, at scenarios with several *successive* instance migrations interwoven with several (positive or negative) user feedbacks. Finally, another challenge is to reduce the information that has to be kept to support this scenario, specifically to avoid having to keep negative information of any kind, and to devise compact representations for the possibilities.

The contribution of this work lies in (a) formalizing the notion of (*possible*) *specificity*, (b) providing principles, rules, and algorithms for computing possibilities after schema migrations, (c) describing a flexible specificity-aware curation process, and (d) reporting experimental results (over real and synthetic datasets) that demonstrate the feasibility of the approach. We could say that from a more general perspective, our work contributes in enriching the lifecycle of Semantic Web data with quality management, appropriate for scenarios where ontologies evolve frequently and independently from instance descriptions. As a consequence, this allows adopting iterative and agile ontology modeling approaches, appropriate for open environments like Linked Open Data.

Note that though there are several works and approaches for dealing with the validity of data during migration [19, 16, 22], there is no work for managing their specificity and quality while ontologies evolve.

Here, we would like to mention that an overview of the main principles of this paper (demonstrated through examples) and an overview of the functionality of the prototype system, all from the perspective of *digital preservation*, is given in our conference paper [29].

The rest of this paper is organized as follows. Section 2 gives the required background information and notations. Section 3 introduces what we call *TFP-partition*

which is fundamental in our approach. Section 4 formalizes the migration problem as a *transition* between TFP-partitions, and proposes the basic *postulates* that should govern such transitions. Based on the above formalization and postulates, Section 5 provides an algorithm for computing the set of possible instance descriptions of an RDF/S KB after a migration, for the case of *backwards compatible* schema evolution, while Section 6 considers the same problem for the more general case of *non-backwards compatible* schema evolution. Section 7 discusses the effect of *successive* migrations. Section 8 presents methods for *ranking possibilities*. Section 9 describes the specificity lifecycle management process (for exploiting the computed possibilities). Section 10 provides experimental results over both real and synthetic data sets, and discusses storage approaches. Section 11 discusses related work. Finally, Section 12 concludes the paper and identifies issues for further research. Appendix A provides a list with the symbols used in the paper and Appendix B provides an example of applying the algorithm for backwards compatible migration. An extra electronic Appendix of this paper⁴ provides proofs and supplementary material.

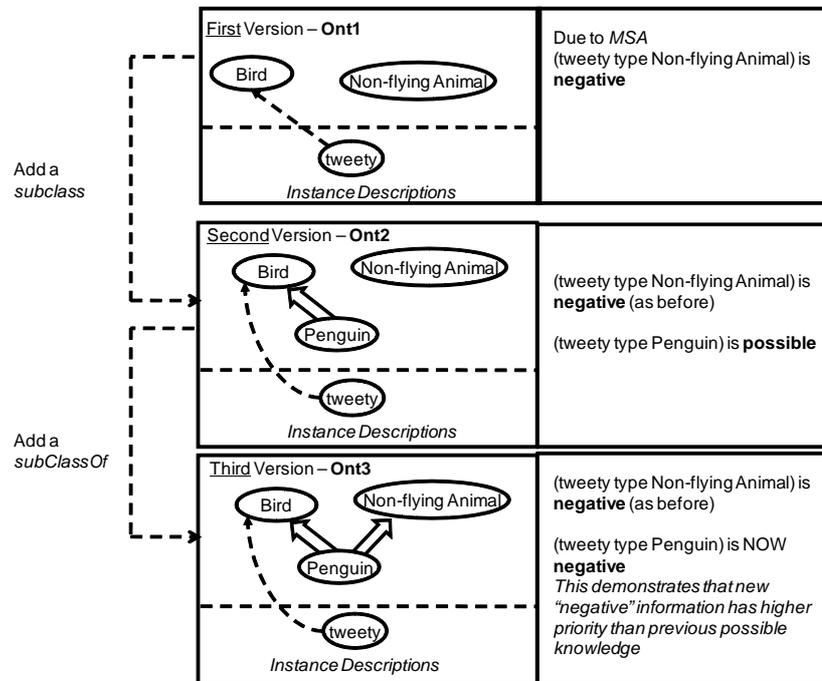


Fig. 2. New “negative information has higher priority than previous possible knowledge

2 Approach and Background

2.1 The Curation Process Model

Apart from identifying the information that could be further specialized (as we discussed just before), we would like to *aid* making it as specific as possible. Therefore, we should support flexible and *interactive* processes for managing the computed

⁴ Available at http://www.ics.forth.gr/~analyti/Papers_2/CurDelta_extra_appendix.pdf

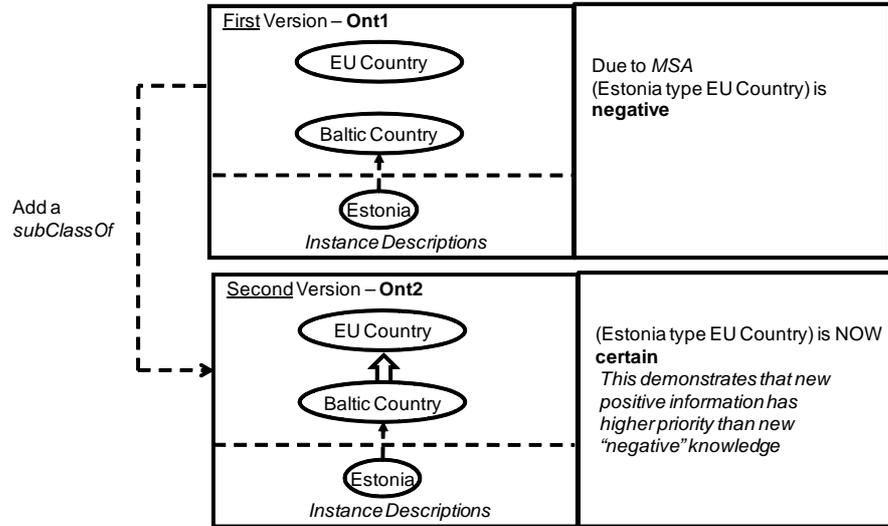


Fig. 3. New positive information has higher priority than the new “negative knowledge

possibilities, where the user will be able to either *accept* or *reject* the computed *recommendations*, and eventually update the knowledge base reaching to a state where the *MSA* holds (at least for those resources for which this is possible). The *ranking* of possibilities is important for designing user-friendly interaction schemes, since we may have a high number of recommendations. Essentially, we propose a process like the one sketched in the right part of Figure 4. Specifically, assume that the user selects some instances then the system displays ranked all or some of the possible instance descriptions for the selected instances. The user accepts or rejects these instance descriptions and the system updates appropriately the KB and its possible part. Note that the possible part of the KB is stored explicitly and separately. In our toy example, this means that we can rank the possible classes for *c1*, so that if the user is prompted to select a possible class for *c1*, then *Diesel* and *Ecological* will be the first classes to be displayed. If the user rejects the class *Ecological*, then all its subclasses will be rejected from the possible classes (and this reduces the effort required for reaching a state where the *MSA* holds).

2.2 On MSA, RDF, and Open/Closed World Assumptions

To avoid misunderstandings, we should clarify from the beginning that we do not violate the Open World Assumption of RDF/S. It is the *MSA* that allows us to infer the negative knowledge of the previous example. Only if we explicitly make the *MSA* we can then exploit it (in a Closed World Assumption manner) in the context of ontology evolution for formalizing the way possibilities are defined. To clarify that we can also capture the Open World Assumption of RDF, suppose that we start from an RDF/S KB for which we know nothing regarding its completeness or specificity. We can capture this case by considering that every instance description, that can be formed using the ontology and is not certain, is possible. That is, the set of “negative assertions” is empty (in our example that would mean that *Adult* *can be* a possible class for *c1* and *Car* *can be* a possible class for *p1*). Then, we can still use and exploit our machinery when we migrate our descriptions to subsequent schema versions, and the steps of the life cycle that we propose can produce the “negative” statements. If

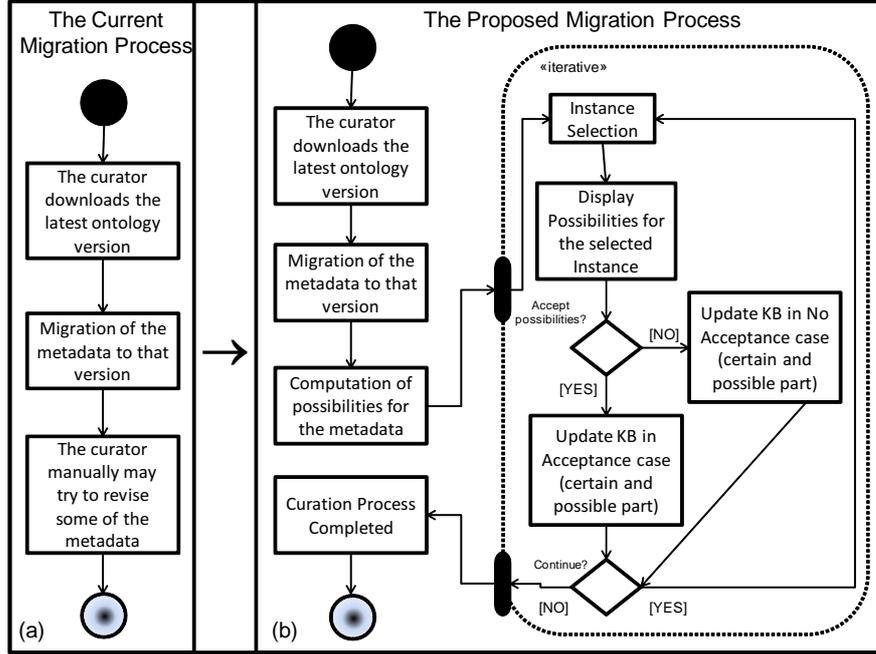


Fig. 4. Current and Proposed Migration Process

however one knows that one particular RDF/S KB is complete (regarding specificity), which can be true in the context of curated knowledge bases, then he can “apply” the *MSA*. This means that every instance description, that can be formed using the ontology and is not certain, is negative. Thus, the set of possibilities is empty (in our example that would mean that *Adult* is not a possible class for *c1* and *Car* is not a possible class for *p1*).

As we shall see later, the notion of TFP-partition allows modeling all such cases, and in Section 3.1 will show how exactly this can be done.

2.3 RDF/S Knowledge Bases

This section introduces notions and notations about RDF/S that shall be used in the sequel. Let URI be the set of URI references and LIT be the set of plain and typed literals. In our framework, an RDF/S Knowledge Base (KB) is defined by a set of RDF triples of the form $(subject\ predicate\ object)$, where $subject, predicate \in URI$ and $object \in URI \cup LIT$.

Let \mathcal{T} be the set of all possible triples that can be constructed from a countably infinite set of URIs, as well as literals (e.g. strings, integers, float numbers) [11]. Then, an RDF/S KB (for short KB) can be seen as a finite subset K of \mathcal{T} , i.e. $K \subseteq \mathcal{T}$. Apart from the explicitly specified triples of a KB K , other triples can be inferred based on the RDF/S semantics [13]. For this reason, we introduce the notion of closure.

The *closure* of a KB K , denoted by $\mathcal{C}(K)$, is the set of all triples that either are explicitly asserted or can be inferred from K based on RDFS-entailment of the RDF/S semantics [13]. Essentially, the following derivation rules are used:

- (i) if $(c_1\ subClassOf\ c_2)$ and $(c_2\ subClassOf\ c_3)$ then $(c_1\ subClassOf\ c_3)$,
- (ii) if $(pr_1\ subPropertyOf\ pr_2)$ and $(pr_2\ subPropertyOf\ pr_3)$ then $(pr_1\ subPropertyOf\ pr_3)$,
- (iii) if $(o\ type\ c_1)$ and $(c_1\ subClassOf\ c_2)$ then $(o\ type\ c_2)$,

- (iv) if $(o \text{ pr}_1 o')$ and $(\text{pr}_1 \text{ subPropertyOf } \text{pr}_2)$ then $(o \text{ pr}_2 o')$,
- (v) if $(o \text{ pr } o')$ and $(\text{pr } \text{domain } c)$ then $(o \text{ type } c)$, and
- (vi) if $(o \text{ pr } o')$ and $(\text{pr } \text{range } c)$ then $(o' \text{ type } c)$.

Def. 1 Let K be a KB. We define the *schema* of K as the tuple $\Gamma_K = \langle C_K, Pr_K, \text{domain}, \text{range}, \leq_{cl}^*, \leq_{pr}^* \rangle$, where⁵:

- C_K is the set of classes of $\mathcal{C}(K)$,
- Pr_K is the set of properties of $\mathcal{C}(K)$,
- *domain* is a total function $\text{domain} : Pr_K \rightarrow C_K$ that maps a property in Pr_K to its domain,
- *range* is a total function $\text{range} : Pr_K \rightarrow C_K$ that maps a property in Pr_K to its range,
- \leq_{cl}^* is the *subClassOf* relation between C_K , and
- \leq_{pr}^* is the *subPropertyOf* relation between Pr_K . □

Below we introduce notations for the *resources* of K , the *instances* of K , and the *instances of a class* $c \in C_K$:

$$\begin{aligned} Res_K &= \{o \mid (o \text{ type } \text{Resource}) \in \mathcal{C}(K)\} \\ Inst_K &= Res_K \setminus (C_K \cup Pr_K) \\ inst_K(c) &= \{o \mid (o \text{ type } c) \in \mathcal{C}(K)\} \end{aligned}$$

Def. 2 (Valid KB)

We consider a KB K to be *valid* if:

- (i) the relations \leq_{cl}^* and \leq_{pr}^* are acyclic,
- (ii) if $\text{pr} \leq_{pr}^* \text{pr}'$ then $\text{domain}(\text{pr}) \leq_{cl}^* \text{domain}(\text{pr}')$ and $\text{range}(\text{pr}) \leq_{cl}^* \text{range}(\text{pr}')$. □

In this paper, we consider only valid KBs. The triples of \mathcal{T} can be partitioned to *schema* and *instance* triples, as shown in Table 1. The triples that are not schema triples, according to that table, are instance triples. The latter can be further partitioned to *class instance triples* (having the form $(o \text{ type } c)$) and *property instance triples* (having the form $(o \text{ pr } o')$).

Schema Triples	
triple	abbreviation
$c \text{ type } \text{Class}$	
$c \text{ subClassOf } c'$	$c \leq_{cl} c'$
$\text{pr } \text{type } \text{Property}$	
$\text{pr } \text{subPropertyOf } \text{pr}'$	$\text{pr} \leq_{pr} \text{pr}'$
$\text{pr } \text{domain } c''$	$\text{domain}(\text{pr}) = c''$
$\text{pr } \text{range } c''$	$\text{range}(\text{pr}) = c''$
Instance Triples	
$o \text{ type } c$	
$o \text{ pr } o'$	

Table 1. Schema and Instance Triples

⁵ Note that according to RDF/S semantics [13], \leq_{cl}^* and \leq_{pr}^* are reflexive and transitive relations.

Example 1. Figure 5 illustrates two KBs K and K' , each consisting of schema triples and instance triples. In this example the instances are shown at the bottom of each KB and we have $Inst_K = Inst_{K'} = \{\text{Fiat_1}, \text{BMW_1}, \text{Bob}, \text{Alice}, \text{Computer Science Department}, \text{FORTH}\}$. \square

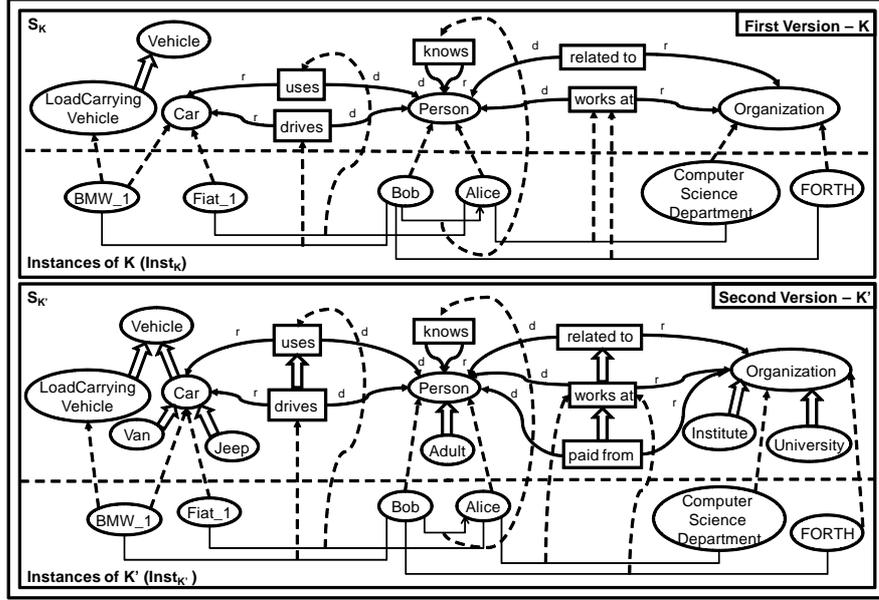


Fig. 5. An instance migration scenario

Classes and class instances are depicted by ovals. Properties are depicted by rectangles and the letters “d” and “r” are used to denote the domain and the range of a property. Fat arrows denote subClassOf/subPropertyOf relationships, while dashed arrows denote instanceOf relationships.

3 The Notion of TFP-partition

In this Section, we introduce the notion of TFP-partition (True-False-Possible partition) of a KB for capturing various application-specific assumptions about the specificity of its descriptions. At first, we define the set of *cartesian instance triples* of a KB K , denoted by B_K (where “B” comes from Base). Intuitively, it includes all instance triples that can be expressed using the schema of K .

Def. 3 (Cartesian Instance Triples)

Given a KB K , the set of *cartesian instance triples* of K , denoted by B_K , is the union of the class instance triples in $(Inst_K \cap URI) \times \{type\} \times C_K$ and the property instance triples in $(Inst_K \cap URI) \times Pr_K \times Inst_K$. \square

Given a KB K , we can distinguish its set of *schema triples* S_K and its set of *instance triples* I_K , i.e. $K = S_K \cup I_K$. However, for migration purposes, we need to consider only instance triples in B_K , i.e. those in $I_K \cap B_K$. Note that the remaining

instance triples, e.g. property instances that connect classes, are not interesting for migration purposes, and we ignore them. For notational simplicity we shall hereafter use I_K to denote $I_K \cap B_K$, and we shall use $K = (S_K, I_K)$. We define $\mathcal{C}_i(K)$ as the set of explicit and inferred instance triples, specifically $\mathcal{C}_i(K) = \mathcal{C}(K) \cap B_K$. Clearly, it holds: $I_K \subseteq \mathcal{C}_i(K) \subseteq B_K$.

We shall write $K = (S_K, I_K)$. We define $\mathcal{C}_i(K)$ as the set of explicit and inferred instance triples, specifically $\mathcal{C}_i(K) = \mathcal{C}(K) \cap B_K$. Clearly, it holds: $I_K \subseteq \mathcal{C}_i(K) \subseteq B_K$.

In the Example 1, we have $K = (S_K, I_K)$, $K' = (S_{K'}, I_{K'})$ and $I_K = I_{K'}$, e.g. $\{(\text{Bob drives BMW}_1), (\text{Fiat}_1 \text{ type Car})\} \subset I_K = I_{K'}$. In addition we have $(\text{domain}(\text{drives}) = \text{Person}) \in S_K$, $(\text{Car} \leq_{cl}^* \text{Vehicle}) \in S_{K'}$, while $(\text{Car} \leq_{cl}^* \text{Vehicle}) \notin S_K$.

Below, we define the notion of *valid to add property instance* in order to indicate which of the property instance triples of B_K are considered invalid.

Def. 4 (Valid to Add Property Instance)

We shall call a property instance triple $(o \text{ pr } o') \in B_K$ *valid to add*, for short *valid*, if it satisfies the constraint:

$$(o \in \text{inst}_K(\text{domain}(pr))) \wedge (o' \in \text{inst}_K(\text{range}(pr))) \quad (1)$$

□

The motivation is the following: if we add to K a property instance triple $(o \text{ pr } o')$ that does not satisfy expression (1) of Def. 4, then at least one new class instance triple would be added to $\mathcal{C}_i(K)$, due to derivation rules (v) and (vi) of RDF/S semantics (see Section 2). Furthermore, several Semantic Web Repositories for well justified reasons do not support derivation rules (v) and (vi), and consequently do not accept the addition of property instance triples that do not satisfy expression (1). To simplify notation, hereafter we shall use $\text{valid}(o, pr, o', K)$ to denote expression (1). We shall also use $\text{Invalid}(K)$ to refer to the invalid property instance triples of B_K , i.e.

$$\text{Invalid}(K) = \{(o \text{ pr } o') \in B_K \mid \neg \text{valid}(o, pr, o', K)\}$$

Let us also introduce some auxiliary notations. We will define the *SubTriples* of an instance triple as follows:

$$\begin{aligned} \text{SubTriples}((o \text{ type } c)) &= \{(o \text{ type } c') \mid c' \leq_{cl}^* c\} \\ \text{SubTriples}((o \text{ pr } o')) &= \{(o \text{ pr } o') \mid pr' \leq_{pr}^* pr\} \end{aligned}$$

If A is a set of instance triples, we define: $\text{SubTriples}(A) = \bigcup_{t \in A} \text{SubTriples}(t)$. Given two triples t and t' , we shall write: $t \leq t'$ iff $t \in \text{SubTriples}(t')$.

Now, we introduce the notion of TFP-partition (from *True False Possible*) which is fundamental for our work. The main idea is to partition the set of cartesian instance triples B_K into three pairwise disjoint subsets: *True*, *False*, and *Possible* instance triples:

- the first comprises I_K and the inferred instance triples (i.e. $\mathcal{C}_i(K)$),
- the second comprises instance triples which are not true (denoted by F_K), and
- the last comprises instance triples (denoted by P_K) that are possible.

Def. 5 (TFP-partition)

An *TFP-partition* of B_K is a three-fold partition of B_K , denoted by $(\mathcal{C}_i(K), F_K, P_K)$, that satisfies the following:

- (i) $\mathcal{C}_i(K) = \mathcal{C}(K) \cap B_K$.
- (ii) F_K is a *lower set*⁶ wrt \leq , i.e. $SubTriples(F_K) = F_K$.
- (iii) If an element of B_K is not valid then it belongs to F_K , i.e. $Invalid(K) \subseteq F_K$. \square

Note that $\mathcal{C}_i(K)$ is an *upper set*⁷ wrt \leq , as consequence of the derivation rules (i) and (ii) of RDF/S semantics (see Section 2) and the fact that $\mathcal{C}(K)$ is closed with respect to the closure operator \mathcal{C} . Def. 5(ii) actually says that: (a) if $(o \text{ type } c_1) \in F_K$ and $c_2 \leq_{cl}^* c_1$ then it should hold $(o \text{ type } c_2) \in F_K$, and (b) if $(o \text{ pr}_1 \text{ o}') \in F_K$ and $pr_2 \leq_{pr}^* pr_1$ then it should be $(o \text{ pr}_2 \text{ o}') \in F_K$. It is not hard to see that it would be “irrational” if these two were not satisfied.

It follows from Def. 5(ii) and the fact that a TFP-partition is a partition that it does not exist $m \in F_K$ and $p \in P_K$ such that $p \leq m$.

Clearly, the triples in P_K fall into the following two categories:

1. class instance triples in $((Inst_K \cap URI) \times \{type\} \times C_K)$,
2. *valid to add* property instance triples, i.e. property instance triples whose addition in K would not add any inferred class instance triple to $\mathcal{C}_i(K)$.

Figure 6 illustrates some the above notions and notations using Venn Diagrams.

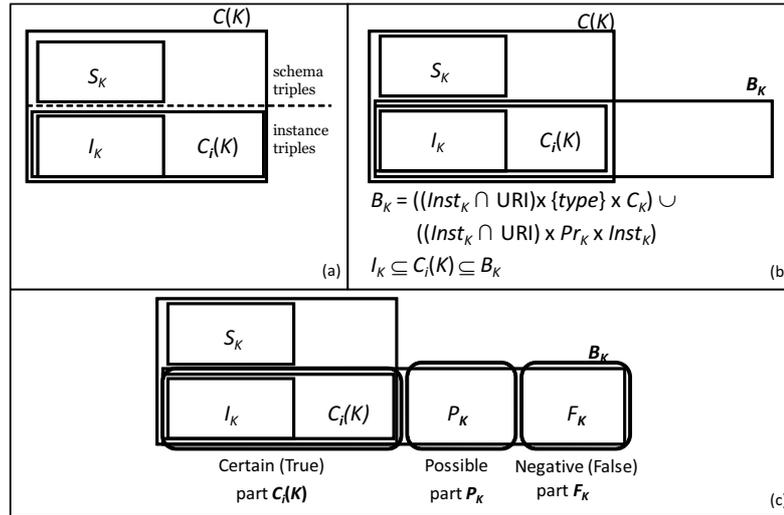


Fig. 6. TFP-partition illustrated with Venn Diagrams

Two lemmas follow.

Lemma 1 ($P_K \cup \mathcal{C}_i(K)$ is an Upper Set wrt \leq)

1. If $(o \text{ type } c_2) \in P_K$ and $c_2 \leq_{cl}^* c_1$ then $(o \text{ type } c_1) \in (P_K \cup \mathcal{C}_i(K))$.
2. If $(o \text{ pr}_2 \text{ o}') \in P_K$ and $pr_2 \leq_{pr}^* pr_1$ then $(o \text{ pr}_1 \text{ o}') \in (P_K \cup \mathcal{C}_i(K))$. \square

⁶ A *lower set* (else called *downward closed set*) is a subset Y of a given partially ordered set (X, \leq) s.t. for all elements x and y , if $x \leq y$ and y is an element of Y , then x is also in Y .

⁷ *Upper set* is the dual notion of lower set.

Lemma 2 (P_K is interval-based wrt \leq)

1. If $c_1 \leq_{cl}^* c_2 \leq_{cl}^* c_3$ and $(o \text{ type } c_1), (o \text{ type } c_3) \in P_K$ then $(o \text{ type } c_2) \in P_K$.
2. If $pr_1 \leq_{pr}^* pr_2 \leq_{pr}^* pr_3$ and $(o \text{ pr}_1 \text{ } o'), (o \text{ pr}_3 \text{ } o') \in P_K$ then $(o \text{ pr}_2 \text{ } o') \in P_K$. \square

Lemma 2 states that if t and t' belong to P_K then all t_x in the interval $[t, t']$ (defined over the \leq partial order) belong to P_K , too. This lemma can be exploited for specifying compact (interval-based) representations of P_K (more at Section 10.2).

3.1 Capturing Application Hypotheses/Scenarios with TFP-partitions

[MSA case].

Consider a KB K that satisfies the *Maximum Specificity Assumption* (MSA), i.e. all instances have been described with the most specific classes or properties of the schema that hold in the application domain. We can capture this case with a TFP-partition in which $P_K = \emptyset$. Since a TFP-partition is a partition, and $P_K = \emptyset$, it follows that all elements of B_K that are not in $\mathcal{C}_i(K)$, should be considered as false, i.e. $F_K = B_K \setminus \mathcal{C}_i(K)$, and thus the TFP-partition of B_K is:

$$(\mathcal{C}_i(K), F_K, P_K) = (\mathcal{C}_i(K), B_K \setminus \mathcal{C}_i(K), \emptyset)$$

It is not hard to verify that the above is a TFP-partition. Indeed (ii) and (iii) of Def. 5 are satisfied. Regarding (ii), it holds that $SubTriples(B_K \setminus \mathcal{C}_i(K)) = B_K \setminus \mathcal{C}_i(K)$, because if $t \in B_K \setminus \mathcal{C}_i(K)$ and $t' \leq t$, then $t' \in B_K \setminus \mathcal{C}_i(K)$, due to derivation rules (iii) and (iv) of RDF/S semantics (see Section 2). Regarding (iii), note that $\mathcal{C}_i(K) \cap Invalid(K) = \emptyset$, due to derivation rules (v) and (vi) of RDF/S semantics, thus, $Invalid(K) \subseteq B_K \setminus \mathcal{C}_i(K)$.

[Open World case].

Now consider the other extreme case where the MSA does not hold for any instance. For example, consider that we start from a KB for which we know nothing regarding its completeness or specificity, and we want to consider that every valid instance triple that can be formed using the ontology and is not certain, it is possible. We can capture this knowledge by a TFP-partition whose F_K contains only the invalid instance triples of B_K , i.e. $F_K = Invalid(K)$. Since a TFP-partition is a partition of B_K , it follows that $P_K = B_K \setminus (\mathcal{C}_i(K) \cup F_K) = B_K \setminus (\mathcal{C}_i(K) \cup Invalid(K))$, meaning that P_K includes all valid instance triples of B_K that do not belong to $\mathcal{C}_i(K)$, and hence the TFP-partition of B_K is:

$$(\mathcal{C}_i(K), F_K, P_K) = (\mathcal{C}_i(K), Invalid(K), B_K \setminus (\mathcal{C}_i(K) \cup Invalid(K)))$$

It is not hard to verify that the above is a TFP-partition⁸.

[Mixed case].

Apart from the above (extreme) cases, TFP-partitions can be used to capture various application specific assumptions, e.g. cases where we know that MSA holds only for a *part* of the KB.

⁸ Def. 5(ii) is satisfied because if it holds that $\neg valid(o, pr, o', K)$ and $pr' \leq_{pr}^* pr$ then certainly it holds that $\neg valid(o, pr', o', K)$, due to (ii) of Def. 2. Thus, $SubTriples(Invalid(K)) = Invalid(K)$. Def. 5(iii) is obvious.

4 Transitioning from one TFP-partition to Another

Here we will formalize schema migrations, by conceiving them as *transitions* between TFP-partitions which should be governed by a few fundamental postulates. For reasons of readability, we will first focus on migrations to *backwards compatible* schemas.

Def. 6 (Backwards Compatibility)

Let S and S' be two sets of schema triples. S' is *backwards compatible* with S , denoted by $S \sqsubseteq S'$, iff $\mathcal{C}(S) \subseteq \mathcal{C}(S')$. \square

For example, for the KBs of Figure 5 we have $S_K \sqsubseteq S_{K'}$. Note that we can check if a set of schema triples S' is backwards compatible with another set of schema triples S by computing the Delta function $\Delta_d(S \rightarrow S')$, defined as follows [32, 33]:

$$\Delta_d(S \rightarrow S') = \{Add(t) \mid t \in S' \setminus \mathcal{C}(S)\} \cup \{Del(t) \mid t \in S \setminus \mathcal{C}(S')\}$$

Note that $Add(t)$ and $Del(t)$ are not functions on t but strings, where t is replaced by the appropriate schema triple.

Prop. 1 $S \sqsubseteq S'$ iff $\Delta_d(S \rightarrow S')$ contains only *Add* operations. \square

Consider that we want to migrate the instance triples of a KB $K = (S_K, I_K)$, to a schema $S_{K'} \supseteq S_K$, reaching to a KB $K' = (S_{K'}, I_{K'})$. It is not hard to see that it holds $B_{K'} \supseteq B_K$. Since $S_{K'}$ is backwards compatible with S_K , every b that belongs to B_K certainly belongs to $B_{K'}$. The superset relationship between B_K and $B_{K'}$ can be strict (i.e. $B_{K'} \supset B_K$) because $S_{K'}$ can contain new elements (classes or properties) that could be used for generating instance triples to be added to $B_{K'}$.

Prop. 2 If $S_K \sqsubseteq S_{K'}$ then $B_{K'} \supseteq B_K$. \square

Suppose that we know the TFP-partition of K , i.e. $(\mathcal{C}_i(K), F_K, P_K)$. Our objective is to define the new TFP-partition, i.e. we want to define the transition:

$$(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$$

By migrating I_K to $S_{K'}$, we can get $\mathcal{C}(K')$, and consequently $\mathcal{C}_i(K')$. The rising question is how $F_{K'}$ and $P_{K'}$ are defined.

As we mentioned in the introductory section, we can have conflicts among:

- (a) new positive knowledge inferable from the instance triples and the new schema,
- (b) new “negative” information inferable from the past negative instance triples and the new schema, and
- (c) the previously computed possibilities (possible refinements).

*We resolve such conflicts by considering that:
 (a) has higher priority than (b), and
 (b) has higher priority than (c).*

These priorities can be formally expressed by two postulates that concern the transition from one TFP-partition to another.

Def. 7 (TFP-partition Evolution Postulates)

A transition $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$ is *consistent* if the following postulates are satisfied:

(II1) $\mathcal{C}_i(K')$ does not depend on F_K or P_K .

Priority of the positive knowledge inferrable from the instance triples and the new schema.

(II2) $F_K \cap P_{K'} = \emptyset$.

Past negative information cannot become possible. \square

Postulate II1 gives priority to the new positive knowledge over past negative or possible knowledge. It is consistent with (and reminiscent of) the principle “Recent knowledge prevails the old one” (also, called “Principle of Success” [1] and “Primacy of New Information” [6]).

Postulate II2 says that past negative information cannot become possible. This implies, as stated and proved by the Prop. 3 that follows, that past negative information is *preserved* as long as it does not contradict with the new positive knowledge.

Prop. 3 (Inertia Rule for Negatives)

In the context of a *consistent* transition $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, it follows that: $F_K \cap P_{K'} = \emptyset$ (i.e. II2) iff $(F_K \setminus \mathcal{C}_i(K')) \subseteq F_{K'}$. \square

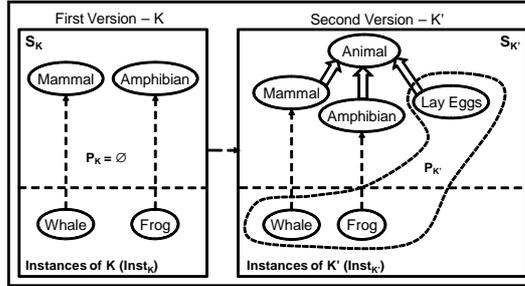


Fig. 7. Motivating example for Postulate II2

Below, we present also an example justifying Postulate II2.

Example 2. Consider Figure 7. It holds that $I_K = I_{K'} = \{(\text{Whale type Mammal}), (\text{Frog type Amphibian})\}$. Assuming that the *MSA* holds for KB K , it follows that $F_K = \{(\text{Whale type Amphibian}), (\text{Frog type Mammal})\}$. In $S_{K'}$, two new classes are added *Animal* and *Lay Eggs*, as well as the relationships $\text{Mammal} \leq_{cl} \text{Animal}$, $\text{Amphibian} \leq_{cl} \text{Animal}$, and $\text{Lay Eggs} \leq_{cl} \text{Animal}$. Obviously, we want $F_{K'} = F_K$ and $P_{K'} = \{(\text{Whale type Lay Eggs}), (\text{Frog type Lay Eggs})\}$. Note that it holds $F_K \cap P_{K'} = \emptyset$. \square

Based on the above postulates, we will specify how a transition can be performed. At first we introduce some notational conventions.

Notational conventions: In the context of two KBs K and K' , and unless otherwise indicated, with c we will denote a class in both C_K and $C_{K'}$, while with c' we will denote a class in $C_{K'}$ (that could however belong also to C_K). Analogously, with pr we will denote a property in both Pr_K and $Pr_{K'}$, while with pr' we will denote a property in $Pr_{K'}$ (that could belong also to Pr_K). In addition, with \leq_{cl}^* and \leq_{pr}^* , we will refer to the relations of the new schema $S_{K'}$. Further, we consider that the \leq_{cl}^* and \leq_{pr}^* relations, used in defining $SubTriples(A)$, refer to the relations of the new schema $S_{K'}$.

Def. 8 (Derivation of Negatives at a Transition)

Consider a TFP-partition $(\mathcal{C}_i(K), F_K, P_K)$ based on a schema S_K and suppose that we want to define the TFP-partition after migrating I_K to a backwards compatible schema $S_{K'}$. Consider deriving $F_{K'}$ using the following rules:

- (R1) If $(o \text{ type } c) \in F_K$, $c' \leq_{cl}^* c$, and $(o \text{ type } c) \notin \mathcal{C}_i(K')$ then $(o \text{ type } c') \in F_{K'}$.
- (R2) If $(o \text{ pr } o') \in F_K$, $pr' \leq_{pr}^* pr$, and $(o \text{ pr } o') \notin \mathcal{C}_i(K')$ then $(o \text{ pr } o') \in F_{K'}$.
- (R3) If $(o \text{ pr } o') \in B_{K'}$ and $\neg \text{valid}(o, pr', o', K')$ then $(o \text{ pr } o') \in F_{K'}$. \square

Essentially, the above rules produce the following set of instance triples:

$$F_{K'} = \text{Invalid}(K') \cup \text{SubTriples}(F_K \setminus \mathcal{C}_i(K'))$$

It is not hard to see that the $F_{K'}$ derived by the above rules satisfies the following: $F_{K'}$ and $\mathcal{C}_i(K')$ are disjoint, $F_{K'}$ is a lower set, and $F_{K'}$ contains all invalid property instance triples of $B_{K'}$.

Prop. 4 (Derivation of Negatives at a Transition)

The negative part of the target TFP-partition of a consistent transition, always satisfies the rules of Def. 8. \square

Having derived $\mathcal{C}_i(K')$ and $F_{K'}$, we can now derive $P_{K'}$, and thus obtain the TFP-partition of $B_{K'}$, by excluding from $B_{K'}$ all triples in $\mathcal{C}_i(K')$ or in $F_{K'}$.

Def. 9 (Deriving $P_{K'}$ through $B_{K'}$, $\mathcal{C}_i(K')$ and $F_{K'}$) If we know $\mathcal{C}_i(K')$ and $F_{K'}$, then we can derive $P_{K'}$ as follows: $P_{K'} = B_{K'} \setminus (F_{K'} \cup \mathcal{C}_i(K'))$. \square

We can now state and prove the following important proposition.

Prop. 5 (Transition Correctness)

The derivation of $F_{K'}$ using the rules of Def. 8, and the derivation of $P_{K'}$ by Def. 9, yields a three-fold partition that is a TFP-partition and

$(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$ is a *consistent transition*, i.e. it respects postulates *II1* and *II2*. We will denote this particular transition by:

$$(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'}) \quad \square$$

The derived TFP-partition, not only makes the transition consistent, but it has the *minimum set of negatives* among all TFP-partitions which constitute consistent transitions (it follows directly from Prop. 4).

Algorithmically, from F_K , K , and $S_{K'}$, we can produce $F_{K'}$ by straightforwardly applying the rules of Def. 8.

5 On Managing TFP-partitions without F -sets

Suppose that we want to migrate the instance triples of a KB $K = (S_K, I_K)$, to a schema $S_{K'}$ such that $S_K \sqsubseteq S_{K'}$. To compute $P_{K'}$, one approach is to apply the method described in the previous section. However, the shortcoming of that approach is that it requires computing the whole TFP-partition and to keep F_K stored, something that could be very space consuming. Therefore, below we investigate whether we can achieve our goal (derive $P_{K'}$) *without having to compute $F_{K'}$, and without having kept F_K* .

We prefer an approach requiring the availability only of P_K , because it is reasonable to assume that $|P_K| < |F_K|$, meaning that a P_K -based approach for computing $P_{K'}$ requires fewer space than a F_K -based approach (like the one described in the previous section). After all, the ideal state of a curated KB is a state that satisfies the *MSA*, and such a state requires keeping only K (since P_K is empty at that case). Instead, a F_K -based approach would require keeping a non-empty (and possibly very large) F_K even if the *MSA* holds. In general, the more we approach the *MSA*, the fewer storage is required for a P_K -based approach, while the opposite holds for the F_K -based approach. Also note that the lifecycle management process that we propose at Section 9 aims at reducing possibilities and approaching a state satisfying *MSA*.

Thus, we introduce what we call *extended KB*.

Def. 10 (Extended KB)

An *extended KB*, for short *eKB*, is a pair $\mathcal{E} = (K, P_K)$, where K is a KB and P_K is a set of possible instance triples s.t. $P_K \cap \mathcal{C}(K) = \emptyset$. \square

In fact, the instance triples of K , and the set P_K , are two parts of the TFP-partition of K . Below, we will try to compute $P_{K'}$ just from K , P_K , and $S_{K'}$. The envisioned process will be *correct* if it gives the same result ($P_{K'}$) as Def. 9.

In general, $P_{K'}$ can be produced by adding and deleting triples to/from P_K , i.e. we can write

$$P_{K'} = (P_K \setminus P_{K_Del}) \cup P_{K_Add}$$

where P_{K_Del} are the elements of P_K that should be deleted from it and P_{K_Add} are the elements that should be added to P_K . Due to priority of new knowledge, we certainly know that

$$P_{K_Del} \supseteq P_K \cap \mathcal{C}_i(K')$$

In fact,

$$P_{K_Del} = (P_K \cap \mathcal{C}_i(K')) \cup (P_K \cap F_{K'})$$

This is because, the only triples that can be deleted from P_K belong to $\mathcal{C}_i(K')$ or $F_{K'}$. Regarding P_{K_Add} , it should contain instance triples that involve the new classes and properties. The following two Propositions indicate the instance triples to be added to P_{K_Add} .

Prop. 6 Suppose we are in the context of $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$.

For a new class $c' \in C_{K'} \setminus C_K$, it holds that: (*o type c'*) $\in P_{K'}$ iff:

(i) $o \in Inst_K \cap URI$, and

(ii) for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that (*o type c*) $\in (\mathcal{C}_i(K') \cup P_K)$, and

(iii) (*o type c'*) $\notin \mathcal{C}_i(K')$. \square

Prop. 7 Suppose we are in the context of $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$.

For a new property $pr' \in Pr_{K'} \setminus Pr_K$, it holds that: (*o pr' o'*) $\in P_{K'}$ iff:

(i) $o \in URI$ and *valid(o, pr', o', K')*, and

(ii) for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that (*o pr o'*) $\in (\mathcal{C}_i(K') \cup P_K)$, and

(iii) (*o pr' o'*) $\notin \mathcal{C}_i(K')$. \square

5.1 Algorithmic Perspective

Below we present Algorithm 1, called *Produce_Possibilities*, which takes as input a KB K , its set of possible instance triples P_K , and a new set of schema triples $S_{K'}$ (s.t. $S_K \sqsubseteq S_{K'}$), where $K' = (S_{K'}, I_K)$ is a (valid) KB, and it produces the set of possible instance triples $P_{K'}$ for the new KB K' .

As the following Proposition states, Algorithm *Produce_Possibilities*($K, P_K, S_{K'}$) is *correct*, i.e. that it produces the same $P_{K'}$ as that defined in Def. 9 without the need of computing $B_{K'}$ and $F_{K'}$.

Prop. 8 If $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, then $P_{K'} = \text{Produce_Possibilities}(K, P_K, S_{K'})$. □

Algorithm 1 *Produce_Possibilities*($K, P_K, S_{K'}$)

Input: a KB K , its set of possible instance triples P_K , and a set of schema triples

$S_{K'}$ s.t. $S_K \sqsubseteq S_{K'}$ and $K' = (S_{K'}, I_K)$ is a (valid) KB

Output: the set of possible instance triples $P_{K'}$ of the KB K'

- (1) $K' = (S_{K'}, I_K)$;
 - (2) $P_{K_Add} = \emptyset$;
 - (3) $P_{K_Del} = \emptyset$;
 - /* FOR CLASS INSTANCES: */
 - (4) $NC = C_{K'} \setminus C_K$; /* new classes appearing in K' */
 - /* PART A: New classes */
 - (5) For all $(c' \in NC)$ do /* for each new class */
 - (6) $P_{K_Add} = P_{K_Add} \cup \{(o \text{ type } c') \mid o \in \text{Inst}_K \cap \text{URI},$
 $\forall c \in C_K \text{ s.t. } c' \leq_{cl}^* c \text{ it holds that } (o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K), \text{ and}$
 $(o \text{ type } c') \notin \mathcal{C}_i(K')\}$;
 - /* PART B: Existing classes */
 - (7) For all $(c_1 \in C_K)$ do
 - /* Moving class instance triples from P_K to $F_{K'}$ due to Rule R1. */
 - (8) $P_{K_Del} = P_{K_Del} \cup \{(o \text{ type } c_1) \in P_K \mid c_2 \in C_K, c_1 \leq_{cl}^* c_2, \text{ and}$
 $(o \text{ type } c_2) \notin (P_K \cup \mathcal{C}_i(K'))\}$;
 - /* FOR PROPERTIES: */
 - (9) $NP = Pr_{K'} \setminus Pr_K$; /* new properties appearing in K' */
 - /* PART C: New properties */
 - (10) For all $(pr' \in NP)$ do /* for each new property */
 - (11) $P_{K_Add} = P_{K_Add} \cup \{(o \text{ pr}' o') \mid o \in \text{URI}, \text{ valid}(o, pr', o', K'),$
 $\forall pr \in Pr_K \text{ s.t. } pr' \leq_{pr}^* pr \text{ it holds that } (o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K), \text{ and}$
 $((o \text{ pr}' o') \notin \mathcal{C}_i(K'))\}$;
 - /* PART D: Existing properties */
 - (12) For all $(pr_1 \in Pr_K)$ do
 - /* Moving property instance triples from P_K to $F_{K'}$ due to Rule R2. */
 - (13) $P_{K_Del} = P_{K_Del} \cup \{(o \text{ pr}_1 o') \in P_K \mid pr_2 \in Pr_K, pr_1 \leq_{pr}^* pr_2, \text{ and}$
 $(o \text{ pr}_2 o') \notin (P_K \cup \mathcal{C}_i(K'))\}$;
 - (14) $P_{K_Del} = P_{K_Del} \cup (P_K \cap \mathcal{C}_i(K'))$;
 - (15) $P_{K'} = P_K \setminus P_{K_Del}$;
 - (16) $P_{K'} = P_{K'} \cup P_{K_Add}$;
 - (17) Return $P_{K'}$;
-

Part A (lines 5-6) of Algorithm 1 follows from Prop. 6 and concerns new classes $c' \in NC$. If a new class c' has superclasses c (which are existing classes in K), we have to check if P_K or $\mathcal{C}_i(K')$ includes triples of the form $(o \text{ type } c)$ for each superclass c , where $o \in Inst_K \cap URI$. Only if this is true and $(o \text{ type } c')$ does not belong to $\mathcal{C}_i(K')$, we can safely add triples of the form $(o \text{ type } c')$ to P_{K_Add} . If a new class c' has no superclasses that are existing classes in K , we just have to check whether a triple of the form $(o \text{ type } c')$ does not belong to $\mathcal{C}_i(K')$ before adding it to P_{K_Add} .

Part B (lines 7-8), concerns existing classes $c_1 \in C_K$, for which we have to add to P_{K_Del} all those triples $(o \text{ type } c_1) \in P_K$, if there is a class $c_2 \in C_K$ such that $c_1 \leq_{cl}^* c_2$ in K' and $(o \text{ type } c_2) \notin (\mathcal{C}_i(K') \cup P_K)$. This is because, since $S_K \sqsubseteq S_{K'}$, it holds that $(o \text{ type } c_2) \notin \mathcal{C}_i(K)$. Thus, $(o \text{ type } c_2) \in F_K$. Therefore, it follows from Rule R1 of Prop. 4 that $(o \text{ type } c_1) \in F_{K'}$.

Part C (lines 10-11) follows from Prop. 7 and concerns new properties $pr' \in NP$. If a new property pr' has superproperties pr (which are existing properties in K), we have to check if P_K or $\mathcal{C}_i(K')$ includes triples of the form $(o \text{ pr } o')$ for each superproperty pr . Only if this is true, it holds that $valid(o, pr', o', K')$, and $(o \text{ pr } o')$ does not belong to $\mathcal{C}_i(K')$, we can safely add triples of the form $(o \text{ pr } o')$ to P_{K_Add} . If a new property has no superproperties that are existing properties in K , we just have to check whether it holds $valid(o, pr', o', K')$ and $(o \text{ pr } o')$ does not belong to $\mathcal{C}_i(K')$ before adding $(o \text{ pr } o')$ to P_{K_Add} .

Part D (lines 12-13) concerns existing properties $pr_1 \in Pr_K$, for which we have to add to P_{K_Del} all those triples $(o \text{ pr}_1 \text{ } o') \in P_K$, if there is a property $pr_2 \in Pr_K$ such that $pr_1 \leq_{pr}^* pr_2$ in K' and $(o \text{ pr}_2 \text{ } o') \notin (\mathcal{C}_i(K') \cup P_K)$. This is because, since $S_K \sqsubseteq S_{K'}$, it holds that $(o \text{ pr}_2 \text{ } o') \notin \mathcal{C}_i(K)$. Thus, $(o \text{ pr}_2 \text{ } o') \in F_K$. Therefore, it follows from Rule R2 of Prop. 4 that $(o \text{ pr}_1 \text{ } o') \in F_{K'}$.

In line 14 of Algorithm 1, we add to P_{K_Del} the set $P_K \cap \mathcal{C}_i(K')$, because all instance triples that belong to $\mathcal{C}_i(K')$ have to be removed from P_K . At the end (lines 15-16), we have to update P_K by adding to it the P_{K_Add} set and by removing from it the P_{K_Del} set. Then, we return $P_{K'}$. The execution order of the above parts A, B, C, and D does not matter since in any order the output result $P_{K'}$ is the same.

Examples of applying the algorithm are given in Appendix B.

Prop. 9 The time complexity of Algorithm 1 is $O(|Inst_K|^2 * M^2 * (|K'|^2 + |P_K|))$, where $M = \max(|C_{K'}|, |Pr_{K'}|)$. \square

6 Non-backwards Compatible Schema Evolution

Here we consider the case where the next schema version $S_{K'}$ of a KB K' , is not backwards compatible with S_K , i.e. $S_K \not\sqsubseteq S_{K'}$ (and consequently $\mathcal{C}(S_K) \not\subseteq \mathcal{C}(S_{K'})$), however the instance triples in I_K are transferred to K' , i.e. $K' = (S_{K'}, I_K)$, and K' is a (valid) KB.⁹ The changes between $S_{K'}$ and S_K may include deletions of classes, deletions of properties, changes in the subclassOf/subPropertyOf relations, or changes in the domain and range of properties. The elements of I_K refer to (usually leaf) classes and properties of S_K which are preserved in K' automatically by the semantics of RDF/S¹⁰ [13].

⁹ We discuss the notion of *validity* at Sect. 11 along with the discussion of the related works that focus on that issue.

¹⁰ If $(o \text{ type } c) \in K'$ then $(c \text{ type } Class) \in \mathcal{C}(K')$ and if $(o \text{ pr } o') \in K'$ then $(pr \text{ type } Property) \in \mathcal{C}(K')$.

Let $K' = (S_{K'}, I_K)$ be a KB. Obviously, it may hold that $B_K \not\subseteq B_{K'}$. For example, if (*o type c*) $\in B_K$ and $c \in C_K \setminus C_{K'}$ then (*o type c*) $\notin B_{K'}$.

Our goal is to describe how the rules that are used to derive $F_{K'}$ are modified. In brief, the following remain the same:

Def. 7 defining the postulates *II1* and *II2*,

Def. 8 and Prop. 4 regarding the rules *R1*, *R2*, and *R3*, and

Def. 9 defining $P_{K'}$.

Prop. 3 essentially remains the same, but we have to replace $F_K \setminus C_i(K') \subseteq F_{K'}$ by $(F_K \setminus C_i(K')) \cap B_{K'} \subseteq F_{K'}$. This is because there may exist an instance triple $t \in F_K$ but $t \notin B_{K'}$, due to deletions of classes and properties in $S_{K'}$.

However, and regarding the transition postulates, we have to add a *new postulate II3* applicable only to non-backwards compatible schema evolution case. Postulate *II3* expresses that if a triple $t \in B_{K'}$ that existed in $C_i(K)$, does not exist in $C_i(K')$ then t should go to $F_{K'}$.

Def. 11 (NBC Evolution Postulates)¹¹

A transition $(C_i(K), F_K, P_K) \rightsquigarrow (C_i(K'), F_{K'}, P_{K'})$ is *consistent* if, apart from the postulates *II1* and *II2* of Def. 7, the following postulate is satisfied.

(*II3*) All elements $t \in B_{K'}$ s.t. $t \in C_i(K) \setminus C_i(K')$ are placed to $F_{K'}$. □

Note that postulate *II3* is not needed in the backwards compatible schema evolution case, because in this case it holds that $C_i(K) \subseteq C_i(K')$.

Example 3. Consider the scenario shown in Figure 8, where $S_K \not\subseteq S_{K'}$. Suppose that $\{(\text{John type Full-time Permanent Employee}), (\text{John type Technical Staff}), (\text{John type Professor})\} \in P_K$. Note that the specialization relationship $\text{University Employee} \leq_{cl} \text{Permanent Employee}$, which exists in S_K , has been deleted in $S_{K'}$. Note that $(\text{John type Permanent Employee}) \in C_i(K)$, while $(\text{John type Permanent Employee}) \notin C_i(K')$. According to postulate *II3*, $(\text{John type Permanent Employee})$ should go to $F_{K'}$. □

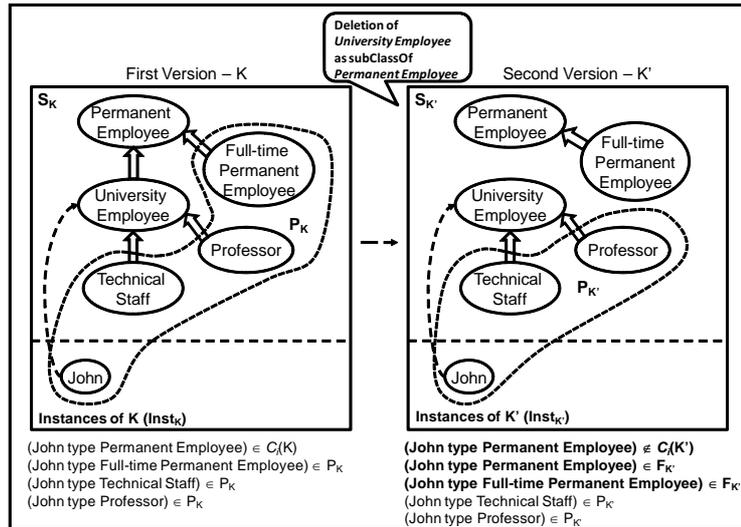


Fig. 8. Certain Modification Inheritance Rules

¹¹ NBC stands for non-backwards compatible.

The following proposition provides an equivalent form of postulate *II3*.

Prop. 10 In the context of a transition $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, it follows that: $(\mathcal{C}_i(K) \setminus \mathcal{C}_i(K')) \cap B_{K'} \subseteq F_{K'}$ (*II3*) iff $\mathcal{C}_i(K) \cap P_{K'} = \emptyset$. \square

Based on *II3*, we provide the following rules that produce (additionally to rules *R1*, *R2*, and *R3*) elements of $F_{K'}$, for the classes and properties that exist in K' .

Prop. 11 (Certain Modification Inheritance Rules) The negative part of the target TFP-partition of a consistent transition, always satisfies the following rules:¹²

- (*R4*) If $(o \text{ type } c) \in \mathcal{C}_i(K)$, $c' \leq_{cl}^* c$, and $(o \text{ type } c) \notin \mathcal{C}_i(K')$ then $(o \text{ type } c') \in F_{K'}$.
 - (*R5*) If $(o \text{ pr } o') \in \mathcal{C}_i(K)$, $pr' \leq_{pr}^* pr$, and $(o \text{ pr } o') \notin \mathcal{C}_i(K')$ then $(o \text{ pr } o') \in F_{K'}$.
- \square

Essentially, rules *R1* – *R5* produce the following set of instance triples:

$$F_{K'} = \text{Invalid}(K') \cup \text{SubTriples}((F_K \setminus \mathcal{C}_i(K')) \cap B_{K'}) \cup \\ \text{SubTriples}((\mathcal{C}_i(K) \setminus \mathcal{C}_i(K')) \cap B_{K'})$$

Obviously, a $F_{K'}$ derived by the rules *R1*-*R5* respects the following validity constraints of TFP-partition (Def. 5): $F_{K'}$ and $\mathcal{C}_i(K')$ are disjoint, $F_{K'}$ is a lower set, and $F_{K'}$ contains all invalid property instance triples of $B_{K'}$.

It is easy to see that the derivation of $F_{K'}$ by the rules *R1*, *R2*, and *R3* of Prop. 4 and *R4*, *R5* of Prop. 11, as well as the derivation of $P_{K'}$ by Def. 9, yield a TFP-partition that makes the transition $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$ *consistent*, i.e. it respects postulates *II1* and *II2* of Def. 7 and *II3* of Def. 11. Moreover, and analogously to the backward compatible case, the derived TFP-partition has the *minimum set of negatives* (again we will use \Rightarrow to denote this particular transition).

Example 4. Continuing Example 3, according to Rule *R4*, (**John type Permanent Employee**) should go to $F_{K'}$. Further, according to Rule *R4*, (**John type Full time Permanent Employee**) must be moved from P_K to $F_{K'}$. \square

Now, similarly to the case of backwards compatible schema evolution, we will try to compute $P_{K'}$ from K , P_K , and $S_{K'}$ without having to know or compute neither F_K nor $F_{K'}$. The envisioned process will be *correct* if it gives the same result ($P_{K'}$) as Def. 9.

In general, $P_{K'}$ can be produced by adding and deleting triples to/from P_K , i.e. we can write $P_{K'} = (P_K \setminus P_{K_Del}) \cup P_{K_Add}$, where P_{K_Del} are the elements of P_K that should be deleted from it and P_{K_Add} are the elements that should be added to P_K . In fact, $P_{K_Del} = (P_K \cap \mathcal{C}(K')) \cup (P_K \cap F_{K'}) \cup \{t \in P_K \mid t \notin B_{K'}\}$ ¹³. This is because, the only triples that can be deleted from P_K belong to $\mathcal{C}(K') \cup F_{K'}$ or do not belong to $B_{K'}$. Regarding P_{K_Add} , it should contain instance triples that involve the new classes and properties. The following two Propositions indicate the instance triples to be added to P_{K_Add} .

Prop. 12 The same as Prop. 6 but now applies to KBs $K = (S_K, I_K)$ and $K' = (S_{K'}, I_{K'})$, where S_K and $S_{K'}$ are not necessarily backwards compatible. \square

¹² The relationships \leq_{cl}^* and \leq_{pr}^* hold in K' .

¹³ Note that $t \notin B_{K'}$, if the class c or property pr appearing in t have been deleted in K' .

Prop. 13 The same as Prop. 7 but now applies to KBs $K = (S_K, I_K)$ and $K' = (S_{K'}, I_K)$, where S_K and $S_{K'}$ are not necessarily backwards compatible. \square

Below we present Algorithm 2, called *Produce_Possibilities_{NBC}*, which takes as input a KB K , its set of possible instance triples P_K , and a new set of schema triples $S_{K'}$ s.t. $K' = (S_{K'}, I_K)$ is a (valid) KB, and produces the set of possible instance triples $P_{K'}$ for the new KB K' .

The following Proposition shows that Algorithm *Produce_Possibilities_{NBC}*($K, P_K, S_{K'}$) is *correct*.

Prop. 14 Let $K = (S_K, I_K)$ and $K' = (S_{K'}, I_K)$ be two valid KBs. If $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, then $P_{K'} = \text{Produce_Possibilities}_{NBC}(K, P_K, S_{K'})$. \square

As we can see, the only differences between Algorithm 2 and Algorithm 1 are the new parts C and F (see bold lines).

Note that lines (5-8) of Algorithm 2 and lines (5-8) of Algorithm 1, as well as lines (12-15) of Algorithm 2 and lines (10-13) of Algorithm 1 are the same. However, their proof is different. This is because the second ontology version is not necessarily backwards compatible with the first one and the instance triples (*o type c*₂) or (*o pr*₂ *o'*) may belong to $\mathcal{C}_i(K)$ and not to $\mathcal{C}_i(K')$.

Part C (lines 9-10) concerns class instance triples of the form (*o type c*) that belong to P_K and refer to classes c that have been deleted in K' . These class instance triples have to be added to P_{K_Del} set, as they do not belong to $B_{K'}$.

Part F (lines 16-17) concerns property instance triples of the form (*o pr o'*) that belong to P_K and refer to properties pr that either have been deleted in K' or it holds that $\neg \text{valid}(o, pr, o', K')$, possibly because the domain and/or range of pr has been changed or specialization relationships between classes that involve (directly or indirectly) the domain and/or range of pr have been deleted in $S_{K'}$. These property instance triples have to be added to P_{K_Del} set because they either do not belong to $B_{K'}$ or belong to $\text{Invalid}(K')$.

If a property, pr , that appears in S_K and I_K , is removed in $S_{K'}$ and consequently the statements regarding its domain and range, to restore validity of $K' = (S_{K'}, I_K)$, we consider that the domain and range of pr , in K' , is the class *Resource*. Consider, for instance, Figure 9. The property instance triple $(o_1 pr_1 o_2) \in I_K$ leads to an invalid KB because pr_1 and thus, also its domain and range are deleted in K' . If we consider that the domain and range of pr_1 , in K' , is the class *Resource* then the validity of K' is restored.

Prop. 15 The time complexity of Algorithm 2 is $O(|Inst_K|^2 * M^2 * (S^2 + |P_K|))$, where $S = \max(|K|, |K'|)$ and $M = \max(|C_K|, |Pr_K|, |C_{K'}|, |Pr_{K'}|)$. \square

We would like to note that Algorithm 2 is more general than Algorithm 1 in the sense that it can be applied even in the backwards compatible schema evolution case. However, Algorithm 1 has fewer steps and is more efficient.

Algorithm 2 *Produce_Possibilities_{NBC}(K, P_K, S_{K'})*

Input: a KB K , its set of possible instance triples P_K , and a new set of schema triples

$S_{K'}$ s.t. $K' = (S_{K'}, I_K)$ is a (valid) KB

Output: the set of possible instance triples $P_{K'}$ of the KB K'

```
(1)   $K' = (S_{K'}, I_K)$ ;  
(2)   $P_{K\_Add} = \emptyset$ ;  
(3)   $P_{K\_Del} = \emptyset$ ;  
/* FOR CLASS INSTANCES: */  
(4)   $NC = C_{K'} \setminus C_K$ ; /* new classes appearing in  $K'$  */  
/* PART A: New classes */  
(5)  For all ( $c' \in NC$ ) do /* for each new class */  
(6)     $P_{K\_Add} = P_{K\_Add} \cup \{(o \text{ type } c') \mid o \in Inst_K \cap URI,$   
         $\forall c \in C_K \text{ s.t. } c' \leq_{cl}^* c \text{ it holds that } (o \text{ type } c) \in (C_i(K') \cup P_K), \text{ and}$   
         $(o \text{ type } c') \notin C_i(K')\}$ ;  
/* PART B: Existing classes */  
(7)  For all ( $c_1 \in C_K$ ) do  
/* Moving class instance triples from  $P_K$  to  $F_{K'}$  due to Rule R1 and Rule R4. */  
(8)     $P_{K\_Del} = P_{K\_Del} \cup \{(o \text{ type } c_1) \in P_K \mid c_2 \in C_K, c_1 \leq_{cl}^* c_2, \text{ and}$   
         $(o \text{ type } c_2) \notin (P_K \cup C_i(K'))\}$ ;  
/* PART C: Removing class instance triples from  $P_K$ , due to removed  
classes */  
(9)  For all (( $o \text{ type } c$ )  $\in P_K$ ) do  
(10)  If  $c \notin C_{K'}$  then  $P_{K\_Del} = P_{K\_Del} \cup \{(o \text{ type } c)\}$ ;  
/* FOR PROPERTIES: */  
(11)  $NP = Pr_{K'} \setminus Pr_K$ ; /* new properties appearing in  $K'$  */  
/* PART D: New properties */  
(12) For all ( $pr' \in NP$ ) do /* for each new property */  
(13)     $P_{K\_Add} = P_{K\_Add} \cup \{(o \text{ pr}' o') \mid o \in URI,$   
         $valid(o, pr', o', K'),$   
         $\forall pr \in Pr_K \text{ s.t. } pr' \leq_{pr}^* pr \text{ it holds that } (o \text{ pr } o') \in (C_i(K') \cup P_K), \text{ and}$   
         $(o \text{ pr}' o') \notin C_i(K')\}$ ;  
/* PART E: Existing properties */  
(14) For all ( $pr_1 \in Pr_K$ ) do  
/* Moving property instance triples from  $P_K$  to  $F_{K'}$  due to Rule R2 and R5. */  
(15)     $P_{K\_Del} = P_{K\_Del} \cup \{(o \text{ pr}_1 o') \in P_K \mid pr_2 \in Pr_K, pr_1 \leq_{pr}^* pr_2, \text{ and}$   
         $(o \text{ pr}_2 o') \notin (P_K \cup C_i(K'))\}$ ;  
/* PART F: Removing property instance triples from  $P_K$  due to removed  
properties, subject out of the domain, or object out of range */  
(16) For all (( $o \text{ pr } o'$ )  $\in P_K$ ) do  
(17)  If  $pr \notin Pr_{K'}$  or  $\neg valid(o, pr, o', K')$  then  
         $P_{K\_Del} = P_{K\_Del} \cup \{(o \text{ pr } o')\}$ ;  
/* Moving instance triples from  $P_K$  to  $C_i(K')$ . */  
(18)  $P_{K\_Del} = P_{K\_Del} \cup (P_K \cap C_i(K'))$ ;  
(19)  $P_{K'} = P_K \setminus P_{K\_Del}$ ;  
(20)  $P_{K'} = P_{K'} \cup P_{K\_Add}$ ;  
(21) Return  $P_{K'}$ ;
```

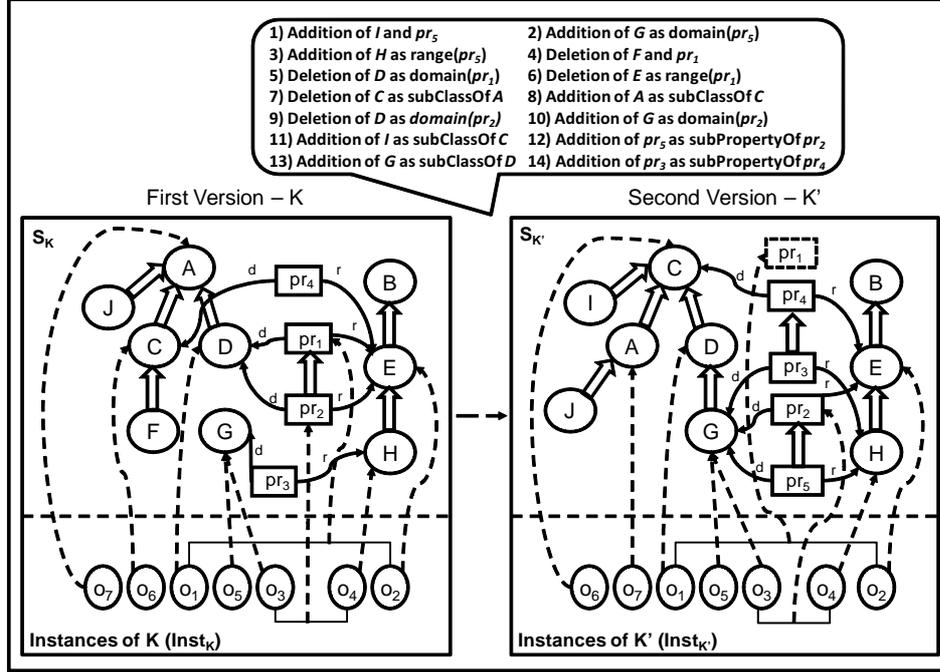


Fig. 9. Non-Backwards Compatible Schema Evolution

7 Sequences of Transitions

In this Section, we investigate what happens in *sequences* of successive transitions.

Let $K_1 = (S_1, I_1)$ be a KB and let S_2, \dots, S_n be a sequence of new schema versions, resulting to KBs $K_i = (S_i, I_1)$, for $i = 2, \dots, n$. Now consider the *one-step* migration of I_1 from the first (S_1) to the last (S_n) schema, i.e. consider the transition:

$$(\mathcal{C}_1, F_1, P_1) \rightsquigarrow (\mathcal{C}_n, F_n, P_n)$$

where \mathcal{C}_1 is based on S_1 and I_1 , and suppose that $P_1 = \emptyset$ ¹⁴.

Now consider a *sequential migrations* scenario where I_1 is migrated to S_2 , then to S_3 , and so on, up to S_n . So, the scenario consists of the following sequence of $n - 1$ transitions:

$$\begin{aligned} (\mathcal{C}_1, F_1, P_1) &\rightsquigarrow (\mathcal{C}'_2, F'_2, P'_2) \\ &\rightsquigarrow \dots \\ &\rightsquigarrow (\mathcal{C}'_{n-1}, F'_{n-1}, P'_{n-1}) \\ &\rightsquigarrow (\mathcal{C}'_n, F'_n, P'_n) \end{aligned}$$

Let us now discuss the relationship between the outcome of the *one-step* migration, i.e. $(\mathcal{C}_n, F_n, P_n)$, with respect to the final outcome of the *sequential migrations*, i.e. with $(\mathcal{C}'_n, F'_n, P'_n)$. The main points are:

(a) $\mathcal{C}_n = \mathcal{C}'_n$, i.e. the certain parts of the resulting partitions are the same. This is

¹⁴ To keep notations simple, here and below we omit the subscripts K from the notations of \mathcal{C} , P and M .

due to postulate *II1* of Def. 7.

(b) The rest parts of the partitions (i.e. F and P) can be different.

Point (b) can be made evident through a small example, like that of Figure 10. The last row (III) shows the one step migration. The first and the second rows show two different sequential migrations that lead to the *same* final schema. Notice that the first sequence gives the same result with the one step migration, however the second does not.

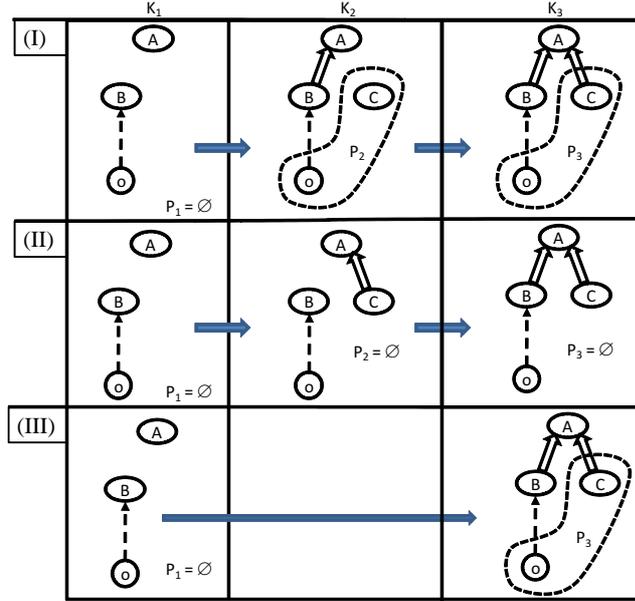


Fig. 10. Different sequences of transitions

We could say that this phenomenon is a kind of “*negative prejudgement, due to lack of the final (forthcoming) knowledge*”. In our case, the negative prejudgement is realized by postulate *II2*.

[*Suggested Policy*]

If no user feedback is expected/given after a migration, then there is no need to compute or store the intermediate P_i . Instead, it is better to compute it between the first and the last schema, and only when needed (i.e. just before the curator starts the lifecycle process that we describe in the next section). In this way, we can bypass the “negative prejudgement” due to lack of the forthcoming knowledge.

Below, we present a proposition that indicates when the *one-step migration* gives the same result with the *sequential migrations* scenario.

Prop. 16 Let $K_1 = (S_1, I_1)$ be a KB and let S_2, \dots, S_n be a sequence of *backwards compatible* schema versions, resulting to KBs $K_i = (S_i, I_1)$, for $i = 2, \dots, n$. Now consider the *one-step* migration of I_1 from the first (S_1) to the last (S_n) schema: $(C_1, F_1, P_1) \rightsquigarrow (C_n, F_n, P_n)$. Additionally, consider the *sequential migrations* scenario, where: $(C_1, F_1, P_1) \rightsquigarrow (C'_2, F'_2, P'_2) \rightsquigarrow \dots \rightsquigarrow (C'_n, F'_n, P'_n)$.

If (a) for all $c_1, c_2 \in C_{K_1}$, it holds that $c_1 \leq_{cl}^* c_2$ in K_1 iff $c_1 \leq_{cl}^* c_2$ in K_n , and (b) for all $pr_1, pr_2 \in Pr_{K_1}$, it holds that $pr_1 \leq_{pr}^* pr_2$ in K_1 iff $pr_1 \leq_{pr}^* pr_2$ in K_n , then $(C_n, F_n, P_n) = (C'_n, F'_n, P'_n)$. \square

We would like to note that Prop. 16 does not hold in the case that the sequence of schema versions is *not backwards compatible* (see Example 5 below). This happens because a class or property that exist in a schema version S_i may have been removed from the next schema version S_{i+1} .

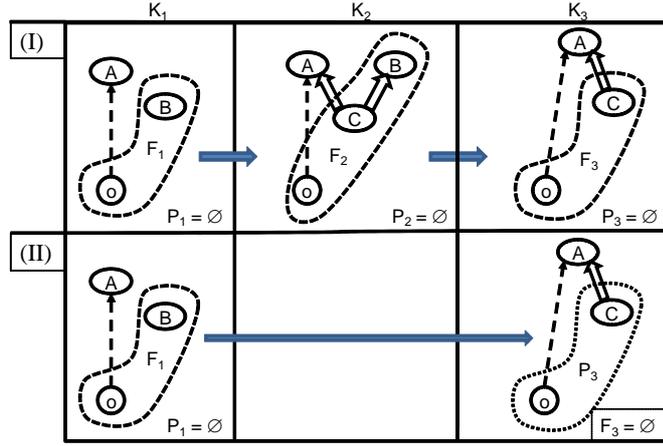


Fig. 11. Different sequences of non-backwards compatible transitions

Example 5. Consider Figure 11. In KB K_1 , there exist two classes A and B and o is an instance of A. Assuming that the *MSA* holds for K_1 , it follows that $P_1 = \emptyset$ and $F_1 = \{(o \text{ type B})\}$. In scenario (I), in KB K_2 , a new class C is added, as well as the relationships $C \leq_{cl} A$ and $C \leq_{cl} B$. From Rule R1, it follows that $F_2 = \{(o \text{ type B}), (o \text{ type C})\}$. In K_3 , class B is removed. Thus, $F_3 = \{(o \text{ type C})\}$. In scenario (II), class B is removed and a new class C is added, along with the relationship $C \leq_{cl} A$. Then, since none of the rules R1 – R5 applies, it follows that $F_3 = \emptyset$ and $P_3 = \{(o \text{ type C})\}$. Note that the non-backwards compatible transitions in scenarios (I) and (II) lead to different results. \square

8 Ranking Possibilities

Here, we discuss methods for ranking the computed possibilities for controlling the amount of possible information that is kept stored, and for aiding the interaction process that will be described later on.

If t is a triple in P_K or $C_i(K)$, we could “score” t according to a *degree of certainty*. A naive approach would be to set $score(t) = 1$ if $t \in C_i(K)$, and $score(t) = 1/2$ if $t \in P_K$. Below we introduce a more sophisticated model for ranking the possible triples. Consider an example, where $PhD_Student \leq_{cl}^* Postgraduate \leq_{cl}^* Student \leq_{cl}^* Person$ and $Manager \leq_{cl}^* Employee$. Assume that John was originally classified to the class *Person*. We can say that *Student* and *Employee* are more probable classes for John than *Postgraduate*, *Manager*, and *PhD_Student*. To this end we propose an extension of P_K that we call *quantified P_K* , such that each triple

is accompanied by a positive integer that is interpreted as follows: the *less* this value is the *more possible* the triple is.

Let $dist_{cl}(c \rightarrow c')$ be the length of the shortest path from class c to class c' comprised from \leq_{cl} relationships (over the reflexive and transitive reduction of \leq_{cl}^*). If there is no path from class c to class c' comprised from \leq_{cl} relationships then $dist_{cl}(c \rightarrow c') = \infty$. If $c \leq_{cl} c_1 \leq_{cl} c_2 \leq_{cl} c_3 \leq_{cl} c'$ and this is the shortest path from class c to class c' comprised from \leq_{cl} relationships, then $dist_{cl}(c \rightarrow c') = 4$. For each element (*o type c*) in Pr_K , the *quantified* Pr_K contains an element (*(o type c), distClass(o, c)*), where:

$$distClass(o, c) = \min\{dist_{cl}(c \rightarrow c') \mid c' \in C_K \text{ and } o \in inst_K(c')\}$$

So, $distClass(o, c)$ is the shortest distance of c from one of the certain classes of o . A measure, similar to $distClass(o, c)$, appears in [23] for measuring the conceptual distance between two concepts.

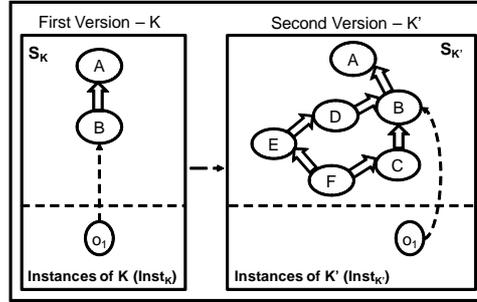


Fig. 12. Ranking possible class instance triples

Example 6. Consider Figure 12. The second schema version $S_{K'}$ is backwards compatible with the first schema version S_K . The new set of classes from K to K' is $\{C, D, E, F\}$. So, according to Algorithm 1, the derived possible class instance triples are: $\{(o_1 \text{ type } C), (o_1 \text{ type } D), (o_1 \text{ type } E), (o_1 \text{ type } F)\}$.

If we rank the possible class instance triples of the instance o_1 , using the above formula, i.e. $distClass(o, c)$, where o corresponds to o_1 and c corresponds to one of the possible classes of o_1 , we get the following quantified possible class instance triples: $\{((o_1 \text{ type } C), 1), ((o_1 \text{ type } D), 1), ((o_1 \text{ type } E), 2), ((o_1 \text{ type } F), 2)\}$.

For example, note that in the case of the possible class instance triple ($o_1 \text{ type } F$), there are two paths from F to B (which is a certain class of o_1). The first one is $F \leq_{cl} C \leq_{cl} B$ and the second one is $F \leq_{cl} E \leq_{cl} D \leq_{cl} B$. Thus, $dist_{cl}(F \rightarrow B)$ is 2. Additionally, there are two paths from F to A (which is also a certain class of o_1). The first one is $F \leq_{cl} C \leq_{cl} B \leq_{cl} A$ and the second one is $F \leq_{cl} E \leq_{cl} D \leq_{cl} B \leq_{cl} A$. Thus, $dist_{cl}(F \rightarrow A)$ is 3. Therefore, $distClass(o_1, F) = \min(\{2, 3\}) = 2$. \square

Similarly, let $dist_{pr}(pr \rightarrow pr')$ be the length of the shortest path from property pr to property pr' comprised from \leq_{pr} relationships (over the reflexive and transitive reduction of \leq_{pr}^* , which is unique in our case because \leq_{pr}^* is finite and acyclic). If there is no path from property pr to property pr' comprised from \leq_{pr} relationships then $dist_{pr}(pr \rightarrow pr') = \infty$. If $pr \leq_{pr} pr_1 \leq_{pr} pr_2 \leq_{pr} pr_3 \leq_{pr} pr'$ and this is the shortest path from pr to pr' comprised from \leq_{pr} relationships, then $dist_{pr}(pr \rightarrow pr') = 4$. We define an auxiliary property pr_{dummy} such that if $pr \in Pr_K$ and there is no $pr' \in Pr_K$

s.t. $pr \leq_{pr} pr'$ then we add $pr \leq_{pr} pr_{\text{dummy}}$. For each element $(o \ pr \ o')$ in P_K , the *quantified* P_K contains an element $((o \ pr \ o'), distProperty(o, pr, o'))$, where:

$$distProperty(o, pr, o') = \begin{cases} \min\{dist_{pr}(pr \rightarrow pr') \mid pr' \in Pr_K, (o \ pr' \ o') \in \mathcal{C}(K)\} \\ \quad \text{if } \exists pr' \in Pr_K \text{ s.t. } (o \ pr' \ o') \in \mathcal{C}(K) \\ dist_{pr}(pr \rightarrow pr_{\text{dummy}}) & \text{otherwise} \end{cases}$$

So, $distProperty(o, pr, o')$ is the shortest distance of pr from one of the properties pr' such that $(o \ pr' \ o') \in \mathcal{C}(K)$. In the case that there is no property pr' such that $(o \ pr' \ o') \in \mathcal{C}(K)$ then $distProperty(o, pr, o')$ is the distance of pr from pr_{dummy} .

The $distClass(o)$ and $distProperty(o, o')$ values can be used for ranking the possible class instance triples and possible property instance triples, respectively.

9 Specificity Life Cycle Management

So far, we have seen principles and algorithms for computing possibilities. Now, we will focus on the management of these possibilities, specifically on the *exploitation* of the computed possibilities.

Suppose a number of instance descriptions that are migrated to a new schema version, and as a consequence our machinery has computed a set of possible instance triples P_K . If o is an instance, let us denote by $posTriples(o)$ all those triples that include o and belong to P_K . Suppose a system that for an instance o shows to the user a set of possible instance triples $U(o)$ such that $U(o) \subseteq posTriples(o)$. The user then decides whether he/she should add one or more than one of these to the certain knowledge base. After that we should also update P_K . Figure 4 (on the right) illustrates the proposed process.

In Figure 4, on the left, we can see the current migration process, where the curator of a KB downloads its latest ontology version, he migrates the instance descriptions to that version, and then he manually tries to revise some of the migrated descriptions. On the right, we can see the proposed migration process. After the migration of the instance descriptions to the latest ontology version, the system computes the possible instance descriptions by executing either Algorithm 1 (when the new schema is backwards compatible with the previous one) or Algorithm 2 (when the new schema is not backwards compatible with the previous one). Then, an iterative procedure starts, where in each iteration, the curator can select a specific instance and then a ranked subset of its possible descriptions is displayed. The curator can accept or reject some or all of these possible descriptions (of the selected instance) and then the *eKB* (its certain and possible part) is updated, analogously. Below, we describe formally the updating of the certain and the possible part of the *eKB*.

What we have to ensure is that the update should: (a) respect the user's request, (b) reduce uncertainty based on what the user was prompted and decided, and (c) yield a valid *eKB* (that respects the constraint of Def. 10). To specify exactly the updating we need to introduce notations for the possible class instance triples and property instance triples of an instance o .

$$\begin{aligned} posTriples^{cl}(o) &= \{(o \ type \ c) \mid (o \ type \ c) \in P_K\} // \text{cl: instance of class} \\ posTriples^{spr}(o) &= \{(o \ pr \ o') \mid (o \ pr \ o') \in P_K\} // \text{spr: subject of property} \\ posTriples^{opr}(o) &= \{(o' \ pr \ o) \mid (o' \ pr \ o) \in P_K\} // \text{opr: object of property} \end{aligned}$$

At Section 3 we defined $SubTriples(\cdot)$. Now we define the $SuperTriples$ of an instance triple, or a set of instance triples A , as follows:

$$\begin{aligned} SuperTriples((o \text{ type } c)) &= \{(o \text{ type } c') \mid c \leq_{cl}^* c'\} \\ SuperTriples((o \text{ pr } o')) &= \{(o \text{ pr}' o') \mid pr \leq_{pr}^* pr'\} \\ SuperTriples(A) &= \bigcup_{t \in A} SuperTriples(t) \end{aligned}$$

[Rejections].

Let o be an instance. If the system shows to the user all instance triples in $posTriples^t(o)$, where $t \in \{pr, spr, opr\}$, and he/she decides that none of these should be added then we should update P_K as follows: $P_K^{up} = P_K \setminus posTriples^t(o)$ where in P_K^{up} , up stands for *updated*. If the system had prompted to the user only a subset of the possibilities, say $U(o)$ (where $U(o) \subseteq posTriples^t(o)$), and the user had decided that none of these should be added then: $P_K^{up} = P_K \setminus SubTriples(U(o))$. Note that except from $U(o)$, the $SubTriples(U(o))$ should be removed from P_K due to (ii) of Def. 5. This ensures that not only $U(o)$ but also all possibilities which are more “specific” than those in $U(o)$ will be excluded. Note that the instance triples in $U(o)$ and thus, the instance triples in $P_K \cap SubTriples(U(o))$, are actually moved to F_K^{up} .

[Acceptances].

If the user has selected some of the suggested possibilities, say $X(o)$, ($X(o) \subseteq U(o)$), to be added, then we should update the certain part of the new KB $K^{up} = K \cup X(o)$, as follows $\mathcal{C}(K^{up}) = \mathcal{C}(K) \cup SuperTriples(X(o))$ (see Prop. 17 below), and then update P_K , accordingly. The latter can be done as follows:

- 1) $P_K^{up} = P_K \setminus SuperTriples(X(o))$
- 2) $P_K^{up} = P_K^{up} \setminus SubTriples(U(o) \setminus SuperTriples(X(o)))$

The first step excludes from P_K also the newly entailed instance triples from K^{up} , i.e. $SuperTriples(X(o))$, (to satisfy the constraint of Def. 10). The second excludes from P_K^{up} the instance triples in $U(o) \setminus SuperTriples(X(o))$ and their specializations. Note that the instance triples in $U(o) \setminus SuperTriples(X(o))$, and thus the instance triples in $P_K \cap SubTriples(U(o) \setminus SuperTriples(X(o)))$, are actually moved to F_K^{up} .

Prop. 17 Let $X \subseteq P_K$. If $K^{up} = K \cup X$ then $\mathcal{C}(K^{up}) = \mathcal{C}(K) \cup SuperTriples(X)$. \square

Example 7. Consider the scenario of Figure 13. P_K and $posTriples^{cl}(\mathbf{John})$ contain the triples: $\{(\mathbf{John} \text{ type } \mathbf{Employee}), (\mathbf{John} \text{ type } \mathbf{Manager}), (\mathbf{John} \text{ type } \mathbf{Student}), (\mathbf{John} \text{ type } \mathbf{Postgraduate}), (\mathbf{John} \text{ type } \mathbf{PhD_Student})\}$ (shown enclosed in a dashed frame)¹⁵. If the system shows to the user all triples in $posTriples^{cl}(\mathbf{John})$ and the user decides that none of these should be added, then P_K should be updated as follows: $P_K^{up} = P_K \setminus posTriples^{cl}(\mathbf{John}) = \emptyset$.

Now suppose that the system had showed to the user only three of the five possible instance triples, such as: $U(\mathbf{John}) = \{(\mathbf{John} \text{ type } \mathbf{Employee}), (\mathbf{John} \text{ type } \mathbf{Student}),$

¹⁵ In the figure, $posCl(\mathbf{John})$ denotes the possible classes of \mathbf{John} .

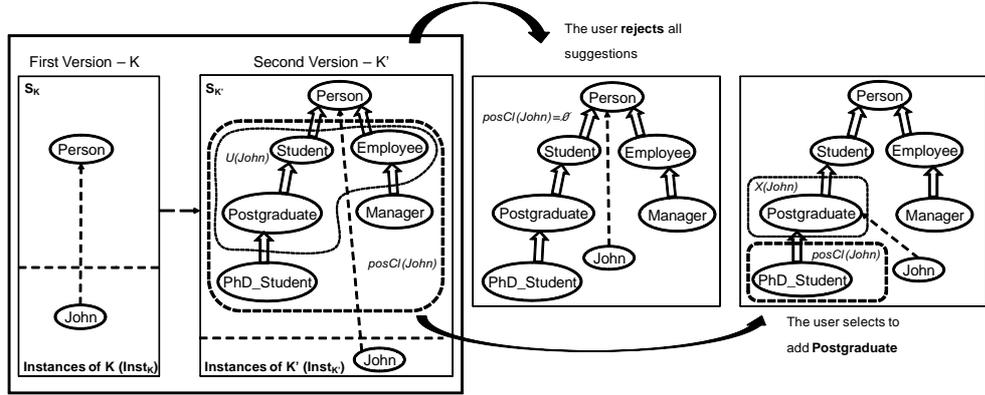


Fig. 13. Example of the Possibility Resolution Process

$(\text{John type Postgraduate})\}$. If the user decides that he/she does not want to add any of these triples to the certain KB then P_K should be updated as follows:

$$P_K^{up} = P_K \setminus \text{SubTriples}(\{(\text{John type Employee}), (\text{John type Student}), (\text{John type Postgraduate})\}) = \emptyset$$

However, if the user decides to add one of these three suggested triples, say the triple $X(\text{John}) = \{(\text{John type Postgraduate})\}$ then the KB has to be updated such that $K^{up} = K \cup \{(\text{John type Postgraduate})\}$. Note that $\text{SubTriples}(X(\text{John})) = \{(\text{John type Postgraduate}), (\text{John type PhD_Student})\}$ and $\text{SuperTriples}(X(\text{John})) = \{(\text{John type Postgraduate}), (\text{John type Student}), (\text{John type Person})\}$.

The possible part P_K has to be updated as follows:

- 1) $P_K^{up} = \{(\text{John type Employee}), (\text{John type Manager}), (\text{John type Student}), (\text{John type Postgraduate}), (\text{John type PhD_Student})\} \setminus \text{SuperTriples}(X(\text{John}))$
 $= \{(\text{John type Employee}), (\text{John type Manager}), (\text{John type PhD_Student})\}$
- 2) $P_K^{up} = P_K^{up} \setminus \text{SubTriples}(U(\text{John}) \setminus \text{SuperTriples}(X(\text{John})))$
 $= P_K^{up} \setminus \text{SubTriples}(\{(\text{John type Employee})\})$
 $= \{(\text{John type PhD_Student})\}$

The updated KB K , i.e. K^{up} , and the updated P_K , i.e. P_K^{up} , can now be given as input to Algorithm 1, in the case that a new set of schema triples $S_{K'}$ is available, for generating the new $P_{K'}$. Note that the set of instance triples that is migrated from S_K to $S_{K'}$ is now $I_{K^{up}}$.

The ranking methods defined at Section 8 can aid the interaction, e.g. at the example of Figure 13, the system will suggest first **Student** and **Employee** as possible classes of **John**. If the user will not select any of these then P_K will be updated and **Postgraduate**, **Manager**, and **PhD_Student** will never be suggested as possible classes (for the instance **John**).

However, the important point is that if the curation process is followed and the curator accepts/rejects the migration-related uncertainties, then the possible part of the KB becomes empty, i.e. no extra storage is required.

10 Implementation and Application Issues

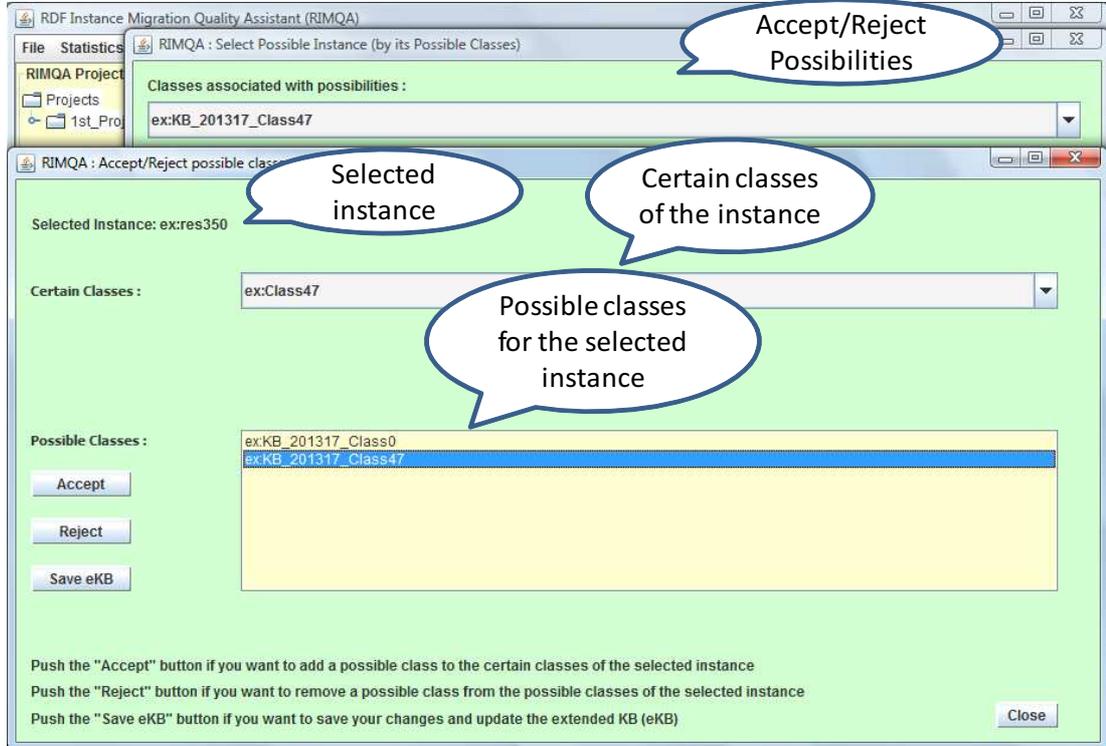


Fig. 14. Selected screenshot of RIMQA

We have implemented a proof-of-concept prototype, called RIMQA (**R**DF **I**nstance **M**igration **Q**uality **A**ssistant) which supports the entire lifecycle that we have proposed¹⁶. A selected screenshot of the system is shown in Figure 14. Below we report some experimental results for demonstrating the feasibility of the approach (Section 10.1), we introduce methods for compacting the representation of possibilities (Section 10.2), and finally we discuss applicability issues (Section 10.3).

10.1 Experimental Evaluation

We have conducted an experimental evaluation whose objectives is to (a) measure the size in triples of the computed $P_{K'}$, and

(b) measure the time required to compute $P_{K'}$ for investigating the applicability of this method to large datasets. Regarding datasets and measurements, we adopt the following methodology. For each dataset we get a sequence of schema versions, specifically a sequence of the form S_{K_0}, \dots, S_{K_n} , where each S_{K_i} is a set of schema triples, and a set of instance triples I_K w.r.t. the first version of the schema, i.e. S_{K_0} . Subsequently, we migrate I_K to each of the subsequent versions of the schema and for each one we compute the corresponding P_{K_i} . Specifically, we do the migrations

¹⁶ More information about the prototype is available through its web page: <http://www.ics.forth.gr/isl/RIMQA>.

$S_{K_{i-1}} \rightarrow S_{K_i}$ for all $i = 1 \dots n$. We use two scenarios: in the first, we consider that $P_{K_{i-1}} = \emptyset$ and thus the computation of P_{K_i} depends only on the current migration, i.e. $S_{K_{i-1}} \rightarrow S_{K_i}$, while in the second $P_{K_{i-1}}$ has been specified from the migration $S_{K_{i-2}} \rightarrow S_{K_{i-1}}$ (or former migrations).

The implementation is written in Java using the SWKM¹⁷ (Semantic Web Knowledge Middleware) main memory model¹⁸. One important implementation detail is that we *do not have to compute the closure* of any of the involved KBs. Instead we check whether a particular triple belongs to the closure and this is done efficiently by exploiting the labeling schemes [5] supported by SWKM.

Real Data Set We used the RDF/S versions of Music Ontology¹⁹ which is a formal framework for dealing with music-related information on the Semantic Web, including editorial, cultural and acoustic information. In our experiments, we used the following successive versions: v.7.08.2007, v.10.08.2007, v.12.08.2007, v.18.09.2007, v.6.12.2007, v.28.07.2008, v.28.10.2008 and v.13.02.2010. Since each version is not backwards compatible with the previous ones (although it is migration compatible), Algorithm 2 is used.

For each version, Table 2 shows the number of classes (i.e. $|C_i|$), the new classes added from S_K to $S_{K'}$ (i.e. $|C_i| - |C_{i-1}|$), the number of properties (i.e. $|Pr_i|$), the new properties added from S_K to $S_{K'}$ (i.e. $|Pr_i| - |Pr_{i-1}|$), the number of explicit instance triples (i.e. $|I_K|$) that are migrated and the number of explicit schema triples (i.e. $|S_K|$). The last line shows the size difference in schema triples between S_K and $S_{K'}$. We can see that version v.18.09.2007 has fewer schema triples than version v.12.08.2007 and version v.28.07.2008 has fewer schema triples than version v.6.12.2007. This means that a number of deletions has taken place from v.12.08.2007 to v.18.09.2007 and from v.6.12.2007 to v.28.07.2008, accordingly.

	Versions of Music Ontology (S_i)								
	v.7 .08 .2007	v.10 .08 .2007	v.12 .08 .2007	v.18 .09 .2007	v.6 .12 .2007	v.28 .07 .2008	v.28 .10 .2008	v.13 .02 .2010	
$ C_i $	113	94	93	93	94	86	124	95	
$ C_i - C_{i-1} $		-19	-1	0	1	-8	38	-29	
$ Pr_i $	147	167	167	167	174	160	160	183	
$ Pr_i - Pr_{i-1} $		20	0	0	7	-14	0	23	
$ I_i $	26	26	26	26	26	26	26	26	
$ S_i $	1059	1266	1266	1259	1269	1115	1153	1302	
$ S_i - S_{i-1} $		207	0	-7	10	-154	38	149	

Table 2. $|C_i|$, $|Pr_i|$, $|I_i|$ and $|S_i|$ for each i Music Ontology version

Number of Possible Triples. Table 3 shows the number of possible triples (i.e. $|P_{K'}|$) produced during the migration from S_K to $S_{K'}$, assuming that $P_K = \emptyset$, while in Table 4 we consider that $P_K \neq \emptyset$ for each migration from S_K to $S_{K'}$. Specifically

¹⁷ <http://athena.ics.forth.gr:9090/SWKM>

¹⁸ All experiments were carried out in an ordinary laptop with processor Pentium(R) Dual-Core CPU T4200 @2.0 Ghz, 2 GB Ram, running Windows Vista.

¹⁹ <http://www.musicontology.com>

	$ P_{K'} $ Sizes (for migrations where $P_K = \emptyset$)						
$ P_{K'} $ changes at each part of Alg. 2	v.7 .08.2007 → v.10 .08.2007	v.10 .08.2007 → v.12 .08.2007	v.12 .08.2007 → v.18 .09.2007	v.18 .09.2007 → v.6 .12.2007	v.6 .12.2007 → v.28 .07.2008	v.28 .07.2008 → v.28 .10.2008	v.28 .10.2008 → v.13 .02.2010
A	0	0	0	+21	+34	+289	+153
D	0	0	0	0	0	0	0
TOTAL	0	0	0	21	34	289	153

Table 3. $|P_{K'}|$ Sizes for Music Ontology migrations. $|P_K| = 0$

	$ P_{K'} $ Sizes (for migrations where $P_K \neq \emptyset$)		
$ P_{K'} $ changes at each part of Alg. 2	v.6.12.2007 → v.28.07.2008 P_K (v.18.09.2007 → v.6.12.2007)	v.28.07.2008 → v.28.10.2008 P_K (v.6.12.2007 → v.28.07.2008)	v.28.10.2008 → v.13.02.2010 P_K (v.28.07.2008 → v.28.10.2008)
$ P_K $	21	34	289
A	+34	+289	+153
B	0	0	0
C	0	-17	-289
D	0	0	0
E	0	0	0
F	0	0	0
Line 18	0	0	0
TOTAL	55	306	153

Table 4. $|P_{K'}|$ Sizes for Music Ontology migrations. $|P_K| \neq 0$

in Table 4, P_K is the one derived by the previous migration. The last lines of these tables show the final $|P_{K'}|$ for each migration $S_K \rightarrow S_{K'}$. As we can see in Table 3, during the migrations v.7.08.2007 → v.10.08.2007, v.10.08.2007 → v.12.08.2007, and v.12.08.2007 → v.18.09.2007, Algorithm 2 does not return any possible triples. Moreover, note that the only possible triples that are produced, are those after the execution of part A of Algorithm 2, i.e. after the update of P_K as a consequence of the new classes added from S_K to $S_{K'}$. Note that in Table 3, only the parts A and D of Algorithm 2 are shown, because the rest parts do not reduce the possible triples, since $P_K = \emptyset$. In Table 4, we can see that possible triples are produced again only by part A and we also have deletions by part C of Algorithm 2.

Execution Times. Tables 5 and 6 show the execution times, corresponding to the scenarios of Table 3 and 4, respectively. We can see that the migration takes just some milliseconds and mainly depends on the number of the new classes (part A) added from S_K to $S_{K'}$.

	$P_{K'}$ Execution Times (for migrations where $P_K = \emptyset$)						
Times for parts of Alg. 2	v.7 .08.2007 → v.10 .08.2007	v.10 .08.2007 → v.12 .08.2007	v.12 .08.2007 → v.18 .09.2007	v.18 .09.2007 → v.6 .12.2007	v.6 .12.2007 → v.28 .07.2008	v.28 .07.2008 → v.28 .10.2008	v.28 .10.2008 → v.13 .02.2010
A	11.0	4.7	0.6	3.3	4.2	88.5	23.8
D	4.7	0.1	0.1	3.9	3.4	1.3	4.2
TOTAL	15.7	4.8	0.7	7.2	7.6	89.8	27.7

Table 5. Execution Times (in msec) for Music Ontology migrations for $|P_K| = 0$

	$P_{K'}$ Execution Times (for migrations where $P_K \neq \emptyset$)		
Times for each part of Alg. 2	v.6.12.2007 → v.28.07.2008 P_K (v.18.09.2007 → v.6.12.2007)	v.28.07.2008 → v.28.10.2008 P_K (v.6.12.2007 → v.28.07.2008)	v.28.10.2008 → v.13.02.2010 P_K (v.28.07.2008 → v.28.10.2008)
A	3.5	167.0	210.9
B	0.6	0.5	1.8
C	0.1	0.5	4.0
D	0.8	0.1	4.0
E	0.2	0.2	0.3
F	0.1	0.1	0.1
LINE 18	0.5	0.6	4.6
TOTAL	5.8	169.0	225.7

Table 6. Execution Times (in *msec*) for Music Ontology migrations for $|P_K| \neq 0$

Synthetic Data Set To conduct experiments over larger datasets and backwards compatible ontologies, we created and used one synthetic data set. Specifically, using the synthetic KB generator described in [27], we created a KB, v.K1, with 100 classes and 300 properties. To obtain a schema whose features resemble those of real ones, the subsumption relation follows a power law distribution. Specifically and accordingly to the metrics used in [28], we set the power-law exponent to 0.5 for the total-degree VR²⁰ function of the property graph and to 1.7 for the PDF²¹ function of the descendants distribution. The depth of the class subsumption hierarchy in the schema is 7. Subsequently, we created three subsequent schemas of v.K1, namely v.K2, v.K3 and v.K4; v.K2 was derived by adding to v.K1 14 new classes as specializations to randomly selected leaf classes of v.K1 and similarly 10 new properties as specializations to randomly selected properties of v.K1; v.K3 was derived by adding to v.K2 14 new classes as specializations to randomly selected leaf classes of v.K2 (where 11 of them existed also in v.K1) and similarly 10 new properties as specializations to randomly selected properties of v.K2 (where 9 of them existed also in v.K1); v.K4 was derived by adding to v.K3 14 new classes as specializations to randomly selected leaf classes of v.K3 (where 13 of them existed also in v.K1) and similarly 10 new properties as specializations to randomly selected properties of v.K3 (where 9 of them existed also in v.K1).

For each class of v.K1 we created 100 instances, while for each property of v.K1 we created 10 property instance triples, among randomly selected instances of the corresponding domain and range classes. Table 7 shows the features of these schemas and the number of instance triples.

Number of Possible Triples. Since each version is backwards compatible with the previous ones, Algorithm 1 is used. Table 8 shows the number of possible triples (i.e. $|P_{K'}|$) produced during the migration from S_K to $S_{K'}$, assuming that $P_K = \emptyset$, while in Table 9 we consider that $P_K \neq \emptyset$ for each migration from S_K to $S_{K'}$. Specifically in Table 9, P_K is the one derived by the previous migration. The last lines of these tables show the final $|P_{K'}|$, for each migration $S_K \rightarrow S_{K'}$.

Note that in Table 8, migration v.K2 → v.K3 and v.K3 → v.K4 produce fewer possible class instance triples than v.K1 → v.K2. This is because some of the added

²⁰ VR stands for Value vs Rank (it measures the relationship between the i^{th} biggest value and its rank i , assuming a descending order).

²¹ PDF stands for Probability Density Function.

	Versions of Synthetic Data (S_i)			
	v.K1	v.K2	v.K3	v.K4
$ C_i $	101	115	129	143
$ C_i - C_{i-1} $		14	14	14
$ Pr_i $	289	299	309	319
$ Pr_i - Pr_{i-1} $		10	10	10
$ I_i $	3790	3790	3790	3790
$ S_i $	1166	1219	1273	1338
$ S_i - S_{i-1} $		53	54	65

Table 7. $|C_i|$, $|Pr_i|$, $|I_i|$ and $|S_i|$ for each i version of the Synthetic data set

$ P_{K'} $ changes at each part of Alg. 1	$ P_{K'} $ Sizes (for migrations where $P_K = \emptyset$)		
	v.K1 \rightarrow v.K2	v.K2 \rightarrow v.K3	v.K3 \rightarrow v.K4
A	+1290	+230	+150
C	+97	+87	+86
TOTAL	1387	317	236

Table 8. $|P_{K'}|$ Sizes for Synthetic Data migrations. $|P_K| = 0$

classes from v.K2 to v.K3 (or from v.K3 to v.K4 respectively) are subclasses of classes added from v.K1 to v.K2 (or from v.K1 to v.K2 and from v.K2 to v.K3 respectively), which have no instances (although the rest classes have instances).

$ P_{K'} $ changes at each part of Alg. 1	$ P_{K'} $ Sizes (for migrations where $P_K \neq \emptyset$)	
	v.K2 \rightarrow v.K3 P_K (v.K1 \rightarrow v.K2)	v.K3 \rightarrow v.K4 P_K (v.K2 \rightarrow v.K3)
$ P_K $	1387	317
A	+1250	+150
B	0	0
C	+97	+86
D	0	0
LINE 14	0	0
TOTAL	2734	553

Table 9. $|P_{K'}|$ Sizes for Synthetic Data migrations. $|P_K| \neq 0$

Execution Times. Tables 10 and 11 show the execution times, corresponding to the scenarios of Table 8 and 9, respectively. We can observe times that range from 4 seconds to 7 minutes and the cost mainly depends on the number of new classes (part A) and on the number of new properties (part C).

Times for parts of Alg. 1	$P_{K'}$ Execution Times (for migrations where $P_K = \emptyset$)		
	v.K1 \rightarrow v.K2	v.K2 \rightarrow v.K3	v.K3 \rightarrow v.K4
A	1.3	3.0	1.7
C	4.8	3.7	2.6
TOTAL	6.1	6.7	4.3

Table 10. Execution Times (in sec) for Synthetic Data migrations for $|P_K| = 0$

	$P_{K'}$ Execution Times (for migrations where $P_K \neq \emptyset$)	
Times for each part of Alg. 1	$v.K2 \rightarrow v.K3$ $P_K (v.K1 \rightarrow v.K2)$	$v.K3 \rightarrow v.K4$ $P_K (v.K2 \rightarrow v.K3)$
A	284.7	52.2
B	1.4	0.03
C	111.9	38.3
D	0.3	0.2
LINE 14	0.04	0.01
TOTAL	398.5	90.74

Table 11. Execution Times (in *sec*) for Synthetic Data migrations for $|P_K| \neq 0$

Scalability Issues and Optimizations

We have seen that parts A and C of Algorithm 1 and parts A and D of Algorithm 2 require most of the time. Below, we analyze their costs. Specifically, in part A of Algorithm 1 and Algorithm 2, for each *new* class or property, we have to make an iteration whose cost is equal to the number of elements of $Inst_K$ multiplied by the number of the superclasses of the new class/property, multiplied by the cost of performing two lookups in closures. Specifically, the lookups, for part A, are the following: (*o type c*) $\in (\mathcal{C}_i(K') \cup P_K)$ and (*o type c'*) $\notin \mathcal{C}_i(K')$, while the lookups for part C are: (*o pr o'*) $\in (\mathcal{C}_i(K') \cup P_K)$ and (*o pr' o'*) $\notin \mathcal{C}_i(K')$.

Therefore, we can say that the cost essentially depends on $Inst_K$, the number of new classes/properties, their depth (specifically their superclasses) and the cost of lookup in closures and in P_K .

First, note that the lookups can be performed fast if the involved closures are materialized and indexed. This is a common practice that is followed for supporting efficient query answering (this means that it is not unrealistic to expect that $\mathcal{C}_i(K')$ has been produced and stored). In the reported times we had not materialized the closures.

Various other optimizations can be devised and applied. For instance, for each $c' \in NC$, we can get each direct superclass $c \in C_K$ of c' and if c has no certain or possible instances then we can ignore c' . The gain is that we avoid unnecessary scans at $Inst_K$.

Even if we cannot ignore a new class c' , we can ignore instances $o \in Inst_K$ if there is a superclass $c \in C_K$ of c' such that (*o type c*) $\notin \mathcal{C}_i(K') \cup P_K$. In this case, not all superclasses $c \in C_K$ of c' need to be checked for satisfying the condition (*o type c*) $\in \mathcal{C}_i(K') \cup P_K$.

The above optimizations are just indicative. One could even change the flow of control of part A. For instance, instead of scanning the entire $Inst_K$ in line (6) of Algorithms 1 and 2 (part A), one could scan only the URIs of the instances of the classes which are superclasses of the new classes. Specifically, instead of the expression $o \in Inst_K \cap URI$, we can have $o \in I_A \cup I_B$, where I_A and I_B are computed as follows: at first, we compute the superclasses of the new classes, i.e. $SupsOfNewClasses = \bigcup_{c' \in NC} \{c \mid c' \leq_{cl}^* c\}$, and then we get their instances in $\mathcal{C}_i(K')$ and P_K , i.e. $I_A = \{o \mid (o \text{ type } c) \in \mathcal{C}_i(K') \text{ and } c \in SupsOfNewClasses\}$, $I_B = \{o \mid (o \text{ type } c) \in P_K \text{ and } c \in SupsOfNewClasses\}$.

The gain from the above change is that now the cost of part A of Algorithms 1 and 2 does not depend on $Inst_K$ but on $|NC|$ multiplied by the average number of instances that a class has plus the average number of possible instances of a

class. Of course, we have to consider the cost of computing *SupOfNewClasses*, I_A , and I_B , which however does not increase the theoretical complexity. Note that such optimization is not needed for line (11) of Algorithm 1 (part C) and line (13) of Algorithm 2 (part D) because $valid(o, pr', o', K')$ should hold.

Other optimizations can be designed depending on the implementation context, e.g. if a triple store is used then the set *SupOfNewClasses* can be computed by a single SPARQL query. The same is true for the sets I_A and I_B .

What we have seen in our experiments is that the computation of possibilities after a migration, using an ordinary laptop and *without* any special optimization, takes at most 7 minutes for the KBs in our datasets. With optimizations like those that we have sketched before, that could be significantly reduced.

Finally, we should say that since migration is not an every day task, the computation of possibilities after a migration is not expected to have any significant scalability problem.

10.2 A Compact Representation for Possibilities

Even if our approach does not require storing any kind of negative information, one may wonder whether the proposed approach is still prohibitively expensive, in terms of space, to apply in large datasets. Two remarks are in order. The first is that the lifecycle process reduces the possibilities that have to be kept. The second point is that we can reduce the size required for the possible instance triples by exploiting various properties that hold.

At first note that if two classes, say c_1 and c_2 , are possible classes for an instance o and it holds $c_1 \leq_{cl}^* c_2$ then all classes between, i.e. all c' such that $c_1 \leq_{cl}^* c' \leq_{cl}^* c_2$, are also possible classes for o (recall Lemma 2(1)). This allows devising storage representations based on *intervals* over the reflexive and transitive reduction of the \leq_{cl}^* relation. For example, if it holds $c_1 \leq_{cl} c_2 \leq_{cl} c_3 \leq_{cl} \dots \leq_{cl} c_{10}$ and all of them are possible classes for o then we can represent them by the interval $[c_1, c_{10}]$.

If for a given instance, o , there are several intervals having a common end (it is more probable to have a common right end), then we could save more space by adopting a more compact representation, e.g. the intervals $[c_1, c_{10}]$ and $[c_2, c_{10}]$ can be represented by $[\{c_1, c_2\}, c_{10}]$. Moreover, if we have a compact representation of the form $[\{c_1, \dots, c_k\}, c]$ and $\{c_1, \dots, c_k\}$ are the leaves of the hierarchy rooted at c then we can even omit their representation and adopt a more declarative method, like $[*, c]$, meaning that all subclasses of c are possible classes for o . In case where a class c is a possible class of an instance o but there are not subclasses of c or superclasses of c that are possible classes of o then the corresponding compact representation is $[c, c]$, indicated by a point interval $[c]$, for short.

We can follow an analogous approach for property instances. If two property instance triples, say $(o \text{ } pr_1 \text{ } o')$ and $(o \text{ } pr_2 \text{ } o')$, are possible property instance triples for two instances o and o' and it holds $pr_1 \leq^* pr_2$ then, for all properties pr' such that $pr_1 \leq_{pr}^* pr' \leq_{pr}^* pr_2$, $(o \text{ } pr' \text{ } o')$ is also a possible property instance triple for o and o' (recall Lemma 2(2)). This again allows compacted representations based on intervals over \leq_{pr} . For example, if it holds $pr_1 \leq_{pr} pr_2 \leq_{pr} pr_3 \leq_{pr} \dots \leq_{pr} pr_{10}$ and $(o \text{ } pr_i \text{ } o')$, for all $i \in \{1, \dots, 10\}$, is a possible property instance triple for o and o' , then we can represent these possible instance triples by the interval $[pr_1, pr_{10}]$. The same storage policy (intervals with the same right end), as in class instance triples above, can be followed in the case of property instance triples.

Figure 15(left) illustrates a data structure for the compact representation, which follows an *object-centered storage policy* (beneficial for the requirements of the life cycle management) regarding classes. As we can see, on the left there is a list of all the instances, lexicographically ordered. In the middle, there is a list of pointers each pointing to an interval on the right part, where all the intervals of the compact representation of possibilities are found. Note that every instance o on the left points to a consecutive list of pointers in the middle and thus, to a list of intervals on the right. For example, instances o_1 and o_2 , in Figure 15(left), point to the same intervals, i.e. the first interval $[*, A]$ and the second interval $[\{B, C, D\}, E]$. Figure 15(right) illustrates the corresponding data structure for properties. Note that two or more pairs of instances may have the same compact representation. For example, (o_1, o_2) and (o_1, o_3) , in Figure 15(right), point to the same two intervals, i.e. $[a]$ and $[\{b, c, d\}, e]$.

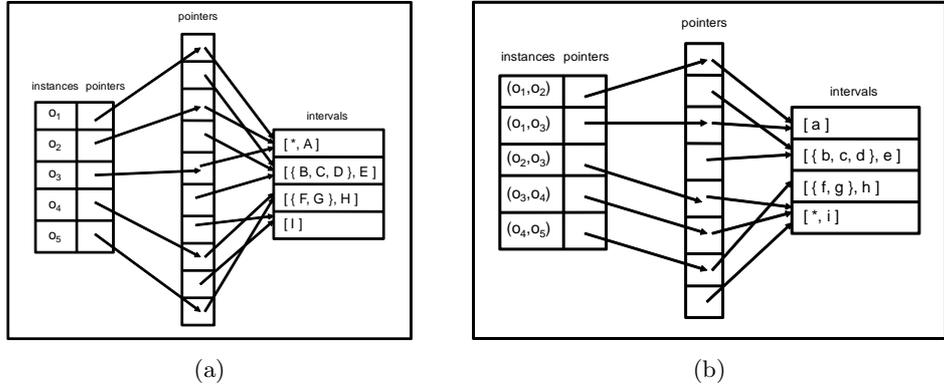


Fig. 15. Object-centered compact storage policy for possible class instance triples (left), and possible property instance triples (right)

We should clarify that the above compression technique is orthogonal and complementary to the various techniques which have been proposed (e.g. [17, 7]) for compressing RDF triples by using ids for each (subject-predicate-object) element of the triples. Specifically, the internal-based method exploits Lemmas 2(1) and 2(2), something that would not be possible with a generic approach for compression. However, the space required by the interval method could be further reduced by using ids for the triple elements.

The main advantage of adopting the compact representation is the space saving that we can achieve. However, if we execute Algorithms 1 and 2 using the compact representation of possibilities, we can see that the cost of looking up a specific possible instance triple in a compact representation is higher than in an explicit representation of all possibilities. Specifically, to decide whether a class instance triple (o type c) belongs to P_K requires locating the instance o and checking every interval which is pointed to by o .

Example 8. Consider Figure 15. Suppose that we want to find out if the class instance triple $(o_1$ type K) belongs to the compact representation of possibilities. Assume that $C \leq_{cl}^* K$ and $K \leq_{cl}^* E$ holds. At first, we find o_1 from the list of instances and then we check the intervals that o_1 points to, i.e. $[*, A]$ and $[\{B, C, D\}, E]$. Thus, we first

check if $K \leq_{cl}^* A$ holds. Since this is not true, we check if $K \leq_{cl}^* E$ and if $(B \leq_{cl}^* K$ or $C \leq_{cl}^* K$ or $D \leq_{cl}^* K)$. Since these conditions hold, $(o_1 \text{ type } K)$ is a possible class instance triple. \square

Let us now comment the time complexity of looking up an instance triple. Let $P_K^{explicit}$ denote an explicit representation of all possibilities ordered lexicographically, and $P_K^{compact}$ its compacted version. The time complexity of looking up an instance triple in $P_K^{explicit}$ (using binary search) is $O(\log_2(|P_K|))$. In $P_K^{compact}$ the cost depends on (a) the number of intervals associated with o , and (b) the cost of subsumption checking. The cost of the later is in $O(|\leq_{cl}|)$ if a plain graph structure is used (and no special index). However, given that most RDF triple stores currently use *labeling schemes* for encoding transitive subsumption relationships, subsumption checking can be performed in $O(1)$ (see [30, 14]).

10.3 Application Issues

The proposed technique can be applied to any ontology-based repository. Just indicatively, CIDOC CRM²² is one ontology (expressed in RDF/S) which is used by several ongoing EU projects, and it is curated (i.e. extended) by an authority (community) that is different from the various communities, curators of repositories, or simple users who keep creating instance descriptions with respect to that ontology. In practice, whenever a new version appears, the available instance descriptions are migrated to that version, and it is worth noting that this ontology has been revised at least 5 times the last two years (one recent version is described at [26]).

Collaborative Curation (or crowdcurating). The proposed technique can also be applied in a *crowdsourcing* setting, where several curators could interact in parallel after the migration of a dataset for curating its specificity. Specifically, each curator could carry out the steps in the area marked «interactive» in the right part of Figure 4. A server stores the extended KB and whenever a curator accepts or rejects one possibility, the changes are applied in the KB of the server in the context of a transaction. In this way, no inconsistency will arise.

[Applicability to Other Domains]

Although we confine ourselves to RDF/S, the results of our work can be applied to any object-oriented conceptual modeling approach that supports classes, inter-class associations, specialization/generalization hierarchies (among classes and among inter-class associations) and instantiation. Below we sketch an application in object-oriented software engineering.

Software Engineering. Our approach can be used also in object-oriented software engineering for upgrading libraries. Commonly, custom software relies on several libraries usually bundled in the form of jars. Most libraries evolve over time and their versions in most cases are backwards compatible. A new version of a library usually offers new subclasses of existing classes which provide improved/diversified functioning while respecting the ADT (Abstract Data Type) of the superclass. If the

²² CIDOC CRM (ISO 21127) is a core ontology describing the underlying semantics of data schemata and structures from all museum disciplines and archives (its RDF representation contains 78 classes and 250 properties from which 7 are literal-valued) (available from <http://www.cidoc-crm.org/>).

new version of a library is backwards compatible with the previous version, replacing the old version with the new version is enough for upgrading a software that depends on that library. However, this does not allow exploiting the new subclasses of the library: the user has to refer to textual descriptions and release notes in order to identify the new classes/features. An IDE (Integrated Development Environment) could adopt our approach for aiding the developer to upgrade his code. Specifically, it could be used for providing suggestions for refinements for those classes that are used by the code, i.e. those objects that instantiate library classes, and this can be done gradually by the lifecycle process.

We should also mention that RDF has been proposed as a data structure for software engineering, specifically for expressing software structure and dependencies²³.

As a brief and very rough example, and assuming a Java library, each Java class corresponds to an RDF class, each public instance variable of a class A with name v and type B corresponds to a property v with $domain(v) = A$ and $range(v) = B$, each object o that instantiates a class A corresponds to a class instance triple (o type A), etc.

10.4 Feedback from Related Communities and Users

We have already showed that the proposed method is feasible in terms of efficiency, and a tool that implements this approach is already available.

Regarding applicability, we have already communicated this work to the organizations participating to the NoE *APARSEN* (Alliance Permanent Access to the Records of Science in Europe, FP7, Proj. No 269977, 2011-2014), and it is referred to the deliverable [20]. The ultimate objective is to plan with the interested partners, a pilot application in an operational setting. We do not foresee any problem or unexpected issues, since this work actually proposed a process for making the curation easy and gradual, and the implemented tool provides the offered functionality in various different ways: starting from a particular instance, class or property, etc. However, its application in a real setting, will probably require adapting it to the processes and systems of the particular data provider/curator. Of course, the developed tool is not a polished end-user product, so the aspect of user interface usability and user-experience has not been considered in depth.

Our experience from a recent national conference of academic libraries²⁴, is that such organizations (libraries, national archives) are not aware of the problem, i.e. of the gradual loss of specificity that happens after each migration. The organizations that participated to that event understood directly the problem and they were interested in learning more about this. So, the key issue is to make such organizations aware of, and then to agree to undertake the responsibility for curating the quality of their metadata.

However, to check how intuitive the notions are, and how usable the current tool is, we selected one person with computer science background, and after making a slide presentation (that took 10 minutes), and a demo of the tool that took 5 minutes, we asked her to use two versions of the CIDOC ontology²⁵ (the first containing 810

²³ For example, there are tools that scan Java bytecode for method calls and create a description of the dependencies between classes and the package/archive encoded in RDF.

Other tools transform Maven POM (Project Object Model) files into RDF.

²⁴ <http://leo.hua.gr/palc2012/?q=en>

²⁵ <http://www.cidoc-crm.org/>

schema triples and the second 1004) and then to migrate and curate a corresponding set of instance triples consisting of 116 instance triples that document the artwork “Avis de Tempete” of Georges Aperghis. From this process, we understood that: (a) the concept and the process is easy to grasp (she understood it from the 10’ presentation), (b) in 5 minutes, she managed to get numerous possible instance triples something that would be impossible without the support of a system.

10.5 Assumptions and Associated Limitations

Single Domain and Range of Properties. In Def. 1 we consider properties which have a single domain and a single range. Note that this is a good practice in conceptual modeling, also in semantic networks, and obligatory in most (or even all) object-oriented systems. We should however note, that RDF/S allows having multiple classes defined as the domain and range of a property. For applying our approach, in such cases, it is suggested to define new classes so that all properties have a single domain (resp. single range) class. In this way, we also avoid some strange inferences. For example, suppose the following three RDF triples (`hasName domain Dog`), (`hasName domain Person`), and (`yannisURI hasName Yann`). From the above, it is inferred (recall the RDF/S derivation rule (v)) that (`yannisURI type Person`) and (`yannisURI, type Dog`)! We avoid such cases if properties have single domain and range classes.

The constraint (ii) of Def. 2, that states that the domain and range of a sub-property should be subclasses of (or the same with) the domain and range of its superproperty, is also classical in object-oriented and semantic modeling and makes our propositions and algorithms more clear and elegant.

Acyclicity. Regarding the acyclicity constraint (Def. 2), note that the construction of cyclic schemas is not a good practice in conceptual modeling. In our task, if a non cyclic schema becomes cyclic, then this could result to counterintuitive behavior (from the curation point of view), e.g. too much past negative triples can become certain, etc. However, note that the acyclicity constraint is not used in any of the proofs, therefore this constraint could be eliminated.

11 Related Work

There are several works on ontology evolution and versioning, for a recent overview see [12]. Below we describe in brief some of these works, although they are not directly related with our proposal. At first, we refer to works related to backwards compatible ontology evolution and then to works concerned with data validity and preservation after ontology evolution.

[*On Backwards Compatible Ontology Evolution*]

Klein et al. [15] propose a versioning mechanism for reducing the problems caused by ontology evolution. They argue that ontology versioning is necessary because changes to ontologies may cause incompatibilities, and drive to situations where the new (changed) ontology cannot be used in place of its previous version. They list a number of artifacts that may depend on an ontology, and thus can become

incompatible after ontological evolution, and data that conforms to an ontology is one of them. When an ontology is changed, data may get a different interpretation or may use terms that do not exist any more. The authors introduce various forms of compatibility and one of them is *backwards compatibility*. In the same direction, Xuan et al. [31] propose a model to deal with the problem of asynchronous ontology versions in the context of a materialized integration system, which is based on the principle of *ontological continuity*, which refers to the permanence of classes, properties, and subsumption. This principle is actually what we call *backwards compatibility*.

[*Ontology Evolution and Data Validity*]

Noy and Klein [19] present an informal discussion on the differences between ontology evolution and database schema evolution, and how structural changes in ontologies affect the preservation of their data instances. They focus on whether instance data can still be accessed through the changed ontology, and they classify the operation effects as information-preserving changes, translatable changes, and information-loss changes.

Now, Flouris et al. [10, 16] focus on the *effects* of a requested change operation, i.e. how the new knowledge base should be after a request for a change, and on its *side-effects* on both schema and instance data, i.e. certain additional actions executed to restore validity. They propose a general-purpose algorithm for determining the effects and side-effects of a requested elementary or complex change operation, and such works can be used to resolve the conflicts.

In addition, Qin and Alturi [22] focus on the validity issue of data instances during ontological evolution. They classify the changes to ontologies into two levels - structural and semantic. Structural changes include changes such as adding/removing classes or properties, modifying the range of a property, etc. Semantic changes mean changes to the semantics presented by the ontology, such as changes to the generalization/descriptiveness of a class or changes to the generalization/restrictiveness of a property. Semantic changes can be further classified into explicit and implicit changes. While explicit semantic changes can be detected based on the structural changes, implicit semantic changes need to be derived from explicit semantic changes by using implication analysis. They propose an algorithm for evaluating the structural validity of a data instance and then another algorithm for evaluating the semantic validity of a data instance.

In our case, since the target schema is fixed, the problem of validity concerns only the instance triples. Mainly there are two approaches to tackle it, and the choice is application specific. According to the first approach (which is actually what we assume in this paper), the set of instance triples should remain the same. In this case if a class c is deleted from S_K we consider that in K' , the instances that were explicitly classified in c , i.e. all $(o \text{ type } c) \in I_K$, are not affected. Thus, c remains in K' as an unconnected class. Now, if a property pr is deleted from S_K (and also the statements regarding its domain and range are deleted from the initial schema), we consider that in K' the property instance triples $(o \text{ } pr \text{ } o') \in I_K$ are not affected, thus, pr remains in K' as an unconnected property (its domain and range is the top class, i.e. *Resource*).

According to the second approach, we can change the set of instance triples so that K' contains schema elements *only* of $S_{K'}$. Obviously this requires updating I_K and it can be done based on the principles of *minimal change* (reminiscent of the

postulates of the AGM theory [1] regarding contraction). For the case of RDF/S this is actually quite straightforward (related work include [19, 16]).

In future work, it would be worth to investigate, in the non-backwards compatible schema evolution case, a possible combination of the “repair” of invalid instance descriptions proposed in [10, 16] with our proposal for computing possible instance descriptions. That would make sense only if instead of having as input the entire new schema we receive as input only the changes between the current and the new schema.

In conclusion, there are several works and approaches for dealing with the validity of data during migration, however there is no work for managing their specificity and quality while ontologies evolve.

Finally, we should mention that we do not focus on supporting *historical queries* regarding how concepts evolved over time, as in [24]. In addition, the way we compute possibilities does not presuppose an explicit modeling of how the individual concepts have evolved over time.

At last, [8] proposes a methodology for adding negative constraints to OWL ontologies, i.e. constraints that define expressions that should not be satisfiable. For example, an expression of “vegetarian pizza having meat as ingredient” should not be satisfiable. Instead, in the current paper, we do not focus on satisfiability. We focus on the uncertainty in the context of migrating sets of instance descriptions to new RDF schemas.

12 Concluding Remarks

The rapid evolution of ontologies requires principles, techniques, and tools for managing the quality of the migrated descriptions, as well as flexible interactive methods for managing the incurred uncertainty. To the best of our knowledge this is the first work that exploits ontology schema evolution for managing the specificity of instance descriptions.

In this work, we use *possible instance triples* for expressing (actually bounding) the uncertainty regarding the specificity of a description caused by its migration to a new schema. We introduce the notion of TFP-partition (True-False-Possible partition) of a KB for capturing various application-specific assumptions about the specificity of its descriptions. Subsequently, we formalize the *migration* to a new schema as a *transition* between two TFP-partitions, which should be governed by few (two for the case of backwards compatible, and three for the case of non-backwards compatible schema evolution) *postulates* which are general (i.e. RDF/S independent). Based on these postulates, we derive rules for carrying out a transition for the case of RDF/S. To reduce the space requirements, we show how transitions can be defined without having to keep any negative information (i.e. the “F” part of a TFP-partition), instead only the certain and the possible part of the KB has to be kept, reaching to what we call extended KB (*eKB*). To further compress the possible part of the eKB, and making it feasible for large data sets, we propose a compact (interval-based) representation (exploiting intrinsic properties of TFP-partitions).

Since the ultimate objective is not just the identification of possibilities, but to aid making the instance descriptions as specific as possible, we propose a *specificity lifecycle management process* that *ranks* the possible instance triples, prompts to the

user a subset of the possible instance triples and we show how the extended KB should be *updated* when the user *accepts* or *rejects* some of them.

To investigate the feasibility of our approach, we designed and developed a prototype system, and we applied our approach on real and synthetic datasets. The measurements demonstrate the feasibility of the approach.

There are several issues for future research. One interesting direction is to generalize our approach to the *XSD*²⁶-typed literal values [21] of property instance triples. Such extension would allow reasoning about the *accuracy* of the migrated descriptions over linearly ordered domains (e.g. as consequence of migrating 32-bit floating numbers to a 64-bit representation).

Acknowledgements Work done in the context of NoE *APARSEN* (Alliance Permanent Access to the Records of Science in Europe, FP7, Proj. No 269977, 2011-2014). Many thanks to Kristina Lantzaki that participated to the user study.

References

1. C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
2. C. Bizer, T. Heath, and T. Berners-Lee. Linked data-the story so far. *International Journal on Semantic Web and Information Systems*, 5(3):1–22, 2009.
3. D. Brickley and R. V. Guha. RDF Vocabulary Description Language 1.0: RDF Schema, W3C Recommendation, February 2004. Available at <http://www.w3.org/TR/rdf-schema/>.
4. Peter Buneman, James Cheney, Wang Chiew Tan, and Stijn Vansummeren. Curated Databases. In *27th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS-2008)*, pages 1–12, 2008.
5. V. Christophides, D. Plexousakis, M. Scholl, and S. Tourtounis. On Labeling Schemes for the Semantic Web. In *12th Intern. World Wide Web Conference (WWW-2003)*, pages 544–555, 2003.
6. M. Dalal. Investigations into a Theory of Knowledge Base Revision: Preliminary Report. In *7th National Conference on Artificial Intelligence (AAAI-1988)*, pages 475–479, 1988.
7. J. D. Fernández, M. A. Martínez-Prieto, and C. Gutierrez. Compact Representation of Large RDF Data Sets for Publishing and Exchange. In *9th International Semantic Web Conference (ISWC-2010)*, pages 193–208, 2010.
8. Sébastien Ferré and Sebastian Rudolph. *Advocatus Diaboli*—Exploratory Enrichment of Ontologies with Negative Constraints. In *18th International Conference on Knowledge Engineering and Knowledge Management (EKAW-2012)*, pages 42–56, 2012.
9. G. Flouris. *On Belief Change and Ontology Evolution*. PhD thesis, Computer Science Department, University of Crete, Greece, 2006.
10. Giorgos Flouris, George Konstantinidis, Grigoris Antoniou, and Vassilis Christophides. Formal Foundations for RDF/S KB Evolution. *International Journal on Knowledge and Information Systems (KAIS)*, 35(1):153–191, 2013.
11. C. Gutierrez, C. Hurtado, and A. Mendelzon. Foundations of Semantic Web Databases. In *23rd ACM Symposium on Principles of Database Systems (PODS-2004)*, pages 95–106, 2004.
12. Michael Hartung, James Terwilliger, and Erhard Rahm. Recent Advances in Schema and Ontology Evolution. *Schema Matching and Mapping*, 2011.

²⁶ XML Schema Definition

13. P. Hayes. RDF Semantics, W3C Recommendation, February 2004. Available at <http://www.w3.org/TR/rdf-mt/>.
14. Ruoming Jin, Yang Xiang, Ning Ruan, and Haixun Wang. Efficiently answering reachability queries on very large directed graphs. In *ACM SIGMOD International Conference on Management of Data (SIGMOD-2008)*, pages 595–608, 2008.
15. Michel C. A. Klein and Dieter Fensel. Ontology Versioning on the Semantic Web. In *First Semantic Web Working Symposium (SWWS-2001)*, pages 75–91, 2001.
16. G. Konstantinidis, G. Flouris, G. Antoniou, and V. Christophides. A Formal Approach for RDF/S Ontology Evolution. In *18th European Conference on Artificial Intelligence (ECAI-2008)*, pages 70–74, Patras, Greece, July 2008.
17. Thomas Neumann and Gerhard Weikum. RDF-3X: a RISC-style Engine for RDF. *Proceedings of the VLDB Endowment (PVLDB)*, 1(1):647–659, 2008.
18. Thomas Neumann and Gerhard Weikum. x-RDF-3X: Fast Querying, High Update Rates, and Consistency for RDF Databases. *Proceedings of the VLDB Endowment (PVLDB)*, 3(1):256–263, 2010.
19. Natalya Fridman Noy and Michel C. A. Klein. Ontology Evolution: Not the Same as Schema Evolution. *Knowledge and Information Systems*, 6(4):428–440, 2004.
20. APARSEN: Network of Excellence (FP7 No 269977) 2011–2014. Deliverable D26.1: Report and Strategy on Annotation, Reputation and Data Quality, 2012. <http://www.alliancepermanentaccess.org/wp-content/uploads/downloads/2012/12/APARSEN-REP-D26\1-01-1\0.pdf>.
21. David Peterson, Shudi (Sandy) Gao, Ashok Malhotra, C. M. Sperberg-McQueen, and Henry S. Thompson. W3C XML Schema Definition Language (XSD) 1.1 Part 2: Datatypes, W3C Working Draft 3, December 2009. Available at <http://www.w3.org/TR/xmlschema11-2/>.
22. Li Qin and Vijayalakshmi Atluri. Evaluating the validity of data instances against ontology evolution over the Semantic Web. *Information & Software Technology*, 51(1):83–97, 2009.
23. Roy Rada, Hafehd Mili, Ellen Bicknell, and Maria Blettner. Development and Application of a Metric on Semantic Nets. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(1):17–30, 1989.
24. Flavio Rizzolo, Yannis Velegarakis, John Mylopoulos, and Sjarhei Bykau. Modeling Concept Evolution: A Historical Perspective. In *28th International Conference on Conceptual Modeling (ER-2009)*, pages 331–345, 2009.
25. Satya S. Sahoo, Wolfgang Halb, Sebastian Hellmann, Kingsley Idehen, Ted Thibodeau Jr, Soren Auer, Juan Sequeda, and Ahmed Ezzat. A Survey of Current Approaches for Mapping of Relational Databases to RDF, 2009. Report by the W3C RDB2RDF Incubator Group. Available at http://www.w3.org/2005/Incubator/rdb2rdf/RDB2RDF_SurveyReport.pdf.
26. M. Theodoridou, Y. Tzitzikas, M. Doerr, Y. Marketakis, and V. Melessanakis. Modeling and querying provenance by extending CIDOC CRM. *Distributed and Parallel Databases*, 27(2):169–210, 2010.
27. Y. Theoharis, G. Georgakopoulos, and V. Christophides. On the Synthetic Generation of Semantic Web Schemas. In *Joint ODBIS & SWDB Workshop on Semantic Web, Ontologies, and Databases. Collocated with VLDB2007*, pages 98–116, 2007.
28. Yannis Theoharis, Yannis Tzitzikas, Dimitris Kotzinos, and Vassilis Christophides. On Graph Features of Semantic Web Schemas. *IEEE Trans. Knowl. Data Eng.*, 20(5):692–702, 2008.
29. Yannis Tzitzikas, Mary Kampouraki, and Anastasia Analyti. Curating the Specificity of Metadata while World Models Evolve. In *9th annual iPres conference on digital preservation (iPres-2012)*, 2012. Available at http://www.ics.forth.gr/~analyti/Local_Papers/CuratingSpecificity_iPres_CRC_pv.pdf.
30. Haixun Wang, Hao He, Jun Yang, Philip S. Yu, and Jeffrey Xu Yu. Dual Labeling: Answering Graph Reachability Queries in Constant Time. In *22nd International Conference on Data Engineering (ICDE-2006)*, 2006.

31. D. N. Xuan, L. Bellatreche, and G. Pierra. A Versioning Management Model for Ontology-Based Data Warehouses. In *8th International Conference on Data Warehousing and Knowledge Discovery (DaWaK 2006)*, pages 195–206, 2006.
32. D. Zeginis, Y. Tzitzikas, and V. Christophides. On the Foundations of Computing Deltas Between RDF Models. In *6th International Semantic Web Conference (ISWC-07)*, pages 637–651. Springer, 2007.
33. D. Zeginis, Y. Tzitzikas, and V. Christophides. On computing deltas of RDF/S knowledge bases. *ACM Transactions on the Web (TWEB)*, 5(3):14, 2011.

A List of Symbols

List of Symbols	
Symbol	Description
$\mathcal{C}(K)$	The <i>closure</i> of a KB K
C_K	The set of classes of $\mathcal{C}(K)$
Pr_K	The set of properties of $\mathcal{C}(K)$
\leq_{cl}^*	The <i>subClassOf</i> relation between classes in C_K
\leq_{pr}^*	The <i>subPropertyOf</i> relation between properties in Pr_K
Res_K	The resources of a KB K
$Inst_K$	The instances of a KB K
$inst_K(c)$	The instances of a class c of a KB K
B_K	The set of <i>cartesian instance triples</i> of a KB K
S_K	The set of <i>schema triples</i> of a KB K
I_K	The set of <i>instance triples</i> of a KB K
$\mathcal{C}_i(K)$	The instance triples of the closure of a KB K
$valid(o, pr, o', K)$	A <i>valid</i> property instance triple of a KB K
$Invalid(K)$	The set of <i>invalid</i> property instance triples of a KB K
$SubTriples((o \text{ type } c))$	The set of <i>subtriples</i> of a class instance triple (<i>o type c</i>) of a KB K
$SubTriples((o \text{ pr } o'))$	The set of <i>subtriples</i> of a property instance triple (<i>o pr o'</i>) of a KB K
$SubTriples(A)$	The set of <i>subtriples</i> of a set of instance triples A of a KB K
F_K	The set of negative(false) instance triples of a KB K
P_K	The set of possible instance triples of a KB K
P_{K_Add}	The set of added possible instance triples in P_K
P_{K_Del}	The set of deleted possible instance triples from P_K
$posTriples^{cl}(o)$	The set of possible class instance triples of an instance o
$posTriples^{spr}(o)$	The set of possible property instance triples of an instance o as subject
$posTriples^{opr}(o)$	The set of possible property instance triples of an instance o as object
$SuperTriples((o \text{ type } c))$	The set of <i>supertriples</i> of a class instance triple (<i>o type c</i>) of a KB K
$SuperTriples((o \text{ pr } o'))$	The set of <i>supertriples</i> of a property instance triple (<i>o pr o'</i>) of a KB K
$SuperTriples(A)$	The set of <i>supertriples</i> of a set of instance triples A of a KB K
K^{up}	The updated certain part of an <i>eKB</i>
P^{up}	The updated possible part of an <i>eKB</i>
$dist_{cl}(c \rightarrow c')$	The length of the shortest path from a class c to a class c'
$dist_{Class}(o, c)$	The shortest distance of c from one of the certain classes of o
$dist_{pr}(pr \rightarrow pr')$	The length of the shortest path from a property pr to a property pr'
$dist_{Property}(o, pr, o')$	The shortest distance of pr from one of the certain properties of (o, o')
$Valid(K)$	The set of <i>valid</i> property instance triples of a KB K
P_K^{comp}	The set of <i>composite possibilities</i>
P_K^{ext}	The set of <i>extended possibilities</i> including <i>atomic</i> and <i>composite possibilities</i>
$Cl(o)$	The set of all certain classes of an instance o
$posCl(o)$	The set of all possible classes of an instance o

Table 12. Symbols and Description

B Example of Applying the Algorithm for Backwards Compatible Migration

Example 9. Consider the example of Figure 5 and suppose that $P_K = \{(\text{Fiat_1 type Vehicle}), (\text{Bob uses BMW_1}), (\text{Alice works at FORTH})\}$. Executing Algorithm 1, step by step, we have that:

- (1) $K' = (S_{K'}, I_K)$;
 - (2) $P_{K_Add} = \emptyset$;
 - (3) $P_{K_Del} = \emptyset$;
 - (4) $NC = \{\text{Van, Jeep, Adult, Institute, University}\}$;
 - (6) $P_{K_Add} = P_{K_Add} \cup \{(\text{Fiat_1 type Van}), (\text{BMW_1 type Van}), (\text{Fiat_1 type Jeep}), (\text{BMW_1 type Jeep}), (\text{Bob type Adult}), (\text{Alice type Adult}), (\text{Computer Science Department type Institute}), (\text{FORTH type Institute}), (\text{Computer Science Department type University}), (\text{FORTH type University})\}$;
 - (8) $P_{K_Del} = P_{K_Del} \cup \{\emptyset\}$;
 - (9) $NP = \{\text{paid from}\}$;
 - (11) $P_{K_Add} = P_{K_Add} \cup \{(\text{Alice paid from Computer Science Department}), (\text{Bob paid from FORTH})\}$;
 - (13) $P_{K_Del} = P_{K_Del} \cup \{(\text{Alice works at FORTH})\}$;
- Note that, $(\text{Alice related to FORTH}) \notin (P_K \cup C_i(K'))$ and $\text{works at } \leq_{cl}$ related to holds in K' . So, we have to move $(\text{Alice works at FORTH})$ from P_K to $F_{K'}$ (due to Rule R2).
- (14) $P_{K_Del} = P_{K_Del} \cup \{(\text{Fiat_1 type Vehicle}), (\text{Bob uses BMW_1})\}$;
- Note that P_{K_Del} is updated by those triples that belong to P_K and now belong to $C_i(K')$. So, we have to remove them from P_K .
- (15) $P_{K'} = P_K \setminus P_{K_Del} = \{(\text{Fiat_1 type Vehicle}), (\text{Bob uses BMW_1}), (\text{Alice works at FORTH})\} \setminus \{(\text{Fiat_1 type Vehicle}), (\text{Bob uses BMW_1}), (\text{Alice works at FORTH})\} = \emptyset$;
 - (16) $P_{K'} = \{(\text{Fiat_1 type Van}), (\text{BMW_1 type Van}), (\text{Fiat_1 type Jeep}), (\text{BMW_1 type Jeep}), (\text{Bob type Adult}), (\text{Alice type Adult}), (\text{Computer Science Department type Institute}), (\text{FORTH type Institute}), (\text{Computer Science Department type University}), (\text{FORTH type University}), (\text{Alice paid from Computer Science Department}), (\text{Bob paid from FORTH})\}$;
 - (17) Return $P_{K'}$;

In order to explain line 8 in part B of Algorithm 1, consider Figure 5. Suppose that we have another version $S_{K''} = S_{K'} \cup \{(\text{Van } \leq_{cl} \text{LoadCarrying Vehicle})\}$, where we have a new specialization relationship, i.e. $\text{Van } \leq_{cl} \text{LoadCarrying Vehicle}$. Then, according to line 8 of Algorithm 1, we have that $P_{K_Del} = P_{K_Del} \cup \{(\text{Fiat_1 type Van})\}$. This is because it holds that $(\text{Fiat_1 type LoadCarrying Vehicle}) \notin P_{K'} \cup C_i(K'')$ and $(\text{Van } \leq_{cl} \text{LoadCarrying Vehicle}) \in S_{K''}$. So, we have to move (Fiat_1 type Van) from $P_{K'}$ to $F_{K''}$ (due to Rule R1). \square

Note that if $(o \text{ type } c) \in C_i(K')$ and $(o \text{ type } c') \notin C_i(K') \cup P_{K'}$, for all $c' \leq_{cl}^* c$ and $c' \neq c$, then the MSA property holds for the class instance triple $(o \text{ type } c)$. Similarly, if $(o \text{ pr } o') \in C_i(K')$ and $(o \text{ pr } o') \notin C_i(K') \cup P_{K'}$, for all $pr' \leq_{pr}^* pr$ and $pr' \neq pr$, then the MSA property holds for the property instance triple $(o \text{ pr } o')$.

Curating the Specificity of Ontological Descriptions under Ontology Evolution: APPENDIX

A Proofs

In this Appendix, we provide the proof of the Lemmas and Propositions, presented in the main paper.

Lemma 1

1. If $(o \text{ type } c_2) \in P_K$ and $c_2 \leq_{cl}^* c_1$ then $(o \text{ type } c_1) \in (P_K \cup \mathcal{C}_i(K))$.
2. If $(o \text{ pr}_2 \ o') \in P_K$ and $pr_2 \leq_{pr}^* pr_1$ then $(o \text{ pr}_1 \ o') \in (P_K \cup \mathcal{C}_i(K))$.

Proof:

1) Let $(o \text{ type } c_2) \in P_K$ and $c_2 \leq_{cl}^* c_1$. We will show that $(o \text{ type } c_1) \in (\mathcal{C}_i(K) \cup P_K)$. Certainly, $o \in Inst_K \cap URI$. Assume that $(o \text{ type } c_1) \notin (\mathcal{C}_i(K) \cup P_K)$. Since $o \in Inst_K \cap URI$ and $c_1 \in C_K$, it follows that $(o \text{ type } c_1) \in B_K$. Therefore, it follows from the definition of TFP-partition (Def. 5) that $(o \text{ type } c_1) \in F_K$. It follows from (ii) of Def. 5 that $(o \text{ type } c_2) \in F_K$. Thus, $(o \text{ type } c_2) \notin P_K$, which is impossible. Therefore, $(o \text{ type } c_1) \in (\mathcal{C}_i(K) \cup P_K)$.

2) Let $(o \text{ pr}_2 \ o') \in P_K$ and $pr_2 \leq_{pr}^* pr_1$. We will show that $(o \text{ pr}_1 \ o') \in (\mathcal{C}_i(K) \cup P_K)$. Since $(o \text{ pr}_2 \ o') \in P_{K'}$, it holds that $o \in Inst_K \cap URI$ and $o' \in Inst_K$. Therefore, $(o \text{ pr}_1 \ o') \in B_K$. Assume that $(o \text{ pr}_1 \ o') \notin (\mathcal{C}_i(K) \cup P_K)$. Therefore, it follows from the definition of TFP-partition (Def. 5) that $(o \text{ pr}_1 \ o') \in F_K$. It follows from (ii) of Def. 5 that $(o \text{ pr}_2 \ o') \in F_K$. Thus, $(o \text{ pr}_2 \ o') \notin P_K$, which is impossible. Therefore, $(o \text{ pr}_1 \ o') \in (\mathcal{C}_i(K) \cup P_K)$. \square

Lemma 2

1. If $c_1 \leq_{cl}^* c_2 \leq_{cl}^* c_3$ and $(o \text{ type } c_1), (o \text{ type } c_3) \in P_K$ then $(o \text{ type } c_2) \in P_K$.
2. If $pr_1 \leq_{pr}^* pr_2 \leq_{pr}^* pr_3$ and $(o \text{ pr}_1 \ o'), (o \text{ pr}_3 \ o') \in P_K$ then $(o \text{ pr}_2 \ o') \in P_K$.

Proof:

1) Let $c_1 \leq_{cl}^* c_2 \leq_{cl}^* c_3$ and $(o \text{ type } c_1), (o \text{ type } c_3) \in P_K$, we will show that $(o \text{ type } c_2) \in P_K$. Note that $o \in Inst_K \cap URI$ and $c_2 \in C_K$. Thus, $(o \text{ type } c_2) \in B_K$. Assume that $(o \text{ type } c_2) \notin P_K$ then $(o \text{ type } c_2) \in F_K$ or $(o \text{ type } c_2) \in \mathcal{C}_i(K)$. If $(o \text{ type } c_2) \in F_K$ then it follows from (ii) of Def. 5 that $(o \text{ type } c_1) \in F_K$, which is impossible. If $(o \text{ type } c_2) \in \mathcal{C}_i(K)$ then it follows from the RDF/S semantics that $(o \text{ type } c_3) \in \mathcal{C}_i(K)$, which is also impossible. Thus, $(o \text{ type } c_2) \in P_K$.

2) Let $pr_1 \leq_{pr}^* pr_2 \leq_{pr}^* pr_3$ and $(o \text{ pr}_1 \ o'), (o \text{ pr}_3 \ o') \in P_K$ we will show that $(o \text{ pr}_2 \ o') \in P_K$. Note that $o \in Inst_K \cap URI$, $o' \in Inst_K$, and $pr_2 \in Pr_K$. Thus, $(o \text{ pr}_2 \ o') \in B_K$. Assume that $(o \text{ pr}_2 \ o') \notin P_K$. Then, $(o \text{ pr}_2 \ o') \in F_K$ or $(o \text{ pr}_2 \ o') \in \mathcal{C}_i(K)$. If $(o \text{ pr}_2 \ o') \in F_K$ then it follows from (ii) of Def. 5 that $(o \text{ pr}_1 \ o') \in F_K$, which is impossible. If $(o \text{ pr}_2 \ o') \in \mathcal{C}_i(K)$ then it follows from the RDF/S semantics that $(o \text{ pr}_3 \ o') \in \mathcal{C}_i(K)$. Therefore, $(o \text{ pr}_2 \ o') \in P_K$. \square

The following Proposition is used in several of our proofs.

Prop. A1. When a KB $K = (S_K, I_K)$ evolves to a new KB $K' = (S_{K'}, I_{K'})$, where $I_K = I_{K'}$, it holds that $Inst_K = Inst_{K'}$. \square

Proof:

It holds that $Inst_K = Res_K \setminus (C_K \cup Pr_K)$. Since $K = (S_K, I_K)$ and from the way that instance triples and schema triples are defined, it follows that $Inst_K = \{o \mid (o \text{ type } c) \in I_K\} \cup \{o, o' \mid (o \text{ pr } o') \in I_K\}$, since $I_K = I_{K'}$. It follows that $Inst_K = Inst_{K'}$. \square

Prop. 1 Let S and S' be two sets of schema triples. It holds that: $S \sqsubseteq S'$ iff $\Delta_d(S \rightarrow S')$ contains only add operations.

Proof:

\Rightarrow) $\Delta_d(S \rightarrow S') = \{Add(t) \mid t \in S' \setminus \mathcal{C}(S)\} \cup \{Del(t) \mid t \in S \setminus \mathcal{C}(S')\}$. We have to prove that $S \setminus \mathcal{C}(S') = \emptyset$. It is known from the *closure definition* that $S \subseteq \mathcal{C}(S)$. $S \sqsubseteq S'$ means that $\mathcal{C}(S) \subseteq \mathcal{C}(S')$, so by transitivity we get that $S \subseteq \mathcal{C}(S')$.

\Leftarrow) $S \setminus \mathcal{C}(S') = \emptyset$, so $S \subseteq \mathcal{C}(S')$. Thus, $\mathcal{C}(S) \subseteq \mathcal{C}(S')$. Therefore, $\mathcal{C}(S) \subseteq \mathcal{C}(S')$. By Definition 6, we get that $S \sqsubseteq S'$. \square

Prop. 2 If $S_K \sqsubseteq S_{K'}$ then $B_K \subseteq B_{K'}$.

Proof: Let $b \in B_K$. If b is a class instance triple of the form (*o type c*) then $o \in Inst_K \cap URI$ and $c \in C_K$. Since $Inst_K = Inst_{K'}$, it follows that $o \in Inst_{K'} \cap URI$. Additionally, since $S_K \sqsubseteq S_{K'}$, it holds that $c \in C_{K'}$. Thus, $b \in B_{K'}$. If b is a property instance triple of the form (*o pr o'*) then $o \in Inst_K \cap URI$, $o' \in Inst_K$, and $pr \in Pr_K$. Thus, $o \in Inst_{K'} \cap URI$, $o' \in Inst_{K'}$ and, since $S_K \sqsubseteq S_{K'}$, it holds that $pr \in Pr_{K'}$. Thus, $b \in B_{K'}$. \square

Prop. 3 In the context of a *consistent transition* $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, it follows that: $F_K \cap P_{K'} = \emptyset$ (i.e. *II2*) iff $(F_K \setminus \mathcal{C}_i(K')) \subseteq F_{K'}$.

Proof:

\Rightarrow) Assume that $F_K \cap P_{K'} = \emptyset$. Let $t \in F_K \setminus \mathcal{C}_i(K')$ and assume that $t \notin F_{K'}$. Since $S_K \sqsubseteq S_{K'}$, it holds that $t \in B_{K'}$. Therefore, based on the definition of TFP-partition (Def. 5) and the fact $t \notin \mathcal{C}_i(K')$, it holds that $t \in P_{K'}$. However, in this case, it holds that $F_K \cap P_{K'} \neq \emptyset$, which is impossible. Therefore, $t \in F_{K'}$. Thus, $(F_K \setminus \mathcal{C}_i(K')) \subseteq F_{K'}$.

\Leftarrow) Assume that $(F_K \setminus \mathcal{C}_i(K')) \subseteq F_{K'}$. Further, assume that it exists $t \in F_K \cap P_{K'}$. Then, based on the definition of TFP-partition (Def. 5), it holds that $t \notin \mathcal{C}_i(K')$. Thus, $t \in F_K \setminus \mathcal{C}_i(K')$. Since $(F_K \setminus \mathcal{C}_i(K')) \subseteq F_{K'}$, it follows that $t \in F_{K'}$. However, this is impossible, since $t \in P_{K'}$. Thus, $F_K \cap P_{K'} = \emptyset$. \square

Prop. 4 The negative part of the target TFP-partition of a consistent transition, always satisfies the rules of Def. 8.

Proof:

If the conditions of the “if part” of rules *R1* and *R2* are satisfied, then according to Prop. 3, (*o type c*) and (*o pr o'*) should be placed in $F_{K'}$. Thus, if rules *R1* and *R2* are not satisfied then $F_{K'}$ cannot be a lower set (and thus cannot be part of a TFP-partition).

If Rule *R3* is not satisfied, then $F_{K'}$ cannot be part of a TFP-partition (since it would not contain the invalid property instance triples). \square

Prop. 5 The derivation of $F_{K'}$ using the rules of Def. 8, and the derivation of $P_{K'}$ by Def. 9, yields a three-fold partition that is a TFP-partition and $(\mathcal{C}_i(K), F_K, P_K) \rightsquigarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$ is a *consistent transition*, i.e. it respects postulates *II1* and *II2*. We will denote this particular transition by: $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$.

Proof: The sets $\mathcal{C}_i(K')$, $F_{K'}$, and $P_{K'}$ are pairwise disjoint by construction. $F_{K'}$ is lower set by construction, and contains all invalid property instance triples of $B_{K'}$. Therefore, $(\mathcal{C}_i(K'), F_{K'}, P_{K'})$ is a TFP-partition.

Postulate *II1* is satisfied by construction. What is left to prove is that *II2* is satisfied (past negative information cannot become possible), i.e. that $F_K \cap P_{K'} = \emptyset$. Suppose this is not true, i.e. suppose there exist a t such that $t \in F_K \cap P_{K'}$. Then, $t \in F_K$ and $t \notin \mathcal{C}_i(K')$. Thus, due to Rule *R1* and Rule *R2* of Prop. 4, t (and its subtriples) would belong to $F_{K'}$. Thus, t cannot belong to $P_{K'}$ (since $P_{K'}$ as defined by Def. 9 excludes all elements of $F_{K'}$). \square

Prop. 6 Suppose we are in the context of $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$. For a new class $c' \in C_{K'} \setminus C_K$, it holds that: $(o \text{ type } c') \in P_{K'}$ iff:

- (i) $o \in \text{Inst}_K \cap \text{URI}$, and
- (ii) for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K)$, and
- (iii) $(o \text{ type } c') \notin \mathcal{C}_i(K')$.

Proof:

\Rightarrow) Let $(o \text{ type } c') \in P_{K'}$. Certainly, $o \in \text{Inst}_K \cap \text{URI}$. We will show that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K)$ and $(o \text{ type } c') \notin \mathcal{C}_i(K')$. Assume that it exists $c \in C_K$ s.t. $(c' \leq_{cl}^* c$ and $(o \text{ type } c) \notin (\mathcal{C}_i(K') \cup P_K))$ or $(o \text{ type } c') \in \mathcal{C}_i(K')$. If it exists $c \in C_K$ s.t. $(c' \leq_{cl}^* c$ and $(o \text{ type } c) \notin (\mathcal{C}_i(K') \cup P_K))$ then since $\mathcal{C}_i(K) \subseteq \mathcal{C}_i(K')$, it follows that $(o \text{ type } c) \notin (\mathcal{C}_i(K) \cup P_K)$. Since $o \in \text{Inst}_K \cap \text{URI}$ and $c \in C_K$, it follows that $(o \text{ type } c) \in B_K$. Therefore, it follows from Def. 5 that $(o \text{ type } c) \in F_K$. It follows from Rule R1 of Prop. 4 that $(o \text{ type } c') \in F_{K'}$. Thus, $(o \text{ type } c') \notin P_{K'}$, which is impossible. If $(o \text{ type } c') \in \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ type } c') \notin P_{K'}$, which is impossible.

\Leftarrow) Assume that $o \in \text{Inst}_K \cap \text{URI}$ and that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K)$. Additionally, assume that it holds $(o \text{ type } c') \notin \mathcal{C}_i(K')$. We will show that $(o \text{ type } c') \in P_{K'}$. It follows from Def. 5 that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in B_K \setminus F_K$ or $(o \text{ type } c) \in \mathcal{C}_i(K')$. It follows that $(o \text{ type } c') \notin F_{K'}$ (note Rule R1 of Prop. 4 does not apply). Since $o \in \text{Inst}_K \cap \text{URI}$, $c' \in C_{K'}$, and $\text{Inst}_K = \text{Inst}_{K'}$, it follows that $(o \text{ type } c') \in B_{K'}$. Now, since $(o \text{ type } c') \notin \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ type } c') \in P_{K'}$. \square

Prop. 7 Suppose we are in the context of $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$. For a new property $pr' \in Pr_{K'} \setminus Pr_K$, it holds that: $(o \text{ pr } o') \in P_{K'}$ iff:

- (i) $o \in \text{URI}$ and $\text{valid}(o, pr', o', K')$, and
- (ii) for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that $(o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K)$, and
- (iii) $(o \text{ pr } o') \notin \mathcal{C}_i(K')$.

Proof:

\Rightarrow) Let $(o \text{ pr } o') \in P_{K'}$. Since $(o \text{ pr } o') \in P_{K'}$, it follows that $\text{valid}(o, pr', o', K')$. Additionally, we have $o \in \text{URI}$. We will show that for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that (i) $(o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K)$ and (ii) $(o \text{ pr } o') \notin \mathcal{C}_i(K')$. Assume that it exists $pr \in Pr_K$ s.t. $(pr' \leq_{pr}^* pr$ and $(o \text{ pr } o') \notin (\mathcal{C}_i(K') \cup P_K))$, or (ii) $(o \text{ pr } o') \in \mathcal{C}_i(K')$. If it exists $pr \in Pr_K$ s.t. $(pr' \leq_{pr}^* pr$ and $(o \text{ pr } o') \notin (\mathcal{C}_i(K') \cup P_K))$ then, since $\mathcal{C}_i(K) \subseteq \mathcal{C}_i(K')$, it follows that $(pr' \leq_{pr}^* pr$ and $(o \text{ pr } o') \notin (\mathcal{C}_i(K) \cup P_K))$. Obviously, it holds that $(o \text{ pr } o') \in B_K$. Then, it follows from Def. 5 that $(o \text{ pr } o') \in F_K$. Therefore, it follows from Rule R2 of Prop. 4 that $(o \text{ pr } o') \in F_{K'}$. Thus, $(o \text{ pr } o') \notin P_{K'}$, which is impossible. If $(o \text{ pr } o') \in \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ pr } o') \notin P_{K'}$, which is impossible.

\Leftarrow) Assume that (i) $o \in \text{URI}$ and $\text{valid}(o, pr', o', K')$, (ii) for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that $(o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K)$, and (iii) $(o \text{ pr } o') \notin \mathcal{C}_i(K')$. We will show that $(o \text{ pr } o') \in P_{K'}$. It follows from Def. 5 that for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that $(o \text{ pr } o') \in B_K \setminus F_K$ or $(o \text{ pr } o') \in \mathcal{C}_i(K')$. It follows that $(o \text{ pr } o') \notin F_{K'}$ (note that Rule R2 and Rule R3 of Prop. 4 do not apply). It holds that $o \in \text{Inst}_{K'} \cap \text{URI}$, $o' \in \text{Inst}_{K'}$, $pr \in Pr_{K'}$. Thus, $(o \text{ pr } o') \in B_{K'}$. Since $(o \text{ pr } o') \notin \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ pr } o') \in P_{K'}$. \square

Prop. 8 If $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, then $P_{K'} = \text{Produce_Possibilities}(K, P_K, S_{K'})$.

Proof: Initially, $P_{K_Add} = \emptyset$ and $P_{K_Del} = \emptyset$. For each new class $c' \in NC$, Algorithm $\text{Produce_Possibilities}(K, P_K, S_{K'})$ inserts to P_{K_Add} exactly the class instance triples in-

icated by Prop. 6. For each old class $c_1 \in C_K$, if it exists $c_2 \in C_K$ s.t. $c_1 \leq_{cl}^* c_2$ and (*o type* c_2) $\notin (P_K \cup C_i(K'))$ then all class instance triples (*o type* c_1) $\in P_K$ are added to P_{K_Del} . Since $S_K \sqsubseteq S_{K'}$, it holds that (*o type* c_2) $\notin C_i(K)$. Thus, (*o type* c_2) $\in B_K \setminus (P_K \cup C_i(K)) = F_K$ and (*o type* c_2) $\notin C_i(K')$. Therefore, due to Rule R1 of Prop. 4, it holds that (*o type* c_1) $\in F_{K'}$.

Similarly, for each new property $pr' \in NP$, Algorithm *Produce_Possibilities*($K, P_K, S_{K'}$), inserts to P_{K_Add} exactly the property instance triples indicated by Prop. 7. For each old property $pr_1 \in Pr_K$, if it exists $pr_2 \in Pr_K$ s.t. $pr_1 \leq_{pr}^* pr_2$ and (*o pr* o') $\notin (P_K \cup C_i(K'))$ then all class instance triples (*o pr* o') $\in P_K$ are added to P_{K_Del} . Since (*o pr* o') $\in P_K$, it holds that (*o pr* o') $\in B_K$. Additionally, since $S_K \sqsubseteq S_{K'}$, it holds that (*o pr* o') $\notin C_i(K)$. Therefore, (*o pr* o') $\in B_K \setminus (P_K \cup C_i(K)) = F_K$ and (*o pr* o') $\notin C_i(K')$. Thus, due to Rule R2 of Prop. 4, it holds that (*o pr* o') $\in F_{K'}$.

All instance triples in $P_K \cap C_i(K')$ are moved to P_{K_Del} . Finally, $P_{K'} = (P_K \setminus P_{K_Del}) \cup P_{K_Add}$. \square

Prop. 9 The time complexity of Algorithm 1 is $O(|Inst_K|^2 * M^2 * (|K'|^2 + |P_K|))$, where $M = \max(|C_{K'}|, |Pr_{K'}|)$.

Proof: First, we will prove the following Lemma.

Lemma: Let K be a KB. Then, (i) the size complexity of $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|) * |K|)$ and (ii) the time complexity of $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|)^2 * |K|^2)$.

Proof:

(i) First note that $K \subseteq \mathcal{C}(K)$. Let T_{rdfs} denote the RDF and RDFS axiomatic triples [13], except the ones that contain *rdfs:li* terms, for $i \in \{1, 2, \dots\}$. It holds that $T_{rdfs} \subseteq \mathcal{C}(K)$. The inference rules used for the derivation of $\mathcal{C}(K)$ are the following:

- (1) If $(c \leq_{cl} c') \in K \cup T_{rdfs}$ then $(c \leq_{cl}^* c') \in \mathcal{C}(K)$.
- (2) If $(c \text{ type } Class) \in \mathcal{C}(K)$ then $(c \leq_{cl}^* c) \in \mathcal{C}(K)$.
- (3) If $(c_1 \leq_{cl}^* c_2) \in \mathcal{C}(K)$ and $(c_2 \leq_{cl}^* c_3) \in \mathcal{C}(K)$ then $(c_1 \leq_{cl}^* c_3) \in \mathcal{C}(K)$.
- (4) If $(c \text{ type } Class) \in \mathcal{C}(K)$ then $(c \leq_{cl}^* Resource) \in \mathcal{C}(K)$.
- (5) If $(o \text{ type } c) \in K \cup T_{rdfs}$ then $(o \text{ type } c) \in \mathcal{C}(K)$.
- (6) If $(o \text{ type } c) \in \mathcal{C}(K)$ and $(c \leq_{cl}^* c') \in \mathcal{C}(K)$ then $(o \text{ type } c') \in \mathcal{C}(K)$.
- (7) If $(o \text{ pr } o') \in \mathcal{C}(K)$ and $(pr \text{ domain } c) \in \mathcal{C}(K)$ then $(o \text{ type } c) \in \mathcal{C}(K)$.
- (8) If $(o \text{ pr } o') \in \mathcal{C}(K)$ and $(pr \text{ range } c) \in \mathcal{C}(K)$ then $(o' \text{ type } c) \in \mathcal{C}(K)$.
- (9) If $(pr \leq_{pr} pr') \in K \cup T_{rdfs}$ then $(pr \leq_{pr}^* pr') \in \mathcal{C}(K)$.
- (10) If $(pr \text{ type } Property) \in \mathcal{C}(K)$ then $(pr \leq_{pr}^* pr) \in \mathcal{C}(K)$.
- (11) If $(pr_1 \leq_{pr}^* pr_2) \in \mathcal{C}(K)$ and $(pr_2 \leq_{pr}^* pr_3) \in \mathcal{C}(K)$ then $(pr_1 \leq_{pr}^* pr_3) \in \mathcal{C}(K)$.
- (12) If $(o \text{ pr } o') \in K \cup T_{rdfs}$ and $(pr \leq_{pr}^* pr') \in \mathcal{C}(K)$ then $(o \text{ pr } o') \in \mathcal{C}(K)$.
- (13) If $(o \text{ pr } o') \in K \cup T_{rdfs}$ then $(o \text{ type } Resource), (pr \text{ type } Resource), (o' \text{ type } Resource) \in \mathcal{C}(K)$.
- (14) If $(o \text{ pr } o') \in K \cup T_{rdfs}$ then $(pr \text{ type } Property) \in \mathcal{C}(K)$.
- (15) If $(c \leq_{cl} c') \in K \cup T_{rdfs}$ then $(c \text{ type } Class), (c' \text{ type } Class) \in \mathcal{C}(K)$.
- (16) If $(pr \leq_{pr} pr') \in K \cup T_{rdfs}$ then $(pr \text{ type } Property), (pr' \text{ type } Property) \in \mathcal{C}(K)$.
- (17) If $(pr \text{ rdfs:domain } c) \in K \cup T_{rdfs}$ then $(pr \text{ type } Property), (c \text{ type } rdfs:Class) \in \mathcal{C}(K)$.
- (18) If $(pr \text{ range } c) \in K \cup T_{rdfs}$ then $(pr \text{ type } Property), (c \text{ type } rdfs:Class) \in \mathcal{C}(K)$.

The number of $(c \leq_{cl}^* c')$ triples in $\mathcal{C}(K)$ due to rules (1-4) is in $O(|C_K|^2)$. The number of (*o type* c') triples in $\mathcal{C}(K)$ due to rules (5-8) is in $O(|K| * |C_K|)$. The number of $(pr \leq_{pr}^* pr')$ triples in $\mathcal{C}(K)$ due to rules (9-11) is in $O(|Pr_K|^2)$. The number of (*o pr* o') triples in $\mathcal{C}(K)$ due to rule (12) is in $O(|K| * |Pr_K|)$. The size complexity due to rules (13-18) is $O(|K|)$. Thus, the total size complexity of $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|) * |K|)$.

(ii) Based on the form of the inference rules i(1)-i(18) and the form of the KBs that we consider, it follows that we can compute $\mathcal{C}(K)$ if (i) we execute the rules i(1) and i(3) until fixpoint, (ii) we execute the rules i(9) and i(11) until fixpoint, (iii) we execute the rules

i(13)-i(18), (iv) we execute the rules i(5) and i(12), (v) we execute the rules i(7)-i(8), (vi) we execute the rule i(4), (vii) we execute the rule i(6), and finally (viii) we execute the rules i(2) and i(10). The time complexity of the steps (i) and (ii) is $O(|K|^2)$, because the time complexity to compute the transitive closure of \leq_{cl} and \leq_{pr} is $O(|K|^2)$. Based on the fact that the size complexity of $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|) * |K|)$, the time complexity of the rest of the steps is $O(\max(|C_K|, |Pr_K|)^2 * |K|^2)$.

End Of Lemma

Since $S_K \sqsubseteq S_{K'}$, it holds that $|K| \leq |K'|$, $|C_K| \leq |C_{K'}|$, and $|Pr_K| \leq |Pr_{K'}|$, assuming that the representation of S_K and $S_{K'}$ is *redundancy-free*, i.e. only the reflexive and transitive reduction of the *subClassOf* and *subPropertyOf* relationships are stated in K and K' . Now, we will compute the time complexity of lines (5-6) of the Algorithm. The time for the tests (*o type c*) $\in (C_i(K') \cup P_k)$ in line (6) is $O(|NC| * |Inst_K| * |C_K| * (|K'|^2 + |P_K|))$. The time for the tests in line (6) (*o type c'*) $\notin C_i(K')$ is $O(|NC| * |Inst_K| * |K'|^2)$. The time for computing $Inst_K \cap URI$ in line (6), is $O(|Inst_K|)$. The time for computing $\{c \mid c' \leq_{cl}^* c\}$, for all $c' \in NC$, is $O(|NC| * |K'|^2)$ and the time of finding which of the elements of these sets belong to C_K is $O(|NC| * |C_{K'}| * |K|) = O(|NC| * |K'| * |K|)$. Thus, the total time complexity for the lines (5-6) is: $O(|NC| * |Inst_K| * |C_K| * (|K'|^2 + |P_K|))$.

Now, we will compute the time complexity of lines (7-8). The time for computing $\{c_2 \mid c_1 \leq_{cl}^* c_2\}$, for all $c_1 \in C_K$, is $O(|C_K| * |K'|^2)$ and the time of finding which of the elements of these sets belong to C_K is $O(|C_K| * |C_{K'}| * |K|)$. The time for checking which of the instance triples (*o type c2*) $\notin P_K \cup C_i(K')$ is $|C_K| * |Inst_K| * |C_K| * (|K'|^2 + |P_K|)$. The time for checking which of the instance triples (*o type c1*) $\in P_K$ is $|C_K| * |Inst_K| * |P_K|$. Thus, the time complexity of lines (7-8) is $|C_K| * |Inst_K| * |C_K| * (|K'|^2 + |P_K|)$.

Now, we will compute the time complexity of lines (10-11). The time for the tests (*o pr o'*) $\in (C_i(K') \cup P_k)$ in line (11) is $O(|NP| * |Inst_K| * |Inst_K| * |Pr_K| * (|K'|^2 + |P_K|))$. The time for the tests in line (11) (*o pr' o'*) $\notin C_i(K')$ is $O(|NP| * |Inst_K| * |Inst_K| * |K'|^2)$. The time for computing $\{pr \mid pr' \leq_{pr}^* pr\}$, for all $pr' \in NP$, is $O(|NP| * |K'|^2)$ and the time of finding which of the elements of these sets belong to Pr_K is $O(|NP| * |Pr_{K'}| * |K|) = O(|NP| * |K'| * |K|)$. Thus, the total complexity time for the lines (10-11) is: $O(|NP| * |Inst_K|^2 * |Pr_K| * (|K'|^2 + |P_K|))$.

Now, we will compute the time complexity of lines (12-13). The time for computing $\{pr_2 \mid pr_1 \leq_{pr}^* pr_2\}$, for all $pr_1 \in Pr_K$, is $O(|Pr_K| * |K'|^2)$ and the time of finding which of the elements of these sets belong to Pr_K is $O(|Pr_K| * |Pr_{K'}| * |K|)$. The time for checking which of the instance triples (*o pr2 o'*) $\notin P_K \cup C_i(K')$ is $O(|Pr_K| * |Inst_K|^2 * |Pr_K| * (|K'|^2 + |P_K|))$. The time for checking which of the instance triples (*o pr1 o'*) $\in P_K$ is $O(|Pr_K| * |Inst_K|^2 * |P_K|)$. Thus, the total time complexity of lines (12-13) is $O(|Pr_K| * |Inst_K|^2 * |Pr_K| * (|K'|^2 + |P_K|))$.

The time complexity of line (14) is $O(|P_K| * |K'|^2)$. Now, we will compute the time complexity of line (15). The size of P_{K_Del} is $O(|Inst_K| * |C_K| + |Inst_K|^2 * |Pr_K|) = O(|Inst_K|^2 * |K|)$. Thus, the time complexity of line (15) is $O(|P_K| * |Inst_K|^2 * |K|)$.

As shown in Lemma, the time complexity of computing $\mathcal{C}(K)$ and $\mathcal{C}(K')$ is $O(M^2 * |K'|^2)$.

Therefore, the total time complexity of Algorithm 1 is $O(|Inst_K|^2 * M^2 * (|K'|^2 + |P_K|))$.

□

Prop. 10 In the context of a transition $(C_i(K), F_K, P_K) \rightsquigarrow (C_i(K'), F_{K'}, P_{K'})$, it follows that: $(C_i(K) \setminus C_i(K')) \cap B_{K'} \subseteq F_{K'}$ (II3) iff $C_i(K) \cap P_{K'} = \emptyset$.

Proof:

\Rightarrow) Assume that $(C_i(K) \setminus C_i(K')) \cap B_{K'} \subseteq F_{K'}$. Further, assume that it exists an instance triple $t \in C_i(K) \cap P_{K'}$. Since $t \in P_{K'}$, it follows that $t \in B_{K'}$. Further, from the Definition of TFP-partition (Def. 5), it follows that $t \notin C_i(K')$. Thus, $t \in (C_i(K) \setminus C_i(K')) \cap B_{K'}$. Therefore, $t \in F_{K'}$, which is impossible. Thus, $C_i(K) \cap P_{K'} = \emptyset$.

\Leftarrow) Assume that $C_i(K) \cap P_{K'} = \emptyset$ and let $t \in (C_i(K) \setminus C_i(K')) \cap B_{K'}$. We will show that $t \in F_{K'}$. Assume that $t \notin F_{K'}$. Then, from the Definition of TFP-partition (Def. 5), it follows

that $t \in P_{K'}$. Then, $t \in \mathcal{C}_i(K) \cap P_{K'}$. However, this is impossible since $\mathcal{C}_i(K) \cap P_{K'} = \emptyset$. Thus, $t \in F_{K'}$. \square

Prop. 11 The negative part of the target TFP-partition of a consistent transition, always satisfies the following rules:²⁷

- (R4) If $(o \text{ type } c) \in \mathcal{C}_i(K)$, $c' \leq_{cl}^* c$, and $(o \text{ type } c) \notin \mathcal{C}_i(K')$ then $(o \text{ type } c') \in F_{K'}$.
(R5) If $(o \text{ pr } o') \in \mathcal{C}_i(K)$, $pr' \leq_{pr}^* pr$, and $(o \text{ pr } o') \notin \mathcal{C}_i(K')$ then $(o \text{ pr } o') \in F_{K'}$. \square

Proof:

If the conditions of the “if part” of the above rules are satisfied, then according to postulate *I3* of Def. 11, $(o \text{ type } c)$ and $(o \text{ pr } o')$ should be placed in $F_{K'}$. If *Rule R4* is not satisfied then $F_{K'}$ is not a lower set (and thus cannot be part of a TFP-partition). The same for *Rule R5*. \square

Prop. 12 The same as Prop. 6 but now applies to KBs $K = (S_K, I_K)$ and $K' = (S_{K'}, I_K)$, where S_K and $S_{K'}$ are not necessarily backwards compatible.

Proof:

\Rightarrow) Let $(o \text{ type } c') \in P_{K'}$. Certainly, $o \in \text{Inst}_K \cap \text{URI}$. We will show that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K)$ and $(o \text{ type } c') \notin \mathcal{C}_i(K')$. Assume that it exists $c \in C_K$ s.t. $(c' \leq_{cl}^* c$ and $(o \text{ type } c) \notin (\mathcal{C}_i(K') \cup P_K))$ or $(o \text{ type } c') \in \mathcal{C}_i(K')$. If it exists $c \in C_K$ s.t. $(c' \leq_{cl}^* c$ and $(o \text{ type } c) \notin (\mathcal{C}_i(K') \cup P_K))$ then it follows that $(o \text{ type } c) \notin (\mathcal{C}_i(K) \cup P_K)$ or $((o \text{ type } c) \in \mathcal{C}_i(K)$ and $(o \text{ type } c) \notin \mathcal{C}_i(K'))$. Since $o \in \text{Inst}_K \cap \text{URI}$ and $c \in C_K$, it follows that $(o \text{ type } c) \in B_K$. If $(o \text{ type } c) \notin (\mathcal{C}_i(K) \cup P_K)$ then it follows from Def. 5 that $(o \text{ type } c) \in F_K$. Therefore, it follows that $(o \text{ type } c') \in F_{K'}$ (see Rule R1 of Prop. 4). If $(o \text{ type } c) \in \mathcal{C}_i(K)$ and $(o \text{ type } c) \notin \mathcal{C}_i(K')$ then from Prop. 11 (Rule R4) it follows that $(o \text{ type } c') \in F_{K'}$. Thus, $(o \text{ type } c') \notin P_{K'}$, which is impossible. If $(o \text{ type } c') \in \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ type } c') \notin P_{K'}$, which is impossible.

\Leftarrow) Assume that $o \in \text{Inst}_K \cap \text{URI}$ and that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in (\mathcal{C}_i(K') \cup P_K)$. Additionally, assume that it holds $(o \text{ type } c') \notin \mathcal{C}_i(K')$. We will show that $(o \text{ type } c') \in P_{K'}$. It follows from Def. 5 that for all $c \in C_K$ s.t. $c' \leq_{cl}^* c$, it holds that $(o \text{ type } c) \in B_K \setminus F_K$ or $(o \text{ type } c) \in \mathcal{C}_i(K')$. It follows $(o \text{ type } c') \notin F_{K'}$ (note that Rule R1 of Prop. 4 and Rule R4 of Prop. 11 do not apply). Since $o \in \text{Inst}_K \cap \text{URI}$, $c' \in C_{K'}$, and $\text{Inst}_K = \text{Inst}_{K'}$, it follows that $(o \text{ type } c') \in B_{K'}$. Now, since $(o \text{ type } c') \notin \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ type } c') \in P_{K'}$. \square

Prop. 13 The same as Prop. 7 but now applies to KBs $K = (S_K, I_K)$ and $K' = (S_{K'}, I_K)$, where S_K and $S_{K'}$ are not necessarily backwards compatible.

Proof:

\Rightarrow) Let $(o \text{ pr } o') \in P_{K'}$. Since $(o \text{ pr } o') \in P_{K'}$, it follows that $\text{valid}(o, pr', o', K')$. We will show that for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that (i) $(o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K)$ and (ii) $(o \text{ pr } o') \notin \mathcal{C}_i(K')$. Assume that it exists $pr \in Pr_K$ s.t. (i) $(pr' \leq_{pr}^* pr$ and $(o \text{ pr } o') \notin (\mathcal{C}_i(K') \cup P_K))$, or (ii) $(o \text{ pr } o') \in \mathcal{C}_i(K')$. It follows that $pr' \leq_{pr}^* pr$, and $(o \text{ pr } o') \notin (\mathcal{C}_i(K) \cup P_K)$ or $((o \text{ pr } o') \in \mathcal{C}_i(K)$ and $(o \text{ pr } o') \notin \mathcal{C}_i(K'))$. Obviously, $(o \text{ pr } o') \in B_K$. If $(o \text{ pr } o') \notin (\mathcal{C}_i(K) \cup P_K)$ then it follows from Def. 5 that $(o \text{ pr } o') \in F_K$. Therefore, since $(o \text{ pr } o') \notin \mathcal{C}_i(K')$, it follows that $(o \text{ pr } o') \in F_{K'}$ (see Rule R2 of Prop. 4). If $(o \text{ pr } o') \in \mathcal{C}_i(K)$ and $(o \text{ pr } o') \notin \mathcal{C}_i(K')$ then it follows from Prop. 11 (Rule R5) that $(o \text{ pr } o') \in F_{K'}$. Thus, $(o \text{ pr } o') \notin P_{K'}$, which is impossible.

\Leftarrow) Assume that (i) $o \in \text{URI}$ and $\text{valid}(o, pr', o', K')$, (ii) for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that $(o \text{ pr } o') \in (\mathcal{C}_i(K') \cup P_K)$, and (iii) $(o \text{ pr } o') \notin \mathcal{C}_i(K')$. We will show that

²⁷ The relationships \leq_{cl}^* and \leq_{pr}^* hold in K' .

$(o \text{ } pr' \text{ } o') \in P_{K'}$. It follows from Def. 5 that for all $pr \in Pr_K$ s.t. $pr' \leq_{pr}^* pr$, it holds that $(o \text{ } pr \text{ } o') \in (B_K \setminus F_K)$ or $(o \text{ } pr \text{ } o') \in \mathcal{C}_i(K')$. It follows that $(o \text{ } pr' \text{ } o') \notin F_{K'}$ (note that Rule R2 and Rule R3 of Prop. 4 and Rule R5 of Prop. 11 do not apply). It holds $o \in Inst_{K'} \cap URI$, $o' \in Inst_{K'}$, and $pr \in Pr_{K'}$. Thus, $(o \text{ } pr' \text{ } o') \in B_{K'}$. Since $(o \text{ } pr' \text{ } o') \notin \mathcal{C}_i(K')$, it follows from Def. 5 that $(o \text{ } pr' \text{ } o') \in P_{K'}$. \square

Prop. 14 Let $K = (S_K, I_K)$ and $K' = (S_{K'}, I_K)$ be two (valid) KBs. If $(\mathcal{C}_i(K), F_K, P_K) \Rightarrow (\mathcal{C}_i(K'), F_{K'}, P_{K'})$, then $P_{K'} = Produce_Possibilities_{NBC}(K, P_K, S_{K'})$.

Proof: Initially, $P_{K_Add} = \emptyset$ and $P_{K_Del} = \emptyset$. For each new class $c' \in NC$, Algorithm $Produce_Possibilities_{NBC}(K, P_K, S_{K'})$, inserts to P_{K_Add} exactly the class instance triples indicated by Prop. 12.

For each old class $c_1 \in C_K$, if it exists $c_2 \in C_{K'}$ s.t. $c_1 \leq_{c_1}^* c_2$ and $(o \text{ } type \text{ } c_2) \notin (P_K \cup \mathcal{C}_i(K'))$ then all class instance triples $(o \text{ } type \text{ } c_1) \in P_K$ are added to P_{K_Del} . Note that $(o \text{ } type \text{ } c_2) \notin (\mathcal{C}_i(K) \cup P_K)$ or $((o \text{ } type \text{ } c_2) \in \mathcal{C}_i(K)$ and $(o \text{ } type \text{ } c_2) \notin \mathcal{C}_i(K'))$. Since $o \in Inst_K \cap URI$ and $c \in C_K$, it follows that $(o \text{ } type \text{ } c_2) \in B_K$. If $(o \text{ } type \text{ } c_2) \notin (\mathcal{C}_i(K) \cup P_K)$ then it follows from Def. 5 that $(o \text{ } type \text{ } c_2) \in F_K$. Therefore, since $(o \text{ } type \text{ } c_2) \notin \mathcal{C}_i(K')$, it follows that $(o \text{ } type \text{ } c_1) \in F_{K'}$ (see Rule R1 of Prop. 4). If $(o \text{ } type \text{ } c_2) \in \mathcal{C}_i(K)$ and $(o \text{ } type \text{ } c_2) \notin \mathcal{C}_i(K')$ then from Prop. 11 (Rule R4) it follows that $(o \text{ } type \text{ } c_1) \in F_{K'}$.

Additionally, it adds to P_{K_Del} , all class instance triples $(o \text{ } type \text{ } c) \in P_K$, where $c \notin C_{K'}$, because these instance triples do not belong to $B_{K'}$.

Similarly, for each new property $pr' \in NP$, Algorithm $Produce_Possibilities_{NBC}(K, P_K, S_{K'})$, inserts to P_{K_Add} exactly the property instance triples indicated by Prop. 13. For each old property $pr_1 \in Pr_K$, if it exists $pr_2 \in Pr_{K'}$ s.t. $pr_1 \leq_{pr}^* pr_2$, and $(o \text{ } pr_2 \text{ } o') \notin (P_K \cup \mathcal{C}_i(K'))$ then all class instance triples $(o \text{ } pr_1 \text{ } o') \in P_K$ are added to P_{K_Del} . Note that $(o \text{ } pr_2 \text{ } o') \notin (\mathcal{C}_i(K) \cup P_K)$ or $((o \text{ } pr_2 \text{ } o') \in \mathcal{C}_i(K)$ and $(o \text{ } pr_2 \text{ } o') \notin \mathcal{C}_i(K'))$. Obviously, $(o \text{ } pr_2 \text{ } o') \in B_K$. If $(o \text{ } pr_2 \text{ } o') \notin (\mathcal{C}_i(K) \cup P_K)$ then it follows from Def. 5 that $(o \text{ } pr_2 \text{ } o') \in F_K$. Therefore, since $(o \text{ } pr_2 \text{ } o') \notin \mathcal{C}_i(K')$, it follows that $(o \text{ } pr_1 \text{ } o') \in F_{K'}$ (see Rule R2 of Prop. 4). If $(o \text{ } pr_2 \text{ } o') \in \mathcal{C}_i(K)$ and $(o \text{ } pr_2 \text{ } o') \notin \mathcal{C}_i(K')$ then it follows from Prop. 11 (Rule R5) that $(o \text{ } pr_1 \text{ } o') \in F_{K'}$.

Additionally, it adds to P_{K_Del} all property instance triples $(o \text{ } pr \text{ } o') \in P_K$, where $pr \notin Pr_{K'}$ or $\neg valid(o, pr, o', K')$. This is because, in the first case they do not belong to $B_{K'}$ and, in the second case, they belong to $F_{K'}$ (see Rule R3 of Prop. 4).

All instance triples in $P_K \cap \mathcal{C}_i(K')$ are moved to P_{K_Del} . Finally, $P_{K'} = (P_K \setminus P_{K_Del}) \cup P_{K_Add}$. \square

Prop. 15 The time complexity of Algorithm 2 is $O(|Inst_K|^2 * M^2 * (S^2 + |P_K|))$, where $S = \max(|K|, |K'|)$ and $M = \max\{|C_K|, |Pr_K|, |C_{K'}|, |Pr_{K'}|\}$.

Proof: In the proof of Prop. 9, we have shown that if K is a KB then the size complexity of $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|) * |K|)$ and (ii) the time complexity of computing $\mathcal{C}(K)$ is in $O(\max(|C_K|, |Pr_K|)^2 * |K|^2)$.

Here, we will provide the complexity of the parts that Algorithm 2 differs from Algorithm 1. The complexity of the parts that Algorithm 2 is the same with Algorithm 1 is provided in the proof of Prop. 9, where we replace $|K|$ and $|K'|$ by S . The time complexity of lines (9-10) is $|P_K| * S^2$. The time complexity of lines (16-17) is $|P_K| * S^2$. The time complexity to compute $\mathcal{C}(K)$ and $\mathcal{C}(K')$ is $O(M^2 * S^2)$. Thus, the total complexity of Algorithm 2 is $O(|Inst_K|^2 * M^2 * (S^2 + |P_K|))$. \square

Prop. 16 Let $K_1 = (S_1, I_1)$ be a KB and let S_2, \dots, S_n be a sequence of backwards compatible schema versions, resulting to KBs $K_i = (S_i, I_1)$, for $i = 2, \dots, n$. Now consider the *one-step* migration of I_1 from the first (S_1) to the last (S_n) schema: $(\mathcal{C}_1, F_1, P_1) \rightsquigarrow (\mathcal{C}_n, F_n, P_n)$. Additionally, consider the *sequential migrations* scenario, where: $(\mathcal{C}_1, F_1, P_1) \rightsquigarrow (\mathcal{C}'_2, F'_2, P'_2)$

$\rightsquigarrow \dots \rightsquigarrow (\mathcal{C}'_n, F'_n, P'_n)$. If (a) for all $c_1, c_2 \in C_{K_1}$, it holds that $c_1 \leq_{cl}^* c_2$ in K_1 iff $c_1 \leq_{cl}^* c_2$ in K_n , and (b) for all $pr_1, pr_2 \in Pr_{K_1}$, it holds that $pr_1 \leq_{pr}^* pr_2$ in K_1 iff $pr_1 \leq_{pr}^* pr_2$ in K_n , then $(\mathcal{C}_n, F_n, P_n) = (\mathcal{C}'_n, F'_n, P'_n)$.

Proof: Obviously, $\mathcal{C}_n = \mathcal{C}'_n$. Let $(o \text{ type } c) \in F'_n$. Then, there is a class $c' \in C_{K_1} \cap C_{K_i}$ s.t. $c \leq_{cl}^* c'$ in K_i and $(o \text{ type } c') \in F_1 \cap F'_i$. Assume that $(o \text{ type } c) \notin F_n$. Then, $(o \text{ type } c) \in \mathcal{C}_n \cup P_n$. If $(o \text{ type } c) \in \mathcal{C}_n$ then this is impossible because $\mathcal{C}_n = \mathcal{C}'_n$. If $(o \text{ type } c) \in P_n$ then $(o \text{ type } c') \in P_n \cup \mathcal{C}_n$, due to Lemma 1(1). If $(o \text{ type } c') \in P_n$ then this is also impossible, because $(o \text{ type } c') \in F_1$ and due to Postulate *II2* of Def. 7. If $(o \text{ type } c') \in \mathcal{C}_n$ then $c' \in C_{K_1} \cap C_{K_n}$. It holds that it exists $c_1 \in C_{K_1}$ s.t. $c_1 \leq_{cl}^* c'$ in K_n and (a) there is $(o \text{ pr } o') \in I_1$ and $domain(pr) \leq_{cl}^* c_1$ in K_1 , or (b) there is $(o' \text{ pr } o) \in I_1$ and $range(pr) \leq_{cl}^* c_1$ in K_1 , or (c) there is $(o \text{ type } c_1) \in I_1$. Since $c_1, c' \in C_{K_1} \cap C_{K_n}$, it follows from the assumption of Prop. 16(a) that $c_1 \leq_{cl}^* c'$ in K_1 . Thus, $(o \text{ type } c') \in \mathcal{C}_1$, which is impossible since $(o \text{ type } c') \in F_1$. Thus, $(o \text{ type } c) \in F_n$.

Let $(o \text{ type } c) \in F_n$. Then, there is a class $c' \in C_{K_1} \cap C_{K_n}$ s.t. $c \leq_{cl}^* c'$ in K_n and $(o \text{ type } c') \in F_1 \cap F_n$. Assume that $(o \text{ type } c) \notin F'_n$. Then, $(o \text{ type } c) \in \mathcal{C}'_n \cup P'_n$. If $(o \text{ type } c) \in \mathcal{C}'_n$ then this is impossible because $\mathcal{C}_n = \mathcal{C}'_n$. If $(o \text{ type } c) \in P'_n$ then $(o \text{ type } c') \in P'_n \cup \mathcal{C}'_n$, due to Lemma 1(1). If $(o \text{ type } c') \in P'_n$, this is impossible due to the fact $(o \text{ type } c') \in F_1$ and postulate *II2*. If $(o \text{ type } c') \in \mathcal{C}'_n$ then $c' \in C_{K_1} \cap C_{K_n}$. It holds that it exists $c_1 \in C_{K_1}$ s.t. $c_1 \leq_{cl}^* c'$ in K_n and (a) there is $(o \text{ pr } o') \in I_1$ and $domain(pr) \leq_{cl}^* c_1$ in K_1 , or (b) there is $(o' \text{ pr } o) \in I_1$ and $range(pr) \leq_{cl}^* c_1$ in K_1 , or (c) there is $(o \text{ type } c_1) \in I_1$. Since $c_1, c' \in C_{K_1} \cap C_{K_n}$, it follows from the assumption of Prop. 16(a) that $c_1 \leq_{cl}^* c'$ in K_1 . Thus, $(o \text{ type } c') \in \mathcal{C}_1$, which is impossible since $(o \text{ type } c') \in F_1$. Thus, $(o \text{ type } c) \in F'_n$.

Let $(o \text{ pr } o') \in F'_n$. Then, there is a property $pr' \in Pr_{K_1} \cap Pr_{K_i}$ s.t. $pr \leq_{pr}^* pr'$ in K_i and $(o \text{ pr } o') \in F_1 \cap F'_i$. Assume that $(o \text{ pr } o') \notin F_n$. Then, $(o \text{ pr } o') \in \mathcal{C}_n \cup P_n$. If $(o \text{ pr } o') \in \mathcal{C}_n$ then this is impossible because $\mathcal{C}_n = \mathcal{C}'_n$. If $(o \text{ pr } o') \in P_n$ then $(o \text{ pr } o') \in P_n \cup \mathcal{C}_n$, due to Lemma 1(2). If $(o \text{ pr } o') \in P_n$ then this is also impossible, because $(o \text{ pr } o') \in F_1$ and due to Postulate *II2* of Def. 7. If $(o \text{ pr } o') \in \mathcal{C}_n$ then $pr' \in Pr_{K_1} \cap Pr_{K_n}$. It holds that it exists $pr_1 \in Pr_{K_1}$ s.t. $pr_1 \leq_{pr}^* pr'$ in K_n and $(o \text{ pr } o') \in I_1$. Since $pr_1, pr' \in Pr_{K_1} \cap Pr_{K_n}$, it follows from the assumption of Prop. 16(b) that $pr_1 \leq_{cl}^* pr'$ in K_1 . Thus, $(o \text{ pr } o') \in \mathcal{C}_1$, which is impossible since $(o \text{ pr } o') \in F_1$. Thus, $(o \text{ pr } o') \in F_n$.

Let $(o \text{ pr } o') \in F_n$. Then, there is a property $pr' \in Pr_{K_1} \cap Pr_{K_n}$ s.t. $pr \leq_{pr}^* pr'$ in K_n and $(o \text{ pr } o') \in F_1 \cap F_n$. Assume that $(o \text{ pr } o') \notin F'_n$. Then, $(o \text{ pr } o') \in \mathcal{C}'_n \cup P'_n$. If $(o \text{ pr } o') \in \mathcal{C}'_n$ then this is impossible because $\mathcal{C}_n = \mathcal{C}'_n$. If $(o \text{ pr } o') \in P'_n$ then $(o \text{ pr } o') \in P'_n \cup \mathcal{C}'_n$, due to Lemma 1(2). If $(o \text{ pr } o') \in P'_n$, this is impossible due to the fact $(o \text{ pr } o') \in F_1$ and postulate *II2*. If $(o \text{ pr } o') \in \mathcal{C}'_n$ then $pr' \in Pr_{K_1} \cap Pr_{K_n}$. It holds that it exists $pr_1 \in Pr_{K_1}$ s.t. $pr_1 \leq_{pr}^* pr'$ in K_n and $(o \text{ pr } o') \in I_1$. Since $pr_1, pr' \in Pr_{K_1} \cap Pr_{K_n}$, it follows from the assumption of Prop. 16(b) that $pr_1 \leq_{cl}^* pr'$ in K_1 . Thus, $(o \text{ pr } o') \in \mathcal{C}_1$, which is impossible since $(o \text{ pr } o') \in F_1$. Thus, $(o \text{ pr } o') \in F'_n$.

Therefore, $F_n = F'_n$. The fact that $P_n = P'_n$ follows immediately from the definition of TFP-partition (Def. 5). \square

Prop. 17: Let $X \subseteq P_K$. If $K^{up} = K \cup X$ then $\mathcal{C}(K^{up}) = \mathcal{C}(K) \cup SuperTriples(X)$.

Proof: If $(o \text{ type } c) \in X$, for $c \in C_K$, then the newly derived triples in $\mathcal{C}(K^{up})$, due to this newly added RDF triple to K^{up} , are $\{(o \text{ type } c') \mid c \leq_{cl}^* c'\}$. If $(o \text{ pr } o') \in X$, for $pr \in Pr_K$, then the newly derived triples in $\mathcal{C}(K^{up})$, due to this newly added RDF triple to K^{up} , are $\{(o \text{ pr } o') \mid pr \leq_{pr}^* pr'\} \cup \{(o \text{ type } domain(pr')) \mid pr \leq_{pr}^* pr'\} \cup \{(o' \text{ type } range(pr')) \mid pr \leq_{pr}^* pr'\}$. We will show that $\{(o \text{ type } domain(pr')) \mid pr \leq_{pr}^* pr'\} \cup \{(o' \text{ type } range(pr')) \mid pr \leq_{pr}^* pr'\} \subseteq \mathcal{C}(K)$. Since $(o \text{ pr } o') \in P_K$, it holds from (iii) of Def. 5 that $(o \text{ type } domain(pr)) \in \mathcal{C}(K)$ and $(o' \text{ type } range(pr)) \in \mathcal{C}(K)$. If it exists pr' s.t. $pr \leq_{pr}^* pr'$ then, since K is a valid KB, it holds that (i) $domain(pr) \leq_{cl}^* domain(pr')$ and (ii) $range(pr) \leq_{cl}^* range(pr')$. Therefore, $(o \text{ type } domain(pr')) \in \mathcal{C}(K)$ and $(o \text{ type } range(pr')) \in \mathcal{C}(K)$. \square

B Capturing Application Scenarios with TFP-partitions

TFP-partitions can be used to capture various application specific assumptions, e.g. cases where we know that *MSA* holds only for a *part* of the KB. For instance, suppose we know that the *MSA* holds only for the class instance triples of a subset O of the instances of K , i.e. $O \subseteq \text{Inst}_K$. We can capture this assumption with a TFP-partition having $F_K = \text{Invalid}(K) \cup \{(o \text{ type } c) \in B_K \setminus \mathcal{C}_i(K) \mid o \in O\}$.

Analogously, we can express explicitly assumptions regarding the specificity of property instance triples. Specifically, suppose we know that the *MSA* holds only for the property instance triples whose *subject* belongs to a subset O of the instances of K . We can capture this assumption with a TFP-partition having $F_K = \text{Invalid}(K) \cup \{(o \text{ pr } o') \in B_K \setminus \mathcal{C}_i(K) \mid o \in O\}$.

If we know that *MSA* holds only for the property instance triples whose *object* belongs to a subset O of the instances of K , then we can capture this assumption with a TFP-partition having $F_K = \text{Invalid}(K) \cup \{(o' \text{ pr } o) \in B_K \setminus \mathcal{C}_i(K) \mid o \in O\}$.

If we know the *MSA* holds only for the property instance triples whose both *subject* and *object* belong to a subset O of the instances of K , then we can set $F_K = \text{Invalid}(K) \cup \{(o \text{ pr } o') \in B_K \setminus \mathcal{C}_i(K) \mid o, o' \in O\}$.

Finally, in the case that more than one of the previous conditions hold then F_K is the union of their corresponding “negative statements”. Note that in all cases it holds $\text{SubTriples}(F_K) = F_K$ and $\text{Invalid}(K) \subseteq F_K$.

C Algorithm for Deriving the Negatives

Alg. 3 shows how from F_K , K , and $S_{K'}$, we can produce $F_{K'}$ by applying the rules of Def. 8.

Algorithm 3 $M_K\text{Update}(K, F_K, S_{K'})$

Input: A KB K , F_K , and a set of schema triples $S_{K'}$ s.t. $S_K \sqsubseteq S_{K'}$

Output: $F_{K'}$

- (1) $K' = (S_{K'}, I_K)$;
 - (2) $F_{K'} = \{\}$;
 - /* Part A: Class Instances */
 - /* Rule R1 */
 - (3) For all $((o \text{ type } c) \in F_K \text{ and } (o \text{ type } c) \notin \mathcal{C}_i(K'))$ do
 - (4) For all $(c' \in \mathcal{C}_{K'} \text{ and } c' \leq_{cl}^* c)$ do
 - (5) $F_{K'} = F_{K'} \cup \{(o \text{ type } c')\}$;
 - /* Part B: Property Instances */
 - /* Rule R2 */
 - (6) For all $((o \text{ pr } o') \in F_K \text{ and } (o \text{ pr } o') \notin \mathcal{C}_i(K'))$ do
 - (7) For all $(pr' \in \text{Pr}_{K'} \text{ and } pr' \leq_{pr}^* pr)$ do
 - (8) $F_{K'} = F_{K'} \cup \{(o \text{ pr } o')\}$;
 - /* Rule R3 */
 - (9) For all $((o \text{ pr } o')$ s.t. $o \in \text{Inst}_{K'} \cap \text{URI}$, $o' \in \text{Inst}_{K'}$, $pr \in \text{Pr}_{K'}$, and $\neg \text{valid}(o, pr, o', K')$) do
 - (10) $F_{K'} = F_{K'} \cup \{(o \text{ pr } o')\}$;
 - (11) Return $F_{K'}$;
-

D Example of a Backwards Compatible Migration

Example 10. Consider, for instance, Figure 5. Assume that K satisfies the *MSA*. In K , **Fiat_1** is not a **Person**, i.e. $(\text{Fiat_1 type Person}) \in F_K$. In K' , we have that $(\text{Fiat_1 type Person}) \notin C_i(K')$. Thus, we have that $(\text{Fiat_1 type Person}) \in F_{K'}$ (from Rule *R1*). In K , **Bob** is not a **Car**, i.e. $(\text{Bob type Car}) \in F_K$. In K' , a new class is inserted that specializes the class **Car**, called **Van**. We have that $(\text{Bob type Car}) \notin C_i(K')$. So, $(\text{Bob type Van}) \in F_{K'}$ (from Rule *R1*).

In K , **Alice** does not **drive** a **BMW_1**, i.e. $(\text{Alice drives BMW_1}) \in F_K$. In K' , we have that $(\text{Alice drives BMW_1}) \notin C_i(K')$. Thus, we have that $(\text{Alice drives BMW_1}) \in F_{K'}$ (from Rule *R2*). In K , **Alice** does not **work at** **FORTH**, i.e. $(\text{Alice works at FORTH}) \in F_K$. In K' , a new property is inserted that specializes the property **works at**, called **paid from**. **Alice** belongs to the instances of the domain of property **paid from** and **FORTH** belongs to the instances of the range of property **paid from**. We have that $(\text{Alice works at FORTH}) \notin C_i(K')$. So, $(\text{Alice paid from FORTH}) \in F_{K'}$ (from Rule *R2*).

Additionally, note that $(\text{Fiat_1 paid from FORTH}) \in F_{K'}$ (due to Rule *R3*). This is because $(\text{Fiat_1 type Person}) \notin C_i(K')$, and thus it holds $\neg \text{valid}(\text{Fiat_1, paid from, FORTH, } K')$. \square

E Example of Applying the Algorithm for Non-Backwards Compatible Migration

Example 11. Consider the example of Figure 9. Dashed rectangles denote deleted properties in $S_{K'}$, which appear in the instance triples of $K' = (S_{K'}, I_K)$. Note that, even though property pr_1 is removed in $S_{K'}$, it still exists in K' as an unconnected element because the instance triple $(o_1 \text{ pr}_1 o_2) \in I_K$, and thus $(\text{pr}_1 \text{ type Property}) \in C(K')$, according to the RDF/S semantics [13]. The domain and range of pr_1 is the class *Resource*.

Suppose that $P_K = \{(o_6 \text{ type F}), (o_6 \text{ type G}), (o_6 \text{ type J}), (o_7 \text{ type C}), (o_7 \text{ type D}), (o_7 \text{ type F}), (o_7 \text{ type G}), (o_1 \text{ pr}_2 o_2), (o_5 \text{ pr}_3 o_4)\}$.

Executing Algorithm 2, step by step, we have that:

- (1) $K' = (S_{K'}, I_K)$;
 - (2) $P_{K_Add} = \emptyset$;
 - (3) $P_{K_Del} = \emptyset$;
 - (4) $NC = \{\mathbf{I}\}$;
 - (6) $P_{K_Add} = P_{K_Add} \cup \{(o_1 \text{ type I}), (o_3 \text{ type I}), (o_5 \text{ type I}), (o_6 \text{ type I}), (o_7 \text{ type I})\}$;
 - (8) $P_{K_Del} = P_{K_Del} \cup \{(o_6 \text{ type G}), (o_6 \text{ type J})\}$;
- Note that $\{(o_6 \text{ type D}), (o_6 \text{ type A})\} \notin (P_K \cup C_i(K'))$ and $\mathbf{G} \leq_{cl} \mathbf{D}$, $\mathbf{J} \leq_{cl} \mathbf{A}$ hold in K' . So, $\{(o_6 \text{ type G}), (o_6 \text{ type J})\}$ are moved from P_K to $F_{K'}$. In fact, $(o_6 \text{ type G})$ is moved to $F_{K'}$ due to Rule *R1* and $(o_6 \text{ type J})$ is moved to $F_{K'}$ due to Rule *R4*.
- (10) $P_{K_Del} = P_{K_Del} \cup \{(o_6 \text{ type F}), (o_7 \text{ type F})\}$;
- Note that class *F* is removed in $S_{K'}$.
- (11) $NP = \{\text{pr}_5\}$;
 - (13) $P_{K_Add} = P_{K_Add} \cup \{(o_3 \text{ pr}_5 o_4)\}$;
 - (15) $P_{K_Del} = P_{K_Del} \cup \{(o_5 \text{ pr}_3 o_4)\}$;
- Note that $(o_5 \text{ pr}_4 o_4) \notin (P_K \cup C_i(K'))$ and $\text{pr}_3 \leq_{pr} \text{pr}_4$ holds in K' . So, $(o_5 \text{ pr}_3 o_4)$ is moved from P_K to $F_{K'}$ (due to Rule *R2*).
- (17) $P_{K_Del} = P_{K_Del} \cup \{(o_1 \text{ pr}_2 o_2)\}$;
- Note that in K' , $\text{domain}(\text{pr}_2) = \mathbf{G}$ and $o_1 \notin \text{inst}_{K'}(\text{domain}(\text{pr}_2))$.
- (18) $P_{K_Del} = P_{K_Del} \cup \{(o_7 \text{ type C})\}$;
- Note that P_{K_Del} is updated by those triples that belong to P_K and now belong to $C(K')$. So, we have to remove them from P_K .

(19) $P_{K'} = P_K \setminus P_{K_Del} =$
 $\{(o_6 \text{ type F}), (o_6 \text{ type G}), (o_6 \text{ type J}), (o_7 \text{ type C}), (o_7 \text{ type D}),$
 $(o_7 \text{ type F}), (o_7 \text{ type G}), (o_1 \text{ pr}_2 \ o_2), (o_5 \text{ pr}_3 \ o_4)\} \setminus$
 $\{(o_6 \text{ type F}), (o_6 \text{ type G}), (o_6 \text{ type J}), (o_7 \text{ type F}), (o_7 \text{ type C}), (o_1 \text{ pr}_2 \ o_2),$
 $(o_5 \text{ pr}_3 \ o_4)\} = \{(o_7 \text{ type D}), (o_7 \text{ type G})\};$
(20) $P_{K'} = P_{K'} \cup P_{K_Add} =$
 $\{(o_7 \text{ type D}), (o_7 \text{ type G})\} \cup \{(o_1 \text{ type I}), (o_3 \text{ type I}), (o_5 \text{ type I}), (o_6 \text{ type I}),$
 $(o_7 \text{ type I}), (o_3 \text{ pr}_5 \ o_4)\};$
(21) Return $P_{K'}$; □

F Composite Possibilities

So far P_K contains individual suggestions for a KB K . Here, we describe how we can extend P_K and reach a P_K^{ext} , called *extended set of possibilities*, that contains also *composite suggestions*, where a composite suggestion is a set of instance triples, aiming at further aiding the curator.

Example 12. To grasp the idea, consider the scenario of Figure 16, where $(\text{HKR6263 type Car})$ is a possible class instance triple in K' . The proposed *composite possibility* in K' is: $\{(\text{HKR6263 type Car}), (\text{Mary drives HKR6263})\}$. This means that if the user decides to add the class instance triple $(\text{HKR6263 type Car})$ to K' then the property instance triple $(\text{Mary drives HKR6263})$ can be added to K' . However, if the user does not add HKR6263 as an instance of Car , then $(\text{Mary drives HKR6263})$ would not be added to K' , since this triple is an element of $\text{Invalid}(K')$, and it belongs to $F_{K'}$. This is the motivation for defining composite possibilities. Regarding interaction, if the user accepts the proposed composite possibility, both instance triples $\{(\text{HKR6263 type Car})$ and $(\text{Mary drives HKR6263})\}$ will be added to K' . □

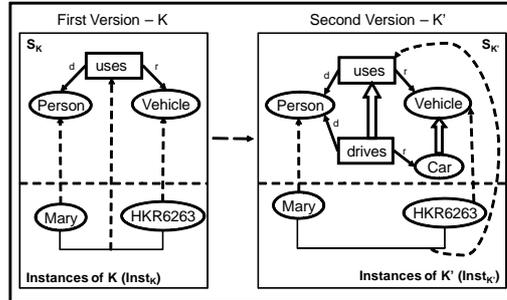


Fig. 16. First scenario for *Extended Possibilities*

To define precisely composite suggestions, we start by an auxiliary definition.

Def. 12 (Valid Property Set)

We define the *valid property set* of K , denoted by $\text{Valid}(K)$, as follows:

$$\text{Valid}(K) = \{(o \text{ pr } o') \in B_K \mid \text{valid}(o, \text{pr}, o', K) \text{ holds}\}. \quad \square$$

The set of composite possibilities P_K^{ext} is not a set of triples like those in P_K , but a family of sets containing:

- one singleton $\{t\}$, for each $t \in P_K$, and
- triple sets of the form $\{h, t\}$ or $\{h_1, t, h_2\}$, where h, h_1, h_2 are class instance triples of P_K and t is a property instance triple that (a) it does not belong to $\text{Valid}(K)$, and (b) t would belong to $\text{Valid}(K)$ (and hence it could be added to K) if h, h_1, h_2 were added to K .

In other words, the non-singleton elements of P_K^{ext} consist of an *hypothesis* that is already possible (i.e. h, h_1, h_2), and a *consequence* of that hypothesis (i.e. t can be added to K , if h, h_1, h_2 were added to K).

We can define the *extended set of possibilities* in general terms, after first defining an ordering over triple sets.

Def. 13 Let s, s' be two sets of instance triples. We define $s \preceq s'$ iff

- (i) they contain the same property instance triples and
- (ii) for all class instance triples $t \in s$, it exists $t' \in s'$ s.t. $t \in \text{SubTriples}(t')$. □

Def. 14 (Extended Set of Possibilities) Let K be a KB. We define the *extended set of possibilities* as follows:

$$P_K^{\text{ext}} = \{\{p\} \mid p \in P_K\} \cup P_K^{\text{comp}}, \text{ where}$$

$$P_K^{\text{comp}} = \text{maximal}_{\preceq}(\{s \cup \{t\} \mid s \subseteq P_K, t \notin \text{Valid}(K), t \in \text{Valid}(K \cup s), \text{ and}$$

$$t \notin \text{Valid}(K \cup s') \forall s' \subset s\}) \square$$

Let $Cl(o) = \{c \mid (o \text{ type } c) \in C_i(K)\}$ and $posCl(o) = \{c \mid (o \text{ type } c) \in P_K\}$. By taking into account Def. 12, we can reach to the following:

Prop. 18 (Composite possibilities) Let K be a KB. It holds that:

$$P_K^{\text{comp}} = P_1 \cup P_2 \cup P_3, \text{ where:}$$

$$P_1 = \{\{h_1, t\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_1 = (o \text{ type } c_1) \in P_K, \text{ where } \text{domain}(pr) = c_1, \text{ range}(pr) \in Cl(o')\}$$

$$P_2 = \{\{t, h_2\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_2 = (o' \text{ type } c_2) \in P_K, \text{ where } \text{range}(pr) = c_2, \text{ domain}(pr) \in Cl(o)\}$$

$$P_3 = \{\{h_1, t, h_2\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_1 = (o \text{ type } c_1) \in P_K \text{ where } \text{domain}(pr) = c_1,$$

$$h_2 = (o' \text{ type } c_2) \in P_K \text{ where } \text{range}(pr) = c_2\} \quad \square$$

Proof: First, we will prove the following lemma.

Lemma: It holds that $\{s \cup \{t\} \mid s \subseteq P_K, t \notin \text{Valid}(K), t \in \text{Valid}(K \cup s), \text{ and } t \notin \text{Valid}(K \cup s') \forall s' \subset s\} = P'_1 \cup P'_2 \cup P'_3$ where

$$P_K^{\text{comp}} = P_1 \cup P_2 \cup P_3, \text{ where:}$$

$$P'_1 = \{\{h_1, t\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_1 = (o \text{ type } c_1) \in P_K, \text{ where } c_1 \leq_{cl}^* \text{domain}(pr), \text{ range}(pr) \in Cl(o')\}$$

$$P'_2 = \{\{t, h_2\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_2 = (o' \text{ type } c_2) \in P_K, \text{ where } c_2 \leq_{cl}^* \text{range}(pr), \text{ domain}(pr) \in Cl(o)\}$$

$$P'_3 = \{\{h_1, t, h_2\} \mid t = (o \text{ pr } o') \notin \text{Valid}(K) \text{ where } o, o' \in \text{Inst}_K, pr \in Pr_K,$$

$$h_1 = (o \text{ type } c_1) \in P_K \text{ where } c_1 \leq_{cl}^* \text{domain}(pr),$$

$$h_2 = (o' \text{ type } c_2) \in P_K \text{ where } c_2 \leq_{cl}^* \text{range}(pr),$$

$$\text{domain}(pr) \notin Cl(o), \text{ and } \text{range}(pr) \notin Cl(o')\}$$

Proof:

\Leftarrow Let $x \in P'_1 \cup P'_2 \cup P'_3$. Then, $x = s \cup \{t\}$, where $t = (o \text{ pr } o') \notin \text{Valid}(K)$, for $o, o' \in \text{Inst}_K$ and $pr \in Pr_K$, and (i) $s = \{(o \text{ type } c_1)\}$, where $c_1 \leq_{cl}^* \text{domain}(pr)$, if $x \in P'_1$, (ii) $s = \{(o' \text{ type } c_2)\}$, where $c_2 \leq_{cl}^* \text{range}(pr)$, if $x \in P'_2$, and (iii) $s = \{(o \text{ type } c_1), (o' \text{ type } c_2)\}$,

where $c_1 \leq_{cl}^* domain(pr)$ and $c_2 \leq_{cl}^* range(pr)$, if $x \in P_3'$. In all cases $s \subseteq P_K$. Since $valid(o, pr, o', K \cup s)$, it follows that $t \in Valid(K \cup s)$. It is easy to see that for all $s' \subset s$, it holds that $t \notin Valid(K \cup s')$, because it does not hold $valid(o, pr, o', K \cup s')$. Thus, $x \in \{s \cup \{t\} \mid s \subseteq P_K, t \notin Valid(K), t \in Valid(K \cup s), \text{ and } t \notin Valid(K \cup s') \forall s' \subset s\}$.
 \Rightarrow Let $x \in \{s \cup \{t\} \mid s \subseteq P_K, t \notin Valid(K), t \in Valid(K \cup s), \text{ and } t \notin Valid(K \cup s') \forall s' \subset s\}$. We will show that $x \in P_1' \cup P_2' \cup P_3'$. Let $x = s \cup \{t\}$. Since $s \subseteq P_K, t \notin Valid(K)$ and $t \in Valid(K \cup s)$, it follows that $t = (o \text{ pr } o')$, where $o, o' \in Inst_K$ and $pr \in Pr_K$. Since $t \notin Valid(K \cup s') \forall s' \subset s$, it follows that (i) if $range(pr) \in Cl(o')$ then $s = \{h_1\}$, where $h_1 = (o \text{ type } c_1) \in P_K$ and $c_1 \leq_{cl}^* domain(pr)$, (ii) if $domain(pr) \in Cl(o)$ then $s = \{h_2\}$, where $h_2 = (o' \text{ type } c_2) \in P_K$ and $c_2 \leq_{cl}^* range(pr)$, and (iii) if $domain(pr) \notin Cl(o)$ and $range(pr) \notin Cl(o')$ then $s = \{h_1, h_2\}$, where $h_1 = (o \text{ type } c_1) \in P_K, h_2 = (o' \text{ type } c_2) \in P_K, c_1 \leq_{cl}^* domain(pr)$, and $c_2 \leq_{cl}^* range(pr)$. Therefore, it follows that $x \in P_1' \cup P_2' \cup P_3'$.
End of Lemma

From the above Lemma and Lemma 1(1), Proposition 18 now follows immediately. \square

It follows that each *composite possibility* of P_K^{comp} contains either two or three instance triples.

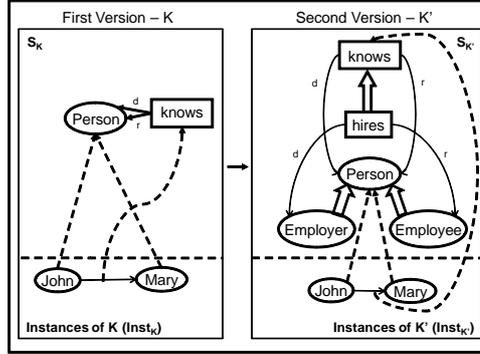


Fig. 17. Second scenario for *Extended Possibilities*

Example 13. The *composite possibility* of the previous example (Figure 16) contains a pair of instance triples. An example with three instance triples is the scenario of Figure 17. Suppose that **Employer** and **Employee** are new classes in K' and **hires** is a new property in K' that specializes the property **knows** in K . Suppose $(John \text{ knows } Mary)$ is a property instance triple included in K . In this case, we have in K' : $posCl(John) = \{Employer, Employee\}$ and $posCl(Mary) = \{Employer, Employee\}$, and $P_{K'}^{ext}$ contains the following *composite possibility*: $\{(John \text{ type } Employer), (John \text{ hires } Mary), (Mary \text{ type } Employee)\}$. Thus, if the user accepts the composite possibility then all instance triples $(John \text{ type } Employer)$, $(Mary \text{ type } Employee)$, and $(John \text{ hires } Mary)$ are added to K' . \square

At last we should stress that there is no need to store composite possibilities as they can be computed on demand.

G RIMQA: Screendumps

Some screendumps of the prototype system, called RIMQA (**R**DF **I**nstance **M**igration **Q**uality **A**ssistant)²⁸, are shown at Figures 18, 19, and 20.

²⁸ <http://www.ics.forth.gr/isl/RIMQA>

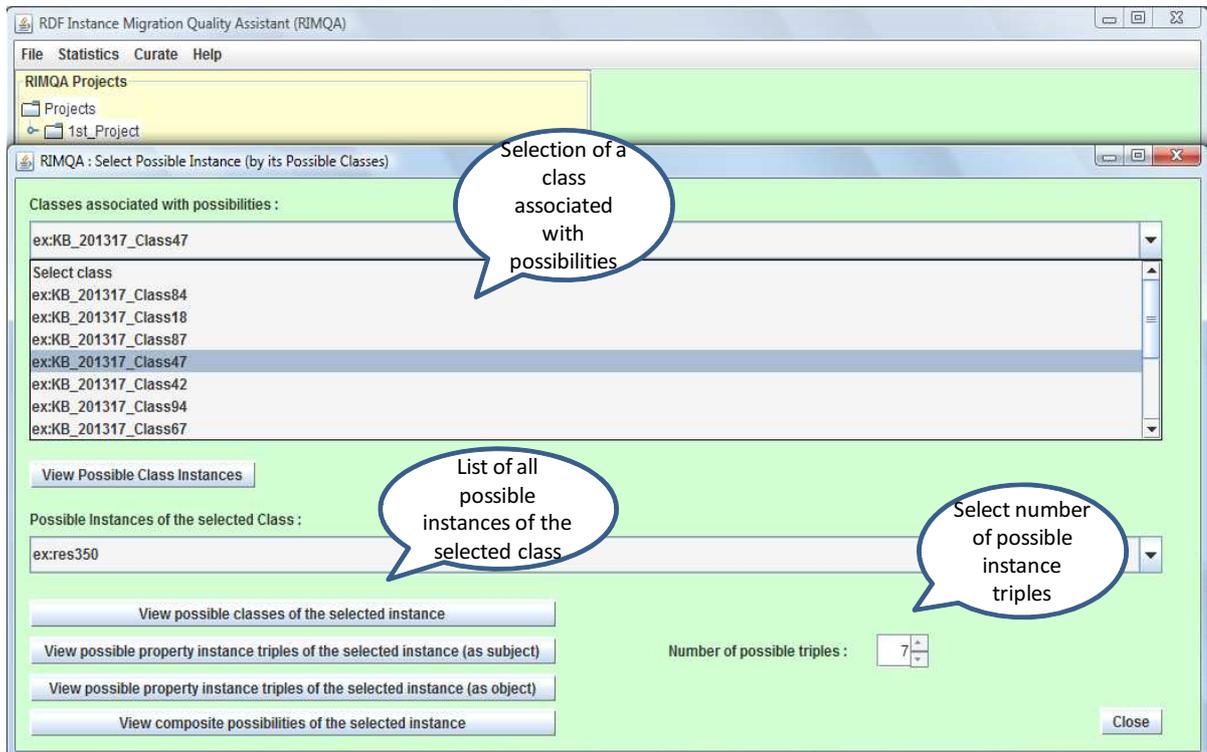


Fig. 18. First selected screenshot of RIMQA

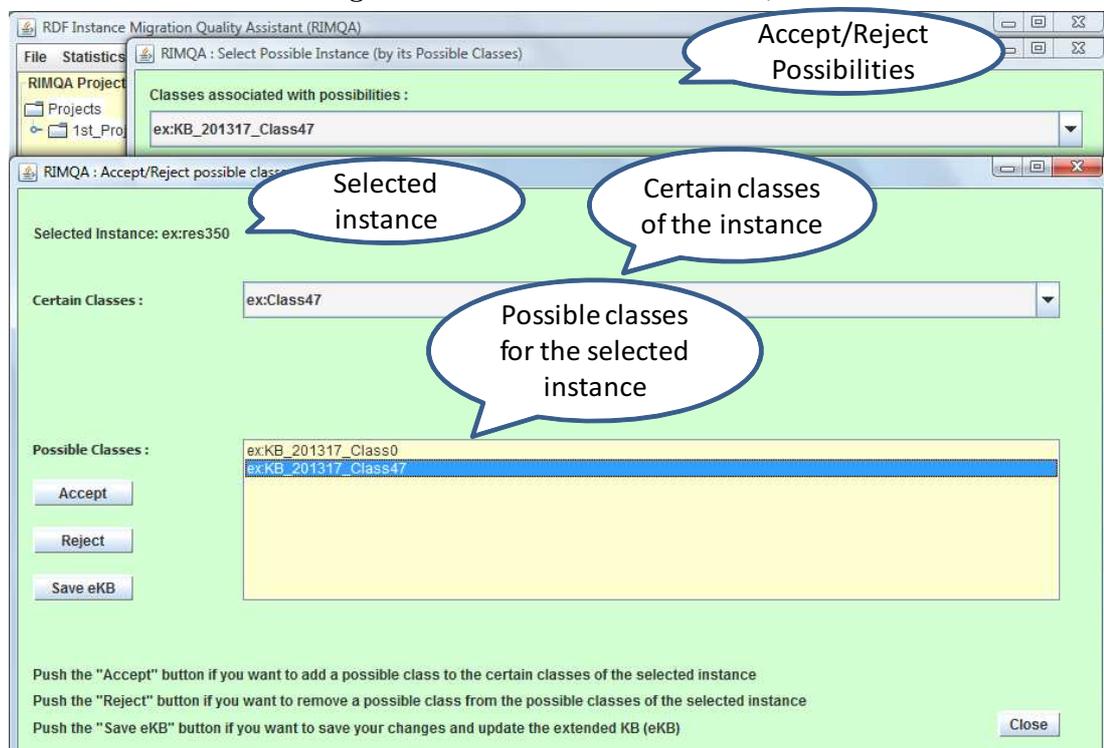


Fig. 19. Second selected screenshot of RIMQA

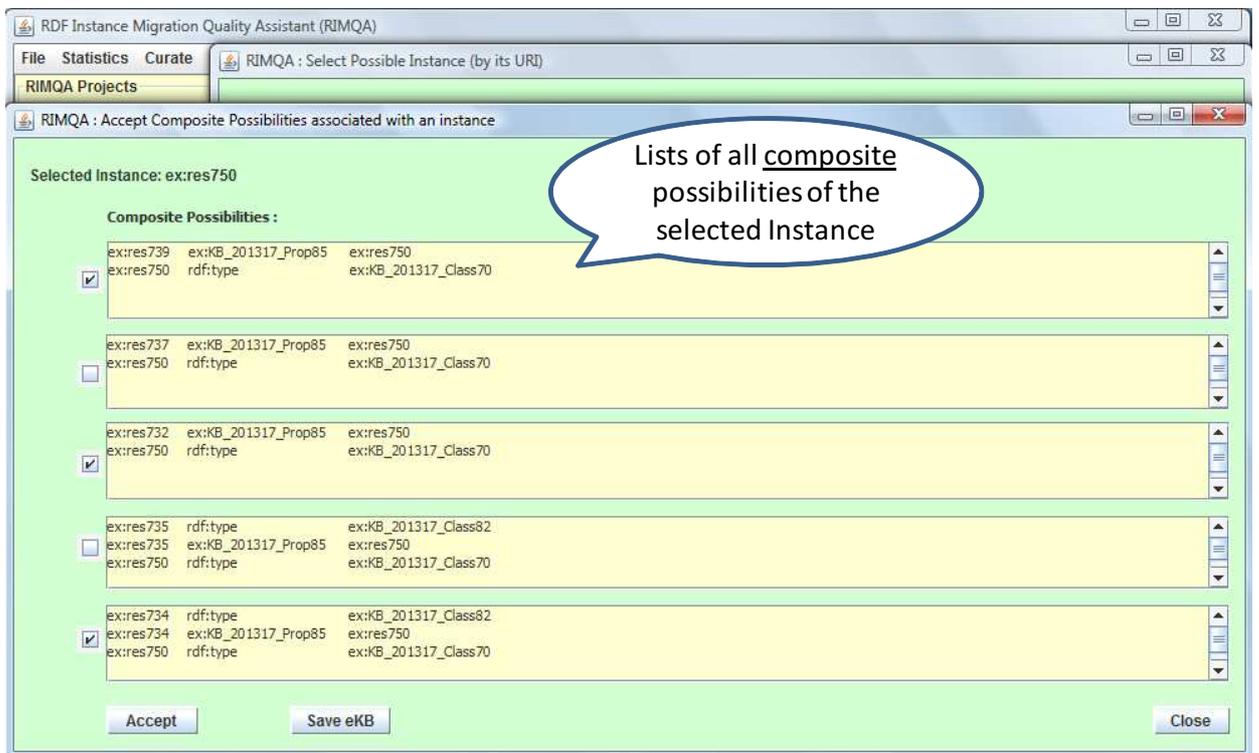


Fig. 20. Third selected screenshot of RIMQA