

# Specification and Incremental Maintenance of Linked Data Mashup Views

Vânia M. P. Vidal<sup>1</sup>, Marco A. Casanova<sup>2</sup>, Narciso Arruda<sup>1</sup>, Mariano Roberval<sup>1</sup>, Luiz Paes Leme<sup>3</sup>, Gisele Lopes<sup>2</sup>, Chiara Renso<sup>4</sup>

<sup>1</sup> Federal University of Ceará, Fortaleza, CE, Brazil  
{vvidal, narciso, mariano}@lia.ufc.br

<sup>2</sup> Department of Informatics – Pontifical Catholic University of Rio de Janeiro, RJ, Brazil  
{casanova, grlopes}@inf.puc-rio.br

<sup>3</sup> Fluminense Federal University, Niteroi, RJ, Brazil  
lapaesleme@ic.uff.br

<sup>4</sup> ISTI Institute of National Research Council, Pisa, Italy  
chiara.renso@isti.cnr.it

**Abstract.** The Linked Data initiative promotes the publication of previously isolated databases as interlinked RDF datasets, thereby creating a global scale data space, known as the Web of Data. Linked Data Mashup applications, which consume data from the multiple Linked Data sources in the Web of Data, are confronted with the challenge of obtaining a homogenized view of this global data space, called a Linked Data Mashup view. This paper proposes an ontology-based framework for formally specifying Linked Data Mashup views, and a strategy for the incremental maintenance of such views, based on their specifications.

**Keywords:** Data Mashup application, RDF dataset interlinking, Linked Data, View Maintenance.

## 1 Introduction

The Linked Data initiative [2] brought new opportunities for building the next generation of Semantic Mashup applications [8]. By exposing previously isolated datasets as data graphs, which can be interlinked and integrated with other datasets, Linked Data allows creating a global-scale interlinked data space, known as the Web of Data. The success of the Linked Data initiative is mainly due to the adoption of known Web standards, such as Web infrastructure standards (URIs and HTTP), Semantic Web standards (RDF and RDFS) and vocabularies, which facilitate the deployment of Linked Data sources. Thanks to this, an unprecedented amount of Linked Data sources were recently produced and continues growing fast.

A *linked data mashup* is an (web) application that offers new functionality by combining, aggregating, and transforming data available on the *web of Data* [11], [19]. Powered by tools and technologies having been developed by the semantic web community, there are various applications domains building applications with Linked

Data mashups [11]. A simple example of LD Mashup is *BBC Music* [13] that integrates data from two Linked Data sources *DBpedia* [3] and *MusicBrainz* [20].

Linked Data Mashup (LDM) applications are confronted with the challenge of obtaining a homogenized view of this global data space, which we call a *Linked Data Mashup view (LDM view)*. The creation of a LDM view is a complex task which involves four major challenges: (1) selection of the Linked Data sources that are relevant for the application; (2) extraction and translation of data from different, possibly heterogeneous Linked Data sources to a common vocabulary; (3) identification of links between resources in different Linked Data sources; (4) combination and fusion of multiple representations of the same real-world object into a single representation and resolution of data inconsistencies to improve the quality of the data.

To be useful, a LDM view must be continuously maintained to reflect dynamic data sources updates. Basically, there are two strategies for materialized view maintenance. *Re-materialization* re-computes view data at pre-established times, whereas *incremental maintenance* periodically modifies part of the view data to reflect updates to the database. It has been shown that incremental maintenance generally outperforms full view recomputation [1], [7], [9], [15].

In this paper, we investigate the problem of incremental maintenance of LDM view. First, we propose an ontology-based framework for formally specifying LDM views. In our framework, a LDM view is specified with the help of exported views, sameAs linkset views, data fusion rules and a normalization function. The LDM view specification is used to automatically materialize the mashup view. Then, we propose a strategy that uses the LDM view specification for incrementally maintain the mashup view materialization.

The incremental maintenance strategy of LDM views is the major contribution of this paper. The strategy addresses the problem of dealing with the combination and fusion of multiple representations of the same real-world object when the Linked Data sources are updated. As discussed in the related work section, this problem received little attention and yet poses new challenges due to very nature of mashup data.

The paper is organized as follows. Section 2 presents the framework to create LDM views. Section 3 discusses how to maintain LDM Views. Section 4 reviews related work. Finally, Section 5 contains the conclusions.

## 2 Ontology-Based Framework for LDM View Specification

### 2.1 Overview

In this section, we discuss a three level ontology-based framework, as summarized in Figure 1, to formally specify LDM views. In the Mashup View Layer, the mashup view ontology OD specifies the concepts of the mashup application (i.e., the *conceptual model*), which is the common vocabulary for integrating data exported by the Linked Data sources.

In the Web of Data Layer, each data source  $S_i$  is described by a *source ontology*  $O_{S_i}$ , published on the Web according to the Linked Data principles. These source ontologies are depicted in the Web of Data layer in Figure 1.

Each Linked Data source  $S_i$  exports one or more views. Each such view  $E_i$  has an ontology  $O_{E_i}$  and a set of rules  $M_{E_i}$  that map concepts of  $O_{S_i}$  to concepts of  $O_D$ . The vocabulary of  $O_{E_i}$  is the subset of the vocabulary of  $O_D$  whose terms occur in the head of the rules in  $M_{E_i}$ . The exported ontologies are depicted in the Exported Views and Linkset Views Layer in Figure 1.

We consider two types of sameAs links: *exported sameAs links*, which are exported by a Linked Data source, and *mashup sameAs links*, which are automatically created based on a sameAs linkset view specification [5] specifically defined for the mashup application.

As detailed in the following subsections, a *LDM view specification* is an n-tuple  $\lambda = (D, O_D, E_1, \dots, E_n, IL_1, \dots, IL_m, ML_1, \dots, ML_p, \mu, \eta, Q)$ , where:

- $D$  is the *name* of the mashup view;
- $O_D$  is the *mashup view ontology*;
- $E_1, \dots, E_n$  are *exported view specifications* with ontologies  $O_{E_1}, \dots, O_{E_n}$ , whose vocabularies are subsets of the vocabulary of  $O_D$ ;
- $IL_1, \dots, IL_m$  are *imported sameAs linkset view specifications* between  $E_1, \dots, E_n$ ;
- $ML_1, \dots, ML_p$  are *mashup sameAs linkset view specifications* between  $E_1, \dots, E_n$ ;
- $\mu$  is a set of *fusion rules* from  $O_{E_1}, \dots, O_{E_n}$  to  $O_D$ ;
- $\eta$  is a *normalization function symbol* whose interpretation defines how to remap IRIs of the exported views to IRIs of the LDM view;
- $Q$  is a set of *quality assessment metrics*, which are used to quantify the quality of the Linked Data sources. This information is required by fusion rules. The specification of quality assessment metric [16] is out of the scope of this work.

The process for generating the LDM view specification  $\lambda$  consists of 5 steps: (1) Modeling of the mashup ontology View. (2) Generation of the exported views specifications. (3) Generation of the imported sameAs linkset view specifications. (4) Generation of the mashup sameAs linkset view specifications. (5) Definition of the normalization function, fusion rules and quality assessment metrics. Steps 2 to 5 are detailed in the following subsections.

In our framework, the materialization of a *LDM view specification*  $\lambda$  is automatically processed based on its specification and consists of four steps:

1. **Materialization of the exported views.** This step translates source data to the exported view vocabulary as specified by the mapping rules in the exported view specification.
2. **Materialization of the imported sameAs linksets views.** Given an imported linkset view  $IL$  over a Linked Data source  $S$ , this step imports sameAs links from  $S$  to the LDM view materialization.
3. **Materialization of the mashup sameAs linksets.** Given a mashup linkset view  $ML$ , this step computes sameAs links based on the specification of  $ML$ .
4. **Materialization of the mashup view:** This step materializes the mashup view by applying the normalization function and the fusion rules to the materialized exported views and the materialized sameAs Linkset views. It includes the combination and fusion of multiple representations of the same real-world object into a single representation and the resolution of data inconsistencies.

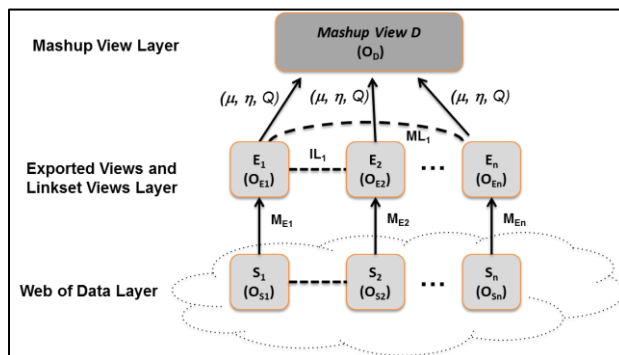


Fig. 1. Three Level Ontology-Based Framework.

## 2.2 Running Example

Throughout the paper, we will adopt a simple example of a mashup application about music, called *DBB\_Music*, which integrates data from two Linked Data sources: *DBpedia* [3] and *MusicBrainz* [20]. Figure 2 depicts, in UML notation, the application ontology *MUSIC\_OWL* for the *DBB\_Music* mashup, which reuses terms from three well-known vocabularies: *FOAF* (Friend of a Friend), *MO* (Music Ontology) and *DC* (Dublin Core). We use the prefix “*moa:*” for the new terms defined in the *MUSIC\_OWL* ontology. *DBpedia* uses the *DBpedia Ontology* [21] which we call *DBpedia\_OWL*, and we use the prefix “*dbpedia-owl:*” to refer to it. *MusicBrainz* uses the *Music Ontology* [17] and we use the prefix “*mo:*” to refer to it.




Fig. 2. MUSIC\_OWL.

## 2.3 Specification and Materialization of Exported Views

An *exported view specification* is a quintuple  $(E, S, OS, OE, ME)$ , where:

- $E$  is the *name* of the view
- $S$  is a Linked Data source
- $O_S$  is the ontology of  $S$
- $O_E$  is the *exported view ontology* such that the terms of its vocabulary occur in the heads of the mapping rules in  $M$
- $M_E$  is a set of mapping rules from  $O_S$  to  $O_E$



In our framework, each linked data source exports a view whose specification is automatically generated considering the source ontology, the application ontology and the mappings between the source ontology and the application ontology. Note that, when an exported view  $E$  is part of a LDM view specification, the vocabulary of the ontology  $O_E$  must be a subset of the vocabulary of the mashup view ontology, as required in Section 2.1. The problem of generating exported views specifications is addressed in our previous work [18] and it is out of the scope of this work.

Consider the mashup application *DBB\_Music* introduced in Section 2.1. An exported view specification, *DBpedia\_EV*, which uses the data source *DBpedia* can be defined based on the mappings rules between *Music\_owl* and *DBpedia\_owl*. Due to space limitation, the mappings rules are omitted here. The vocabulary of the exported view ontology of *DBpedia\_EV* then is:  $\{mo:Record, mo:MusicArtist, mo:Track, foaf:homepage, moa:labelName, dc:title, dc:date, foaf:name, dc:description, mo:track, foaf:maker, foaf:made\}$

An exported view specification, *MusicBrainz\_EV*, which uses the data source *MusicBrainz* can be likewise defined. The vocabulary of the exported view ontology of *MusicBrainz\_EV* is:  $\{mo:Record, mo:MusicArtist, mo:Track, moa:labelName, dc:title, dc:date, foaf:name, foaf:homepage, mo:track\_number, mo:track, foaf:maker, foaf:made\}$

A materialization of an exported view requires translating source data into the exported view vocabulary as specified by the mapping rules. Referring to our case study, Figure 3(b) shows materializations of the exported views *MusicBrainz\_EV* and *DBpedia\_EV* obtained by applying the mappings rules in applying the mappings rules.

## 2.4 Specification and Materialization of Exported sameAs Linkset Views

In our framework, the existing sameAs links between resources of different exported views should also be exported to the mashup view. To create exported links, the user should first specify an exported sameAs linkset view and then materialize the links.

Let  $(T, S_T, O_{S_T}, O_T, \mu_T)$  and  $(U, S_U, O_{S_U}, O_U, \mu_U)$  be two exported views. An *exported sameAs linkset view specification* is a tuple  $(IL, T, U, C, C_{S_T}, C_{S_U})$ , where:

- $IL$  is the *name* of the view;
- $(T, S_T, O_{S_T}, O_T, M_T)$  and  $(U, S_U, O_{S_U}, O_U, M_U)$  are exported view specifications;
- $C$  is a class in both the vocabularies of the exported view ontologies  $O_T$  and  $O_U$ ;
- $C_{S_T}$  is a class in the vocabulary of the data source ontology  $O_{S_T}$  of  $S_T$  such that there is a rule in  $M_T$ , indicating that instances of  $C_{S_T}$  are mapped to instances of  $C$ ;
- $C_{S_U}$  is a class in the vocabulary of the data source ontology  $O_{S_U}$  of  $S_U$  such that there is a rule in  $M_U$ , indicating that that instances of  $C_{S_U}$  are mapped to instances of  $C$ ;

Let  $s_T$  and  $s_U$  be states respectively of  $S_T$  and  $S_U$ . The materialization of  $IL$  in  $s_T$  and  $s_U$  is the set  $IL[s_T, s_U]$  defined as:

$$\begin{aligned} &(t, owl:sameAs, u) \in IL[s_T, s_U] \text{ iff} \\ &(t, owl:sameAs, u) \in s_T \wedge (t, rdf:type, C_{S_T}) \in I[s_T](C_{S_T}) \wedge \\ &(u, rdf:type, C_{S_U}) \in I[s_U](C_{S_U}). \end{aligned}$$

That is,  $IL[s_T, s_U]$  imports sameAs links from  $s_T$  whose subject is in the interpretation in  $s_T$  of the class  $C_{s_T}$  and whose object is in the interpretation in  $s_U$  of the class  $C_{s_U}$ .

Consider, for example, the exported views *MusicBrainz\_EV* and *DBpedia\_EV* introduced in Section 2.3. As shown in Figure 3(a), the data source MusicBrainz contains sameAs links matching instances of classe *mo:Record* with instances of class *dbpedia\_owl:Albums*. It also contains sameAs links matching instances of *mo:MusicArtist* with *dbpedia\_owl:MusicalArtist*. In order to materialize the sameAs links for the instances of the exported views *MusicBrainz\_EV* and *DBpedia\_EV*, two exported linkset views should be specified:

- (IL<sub>1</sub>, MusicBrainz\_EV, DBpedia\_EV, mo:Record, mo:Record, dbpedia\_owl:Album);
- (IL<sub>2</sub>, MusicBrainz\_EV, DBpedia\_EV, mo:MusicArtist, mo:MusicArtist, dbpedia\_owl:MusicalArtist).

Referring to Figure 3(a) of our case study, the sameAs link from *db:artist/Michael\_Jackson* to *dbpedia:Michael\_Jackson* is materialized by IL1, and the sameAs link from *db:record/Thriller* to *dbpedia:Thriller* is materialized by IL2.

## 2.5 Specification and Materialization of Mashup sameAs Linkset Views

In our framework, Mashup sameAs links are inferred by matching property values of resources defined in the exported views. To create mashup links, the user should first specify the mashup sameAs linkset view and then materialize the links. More precisely, a *mashup sameAs linkset view specification* is a tuple  $(ML, T, U, C, p_1, \dots, p_n, \mu)$ , where:

- $ML$  is the *name* of the view;
- $(T, S_T, O_{s_T}, O_T, \mu_T)$  and  $(U, S_U, O_{s_U}, O_U, \mu_U)$  are exported view specifications;
- $C$  is a class in both the vocabularies of the exported view ontologies  $O_T$  and  $O_U$ ;
- $p_1, \dots, p_n$  are properties of class  $C$  in both the vocabularies of the exported view ontologies  $O_T$  and  $O_U$ ;
- $\mu$  is a  $2n$ -relation, called the *match predicate*;

Let  $e_T$  and  $e_U$  be states respectively of  $T$  and  $U$ . The materialization of  $ML$  in  $e_T$  and  $e_U$  is the set  $ML[e_T, e_U]$  defined as:

$$(t, owl:sameAs, u) \in ML[e_T, e_U] \text{ iff there are triples}$$

$$(t, rdf:type, C), (t, P_1, v_1), \dots, (t, P_n, v_n) \in e_T \text{ and}$$

$$(u, rdf:type, C), (u, P_1, w_1), \dots, (u, P_n, w_n) \in e_U$$

$$\text{such that } (v_1, \dots, v_n, w_1, \dots, w_n) \in \mu$$

Consider the exported views *MusicBrainz\_EV* and *DBpedia\_EV* introduced in Section 3.3. Then, sameAs linkset views could be specified for matching instances of the class *mo:Track*. As an example, consider the mashup sameAs view specification for *mo:Track*:  $(ML1, MusicBrainz\_EV, DBpedia\_EV, mo:Track, dc:title, p)$  where the match predicate  $p$  is defined as:  $(v_1, w_1) \in p$  iff  $\sigma(v_k, w_k) \geq \alpha$ , for each  $k=1$ , where  $\sigma$  is the 3-gram distance [14] and  $\alpha = 0.5$ . Referring to our case study, Figure 3(b) shows the sameAs links automatically created using *ML1*.

## 2.6 Specification of the Normalization Function and Fusion Assertions

In this section, we first introduce the concepts of normalization function and fusion assertions, then we describe how the data fusion rules are induced from the fusion assertion.

**Normalization Function.** The LDM view specification includes a *normalization function*, denoted  $\eta$ , that remaps all IRIs in the exported views that are declared to denote the same object, via sameAs links, to one canonical target IRI. The normalization function must satisfy the following axiom (using an infix notation for the sameAs property):

$$(N1) \quad \forall x_1 \forall x_2 (x_1 \text{ sameAs } x_2 \Leftrightarrow \eta(x_1) = \eta(x_2))$$

which says that the normalization function must remap to the same IRI two IRIs  $x_1$  and  $x_2$  iff they are declared to be equivalent via a sameAs statement of the form “ $x_1$  sameAs  $x_2$ ”.

The normalization function partitions the IRIs of the exported views resources into a set of equivalence classes. In the materialization of the LDM view, all IRIs in the same equivalence class are homogenized by grouping all properties of those IRIs into one canonical target IRI. The canonical IRI also have owl:sameAs links pointing to the original IRIs, which makes it possible for applications to refer back to the original data sources on the Web.

Referring to our case study, the equivalence classes induced by the sameAs links in Figure 3(b) are:  $\varepsilon_1 = \{dbpedia:Michael\_Jackson, db:Artist/Michael\_Jackson\}$ ,  $\varepsilon_2 = \{dbpedia:Thriller, db:record/Thriller\}$ ,  $\varepsilon_3 = \{dbpedia:BBM, db:track/BBM\}$ ,  $\varepsilon_4 = \{dbpedia:WBSS\}$ ,  $\varepsilon_5 = \{db:track/WBSS\}$ .

**Data Fusion Assertions and Data Fusion Rules.** In the rest of this section, let  $D$  be a LDM view,  $O_D = (V_D, \Sigma_D)$  be the ontology of  $D$ , and  $\eta$  be the normalization function. Let  $C$  be a class in  $V_D$ . We define  $Props(C, V_D) = \{P \mid P \text{ is a property in } V_D \text{ and } C \text{ is a subclass of the domain of } P\}$ .

In our approach, the user is free to define how to resolve the problem of contradictory attribute values when combining multiple representations of the same real-world object into a single representation (canonical IRI). This is specified with the help of *data fusion property assertions*.

A *data fusion property assertion (FPA)* for a property  $P$  in the context of a class  $C$ , is an expression of form:  $\Psi: P[C] \equiv f/Q$  where  $\Psi$  is the name of the FPA,  $C$  is a class in  $V_D$ ,  $P$  and  $Q$  are properties in  $Props(C, V_D)$ , and  $f$  is a *data fusion function symbol*, which denote functions whose domain is a set of sets of individuals and whose range is a set of individuals.

Data fusion assertions should be regarded as a shorthand notation for a class of data fusion rules. A set  $A$  of data fusion assertions defined over the vocabulary  $V_D$  of  $D$ , induces a set of fusion rules for  $D$  as follows:

- Let  $C$  be a class in  $V_D$  and  $P$  be a property in  $Props(C, V_D)$ . Assume that  $\Psi: P[C] \equiv f/Q$  is in  $A$ , and that  $S_1, \dots, S_n$  are all the exported views of  $D$  whose vocabulary contains class  $C$  and property  $Q$ . Then, the fusion rule for  $P$  in the context  $C$  induced by  $A$  is

$$P(x, u) \leftarrow u = f(v/B[x, v]) \text{ where:}$$

$$B[x,v]=((C(x_1),Q(x_1,v))_{S_1}, x=\eta(x_1)); \dots; ((C(x_n),Q(x_n,v))_{S_n}, x=\eta(x_n))).$$

The subscript  $S_i$  indicates the exported view against which  $(C(x_i), Q(x_i, v))$  will be evaluated. Intuitively,  $\{v/B[x,v]\}$  denotes the set of individuals  $v$  that satisfy  $B[x,v]$ . Then,  $f/v/B[x,v]$  denotes the individual to which  $f$  maps the set  $\{v/B[x,v]\}$ .

- Let  $C$  be a class in  $V_D$  and  $P$  be a property in  $Props(C, V_D)$ . Assume that there is no assertion in  $A$  for property  $P$  in the context of  $C$ , and  $S_1, \dots, S_n$  are all the exported views of  $D$  whose vocabulary contains  $C$  and  $P$ . Then, the fusion rule for  $P$  in the context  $C$  induced by  $A$  is

$$P(x,v) \leftarrow ((C(x_1), P(x_1, v))_{S_1}, x=\eta(x_1)); \dots; ((C(x_n), P(x_n, v))_{S_n}, x=\eta(x_n)))$$

- Let  $C$  be a class in  $V_D$ . Assume that  $S_1, \dots, S_n$  are all the exported views of  $D$  whose vocabulary contains  $C$ . Then, the fusion rule for class  $C$  induced by  $A$  is

$$C(x) \leftarrow ((C(x_1))_{S_1}, x=\eta(x_1)); \dots; ((C(x_n))_{S_n}, x=\eta(x_n))).$$

For example, the *FPA* for the *moa:labelName* property in the context of class *mo:Record*

$$\Psi: moa:labelName[mo:Record] \equiv KeepSingleValueByReputation/moa:labelName$$

induces the following fusion rule:

$$moa:labelName(x,u) \leftarrow u = KeepSingleValueByReputation(v / (mo:Record(y), moa:labelName(y,v))_{DBPedia\_EV}, x=\eta(y); (mo:Record(z), moa:labelName(z,v))_{MusicBrainz\_EV}, x=\eta(z))).$$

The fusion rule creates a triple of the form “ $(x, moa:labelName, u)$ ” in a materialization of the LDM view by taking  $u$  as the most reputable name for  $x$  found in the materialization of the two exported views, *MusicBrainz\_EV* and *DBPedia\_EV*.

When no *FPA* is specified for the *moa:labelName* property in the context of class *mo:Record*, then the default fusion rule

$$moa:labelName(x,v) \leftarrow (mo:Record(y), moa:labelName(y,v))_{DBPedia\_EV}, x=\eta(y); (mo:Record(z), moa:labelName(z,v))_{MusicBrainz\_EV}, x=\eta(z))$$

is induced. The fusion rule creates a triple for all values of *moa:labelName* coming from the exported views.

## 2.7 Materialization of the LD Mashup View

In this section we describe how to materialize a LDM view by applying the normalization function and data fusion rules to the materialized exported views and the materialized sameAs Linkset views. The materialization process includes the combination and fusion of multiple representations of the same real-world object into a single representation and the resolution of data inconsistencies [4].

We begin by introducing the notation that will be used in this section:

- $\lambda = (D, O_D, E_1, \dots, E_n, IL_1, \dots, IL_p, ML_1, \dots, ML_q, \mu, \eta, Q)$  is a LDM view specification.
- $V_D$  is the vocabulary of  $O_D$ .
- $\mathbf{C} = \{ C / C \text{ is a class in } V_D \}$ .
- $\mathbf{P} = \{ P / P \text{ is a property in } V_D \}$ .
- $e_1, \dots, e_n$  are states of  $E_1, \dots, E_n$  and  $l_1, \dots, l_m$  are states of  $IL_1, \dots, IL_p, ML_1, \dots, ML_q$ , with  $m=p+q$ .

- $s = (e_1, \dots, e_n, l_1, \dots, l_m)$
- $IRIs(s) = \{ x / x \text{ is the subject of a triple in a tripleset in } s \}$
- $\varepsilon_1, \dots, \varepsilon_w$  are the equivalence classes induced by  $l_1, \dots, l_m$
- $[\varepsilon_i]$  denotes the IRI that represents  $\varepsilon_i$
- $\eta^s$  denotes the interpretation of  $\eta$  for the  $\varepsilon_1, \dots, \varepsilon_w$
- $I[e](T)$  denotes the set of triples that state  $e$  assigns to a class or property  $T$ , i.e., the interpretation of  $T$  in  $e$ .

The state  $s$  induces a state  $d$  of  $D$  as follows:

- Let  $C$  be a class in  $V_D$ . Assume that  $\mu$  has a fusion rule whose head uses  $C$  and that the rule is of the form

$$C(x) \leftarrow (C(x_1)_{S_1}, x = \eta(x_1)); \dots; (C(x_p)_{S_p}, x = \eta(x_p))$$

Then, the interpretation of  $C$  induced by  $\mu$  in  $s$ , also denoted  $I[s](C)$ , is defined as:

$$(x, \text{rdf:type}, C) \in I[s](C) \text{ iff } ((x, \text{rdf:type}, C) \in I[e_1](C), x = \eta^s(x_1)) \vee \dots \vee \\ (x, \text{rdf:type}, C) \in I[e_n](C), x = \eta^s(x_n)).$$

- Let  $P$  be a property in  $V_D$ . Assume that  $\mu$  has a fusion rule whose head uses  $P$  and that the rule is of the form

$$P(x, v) \leftarrow (C(x_1), P(x_1, v))_{S_1}, x = \eta(x_1); \dots; (C(x_n), P(x_n, v))_{S_n}, x = \eta(x_n).$$

Then, the interpretation of  $P$  induced by  $\mu$  in  $s$ , also denoted  $I[s](P)$ , is defined as:

$$(x, P, v) \in I[s](P) \text{ iff } ((x, \text{rdf:type}, C) \in I[e_1](C), (x, P, v) \in I[e_1](P), x = \eta^s(x_1)) \vee \\ \dots \vee ((x, \text{rdf:type}, C) \in I[e_n](C), (x, P, v) \in I[e_n](P), x = \eta^s(x_n))$$

- Let  $P$  be a property in  $V_D$ . Assume that  $\mu$  has a fusion rule whose head uses  $P$  and that the rule is of the form

$$P(x, u) \leftarrow u = f(v/B[x, v]) \text{ where:}$$

$$B[x, v] = ((C(x_1), Q(x_1, v))_{S_1}, x = \eta(x_1)); \dots; ((C(x_n), Q(x_n, v))_{S_n}, x = \eta(x_n))$$

Then, the interpretation of  $P$  induced by  $\mu$  in  $s$ , also denoted  $I[s](P)$ , is defined as:

$$(x, P, u) \in I[s](P) \text{ iff } u = I[s](f)(\{v / I[s](B[x, v]) = \text{true}\})$$

- If there is more than one fusion rule in  $\mu$  whose head uses  $P$ , then  $I[s](P)$  is the union of all sets of triples as defined on the right-hand sides of the double implications above.
- Recall that  $\mathbf{C}$  is the set of all classes in  $V_D$  and  $\mathbf{P}$  is the set of all properties in  $V_D$  and  $\varepsilon_i$  is an equivalence class induced by  $l_1, \dots, l_m$ . We define the *DataFusion* function as follows:

$$\text{DataFusion}(\varepsilon_i, s) = \cup_{C \in \mathbf{C}} \{ (x, \text{rdf:type}, C) \in I[s](C) / x = [\varepsilon_i] \} \cup \\ \cup_{P \in \mathbf{P}} \{ (x, P, y) \in I[s](P) / x = [\varepsilon_i] \} \cup \\ \cup \{ ([\varepsilon_i], \text{owl:sameAs}, y) / y \in \varepsilon_i \wedge y \neq [\varepsilon_i] \}$$

- We also define the state  $d$  or the *materialization* of  $D$  induced by  $s$  as (recall that  $\varepsilon_1, \dots, \varepsilon_w$  are the equivalence classes induced by  $l_1, \dots, l_m$ ):

$$d = \cup_{i=1, \dots, w} \text{DataFusion}(\varepsilon_i, s).$$

Referring to our case study, Figure 3(c) shows the state of the mashup view computed from the states of the exported views and sameAs linkset views in Figure 3 (b). The mashup view has 5 resources which are computed by applying the *DataFusion* function to equivalence classes  $\varepsilon_1, \dots, \varepsilon_5$  defined in Section 2.6.

### 3 Incremental Maintenance of LDM Views

We now turn to the problem of maintaining a LDM view, when update operations are applied to the Linked Data sources.

In this section, let:

- $\lambda = (D, O_D, E_1, \dots, E_n, IL_1, \dots, IL_p, ML_1, \dots, ML_q, \mu, \eta, Q)$  is a LDM view specification;
- $e_1, \dots, e_n$  are states of  $E_1, \dots, E_n$  and  $l_1, \dots, l_m$  are states of  $IL_1, \dots, IL_p, ML_1, \dots, ML_q$ , with  $m=p+q$ ;
- $s = (e_1, \dots, e_n, l_1, \dots, l_m)$ ;
- $d$  be the state of  $D$  induced by  $s$ .

The incremental view maintenance problem is schematically described by the diagram in Figure 4. The user specifies an update  $u$  against a base data source, which results in new states  $e'_1, \dots, e'_n$  of the exported views and new states  $l'_1, \dots, l'_m$  of the sameAs views. Let  $d'$  be the state of  $D$  induced by  $(e'_1, \dots, e'_n, l'_1, \dots, l'_m)$ . We say that a set of updates  $U_D$  over the state  $d$  correctly maintains  $D$  iff  $U_D(d) = d'$ .




Fig. 4. Incremental view maintenance problem

Figure 5 shows the main components of the architecture we suggest to incrementally maintain the LDM view  $D$ . For each data source  $S$  that exports a view  $E_s$  to  $D$ , there is a *Source View Controller*, with the following functionality:

1. Identify updates on  $S$  that are relevant to  $D$ . This is computed from the specification of the exported view  $E_s$ .
2. For each relevant update  $u$ , create the set:  $U = \{ \langle r, \psi_r \rangle / r \text{ is the IRI of a resource affected by } \mu \text{ and } \psi_r \text{ is the state of } r \text{ in the new state of the exported view } E \}$ .
3. Send  $U$  to the LDM View Controller.

Note that  $R$  can be automatically computed based on the exported view specification, the new and old states of  $S$  and the base update  $u$ . Hence, no access to the mashup view is required. The problem of computing  $R$  is out of the scope of this paper. A similar problem was addressed in a previous work [22].

The *LDM View Controller* receives  $U$  and then performs the incremental maintenance of the exported view  $E_s$  and sameAs views  $IL_1, \dots, IL_p, ML_1, \dots, ML_q$ , followed by the incremental maintenance of mashup view  $D$ . The incremental maintenance of sameAs views has already been addressed in our previous work [5], and the problem of incremental maintenance of exported views is very similar to the problem addressed in [22]. Therefore, we do not address those problems in this paper.

In order to accomplish step 2, i.e. the incremental maintenance of mashup view  $D$ , the *LDM View Controller* executes the procedure  $Effect(p)$  for each pair  $p = \langle r, \psi_r \rangle$  in  $R$  (see Table 1).  $Effect(p)$  applies, for the current state of the mashup view, the required updates to maintain  $D$  considering the new state of  $r$ .




Fig. 5. Suggested platform for LDM view maintenance.

Table 1. Procedure  $Effect(r)$

<p><i>Parameters</i> (as defined above):</p> <ul style="list-style-type: none"> <li><math>\lambda</math>, the LDM view specification</li> <li><math>t = (e'_1, \dots, e'_n, I'_1, \dots, I'_m)</math></li> <li><math>d</math>, the state of mashup view <math>D</math> obtained by applying the fusion rules to <math>s</math></li> </ul> <p><i>Input:</i> pair <math>\langle r, \psi_r \rangle</math></p> <p><b>Step 1.</b> Computes the new state of the resources directly affected by the new state of <math>r</math>.</p> <p>1.1 Compute <math>A</math> the set of directly affected resources:</p> $A_1 := \{x \mid x \in IRIs(t) \wedge y \in IRIs(d) \wedge (y, owl:sameAs, r) \wedge (y, owl:sameAs, x) \in d\};$ $A_2 := \{x \mid x \in IRIs(t) \wedge \eta^r(x) = \eta^r(r)\}$ $A := A_1 \cup A_2;$ <p>1.2. Compute <math>\epsilon_1, \dots, \epsilon_n</math> the new equivalence classes induced by the sameAs links in <math>I'_1, \dots, I'_m</math> for the resources in <math>A</math>;</p> <p>1.3. Compute <math>\Delta_1</math> the new state of the mashup objects for the resources in <math>A</math>;</p> $\Delta_1 := \cup_{i=1, \dots, m} DataFusion(\epsilon_i, t);$ <p>1.4. Compute <math>\Delta_2</math> the old state of the mashup objects for the resources in <math>A</math>;</p> $\Delta_2 := \{(x, P, z) \mid (x, P, z) \in d \wedge y \in A \wedge (x, owl:sameAs, y) \in d\};$ <p>1.5. <math>d = (d - \Delta_2) \cup \Delta_1</math>;</p> <p><b>Step 2.</b> Computes the new state of triples <math>(x, P, y)</math> in <math>d</math> such that <math>y</math> is affected by the updates in step 1.5.</p> <p>2.1. Compute <math>B</math> the set of triples affected by the updates in step 1.</p> $B := \{(x, P, y) \mid (x, P, y) \in d \wedge (y, owl:sameAs, z) \in d \wedge z \in A(r) \wedge y \neq \eta^r(z)\}.$ <p>2.2. Compute <math>\Delta_3</math> the new state of the triples in <math>B</math>;</p> $\Delta_3 := \emptyset;$ <p>For each <math>(x, P, y)</math> in <math>B</math> do {</p> <p>Let <math>Q</math> be the property of the exported views used in the fusion rule of <math>P</math>.</p> $w := \{(u, Q, z) \mid (u, Q, z) \in t \wedge (x, owl:sameAs, u) \in d\};$ $\Delta_3 := \Delta_3 \cup \{(x, P, q) \mid (x, P, q) \in I[w](P)\};$ <p>2.3. <math>d := d - B \cup \Delta_3</math>;</p>
--

To illustrate this strategy, let  $\mu$  be the following update on *DBPedia*:

1. WITH <http://dbpedia.org>
2. DELETE { ?x dc:title "W. B. S. S." }
3. INSERT { ?x dc:title "Wanna Be Startin' Somethin'" }
4. WHERE { ?x rdfs:type dbpedia-owl:Single .
5.         ?x dc:title "W. B. S. S." }

**Phase 1:** (Executed by the *DBPedia* View Controller)

**Step 1:** From the mapping rules we have that  $u$  is relevant to class *mo:Track* of exported view *DBPEDIA\_EV*.

**Step 2:** Considering the state of *DBPedia* in Figure 3(a), the update  $u$  changes the title of the instance *dbpedia:WBSS* to “Wanna Be Startin' Somethin”.

**Step 3:** Send  $R=\{ \text{dbpedia:WBSS}, \psi_{\text{dbpedia:WBSS}} \}$  to the *LDM View Controller*.

**Phase 2:** (Executed by the *LDM View Controller*)

**Step 1:** Update the exported view *DBPEDIA\_EV* and the sameAs linkset views *ML1*. Based on the specification of *ML1*, the resources *dbpedia:WBSS* and *db:track/WBSS* are computed as equivalent. Therefore a new sameAs link is created in *ML1*.

{(dbpedia:WBSS, owl:sameAs, db:track/WBSS)}

**Step 2** The procedure *Effect*(*dbpedia:WBSS*) is executed as follows:

**Step 2.1:**

$A=\{\text{dbpedia:WBSS}, \text{db:track/WBSS}\}$

$\varepsilon_I=\{\text{dbpedia:WBSS}, \text{db:track/WBSS}\}$  is new equivalence class.

$\Delta_1$  is obtained from the fusion of the resources *dbpedia:WBSS*, *db:track/WBSS*.

$\Delta_2$  contains all triples in  $d$  relating to the old state of resources *dbpedia:WBSS* and *db:track/WBSS*.

$d = (d - \Delta_2) \cup \Delta_1$ ;

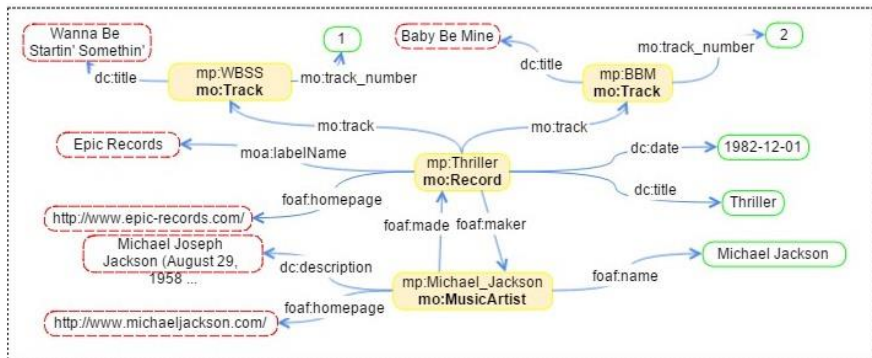
**Step 2.2.**

$B=\{(\text{mp:Thriller}, \text{mo:track}, \text{mp:DB\_WBSS}), (\text{mp:Thriller}, \text{mo:track}, \text{mp:MB\_WBSS})\}$

$\Delta_3=\{(\text{mp:Thriller}, \text{mo:track}, \text{mp:WBSS})\}$

$d := d - B \cup \Delta_3$ .

The new state of the mashup view is shown in Figure 6.



**Fig. 6.** New state of the mashup view

## 4 Related Work

The problem of incremental view maintenance has been extensively studied in the literature. However, for the most part, view is defined over a single data source, e.g., for relational views [6], for object-oriented views [15], for semi-structured views [1], for XLM Views [7], [23], and for RDF views [22]. None of the proposed techniques can be directly applied to data integration views.

The incremental maintenance of data integration views is address in [9], [10]. The focus in those works is the relational views. In [10], an algorithm for the incremental maintenance of outerjoin and match views was developed. In [9], algebraic change propagation algorithms were developed for the maintenance of the outerjoin view. Despite of their important contributions, none of those techniques can be applied to LDM views, since the nature of linked data sources poses new challenges for dealing with the data fusion problem.

Recently, two frameworks were proposed for creating LDM views. The ODClean-Store framework [12] offers Linked Data fusion, dealing with inconsistencies. LDIF - Linked Data Integration Framework [19] implements a mapping language, deals with URIs remapping and uses named graphs for registering data provenance.

To the best of our knowledge, incremental maintenance of LDM views has not yet been addressed in any framework.

## 5 Conclusions

In this paper, we first proposed an ontology-based framework for specifying Linked Data Mashup views. In the framework, a Linked Data mashup view is formally specified with the help of exported views, linkset views, data fusion function and normalization function. The LDM view specification is used to automatically materialize the mashup view. Then, we outlined a strategy that uses the mashup view specification for incrementally maintain the mashup view.

Our strategy addressed the problem of dealing with changes in the set of sameAs links between IRIs from different data sources and with the very question of recomputing mashup property values, in the presence of updates on the data sources.

## 6 References

1. Abiteboul, S., McHugh, J., Rys, M., Vassalos, V., Wiener, J. L.: Incremental Maintenance for Materialized Views over Semistructured Data. In VLDB 1998, pp. 38–49 (1998)
2. Berners-Lee, T.: Linked Data, <http://www.w3.org/DesignIssues/LinkedData.html> (2006)
3. Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Cyganiak, R., Hellmann, S.: DBpedia - A crystallization point for the Web of Data. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, Issue 7, Pages 154–165 (2009)
4. Bleiholder, J., Naumann, F.: Data fusion. *ACM Comput. Surv.* 41(1), 1:1–1:41 (2009)
5. Casanova, M.A., Vidal, V.M.P., Lopes, G.R., Leme, L.A.P.P., Ruback, L.: On Materialized sameAs Linksets. In DEXA 2014 (2014)

6. Ceri, S. and Widom, J.: Deriving productions rules for incremental view maintenance. In VLDB 1991, pp. 577–589 (1991)
7. Dimitrova, K., El-Sayed, M., Rundensteiner, E.A.: Order-sensitive View Maintenance of Materialized XQuery Views. In ER 2003, pp. 144–157 (2003)
8. Endres, B. N.: Semantic Mashups. Springer, Heidelberg (2013)
9. Griffin, T., Libkin, L.: Algebraic change propagation for semijoin and outerjoin queries. In SIGMOD Record, vol. 27(3) (1998)
10. Gupta, A. and Mumick, I.S.: Materialized Views. MIT Press (2000)
11. Hanh, H.H., Tai, N.C., Duy, K.T., Dosam, H., Jason, J.J.: Semantic Information Integration with Linked Data Mashups Approaches. International Journal of Distributed Sensor Networks 2014, Article ID 813875 (2014)
12. Knap, T., Michelfeit, J., Daniel, J., Jerman, P., Rychnovsky, D., Soukup, T., Necasky, M.: ODCleanStore: A Framework for Managing and Providing Integrated Linked Data on the Web. In ISWC 2012, Posters & Demos Track (2012)
13. Kobilarov, G., Scott, T., Raimond, Y., Oliver, S., Sizemore, C., Smethurst, M., Bizer, C., Lee, R.: Media meets semantic web – How the BBC uses DBpedia and Linked Data to make connections. ESWC 2009, Springer-Verlag, Berlin, Heidelberg, pp. 723–737(2009)
14. Kondrak, G.: N-gram similarity and distance. In Proceedings of the Twelfth International Conference on SPIRE, pages 115–126, Buenos Aires, Argentina (2005)
15. Kuno, H. A. and Rundensteiner, E. A.: Incremental Maintenance of Materialized Object-Oriented Views in MultiView: Strategies and Performance Evaluation. In IEEE TDKE, vol. 10, no. 5, pp. 768–792 (1998)
16. Mendes, M., Mühleisen, H., Bizer, C.: Sieve: Linked Data Quality Assessment and Fusion. Invited paper at the LWDM 2012 (2012)
17. Raimond, Y., Abdallah, S., Sandler, M., Giasson, F.: The Music Ontology. In: International Conference on Music Information Retrieval, pp. 417–422 (September 2007)
18. Sacramento, E.R., Vidal, V.M.P., Macedo, J.A.F., Lóscio, B.F., Lopes, F.L.R., Casanova, M.A.: Towards Automatic Generation of Application Ontologies. JIDM 1(3): 535-550 (2010)
19. Schultz, A., Matteini, A., Isele, R., Mendes, P., Bizer, C., Becker, C.: LDIF - A Framework for Large-Scale Linked Data Integration. In WWW2012, Developers Track (2012)
20. Swartz, A.: MusicBrainz: A Semantic Web Service. IEEE Intelligent Systems, vol. 17, no. 1, pp. 76-77 (2002)
21. The DBpedia Ontology (2014), <http://wiki.dbpedia.org/Ontology2014>
22. Vidal, V.M.P., Casanova, M.A., Cardoso, D.S.: Incremental Maintenance of RDF Views of Relational Data. In ODBASE 2013, pp 572-587 (2013)
23. Vidal, V.M.P., Lemos, F.C.L., Araújo, V., Casanova M.A.: A Mapping-Driven Approach for SQL/XML View Maintenance. In ICEIS 2008, pp. 65-73 (2008)