

Transforming OWL 2 RL Schemas to Relational Schemas with Open-world or Closed-world Semantics

Lama Al Khuzayem
Data and Knowledge
Engineering Group
Department of Computing
Imperial College London
laa11@imperial.ac.uk

Yu Liu
Data and Knowledge
Engineering Group
Department of Computing
Imperial College London
yu.liu11@imperial.ac.uk

Peter McBrien
Data and Knowledge
Engineering Group
Department of Computing
Imperial College London
p.mcbrien@imperial.ac.uk

ABSTRACT

Storing and processing Semantic Web knowledge in **relational database management systems (RDBMSs)** is currently a growing interest in both academia and industry. In this paper, we present OWLRel, a database-driven ontology reasoner that supports sound reasoning of OWL 2 RL and can adhere to either the open-world assumption or the closed-world assumption. This is achieved by regarding the mapping of OWL 2 RL schemas to relational schemas as a two-phase process: (1) Convert the OWL 2 RL model into a logical relational model where the operational semantics of constraints are not specified. (2) Implement the OWL 2 RL constraints in the relational model either as triggers achieving an **open-world semantics (OWS)** approach, or as constraints achieving a **closed-world semantics (CWS)** approach.

Keywords

Ontologies, OWL 2 RL, Relational Databases, Database Triggers, Transactional Reasoning, BAV Transformations.

1. INTRODUCTION

With the increasing number of ontologies available on the web, a growing interest has arisen for developing systems that not only store ontologies expressed in the **web ontology language (OWL)** [25] in a **relational database management system (RDBMS)**, but also perform reasoning over the instances to capture the **open-world semantics (OWS)** characteristics of ontologies [8, 27, 18]. Hence, the problem which is the focus of this paper, can be broken down into two parts. Firstly, how can we map the OWL schemas to relational schemas, and secondly how to perform reasoning over the instances held in this relational schema.

Most of the state-of-the-art proposals perform the transformation from knowledge model to data model at a high-

level, which is often referred to as **direct mapping (DM)**. This involves specifying specific mappings from constructs in OWL to constructs in the relational model, or *vice versa*. By contrast, in this paper, we show how to translate schemas expressed in OWL 2 RL (a profile of the most recent version of the web ontology language OWL 2 [26] which aims for rule-based implementations) into relational schemas via an intermediary low-level **hypergraph data model (HDM)** [19].

Using the HDM as an intermediate language has several advantages. First is the benefit of abstracting the high-level constructs in OWL and the relational model to a set of core low-level elemental modelling primitives (nodes, edges, and constraints) [21], making apparent what are the precise differences in the logical semantics of the OWL and relational models. The second benefit is that it has already been shown how to use the HDM to translate between ER, ORM, relational, UML and XML modelling languages [4, 13], and hence the work mapping OWL to relational models presented in this paper will transitively also allow OWL to be mapped to other data models. Thirdly, the **intermodel transformations** described in [4] expressed as BAV mappings [14] are based on using five types of HDM equivalence rules that transform one HDM schema into an equivalent HDM schema, and these equivalence rules may be directly applied to the new task of mapping between the OWL and relational models.

RDBMSs provide a means of storage for ontologies that is able to process ontologies with large number of individuals in a manner more efficient than using Tableau-based reasoners such as Pellet [20] and Fact++ [23]. A drawback, however, is that we are forced to make the **unique name assumption (UNA)**. A challenge we meet in this paper is that normal semantics of an RDBMS is **closed-world semantics (CWS)**, but that of OWL is OWS. We meet this challenge by building a single framework that maps the OWL OWS to either an equivalent OWS in the RDBMS, or a CWS in the RDBMS.

In general, the process of reasoning may be broken down into: classification of the **terminology box (T-Box)** and type-inference over the **assertion box (A-Box)**. The classification of the T-Box is not the concern of this paper, and our implementation work has relied on an OWL reasoner, Pellet, to perform this task. Type-inference, in the context of databases, is the process of relating values for tables/columns to values in other tables/columns. How this is executed will depend on the choice of OWS or CWS. Consid-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

ering the **description logic (DL)** [3] rule $\text{Father} \equiv \text{Man} \sqcap \text{Parent}$, in an OWS, we could deduce the following inference rules:

1. Inserting **John** as an instance of **Father** would then cause **John** to be inserted into both **Man** and **Parent** if not already present.
2. Inserting **John** into **Man** would also cause **John** to be inserted into **Father** if **John** was already in **Parent**.
3. Inserting **John** into **Parent** would also cause **John** to be inserted into **Father** if **John** was already in **Man**.

By contrast, in a CWS:

1. An insert of **John** into **Father** would fail unless **John** was already an instance of both **Man** and **Parent**.
2. An insert of **John** into **Man** would fail if **John** was already present in **Parent** but not present in **Father**.
3. An insert of **John** into **Parent** would fail if **John** was already present in **Man** but not present in **Father**.

A further objective of reasoning in an RDBMS is to achieve **transactional reasoning** [16], where the results of reasoning derived from data changed by the database operations should be available as part of the atomic action of the transaction. To achieve transactional reasoning, two approaches were identified in [16]: **view-based reasoning (VBR)**, where rules are used to derive the result of reasoning as each query is executed over the database, and **trigger-based reasoning (TBR)** where triggers (active rules) are used to materialise the result of reasoning at data insertion time, and queries simply read the materialised views. The advantages and disadvantages of using views or materialised views are well known, and each serve different real world requirements. Specific advantages of the TBR approach include that it is very fast at query processing, and that reasoning is incremental, since insertion of data only requires the change to reasoned results to be computed.

To illustrate the concept of transactional reasoning, let us consider again the DL rule $\text{Father} \equiv \text{Man} \sqcap \text{Parent}$. If **John** was already recorded as a **Man**, then if a transaction added that **John** was a **Parent**, then any query on the result of the transaction should also be able to view that **John** is a **Father**. However, most approaches for storing ontology data in an RDBMS will make the process of reasoning be detached from transaction processing of the data. Specifically, in our example, after the transaction that added **John** was a **Parent**, the database could be queried and find that he was not a **Father** until a separate process of reasoning had derived that fact. Hence, we would say that these approaches do not support transactional reasoning.

The novelty of this paper, which distinguishes it from previous work in the area (reviewed in Section 11) can be summarised as follows:

1. We represent OWL 2 RL axioms and constructs in the HDM.
2. We produce lossless transformations of OWL 2 RL knowledge bases to relational schemas. The transformations are lossless in the sense that all the semantics of the OWL 2 RL model are presented in the relational model.

3. Expressing the transformations as BAV mappings means we have a precise definition of the equivalence between schemas, and can map data back and forth between the OWL 2 RL and relational schemas.
4. We provide a single mapping approach for handling the axioms of OWL 2 RL schemas which can result in either CWS or OWS in the RDBMS.
5. We support sound transactional reasoning of OWL 2 RL ontologies with large A-Boxes.

The remainder of this paper is structured as follows. In Section 2 we detail OWLRel’s architecture and in Section 3 we review the HDM. In Section 4, we show the complete representations of OWL 2 RL constructs in HDM and provide a transformation example followed by explaining the intermodel transformations in Section 5. The process of mapping the resulting HDM schema to an RDBMS is explained in Section 6. We present our novel approach of transforming OWL 2 RL constructs to relational schemas via the HDM under OWS and CWS in Sections 7 and 8 respectively. Section 9 highlights the operational semantics differences between OWS and CWS. We show our system’s evaluation results in Section 10 and review related work in the area in Section 11. Finally, in Section 12, we state our conclusions and elaborate on future directions.

2. OWLREL ARCHITECTURE

OWLRel exploits this division of the OWL reasoning process: classification of the T-Box and type-inference of the A-Box to build a reasoning system in several steps, for which the overall design is illustrated in Figure 1. The process is as follows:

- ① OWLRel uses the OWL API [7] to load an OWL ontology file and then separates the T-Box and the A-Box.
- ② The T-Box is passed into a reasoner for classification. Since this step is conducted only once as a process of building up the database schema, the objective is to have as complete reasoning results as possible. The OWLRel system supports any reasoner that works with OWL API, and our evaluation results in Section 10 are based on using Pellet.
- ③ The fully classified T-Box is mapped to a HDM schema called the HDM-OWL2RL schema. The BAV mapping approach used in this transformation describes the the mappings between schemas on a construct by construct basis, as a pathway of primitive transformation steps applied in sequence. Details of this transformation are found in Section 4.
- ④ OWLRel then applies BAV equivalence rules on the HDM-OWL2RL schema to produce an equivalent HDM schema called HDM-OWL2RL+rel. The purpose of this step is to prepare the HDM schema which will be mapped to a relational model. Details can be found in in Section 5.
- ⑤ A core relational database schema that resembles part of the HDM-OWL2RL+rel schema is created in an RDBMS which contains only tables, columns, and primary keys as discussed in Section 6. The rest of the HDM constraints are treated separately under OWS and CWS as in the following two alternative steps.
- ⑥ Under OWS, OWLRel generates a set of SQL statements to create triggers for the HDM constraints. The OWLRel system uses an external file of template triggers to generate

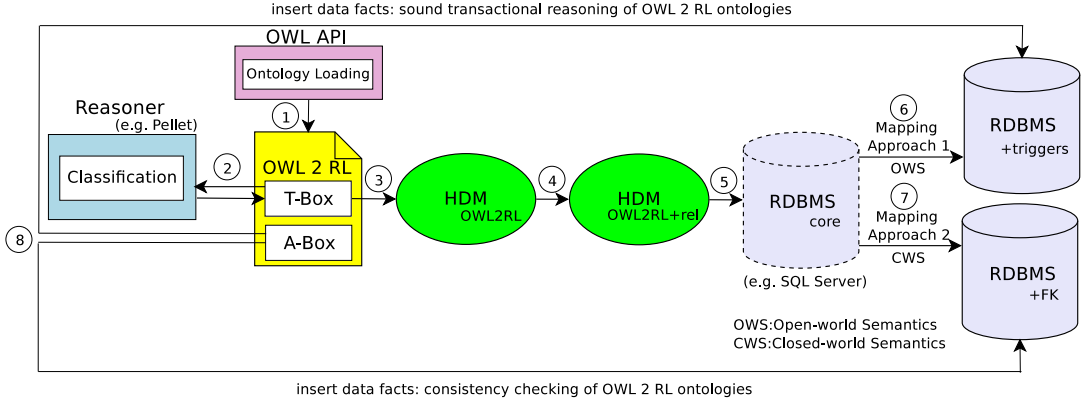


Figure 1: OWL 2 RL to Relational Mapping Approaches

the SQL triggers from the logical triggers which are presented in Section 7.

⑦ Under CWS, OWLRel creates SQL constraints for those unconverted constraints from HDM-OWL2RL+rel as explained in Section 8.

⑧ Once the SQL schema plus triggers or constraints have been created, the A-Box reasoning in an OWS context can take place by inserting data using SQL statements, which results in invoking the triggers that perform type inference and insert derived data into the database. If an inconsistency arises from a particular reasoning step, the transaction in which it occurs will be rolled back by removing any reasoned data as well as the original inserted data. In a CWS scenario, as the A-Box is being inserted to the database, constraints will check the consistency of the database. Hence, any data that violates these constraints will be rejected, rather than performing type inference.

The result of these steps is a stand-alone fully type inferring system in an RDBMS under **OWS**, which produces **sound reasoning** for ontologies expressed in the OWL 2 RL profile or ensuring the **consistency** of the ontologies in a **CWS** setting.

3. HDM OVERVIEW

In this Section, we provide a brief review of the HDM defined in [4, 19]. An HDM schema S is a triple $\langle Nodes, Edges, Cons \rangle$ where:

- $Nodes$ is a set of nodes in the graph such that $Nodes \subseteq \{node: \langle \langle n_n \rangle \mid n_n \in Names \rangle\}$. Given data instance I of the HDM schema, the function $Ext_I(node: \langle \langle n_n \rangle \rangle)$ gives the set of data values associated with $node: \langle \langle n_n \rangle \rangle$.
- $Schemes = Nodes \cup Edges$
- $Edges$ is a set of edges in the graph such that $Edges \subseteq \{edge: \langle \langle n_e, s_1, \dots, s_n \rangle \rangle \mid n_e \in Names \cup \{_ \} \wedge s_1 \in Schemes \wedge \dots \wedge s_n \in Schemes\}$. Note that this defines a hypergraph edge since edges can connect more than two nodes, and also defines a nested graph, since edges can also connect to other edges as follows:
 $\forall I. Ext_I(edge: \langle \langle n_e, s_1, \dots, s_n \rangle \rangle) \subseteq Ext_I(s_1) \times \dots \times Ext_I(s_n)$
- $Cons \subseteq \{c(s_1, \dots, s_n) \mid c \in Funcs \wedge s_1 \in Schemes \wedge \dots \wedge s_n \in Schemes\}$ $Cons$ is a set of boolean-valued functions (constraints) whose variables are members of

$Schemes$ and where the set of functions $Funcs$ forms the HDM constraint language. The set of constraints used in this paper are as follows:

- $cons: \langle \langle \subseteq, s_1, s_2 \rangle \rangle$ is the **inclusion** constraint which states that scheme s_1 is always a subset of scheme s_2 : $\forall I. Ext_I(s_1) \subseteq Ext_I(s_2)$
- $cons: \langle \langle \not\subseteq, s_1, \dots, s_n \rangle \rangle$ is the **exclusion** constraint which states that all the associate schemes are disjoint from each other:
 $\forall I. 1 \leq x \leq n, 1 \leq y \leq n, x \neq y. Ext_I(s_x) \cap Ext_I(s_y) = \emptyset$
- $cons: \langle \langle \cup, s_1, \dots, s_n, s \rangle \rangle$ is the **union** constraint stating scheme s is the union of schemes s_1, \dots, s_n :
 $\forall I. Ext_I(s) = Ext_I(s_1) \cup \dots \cup Ext_I(s_n)$
- $cons: \langle \langle \triangleright, s_1, \dots, s_m, s \rangle \rangle$ is the **mandatory** constraint stating that every combination of values that appears in schemes s_1, \dots, s_m must appear in the edge s connecting those schemes.
- $cons: \langle \langle \triangleleft, s_1, \dots, s_m, s \rangle \rangle$ is the **unique** constraint stating that every combination of values that appears in schemes s_1, \dots, s_m must appear no more than once in the edge s connecting those schemes.
- $cons: \langle \langle \xrightarrow{id}, s_1, s \rangle \rangle$ is the **reflexive** constraint, stating that for any value in s_1 must appear reflexively in the edge s that connects to s_1 , so that $\forall I. Ext_I(s_1) \times Ext_I(s_1) \subseteq Ext_I(s)$
- $cons: \langle \langle |1, s \rangle \rangle$ is the **instance** constraint, stating that only one value may be stored in s such that $\forall I. |Ext_I(s)| = 1$

In addition to referring to schemes directly, constraints may also take joins and projections of schemes as arguments.

To illustrate the use of HDM to describe the semantics of high-level data models, consider the relational schema in Example 3.1 (where primary keys are underlined, and nullable column names are suffixed by a question mark). We will later show how this schema can be derived from the OWL 2 RL ontology listed in Figure 2.

EXAMPLE 3.1. Family Relational Database Schema

Person(id, spouse?)
 Man(id, wife?)
 Woman(id, husband?)
 Parent(id)
 JohnsChildren(id)
 hasChild(parent, child)
 hasParent(child, parent)
 hasGrandParent(grandchild, grandparent)
 hasAncestor(descendant, ancestor)

□

Using the approach from [12], we can translate the relational schema into a HDM schema as follows. For each relational table we create a node, and connect it via HDM edges to nodes created for each column. Hence for the *Woman* table we create:

node:⟨⟨Woman⟩⟩ edge:⟨⟨-, Woman, woman:id⟩⟩
 node:⟨⟨woman:id⟩⟩ edge:⟨⟨-, Woman, husband⟩⟩
 node:⟨⟨husband⟩⟩

Since values in columns cannot exist in isolation from rows in the table, we state that the nodes representing columns have a mandatory association with the edge connecting them to the nodes representing the table:

cons:⟨⟨▷, node:⟨⟨woman:id⟩⟩, edge:⟨⟨-, Woman, woman:id⟩⟩⟩
 cons:⟨⟨▷, node:⟨⟨husband⟩⟩, edge:⟨⟨-, Woman, husband⟩⟩⟩

and since each column of a relation may take only one value, it follows that the node representing table is connected to the same edges by a unique constraint:

cons:⟨⟨◁, node:⟨⟨Woman⟩⟩, edge:⟨⟨-, Woman, woman:id⟩⟩⟩
 cons:⟨⟨◁, node:⟨⟨Woman⟩⟩, edge:⟨⟨-, Woman, husband⟩⟩⟩

If columns are not nullable, such as *woman.id*, we also state that the association of the table to the column edge is mandatory:

cons:⟨⟨▷, node:⟨⟨Woman⟩⟩, edge:⟨⟨-, Woman, woman:id⟩⟩⟩

and if the column is key, then we state that the same edge is reflexive:

cons:⟨⟨ $\overset{id}{\rhd}$, node:⟨⟨Woman⟩⟩, edge:⟨⟨-, Woman, woman:id⟩⟩⟩

which in combination with the unique and mandatory constraints has the consequence that the extent of the node representing the table becomes the set of key values of the table.

Note that you can view the HDM representation of the relational model (HDM nodes connected by HDM edges) as a forest of two-level trees, where the root of each tree is a node representing a table, the leaves of the tree are nodes representing columns, and HDM inclusion constraints exist only between leaf nodes. Other types of HDM constraints may be associated with the edges between nodes.

4. REPRESENTING OWL 2 RL IN HDM

OWL 2 RL is a syntactic subset of OWL 2 DL [26], and OWL 2 DL is an implementation of the DL *SROIQ(D)* with keys added. OWL 2 RL supports almost all OWL 2 axioms except for reflexive object properties and disjoint union expressions, and restricts the usage of some class expressions to make it possible to reason using rule-based engines with a complexity of **P**TIME-complete. Thus, it is well suited for applications that require scalable reasoning without losing too much expressivity.

We now discuss how OWL 2 RL constructs and axioms may be represented in HDM which corresponds to Step ③

in Figure 1. For conciseness, we only discuss some of those OWL 2 RL constructs listed in Tables 1 and 2, which are sufficient to describe how the OWL 2 RL knowledge base illustrated in Figure 2 can be translated into a HDM schema depicted in Figure 3.

All OWL classes are represented as HDM nodes. For example, class *Person* is represented as: node:⟨⟨Person⟩⟩

Object properties are represented as HDM edges with different HDM constraints depending on the type of the **objectProperty**. For example, the *hasSpouse* property is both a **symmetricProperty** and a **functionalProperty** as denoted in rules (26) and (27). In HDM we represent it as follows:

edge:⟨⟨hasSpouse, Person, Person⟩⟩
 cons:⟨⟨⊆, $\pi_{\langle\text{Person}\#2, \text{Person}\#1\rangle}$, ⟨⟨hasSpouse, Person, Person⟩⟩, ⟨⟨hasSpouse, Person, Person⟩⟩⟩
 cons:⟨⟨◁, Person#1, ⟨⟨hasSpouse, Person, Person⟩⟩⟩

Note in the above rule, where the edge *hasSpouse* links node *Person* with itself, we can disambiguate the first and second occurrence of *Person* with #1 and #2.

All OWL 2 RL axioms are represented as HDM constraints. For example, the **subclassOf** and **subPropertyOf** axioms denoted in rules (1) and (21) are represented as inclusion constraints (\subseteq) where the first element is subsumed by the second as follows:

cons:⟨⟨⊆, Man, Person⟩⟩
 cons:⟨⟨⊆, ⟨⟨hasHusband, Woman, Man⟩⟩, ⟨⟨hasSpouse, Person, Person⟩⟩⟩

If a class is a **complementOf** another class as given in rule (3) that is represented as an exclusion constraint (\neg) between the two classes and the union (\cup) of the two classes gives you the class *Thing* as follows:

cons:⟨⟨ \neg , Man, Woman⟩⟩
 cons:⟨⟨ \cup , Thing, Man, Woman⟩⟩

The **someValuesFrom** construct denoted in rule (4) is resembled in HDM as a node connected to the filler node (i.e. class *Person*) with an edge and a mandatory (\triangleright) constraint and making the newly created edge subset of the edge resembling the same property as follows:

node:⟨⟨ \exists hasChild.Person⟩⟩
 edge:⟨⟨hasChildIE, \exists hasChild.Person, Person⟩⟩
 cons:⟨⟨▷, \exists hasChild.Person, ⟨⟨hasChildIE, \exists hasChild.Person, Person⟩⟩⟩
 cons:⟨⟨⊆, ⟨⟨hasChildIE, \exists hasChild.Person, Person⟩⟩, ⟨⟨hasChild, Person, Person⟩⟩⟩

The **hasValues** construct denoted in rule (5) is represented by a node for the individual *John* with an attached cardinality constraint to 1 as follows:

node:⟨⟨John⟩⟩
 cons:⟨⟨|1|, John⟩⟩

Then, similar to representing \exists *hasChild.Person*, we create a node $\langle\langle\exists$ hasParent.John⟩⟩ and connect it via an edge to the node *John* and make it mandatory on the edge. We then represent the **equivalentClass** construct denoted in the same rule by two inclusion constraints (\subseteq) in both directions to make the first class subset of the second, and the second subset of the first as follows:

cons:⟨⟨⊆, \exists hasParent.John, JohnsChildren⟩⟩
 cons:⟨⟨⊆, JohnsChildren, \exists hasParent.John⟩⟩

The fact that one property is the **inverseOf** another as shown in rule (24) is represented by an inclusion constraint between the property and the projection (π) of its inverse in the reverse order as follows:

cons:⟨⟨⊆, $\pi_{\langle\text{hasHusband, Man, Woman}\rangle}$, ⟨⟨hasWife, Man, Woman⟩⟩⟩
 cons:⟨⟨⊆, $\pi_{\langle\text{hasWife, Woman, Man}\rangle}$, ⟨⟨hasHusband, Woman, Man⟩⟩⟩

$\text{Man} \sqsubseteq \text{Person}$	(1)	$\top \sqsubseteq \forall \text{hasWife}^-. \text{Man}$	(11)	$\text{hasHusband} \sqsubseteq \text{hasSpouse}$	(21)
$\text{Woman} \sqsubseteq \text{Person}$	(2)	$\top \sqsubseteq \forall \text{hasSpouse}. \text{Person}$	(12)	$\text{hasParent} \sqsubseteq \text{hasGrandParent}$	(22)
$\text{Man} \equiv \neg \text{Woman}$	(3)	$\top \sqsubseteq \forall \text{hasSpouse}^-. \text{Person}$	(13)	$\text{hasGrandParent} \sqsubseteq \text{hasAncestor}$	(23)
$\exists \text{hasChild}. \text{Person} \sqsubseteq \text{Parent}$	(4)	$\top \sqsubseteq \forall \text{hasParent}. \text{Person}$	(14)	$\text{hasHusband} \equiv \text{hasWife}^-$	(24)
$\text{JohnsChildren} \equiv \exists \text{hasParent}. \{ \text{John} \}$	(5)	$\top \sqsubseteq \forall \text{hasParent}^-. \text{Person}$	(15)	$\text{hasParent} \circ \text{hasParent} \equiv \text{hasGrandParent}$	(25)
$\top \sqsubseteq \forall \text{hasChild}. \text{Person}$	(6)	$\top \sqsubseteq \forall \text{hasAncestor}. \text{Person}$	(16)	$\text{hasSpouse}^- \sqsubseteq \text{hasSpouse}$	(26)
$\top \sqsubseteq \forall \text{hasChild}^-. \text{Person}$	(7)	$\top \sqsubseteq \forall \text{hasAncestor}^-. \text{Person}$	(17)	$\top \sqsubseteq \leq 1 \text{hasSpouse}$	(27)
$\top \sqsubseteq \forall \text{hasHusband}. \text{Man}$	(8)	$\top \sqsubseteq \forall \text{hasGrandParent}. \text{Person}$	(18)	$\top \sqsubseteq \leq 1 \text{hasHusband}$	(28)
$\top \sqsubseteq \forall \text{hasHusband}^-. \text{Woman}$	(9)	$\top \sqsubseteq \forall \text{hasGrandParent}^-. \text{Person}$	(19)	$\top \sqsubseteq \leq 1 \text{hasWife}$	(29)
$\top \sqsubseteq \forall \text{hasWife}. \text{Woman}$	(10)	$\text{hasWife} \sqsubseteq \text{hasSpouse}$	(20)	$\text{hasAncestor} \circ \text{hasAncestor} \sqsubseteq \text{hasAncestor}$	(30)

(a) The T-Box of the OWL 2 RL Family Knowledge Base

$\text{Man}(\text{John})$	(31)	$\text{hasHusband}(\text{Mary}, \text{John})$	(34)
$\text{Woman}(\text{Mary})$	(32)	$\text{hasParent}(\text{Lewis}, \text{Albert})$	(35)
$\text{hasAncestor}(\text{Michael}, \text{Alex})$	(33)	$\text{hasParent}(\text{Albert}, \text{Alex})$	(36)

(b) The A-Box of the OWL 2 RL Family Knowledge Base

Figure 2: A Family Knowledge Base Expressed in OWL 2 RL

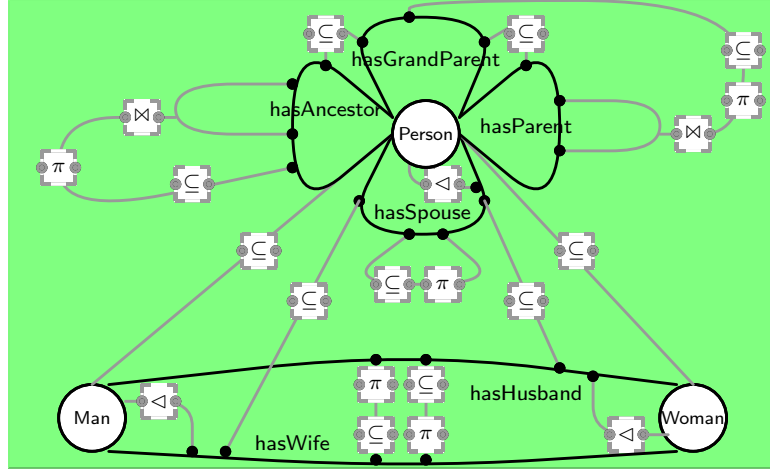


Figure 3: HDM Representation of the OWL 2 RL Family T-Box Rules (1-2) & (8-30)

A **propertyChain** such as the one given in rule (25) is represented by making the projection (π) of joining (\bowtie) the subproperties subset of itself as follows:

edge: $\langle\langle \text{hasGrandParent}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle$

cons: $\langle\langle \sqsubseteq, \pi_{\langle \text{Person}\#1, \text{Person}\#4 \rangle} \langle\langle \text{hasParent}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle$

$\bowtie_{\text{Person}\#2=\text{Person}\#3} \langle\langle \text{hasParent}, \text{Person}\#3, \text{Person}\#4 \rangle\rangle, \langle\langle \text{hasGrandParent}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle \rangle\rangle$

Similarly, a **transitiveProperty** such as the one denoted in rule (30) is represented as follows:

edge: $\langle\langle \text{hasAncestor}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle$

cons: $\langle\langle \sqsubseteq, \pi_{\langle \text{Person}\#1, \text{Person}\#4 \rangle} \langle\langle \text{hasAncestor}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle$

$\bowtie_{\text{Person}\#2=\text{Person}\#3} \langle\langle \text{hasAncestor}, \text{Person}\#3, \text{Person}\#4 \rangle\rangle, \langle\langle \text{hasAncestor}, \text{Person}\#1, \text{Person}\#2 \rangle\rangle \rangle\rangle$

Note that in the HDM diagram, HDM nodes are represented by white circles with thick outlines, and HDM edges are represented by thick black lines. The HDM constraint language is represented by grey dashed boxes connected by grey lines to the nodes and edges to which the constraint applies. Edges pass through black circles in a straight line, hence any edge or constraint applying to an edge meets that edge at an angle.

5. HDM TRANSFORMATIONS

Step ④ in Figure 1 performs a type of normalisation on the HDM-OWL2RL schema, to produce an equivalent HDM-OWL2RL+rel schema that can be directly mapped into a relational schema. This normalisation process is important to overcome the fundamental differences between the two modelling languages. On the one hand, OWL 2 RL is a knowledge model that has the notion of classes, properties. On the other hand, relational, is a key-based data model that has tables, columns, primary key (PK) and foreign key (FK) constraints.

Our three step process converts the HDM graph representing OWL 2 RL into an equivalent graph (in terms of information capacity) that can then be mapped into a relational schema. This uses a set of HDM equivalence mappings presented in [4].

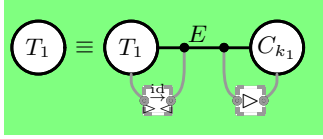
(A) Transform the implied object identifiers (OID) of OWL classes into explicit keys to form PK columns for their respective tables in the relational model.

This can be achieved using a HDM graph equivalence

Table 1: HDM Representations for OWL 2 RL Class Axioms

OWL 2 Construct	DL Syntax	HDM Representation
Class	C	node: $\langle\langle C \rangle\rangle$
subclassOf	$C \sqsubseteq D$	cons: $\langle\langle \sqsubseteq, C, D \rangle\rangle$
equivalentClass	$C \equiv D$	cons: $\langle\langle \sqsubseteq, C, D \rangle\rangle$, cons: $\langle\langle \sqsubseteq, D, C \rangle\rangle$
classDisjointWith	$C \sqcap D \sqsubseteq \perp$	cons: $\langle\langle \not\sqsubseteq, C, D \rangle\rangle$
complementOf	$C \equiv \neg D$	cons: $\langle\langle \not\sqsubseteq, C, D \rangle\rangle$, cons: $\langle\langle \cup, C, D, \text{Thing} \rangle\rangle$
allValuesFrom	$\forall P.D$	node: $\langle\langle \neg D \rangle\rangle$, cons: $\langle\langle \not\sqsubseteq, D, \neg D \rangle\rangle$, cons: $\langle\langle \cup, D, \neg D, \text{Thing} \rangle\rangle$ node: $\langle\langle \exists P.\neg D \rangle\rangle$, edge: $\langle\langle P_IE, \exists P.\neg D, \neg D \rangle\rangle$, cons: $\langle\langle \triangleright, \exists P.\neg D, \langle\langle P_IE, \exists P.\neg D, \neg D \rangle\rangle \rangle\rangle$ node: $\langle\langle \forall P.D \rangle\rangle$, cons: $\langle\langle \not\sqsubseteq, \forall P.D, \exists P.\neg D \rangle\rangle$, cons: $\langle\langle \cup, \forall P.D, \exists P.\neg D, \text{Thing} \rangle\rangle$
someValuesFrom	$\exists P.D$	node: $\langle\langle \exists P.D \rangle\rangle$, edge: $\langle\langle P_IE, \exists P.D, D \rangle\rangle$ cons: $\langle\langle \triangleright, \exists P.D, \langle\langle P_IE, \exists P.D, D \rangle\rangle \rangle\rangle$, cons: $\langle\langle \sqsubseteq, \langle\langle P_IE, \exists P.D, D \rangle\rangle, \langle\langle P, D, D \rangle\rangle \rangle\rangle$
hasValue	$\exists P.\{a\}$	node: $\langle\langle \exists P.a \rangle\rangle$, edge: $\langle\langle P_IE, \exists P.a, a \rangle\rangle$ cons: $\langle\langle \triangleright, \exists P.a, \langle\langle P_IE, \exists P.a, a \rangle\rangle \rangle\rangle$, cons: $\langle\langle \sqsubseteq, \langle\langle P_IE, \exists P.a, a \rangle\rangle, \langle\langle P, D, D \rangle\rangle \rangle\rangle$
oneOf	$\{a_1, \dots, a_n\}$	node: $\langle\langle a_1_a_i_a_n \rangle\rangle$, cons: $\langle\langle \cup, a_1, \dots, a_n, a_1_a_i_a_n \rangle\rangle$
maxCardinality	$\leq nP$	node: $\langle\langle \leq nP \rangle\rangle$, edge: $\langle\langle P_IE, \leq nP, D \rangle\rangle$, cons: $\langle\langle \triangleright^n, \leq nP, \langle\langle P_IE, \leq nP, D \rangle\rangle \rangle\rangle$
intersectionOf	$C \sqcap D$	node: $\langle\langle C - D \rangle\rangle$, cons: $\langle\langle \sqsubseteq, C - D, C \rangle\rangle$, node: $\langle\langle C \cup D \rangle\rangle$, cons: $\langle\langle \cup, C - D, D, C \cup D \rangle\rangle$ cons: $\langle\langle \cup, C, D, C \cup D \rangle\rangle$, cons: $\langle\langle \cup, C - D, D \rangle\rangle$, node: $\langle\langle C \cap D \rangle\rangle$, cons: $\langle\langle \sqsubseteq, C \cap D, C \rangle\rangle$ cons: $\langle\langle \not\sqsubseteq, C \cap D, C - D \rangle\rangle$, cons: $\langle\langle \cup, C \cap D, C - D, C \rangle\rangle$
namedIndividual	a	node: $\langle\langle a \rangle\rangle$, cons: $\langle\langle 1, a \rangle\rangle$

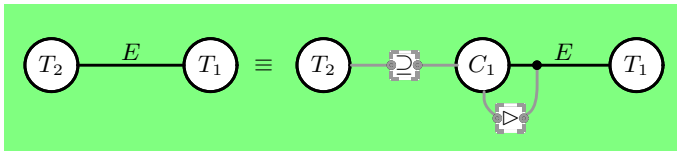
transformation, where the function: $\text{inverse_identity_node_merge}(\langle\langle T_1 \rangle\rangle, \langle\langle C_{k_1} \rangle\rangle)$ will take an existing node $\langle\langle T_1 \rangle\rangle$, and creates a new node $\langle\langle C_{k_1} \rangle\rangle$ connected to $\langle\langle T_1 \rangle\rangle$ by an edge. The edge has constraints that ensure that each instance of $\langle\langle T_1 \rangle\rangle$ appears at least once (\triangleleft), at most once (\triangleright) and reflexively ($\overset{id}{\rightarrow}$) in the edge, so that the contents of $\langle\langle C_{k_1} \rangle\rangle$ must be identical to $\langle\langle T_1 \rangle\rangle$. This is illustrated by:



Applying this step on the HDM OWL 2 RL schema in Figure 3 would be achieved by:
 $\text{inverse_identity_node_merge}(\langle\langle \text{Person} \rangle\rangle, \langle\langle \text{Person: id} \rangle\rangle)$
 $\text{inverse_identity_node_merge}(\langle\langle \text{Woman} \rangle\rangle, \langle\langle \text{Woman: id} \rangle\rangle)$
 $\text{inverse_identity_node_merge}(\langle\langle \text{Man} \rangle\rangle, \langle\langle \text{Man: id} \rangle\rangle)$
 and would result in HDM constructs for tables and their keys as described in Section 3.

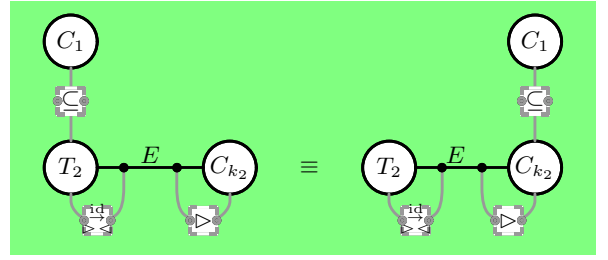
(B) Convert HDM edges representing OWL properties between OWL classes into the HDM equivalent of a column with a foreign key.

This process involves using two HDM equivalences. The first equivalence mapping: $\text{inverse_inclusion_merge}(\langle\langle E, T_1, T_2 \rangle\rangle, \langle\langle C_1 \rangle\rangle)$ uses the $\langle\langle E, T_1, T_2 \rangle\rangle$ edge to identify those members of $\langle\langle T_1 \rangle\rangle$ that participate in the edge, and put them in a new node $\langle\langle C_1 \rangle\rangle$ that is a subset of $\langle\langle T_1 \rangle\rangle$ as illustrated below.



Once this equivalence mapping has been performed, the subset $\langle\langle \sqsubseteq, \langle\langle T_1 \rangle\rangle, \langle\langle C_1 \rangle\rangle \rangle\rangle$ is used in another mapping to be

redirected to node $\langle\langle C_{k_2} \rangle\rangle$ representing the key of table represented by $\langle\langle T_2 \rangle\rangle$. This is generated by a second equivalence mapping $\text{redirect_inclusion_constraint}(\langle\langle \sqsubseteq, \langle\langle T_1 \rangle\rangle, \langle\langle C_1 \rangle\rangle \rangle\rangle, \langle\langle C_{k_2} \rangle\rangle)$



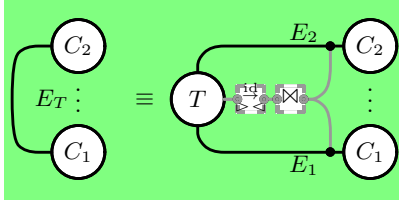
We can apply these two equivalences to the edge: $\langle\langle \text{hasSpouse, Person, Person} \rangle\rangle$ as follows:
 $\text{inverse_inclusion_merge}(\langle\langle \text{hasSpouse, Person, Person} \rangle\rangle, \langle\langle \text{spouse} \rangle\rangle)$
 $\text{redirect_inclusion_constraint}(\langle\langle \sqsubseteq, \text{Person, spouse} \rangle\rangle, \langle\langle \text{Person: id} \rangle\rangle)$
 which would generate a new node $\langle\langle \text{spouse} \rangle\rangle$ which has a mandatory constraint between it and the edge $\langle\langle \text{hasSpouse, Person, Person} \rangle\rangle$ and is a subset of the node $\langle\langle \text{Person} \rangle\rangle$. Applying this process to both ends of the edge: $\langle\langle \text{hasSpouse, Person, Person} \rangle\rangle$:
 $\text{inverse_inclusion_merge}(\langle\langle \text{hasParent, Person\#1, Person\#2} \rangle\rangle, \langle\langle \text{child} \rangle\rangle)$
 $\text{redirect_inclusion_constraint}(\langle\langle \sqsubseteq, \text{Person, child} \rangle\rangle, \langle\langle \text{Person: id} \rangle\rangle)$
 $\text{inverse_inclusion_merge}(\langle\langle \text{hasParent, Person\#2, Person\#1} \rangle\rangle, \langle\langle \text{parent} \rangle\rangle)$
 $\text{redirect_inclusion_constraint}(\langle\langle \sqsubseteq, \text{Person, parent} \rangle\rangle, \langle\langle \text{Person: id} \rangle\rangle)$
 generates two new nodes node: $\langle\langle \text{parent} \rangle\rangle$ and node: $\langle\langle \text{child} \rangle\rangle$ each with a subset constraint with the node representing the key of person.

(C) Represent non-functional properties in OWL as separate tables in the relational model.

The HDM graph equivalence mapping: $\text{identity_edge_merge}(\langle\langle T, C_1, C_2 \rangle\rangle, \langle\langle E_T \rangle\rangle)$ that replaces the edge by a new node $\langle\langle T \rangle\rangle$ connected to the two original nodes via two edges where there are \triangleleft , \triangleright , and $\overset{id}{\rightarrow}$ constraints between the new node and the natural join of the two new edges.

Table 2: HDM Representations for OWL 2 RL Property Axioms

OWL 2 Construct	DL Syntax	HDM Representation
objectProperty	P	edge: $\langle\langle P, C, D \rangle\rangle$
dataProperty	R	$\langle\langle \text{rdf:datatype} \rangle\rangle, \langle\langle R, C, \text{rdf:datatype} \rangle\rangle$
subPropertyOf	$P \sqsubseteq Q$	cons: $\langle\langle \sqsubseteq, P, Q \rangle\rangle$
equivalentProperty	$P \equiv Q$	cons: $\langle\langle \sqsubseteq, P, Q \rangle\rangle, \text{cons:} \langle\langle \sqsubseteq, Q, P \rangle\rangle$
propertyDisjointWith	$P \sqcap Q \sqsubseteq \perp$	cons: $\langle\langle \neg, P, Q \rangle\rangle$
inverseOf	$P \equiv Q^{-}$	cons: $\langle\langle \sqsubseteq, \pi_{\langle C_2, C_1 \rangle} \langle\langle P, C_1, C_2 \rangle\rangle, \langle\langle Q, C_2, C_1 \rangle\rangle \rangle\rangle$ cons: $\langle\langle \sqsubseteq, \pi_{\langle C_1, C_2 \rangle} \langle\langle Q, C_2, C_1 \rangle\rangle, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
symmetricProperty	$P \equiv P^{-}$	edge: $\langle\langle P, C, D \rangle\rangle$ cons: $\langle\langle \sqsubseteq, \pi_{\langle D, C \rangle} \langle\langle P, C, D \rangle\rangle, \langle\langle P, C, D \rangle\rangle \rangle\rangle$
transitiveProperty	$P \circ P \sqsubseteq P$	edge: $\langle\langle P, C_1, C_2 \rangle\rangle$ cons: $\langle\langle \sqsubseteq, \pi_{\langle P\#1.C_1, P\#2.C_2 \rangle} P \bowtie P, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
propertyChain	$P_1 \circ \dots \circ P_n \sqsubseteq P$	edge: $\langle\langle P_1, C_1, C_2 \rangle\rangle, \dots, \text{edge:} \langle\langle P_n, C_n, C_{n+1} \rangle\rangle$ edge: $\langle\langle P, C_1, C_{n+1} \rangle\rangle$ cons: $\langle\langle \sqsubseteq, \pi_{\langle C_1, C_{n+1} \rangle} P_1 \bowtie \dots \bowtie P_n, \langle\langle P, C_1, C_{n+1} \rangle\rangle \rangle\rangle$
functionalProperty	$T \sqsubseteq \leq 1P$	edge: $\langle\langle P, C_1, C_2 \rangle\rangle, \text{cons:} \langle\langle \triangleleft, C_1, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
inverseFunctionalProperty	$T \sqsubseteq \leq 1P^{-}$	edge: $\langle\langle P, C_1, C_2 \rangle\rangle, \text{cons:} \langle\langle \triangleleft, C_2, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
irreflexiveProperty	$T \sqsubseteq \neg \exists P.\text{self}$	edge: $\langle\langle P, C_1, C_2 \rangle\rangle, \text{edge:} \langle\langle Q, D_1, D_2 \rangle\rangle, \text{cons:} \langle\langle \neg, P, Q \rangle\rangle$ cons: $\langle\langle \triangleleft, D_1, Q \rangle\rangle, \text{cons:} \langle\langle \triangleright, D_1, Q \rangle\rangle, \text{cons:} \langle\langle \xrightarrow{\text{id}}, D_1, Q \rangle\rangle$
selfRestriction	$\exists P.\text{self}$	edge: $\langle\langle P, C, D \rangle\rangle, \text{cons:} \langle\langle \triangleright, C, P \rangle\rangle, \text{cons:} \langle\langle \triangleleft, C, P \rangle\rangle, \text{cons:} \langle\langle \xrightarrow{\text{id}}, C, P \rangle\rangle$
key		edge: $\langle\langle P, C, D \rangle\rangle, \text{cons:} \langle\langle \triangleright, C, P \rangle\rangle, \text{cons:} \langle\langle \triangleleft, C, P \rangle\rangle, \text{cons:} \langle\langle \xrightarrow{\text{id}}, C, P \rangle\rangle$



For example, applying the mapping to edge: $\langle\langle \text{hasParent}, \text{child}, \text{parent} \rangle\rangle$ is done by: `inverse_identity_edge_merge($\langle\langle \text{hasParent}, \text{child}, \text{parent} \rangle\rangle, \langle\langle \text{hasParent} \rangle\rangle$)` and would generate HDM objects equivalent to the `hasParent` table in Section 3.

The result of steps (A), (B) and (C) is an HDM graph that is a forest of two-level trees, with subset constraints linking the leaf nodes.

6. BUILDING AN RDB SCHEMA

Step ⑤ in Figure 1 builds a core relational database schema from the HDM OWL2RL+rel schema. By core, we mean it defines tables, columns, and primary keys that allow data to be held, without defining any triggers or constraints that would affect the open-world or closed-world interpretation of the data. In outline, the HDM OWL2RL+rel schema is a forest of two level trees, where root nodes are connected to a number of leaf nodes via edges.

From our methodology in previous section, we can derive three production rules for building a relational schema of the general form $HDM\ patterns \rightsquigarrow Relational\ construct$. For HDM nodes that are roots of the tree (and thus came from OWL 2 RL classes), we map them into a table of the same name:

Class: node: $\langle\langle T \rangle\rangle \rightsquigarrow \text{table:} \langle\langle T \rangle\rangle$

The leaf nodes that have been created by Steps (A) or (B) of the previous section will be mapped into columns:

Property: edge: $\langle\langle \neg, T, C_2 \rangle\rangle, \text{node:} \langle\langle C_2 \rangle\rangle, \text{cons:} \langle\langle \triangleright, C, \langle\langle \neg, T, C \rangle\rangle \rangle\rangle$
 $\rightsquigarrow \text{column:} \langle\langle T, C \rangle\rangle$

Finally, for those edges for which we have defined sufficient constraints to interpret it as the key (*i.e.* resulting from Step (A) above) we can define a primary key of the table.

edge: $\langle\langle \neg, T, C \rangle\rangle, \text{cons:} \langle\langle \triangleright, C, \langle\langle \neg, T, C \rangle\rangle \rangle\rangle, \text{cons:} \langle\langle \triangleright, T, \langle\langle \neg, T, C \rangle\rangle \rangle\rangle,$
cons: $\langle\langle \triangleleft, T, \langle\langle \neg, T, C \rangle\rangle \rangle\rangle, \text{cons:} \langle\langle \xrightarrow{\text{id}}, T, \langle\langle \neg, T, C \rangle\rangle \rangle\rangle$
 $\rightsquigarrow \text{primary_key:} \langle\langle T, C \rangle\rangle$

These rules will result in a relational schema identical to the one shown in Example 3.1.

7. HANDLING CONSTRAINTS IN OWS

We now outline the process of handling the unconverted HDM constraints under OWS which corresponds to Step ⑥ in Figure 1. The approach we follow was inspired by the work of [15, 9] in which we first derive logical triggers over the relational schema that is resulted from the previous section and then implement those logical triggers as SQL physical triggers on a particular target DBMS as will be illustrated in Section 9. The logical triggers are translated from the HDM constructors according to the general form:

$HDM\ construct \rightsquigarrow \text{when event if condition then action.}$

where *event* is the insertion process of a data value into a table. There are two types of event: if *event* is prefixed with $-$ then *condition* and *action* are executed before the insertion, whilst if *event* is prefixed with $+$ then *condition* and *action* are executed after the insertion. SQL **before triggers** (in pl/pgSQL) or **instead of triggers** (in Transact-SQL) are used to implement $-$ events, and **after triggers**, used for $+$ events. The *condition* is a Datalog query over the database, and *action* is either a data tuple to insert, *ignore* (ignoring this insertion that caused the trigger to execute) and *rollback* (rollback the transaction). The logical triggers can be translated into SQL physical triggers following the approach given in [15].

One basic rule deals with the notion that because of the open world nature of reasoning, we might repeatedly infer

the same fact, and thus we have to prevent duplicate updates to a table. This is implemented by the logical triggers:

```
class: node:⟨⟨C⟩⟩
  ~> when  $\neg C(x)$  if  $C(x)$  then ignore
```

This means that when a value is inserted into a table, before the actual insert is done, a check is made to determine if the value is already present in the table, and if so, the insert is ignored.

The logical trigger for **subClassOf** generates a trigger that implies that each insertion to one class will automatically generate the same insertion(s) to its super class(es). Therefore, the consistency of the relations between classes is maintained.

```
subClassOf: cons:⟨⟨ $\subseteq$ , C, D⟩⟩
  ~> when  $^+C(x)$  if true then  $D(x)$ 
```

A similar logical trigger for **subPropertyOf** is generated:

```
subPropertyOf: cons:⟨⟨ $\subseteq$ , P, Q⟩⟩
  ~> when  $^+P(x, y)$  if true then  $Q(x, y)$ 
```

Thus, for rules (1) and (22) we can derive the following logical triggers:

```
when  $^+Man(x)$  then Person(x)
when  $^+hasParent(x, y)$  then hasGrandParent(x, y)
```

The logical trigger for **complementOf** generates a trigger which implies that individuals in class D should not appear in class C and vice versa.

```
complementOf: cons:⟨⟨ $\neg$ , C, D⟩⟩
  cons:⟨⟨ $\cup$ , Thing, D, C⟩⟩
  ~> when  $\neg C(x)$  if  $D(x)$  then rollback
  ~> when  $\neg D(x)$  if  $C(x)$  then rollback
```

Thus, for rule (3) we can derive the following logical triggers:

```
when  $\neg Man(x)$  if Woman(x) then rollback
when  $\neg Woman(x)$  if Man(x) then rollback
```

The logical trigger for **equivalentClass** generates a trigger which implies that each insertion to one class will automatically generate the same insertion(s) to the other class.

```
equivalentClass: cons:⟨⟨ $\subseteq$ , C, D⟩⟩
  cons:⟨⟨ $\subseteq$ , D, C⟩⟩
  ~> when  $^+C(x)$  if true then  $D(x)$ 
  ~> when  $^+D(x)$  if true then  $C(x)$ 
```

The construct **someValuesFrom** $\exists P.D$ defines a set of individuals x that has atleast one tuple like (x, y) in P and y is in D . OWL 2 RL restricts the appearance of a **someValuesFrom** to be only in a subclass expression, so the logical trigger only contains the situation $\exists P.D \sqsubseteq C$:

```
someValuesFrom: cons:⟨⟨ $\subseteq$ ,  $\exists P.D$ , C⟩⟩
  ~> when  $^+P(x, y)$  if  $D(y)$  then  $C(x)$ 
  ~> when  $^+D(y)$  if  $P(x, y)$  then  $C(x)$ 
```

The logical trigger generates a trigger that checks for individuals that satisfy the existensial restriction and inserts them to table C. Thus, for rule (4) we can derive the following logical triggers:

```
when  $^+hasChild(x, y)$  if Person(y) then Parent(x)
when  $^+Person(y)$  if hasChild(x, y) then Parent(x)
```

The **hasValue** construct has two situations in which it may appear (subClass and a superClass) and consequently two logical triggers are generated respectively as follows:

```
hasValue: cons:⟨⟨ $\subseteq$ ,  $\exists P.a$ , C⟩⟩
  ~> when  $^+P(x, a)$  if true then  $C(x)$ 
hasValue: cons:⟨⟨ $\subseteq$ , C,  $\exists P.a$ ⟩⟩
  ~> when  $^+C(x)$  if true then  $P(x, a)$ 
```

If a certain class was a subclass of a **hasValue** expression,

then we should insert (x, a) to the table P whenever there is an insertion of x to C . On the other hand, if the expression is a subset of a class, a tuple (x, a) inserted to P will invoke the trigger which will insert x to the table C . Thus, for rule (5) we can derive the following logical triggers:

```
when  $^+hasParent(x, John)$  if true then JohnsChildren(x)
when  $^+JohnsChildren(x)$  if true then hasChild(x, John)
```

The logical trigger for **inverseOf** keeps the relation of inverse properties which means if P_1 and P_2 are inverse properties and (x, y) is a property instance of property P_1 , then (y, x) has to be an instance of property P_2 .

```
inverseOf: cons:⟨⟨ $\subseteq$ ,  $\pi_{\langle C_2, C_1 \rangle} \langle P_1, C_1, C_2 \rangle$ ,  $\langle P_2, C_2, C_1 \rangle$ ⟩⟩,
  cons:⟨⟨ $\subseteq$ ,  $\pi_{\langle C_1, C_2 \rangle} \langle P_2, C_2, C_1 \rangle$ ,  $\langle P_1, C_1, C_2 \rangle$ ⟩⟩
  ~> when  $^+P_1(x, y)$  if  $\neg P_2(y, x)$  then  $P_2(y, x)$ 
```

In this case, a trigger in table P_1 checks each of its inverse tuples and inserts them to property table P_2 if the inverse tuples do not exist in P_2 .

For a **symmetricProperty**, after inserting a tuple (x, y) to property P , the symmetric tuple (y, x) will be inserted to P by a trigger.

```
symmetricProperty: cons:⟨⟨ $\subseteq$ ,  $\pi_{\langle C_2, C_1 \rangle} \langle P, C_1, C_2 \rangle$ ,  $\langle P, C_1, C_2 \rangle$ ⟩⟩
  ~> when  $^+P(x, y)$  if true then  $P(y, x)$ 
```

For instance, for rule (26) we can derive the following logical trigger:

```
when  $^+hasSpouse(x, y)$  then hasSpouse(y, x)
```

For a **transitiveProperty**, after inserting a tuple (x, y) , it will try to find if tuple (y, z) exists. If so, it will then insert tuple (x, z) . Similarly, it will try to find if tuple (z, x) exists and then insert tuple (z, y) .

```
transitiveProperty: cons:⟨⟨ $\subseteq$ ,  $\pi_{\langle P \# 1, C_1, P \# 2, C_2 \rangle} P \bowtie P$ ,  $\langle P, C_1, C_2 \rangle$ ⟩⟩
  ~> when  $^+P(x, y)$  if  $P(y, z)$  then  $P(x, z)$ 
  if  $P(z, x)$  then  $P(z, y)$ 
```

For example, for rule (30) we can derive the following logical trigger:

```
when  $^+hasAncestor(x, y)$  if hasAncestor(y, z)
  then hasAncestor(x, z)
  if hasAncestor(z, x)
  then hasAncestor(z, y)
```

propertyChain allows for a property to be defined from the concatenation of two or more other properties. The logical triggers are as follows:

```
propertyChain: cons:⟨⟨ $\subseteq$ ,  $\pi_{\langle C_1, C_{n+1} \rangle} P_1 \bowtie \dots \bowtie P_n$ ,  $\langle P, C_1, C_{n+1} \rangle$ ⟩⟩
  ~> when  $^+P_1(x, y)$  if  $P'_{2,n}(y, z)$  then  $P(x, z)$ 
  ~> when  $^+P_n(y, z)$  if  $P'_{1,n-1}(x, y)$  then  $P(x, z)$ 
  ~> when  $^+P_i(p, q)$  if  $P'_{1,i-1}(x, p), P'_{i+1,n}(q, z)$ 
  then  $P(x, z)_{1 < i < n}$ 
```

$*P'_{m,n}(x, y) = \pi_{P_m.domain, P_n.range} \sigma_{P_j.range = P_{j+1}.domain} (P_m \times \dots \times P_n)_{m \leq j < n}$

The first and second situations mean that if there is an insertion to the first or last subchain, the trigger will treat the remaining subchains as a join unit and search data matched inside the unit. The third scenario handles the insertion to the middle subchains by creating two join units and then fetch matching data from them. For example, for rule (25) we can derive the following logical triggers:

```
when  $^+hasParent\#1(x, y)$  if hasParent\#2(y, z)
  then hasGrandParent(x, z)
when  $^+hasParent\#2(y, z)$  if hasParent\#1(x, y)
  then hasGrandParent(x, z)
```


8. HANDLING CONSTRAINTS IN CWS

In this Section, we show an alternative way of handling the unconverted HDM constraints under CWS which corresponds to Step ⑦ in Figure 1.

The basic idea for handling constraints in CSW is we create SQL constraints to check each HDM constraints after data is inserted. We follow the same approach we have used in OWS that, we first generate logical constraints and then show their physical implementation (SQL physical constraints) in Section 9. Logical constraints are translated from the HDM constructors according to productions rules of the general form:

HDM construct \rightsquigarrow **when event if condition then action.**

which is similar to the general form of logical triggers in OWS. The *event* is always happened after the data insertion (denoted by $^+$), since SQL constraints cannot be verified with no data inserted. The *condition* is logical check queries derived from an HDM constraint and the *action* is automatically performed by SQL Server either to *allow* or to *rollback* the insertions. Next, we demonstrate certain logical constraints for handling constraints in CWS.

The logical constraint for **subClassOf** will verify that if the data inserted to one class also in its super class(es).

subClassOf: $\text{cons}:\langle\langle \subseteq, C, D \rangle\rangle$
 \rightsquigarrow **when** $^+C(x)$ **if** $\neg D(x)$ **then rollback**

Similarly, the logical constraint for **subPropertyOf** is to verify all tuples inserted to a property exist in its super properties:

subPropertyOf: $\text{cons}:\langle\langle \subseteq, P, Q \rangle\rangle$
 \rightsquigarrow **when** $^+P(x, y)$ **if** $\neg Q(x, y)$ **then rollback**

For example, for rules (1) and (22) we can derive the following logical constraints:

when $^+\text{Man}(x)$ **if** $\neg\text{Person}(x)$ **then rollback**
when $^+\text{hasParent}(x, y)$ **if** $\neg\text{hasGrandParent}(x, y)$ **then rollback**

The logical constraints for **complementOf** (e.g. $C \equiv \neg D$) which are shown below only allow to insert data to the C if the data is not in the table D, and vice versa:

complementOf: $\text{cons}:\langle\langle \not\subseteq, C, D \rangle\rangle$
 $\text{cons}:\langle\langle \cup, \text{Thing}, D, C \rangle\rangle$
 \rightsquigarrow **when** $^+C(x)$ **if** $D(x)$ **then rollback**
 \rightsquigarrow **when** $^+D(x)$ **if** $C(x)$ **then rollback**

Thus, for rule (3) we can derive the following logical constraints:

when $^-\text{Man}(x)$ **if** $\text{Woman}(x)$ **then rollback**
when $^-\text{Woman}(x)$ **if** $\text{Man}(x)$ **then rollback**

The logical constraint for **equivalentClass** generates a check that verifies that if data is inserted to one table, it is also in the equivalent table of this class.

equivalentClass: $\text{cons}:\langle\langle \subseteq, C, D \rangle\rangle$
 $\text{cons}:\langle\langle \subseteq, D, C \rangle\rangle$
 \rightsquigarrow **when** $^+C(x)$ **if** $\neg D(x)$ **then rollback**
 \rightsquigarrow **when** $^+D(x)$ **if** $\neg C(x)$ **then rollback**

The expression of **someValuesFrom** ($\exists P.D$) only appears in a subclass expression in OWL 2 RL, and a constraint check should be generated to verify x is in the table C, when there are an insertion of (x, y) to the table P and another insertion of y to the table D, of which the logical constraint is shown below:

someValuesFrom: $\text{cons}:\langle\langle \subseteq, \exists P.D, C \rangle\rangle$
 \rightsquigarrow **when** $^+P(x, y), D(y)$ **if** $\neg C(x)$ **then rollback**

For example, for rule (4) we can derive the following logical constraint to check that whether the individuals that satisfy

the existential restriction also exist in the table Parent:

when $^+\text{hasChild}(x, y), \text{Person}(y)$ **if** $\neg\text{Parent}(x)$ **then rollback**

The **hasValue** construct has two situations in which it may appear and consequently two logical constraints are generated respectively as follows:

hasValue: $\text{cons}:\langle\langle \subseteq, \exists P.a, C \rangle\rangle$
 \rightsquigarrow **when** $^+P(x, a)$ **if** $\neg C(x)$ **then rollback**
hasValue: $\text{cons}:\langle\langle \subseteq, C, \exists P.a \rangle\rangle$
 \rightsquigarrow **when** $^+C(x)$ **if** $\neg P(x, a)$ **then rollback**

In the first case, the logical constraint checks if a tuple (x, a) is in the table P, then x should also appear in the table C and vice versa for the second case. Thus, for rule (5) we can derive the following logical constraints:

when $^+\text{hasParent}(x, \text{John})$ **if** $\neg\text{JohnsChildren}(x)$ **then rollback**
when $^+\text{JohnsChildren}(x)$ **if** $\neg\text{hasParent}(x, \text{John})$ **then rollback**

The logical constraint for **inverseOf** checks that if (x, y) is a property instance of property P_1 then (y, x) has to be an instance of property P_2 .

inverseOf: $\text{cons}:\langle\langle \subseteq, \pi_{\langle C_2, C_1 \rangle} \langle\langle P_1, C_1, C_2 \rangle\rangle, \langle\langle P_2, C_2, C_1 \rangle\rangle \rangle\rangle$,
 $\text{cons}:\langle\langle \subseteq, \pi_{\langle C_1, C_2 \rangle} \langle\langle P_2, C_2, C_1 \rangle\rangle, \langle\langle P_1, C_1, C_2 \rangle\rangle \rangle\rangle$
 \rightsquigarrow **when** $^+P_1(x, y)$ **if** $\neg P_2(y, x)$ **then rollback**

For a **symmetricProperty**, we create a similar check which verifies that if a tuple (x, y) is inserted into the table P, its symmetric tuple (y, x) is also in P.

symmetricProperty: $\text{cons}:\langle\langle \subseteq, \pi_{\langle C_2, C_1 \rangle} \langle\langle P, C_1, C_2 \rangle\rangle, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
 \rightsquigarrow **when** $^+P(x, y)$ **if** $\neg P(y, x)$ **then rollback**

For instance, for rule (26) we can derive the following logical constraint:

when $^+\text{hasSpouse}(x, y)$ **if** $\neg\text{hasSpouse}(y, x)$ **then rollback**

For a **transitiveProperty**, a check will verify a tuple (x, z) is in the table P, if there are tuples (x, y) and (y, z) in P:

transitiveProperty: $\text{cons}:\langle\langle \subseteq, \pi_{\langle P\#1.C_1, P\#2.C_2 \rangle} P \bowtie P, \langle\langle P, C_1, C_2 \rangle\rangle \rangle\rangle$
 \rightsquigarrow **when** $^+P(x, y), P(y, z)$ **if** $\neg P(x, z)$ **then rollback**

For example, for rule (30) we can derive the following logical rule:

when $^+\text{hasAncestor}(x, y), \text{hasAncestor}(y, z)$
if $\neg\text{hasAncestor}(x, z)$ **then rollback**

For a **propertyChain**, we generate the following logical constraint:

propertyChain: $\text{cons}:\langle\langle \subseteq, \pi_{\langle C_1, C_{n+1} \rangle} P_1 \bowtie \dots \bowtie P_n, \langle\langle P, C_1, C_{n+1} \rangle\rangle \rangle\rangle$
 \rightsquigarrow **when** $^+P_1(x, y), P'_{2,n}(y, z)$ **if** $\neg P(x, z)$ **then rollback**

$*P'_{m,n}(x, y) = \pi_{P_m.\text{domain}, P_n.\text{range}} \sigma_{P_j.\text{range}=P_{j+1}.\text{domain}} (P_m \times \dots \times P_n)_{m \leq j < n}$

For example, for rule (25) we can derive the following logical rule:

when $^+\text{hasParent}\#1(x, y), \text{hasParent}\#2(y, z)$
if $\neg\text{hasGrandParent}(x, z)$ **then rollback**

9. IMPLEMENTATION OF OWS & CWS

After generating logical triggers and constraints, SQL physical triggers and constraints can be implemented intuitively by SQL statements. Physical trigger generation is based on trigger translation rules introduced in [15].

The implementation from logical constraints to physical constraints is slightly different. If a logical constraint is to check the subsumption relationships between the value of columns (such as **subClassOf**, **subPropertyOf**, **symmetricProperty** and **inverseOf**), foreign keys are used to implement this logical check. For example, the HDM subclass constraint $\text{cons}:\langle\langle \subseteq, \text{Man}, \text{Person} \rangle\rangle$ representing rule (1) can be achieved using a FK, **FK_Man_lisa_Person**, between **Man** and **Person**. Furthermore, HDM constraints resembling **com-**

plementOf, someValuesFrom, transitiveProperty, and propertyChain can be implemented in a CWS approach by writing functions that check if the constraint holds. Next, we show examples of physical triggers and constraints translated from several HDM constructs.

The first example could be a subsumption relationship between properties, such as the rule (23). In the OWS implementation, an after trigger called hasGrandParent_subOf_hasAncestor shown in Figure 5(a) will be created for the table hasGrandParent and it will insert to the table hasAncestor the data inserted to hasGrandParent. However, in a CWS implementation, we can use a constraint of foreign key called FK_hasGrandParent_subOf_hasAncestor from the two columns of hasGrandParent to columns of the table hasAncestor (shown in Figure 5(b)).

Figure 4: Triggers and Constraints for subPropertyOf.

```
CREATE TRIGGER hasGrandParent_subOf_hasAncestor
ON hasGrandParent
AFTER INSERT AS BEGIN
    INSERT INTO hasAncestor
    SELECT grandchild , grandparent
    FROM inserted
    EXCEPT SELECT descendant , ancestor
    FROM hasAncestor
END
```

(a) SQL Trigger for subPropertyOf.

```
ALTER TABLE hasGrandParent
ADD CONSTRAINT FK_hasGrandParent_isa_hasAncestor
FOREIGN KEY (grandchild , grandparent)
REFERENCES hasAncestor(descendant , ancestor)
ALTER TABLE Man
NOCHECK CONSTRAINT FK_hasGrandParent_isa_hasAncestor
```

(b) SQL Constraint for subPropertyOf.

Figure 5: Triggers and Constraints for propertyChain.

```
CREATE TRIGGER hasParent_chain_hasGrandParent
ON hasParent
AFTER INSERT AS BEGIN
    INSERT INTO hasGrandParent
    SELECT p1.child , p2.parent
    FROM hasParent AS p1 JOIN hasParent AS p2
    ON p1.parent = p2.child
    EXCEPT SELECT * FROM hasGrandParent
END
```

(a) SQL Trigger for propertyChain.

```
CREATE FUNCTION dbo.checkPropertyChain ()
RETURNS BIT AS BEGIN
    DECLARE @Exists BIT
    IF NOT EXISTS (
        SELECT p1.child , p2.parent
        FROM hasParent AS p1 JOIN hasParent AS p2
        ON p1.parent = p2.child
        EXCEPT SELECT grandchild , grandparent
        FROM hasGrandParent)
    BEGIN SET @Exists = 1 END
    ELSE BEGIN SET @Exists = 0 END
    RETURN @Exists
END
```

```
ALTER TABLE hasParent
ADD CONSTRAINT CK_propertyChain
CHECK (dbo.checkPropertyChain () = 1)
```

(b) SQL Constraint for propertyChain.

Note that, since the physical constraint check will be performed after data insertions, so we will disable the constraint

check when loading the data and enable it after the update transaction is committed.

Another example could be a new feature of OWL 2 RL which is propertyChain, exemplified by the rule (25). In OWS, we create a trigger to insert the self join values of the table hasParent to the table hasGrandParent. However, the physical constraint is more complex which cannot be implemented by a foreign key. Therefore we create a checking function which verifies the self-joint tuples of the table hasParent (i.e. $\pi_{\text{hasParent}\#1\text{Domain}, \text{hasParent}\#2\text{Range}} \text{hasParent}\#2 \bowtie \text{hasParent}\#1$) are in the table hasGrandParent. The trigger and constraint are shown in Figure 6(a) and Figure 6(b), respectively (Note that in the physical constraint we use the BIT value 1 to denote TRUE).

10. EVALUATION OF OWLREL

In this section, we show the evaluation of our system with regards to OWS and CWS. For evaluating OWLRel under OWS, we considered the completeness, efficiency, and scalability metrics. For evaluating OWLRel's completeness, we have run the 14 queries of the well known **Lehigh University Ontology Benchmark (LUBM)** [6] with original datasets generated from LUBM's A-Box generator, and also more exhaustive A-Boxes generated by SyGENiA [22]. For evaluating the efficiency and scalability, we have run different sizes of LUBM and checked the scalability of the system in terms of data loading and query processing, and compared query processing time with another semantic reasoner, OWLIM-Lite [8]. Under a CWS setting, we ran some checks to guarantee the soundness of our results. In Table 3, we show the time required for processing each step of OWLRel with a total time of 3.40 (min) in OWS and of 3.39 (min) in CWS.

Table 3: OWLRel Performance Report of LUBM.

OWLRel Steps	Time (s)
① & ② Loading and Classification Time:	10.48
③ HDM Transformation Time:	34.48
④ HDM InterModel Transformation Time:	120.92
⑤ Relational Transformation Time:	37.00
⑥ OWS Transformation Time:	0.84
⑦ CWS Transformation Time:	0.29

OWLRel and OWLIM-Lite were tested on a machine with 2 Intel Xeon E5345 with 2.33GHz CPUs and 8GB of memory, which runs a Microsoft SQL Server 2005 database.

10.1 Evaluation Data

LUBM. The LUBM ontology describes concepts in a university domain. It comprises a T-Box which contains several OWL classes, properties and a number of OWL features such as, subClassOf, subPropertyOf, inverseOf, someValuesFrom, intersectionOf and transitiveProperty. Although the T-Box, is quite simple, LUBM contains a number of features that are beyond those permitted by the OWL 2 RL profile. Apart from the T-Box, LUBM contains 14 queries which we number L1-L14 and an A-Box generator to produce A-Boxes with different sizes. In our experiment, we use LUBM(n) to denote the LUBM A-Boxes of n universities. Each university contains about 100,000 individuals and property tuples.

SyGENiA. Only testing the original datasets of the LUBM benchmark would be limited in terms of completeness, since LUBM’s original A-Boxes are not general and exhaustive enough. SyGENiA is able to generate a more complex A-Box for a given query and a T-Box. Thus, we further evaluate our system using another 14 A-Boxes generated by SyGENiA for the 14 queries of LUBM. For each A-Box, we set the number of assertions to be 1000.

10.2 Evaluation of OWLRel under OWS

10.2.1 Completeness of OWLRel

OWLRel shows a 100% completeness level over LUBM for both, the original A-Boxes and the more exhaustive A-Boxes generated by SyGENiA, which means that our system is the better than OWLIM, Minerva [27], HAWK [17] and Sesame [5] mentioned in [22]. Table 4 shows the completeness level of each query processing compared with OWLIM-Lite (The completeness results of OWLIM and Minerva are from [22]). As can be seen, OWLIM cannot process L6, L8 and L10 completely. One reason for the more complete query processing of OWLRel than OWLIM is that we are able to handle the existential qualification, which is not completely supported by OWLIM. Moreover, Minerva’s completeness level is not that high, since it failed to provide complete answers for L5, L6, L7, L8, L10, L12 and L13.

Table 4: Completeness level over SyGENiA LUBM of OWLRel, OWLIM and Minerva.

System	L5	L6	L7	L8	L10	L12	L13
OWLRel	1	1	1	1	1	1	1
OWLIM	1	0.96	1	0.93	0.96	1	1
Minerva	0.89	0.87	0.90	0.76	0.87	0.66	0.24

10.2.2 Efficiency and Scalability of OWLRel

In order to test the scalability of our system and the efficiency of our data loading and query processing, we compared the execution time for answering each query of LUBM and compare it to OWLIM-Lite.

Query Processing. As can be seen from Table 5, we compared our system with OWLIM-Lite in terms of the execution time for each query. The results show that on average OWLRel was faster than OWLIM per query over all different sizes in our experiment. Moreover, the average query processing time of our system was roughly increased linearly when we doubled the size of A-Boxes, which means that OWLRel scaled over the experiment data.

Data Loading. The data loading time of OWLRel for LUBM(5), LUBM(10), LUBM(20) and LUBM(40) is 15(min), 55(min), 97(min) and 187(min), respectively. OWLRel performed quite fast data loading; for example, it was able to insert approximate 356.5 tuples into the database per second for LUBM(40). Moreover, the data loading time also was increased almost linearly when the data size was doubled, which means that OWLRel was also scalable for loading the A-Boxes.

10.3 Evaluation of OWLRel under CWS

Since there is not a good benchmark for evaluating our system in CWS, we just manually verified our constraint checking. For example, considering the rule (3), inserting John to table Woman has been rolled back, since John was

already in table Man. On the other hand, loading the completely reasoned data has not generated any violations.

11. RELATED WORK

In the context of mapping ontological models to the relational model, most proposals (e.g. [24, 2]) suffer from one of these limitations: ignore OWL restrictions that do not have correspondences in the relational model, store the ontology in a fixed schema (adopting a meta-schema approach), not support OWL 2 ontologies, or do not adhere to the OWS characteristics of ontologies.

In the context of reasoning in an RDBMS, RDBMSs are not only capable of processing ontologies with large number of individuals, but also provides many benefits, such as transaction management, security, integrity control, and scalability [1].

McBrien et al. [16], classified the existing methods that support consistency checking over relational data into three types: Application-based reasoners (ABR), VBR, and TBR. VBR systems such as, DLDB2 uses SQL views to achieve type inference, while SQOWL and its extension SQOWL2 [9] are TBR systems like OWLRel that applies SQL triggers to infer new knowledge. ABR systems such as, SOR [10] and OWLIM rely on reasoners to perform type inference outside an RDBMS.

SOR (previously called Minerva) uses a standard tableau-based DL reasoner to perform the classification of the T-Box. Subsequently, it generates rules to perform type inference outside the database then materialises the results of inferences inside the RDBMS which makes query processing fast. Since the type inference process, however, was implemented outside the database, SOR does not support transactional reasoning nor incremental reasoning in an RDBMS as opposed to OWLRel or SQOWL2. Moreover, the current version of SOR only supports OWL 1 DL.

OWLIM-Lite, a sub system of OWLIM, does its reasoning in memory. It performs materialisation while loading the A-Box just like OWLRel and SQOWL2, however, it can not handle existential quantifications which makes OWLRel and SQOWL2 more complete.

DLDB2 and DBOWL store their rules inside the database as views and do not materialise the inferred closure at loading time. This results in very fast loading time, but slow query processing. On the other hand, SQOWL and more recently SQOWL2 compile T-Box rules into DBMS before and after trigger statements providing a forward chaining materialisation approach.

12. CONCLUSIONS & FUTURE WORK

The mapping of ontologies to relational models has been an active area of research during the past decade. In this paper, we have given a complete, lossless transformation of an OWL 2 RL ontology to a relational schema via a HDM under two approaches; OWS using SQL triggers, and CWS using SQL constraints. So far, OWLRel provides faster query processing time than OWLIM with respect to the LUBM benchmark and shows promising scalability results. Future works will be directed towards first, using other benchmarks like UBOM [11] as well as real-world data for conducting exhaustive experiments to assure the quality of our method and improving the scalability of our system to handle billions of assertions. Finally, we will consider performing the

Table 5: LUBM Query Processing time (ms) of OWLRel and OWLIM.

Size	System	L1	L2	L3	L4	L5	L6	L7	L8	L9	L10	L11	L12	L13	L14	Average
LUBM(5)	OWLRel	42	62	21	2	16	6	71	162	253	42	2	7	0	4	49.29
	OWLIM	119	70	1	5	10	98	3	166	222	1	1	8	2	13	51.36
LUBM(10)	OWLRel	79	121	42	4	31	11	22	215	175	81	3	12	0	9	57.5
	OWLIM	118	115	1	6	9	124	3	215	423	1	2	13	1	26	75.5
LUBM(20)	OWLRel	29	376	87	6	69	24	44	61	901	28	7	11	1	17	118.64
	OWLIM	161	172	1	5	6	106	3	463	911	1	2	24	1	56	136.57
LUBM(40)	OWLRel	56	216	33	11	20	10	84	125	2391	55	12	23	2	35	219.5
	OWLIM	246	450	1	5	7	199	2	847	1895	1	1	47	1	118	272.86

mappings as bidirectional i.e., from relational databases with triggers or constraints to an ontology.

13. REFERENCES

- [1] I. Astrova, N. Korda, and A. Kalja. Storing OWL Ontologies in SQL Relational Databases. *IJECSE*, 1(4):242–247, 2007.
- [2] P. Atzeni, P. Del Nostro, and S. Paolozzi. Ontologies and Databases: Going Back and Forth. In *Proceedings of the 4th ODBIS*. Citeseer, 2008.
- [3] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Patel-Schneider. *The Description Logic Handbook*. CUP, 2003.
- [4] M. Boyd and P. McBrien. Comparing and Transforming between Data Models via an Intermediate Hypergraph Data Model. *Journal on Data Semantics*, IV:69–109, 2005.
- [5] J. Broekstra, A. Kampman, and F. Van Harmelen. Sesame: A generic architecture for storing and querying rdf and rdf schema. In *The Semantic Web-ISWC 2002*, pages 54–68. Springer, 2002.
- [6] Y. Guo, Z. Pan, and J. Heflin. LUBM: A benchmark for OWL knowledge base systems. *Web Semantics: Science, Services and Agents on the WWW*, 3(2):158–182, 2005.
- [7] M. Horridge and S. Bechhofer. The OWL API: A Java API for Working with OWL 2 Ontologies. In *OWLED*, volume 529, pages 11–21, 2009.
- [8] A. Kiryakov, D. Ognyanov, and D. Manov. OWLIM—a pragmatic semantic repository for OWL. In *WISE 2005 Workshops*, pages 182–192. Springer, 2005.
- [9] Y. Liu and P. McBrien. SQOWL2: Transactional Type Inference for OWL 2 DL in an RDBMS. In *Description Logics*, pages 779–790, 2013.
- [10] J. Lu, L. Ma, L. Zhang, J.-S. Brunner, C. Wang, Y. Pan, and Y. Yu. SOR: a practical system for ontology storage, reasoning and search. In *Proceedings of the 33rd VLDB*, pages 1402–1405, 2007.
- [11] L. Ma, Y. Yang, Z. Qiu, G. Xie, Y. Pan, and S. Liu. Towards a Complete OWL Ontology Benchmark. In *ESWC*, pages 125–139, 2006.
- [12] P. McBrien and A. Poulouvasilis. Automatic migration and wrapping of database applications a schema transformation approach. In *Proceedings of ER*, 1999.
- [13] P. McBrien and A. Poulouvasilis. A semantic approach to integrating XML and structured data sources. In *Proc. CAiSE’01*, volume 2068 of *LNCS*, pages 330–345. Springer, 2001.
- [14] P. McBrien and A. Poulouvasilis. Data integration by bi-directional schema transformation rules. In *Proc. ICDE’03*, pages 227–238. IEEE, 2003.
- [15] P. McBrien, N. Rizopoulos, and A. Smith. SQOWL: Type Inference in an RDBMS. *ER*, pages 362–376, 2010.
- [16] P. McBrien, N. Rizopoulos, and A. Smith. Type inference methods and performance for data in an RDBMS. In *Proceedings of the 4th International Workshop on SWIM*, SWIM ’12, pages 6:1–6:8, New York, NY, USA, 2012. ACM.
- [17] Z. Pan. HAWK: OWL Repository and Toolkit. Lehigh University, Bethlehem, 2008.
- [18] Z. Pan, X. Zhang, and J. Heflin. DLDB2: A scalable multi-perspective semantic web repository. In *Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT’08. IEEE/WIC/ACM International Conference on*, volume 1, pages 489–495. IEEE, 2008.
- [19] A. Poulouvasilis and P. McBrien. A General Formal Framework for Schema Transformation. *DKE*, 28(1):47–71, 1998.
- [20] E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical owl-dl reasoner. *Web Semantics: science, services and agents on the World Wide Web*, 5(2):51–53, 2007.
- [21] A. Smith, N. Rizopoulos, and P. McBrien. Automated Model Management. *ER*, pages 542–543, 2008.
- [22] G. Stoilos, B. C. Grau, and I. Horrocks. How Incomplete is Your Semantic Web Reasoner? In *AAAI*, 2010.
- [23] D. Tsarkov and I. Horrocks. FaCT++ description logic reasoner: System description. In *Automated reasoning*, pages 292–297. Springer, 2006.
- [24] E. Vysniauskas, L. Nemuraite, R. Butleris, and B. Paradauskas. Reversible Lossless Transformation From OWL 2 Ontologies into Relational Databases. *IJITCA*, 40(4):293–306, 2011.
- [25] W3C. Web Ontology Language Guide, February 2004. <http://www.w3.org/TR/2004/REC-owl-guide-20040210/>.
- [26] W3C. OWL 2 Web Ontology Language New Features and Rationale, June 2009. <http://www.w3.org/TR/2009/WD-owl2-new-features-20090611/>.
- [27] J. Zhou, L. Ma, Q. Liu, L. Zhang, Y. Yu, and Y. Pan. Minerva: A scalable OWL ontology storage and inference system. In *The Semantic Web-ASWC 2006*, pages 429–443. Springer, 2006.