

Converting Instance Checking to Subsumption: A Rethink for Object Queries over Practical Ontologies

Jia Xu¹, Patrick Shironoshita¹, Ubbo Visser², Nigel John¹, Mansur Kabuka¹

¹Department of Electrical and Computer Engineering, University of Miami, Coral Gables, USA

²Department of Computer Science, University of Miami, Coral Gables, USA

Email: j.xu11@umiami.edu, patrick@infotechsoft.com, nigel.john@miami.edu, m.kabuka@miami.edu, visser@cs.miami.edu

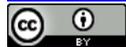
Received 26 October 2014; revised 25 November 2014; accepted 22 December 2014

Academic Editor: Zhongzhi Shi, Institute of Computing Technology, CAS, China

Copyright © 2015 by authors and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Efficiently querying Description Logic (DL) ontologies is becoming a vital task in various data-intensive DL applications. Considered as a basic service for answering object queries over DL ontologies, instance checking can be realized by using the most specific concept (MSC) method, which converts instance checking into subsumption problems. This method, however, loses its simplicity and efficiency when applied to large and complex ontologies, as it tends to generate very large MSCs that could lead to intractable reasoning. In this paper, we propose a revision to this MSC method for DL *SHI*, allowing it to generate much simpler and smaller concepts that are specific enough to answer a given query. With independence between computed MSCs, scalability for query answering can also be achieved by distributing and parallelizing the computations. An empirical evaluation shows the efficacy of our revised MSC method and the significant efficiency achieved when using it for answering object queries.

Keywords

Description Logic, Ontology, Object Query, *SHI*, Most Specific Concept

1. Introduction

Description logics (DLs) play an ever growing role in providing a formal and semantic-rich way to model and

represent structured data in various applications, including semantic web, healthcare and biomedical research, etc. [1]. A knowledge base in description logic, usually referred to as a DL ontology, consists of an *assertional* component (ABox \mathcal{A}) for data description, where *individuals* (single objects) are introduced and their mutual relationships are described using assertional axioms. Semantic meaning of the ABox data can then be unambiguously specified by a *terminological* component (TBox \mathcal{T}) of the DL ontology, where abstract *concepts* and *roles* (binary relations) of the application domain are properly defined.

In various applications of description logics, one of the core tasks for DL systems is to provide an efficient way to manage and query the assertional knowledge (*i.e.* ABox data) in a DL ontology, especially for those data-intensive applications; and DL systems are expected to scale well with respect to (w.r.t.) the fast growing ABox data, in settings such as semantic webs or biomedical systems. The most basic reasoning service provided by existing DL systems for retrieving objects from ontology ABoxes is *instance checking*, which tests whether an individual is a member of a given concept. Instance retrieval (*i.e.* retrieve all instances of a given concept) then can be realized by performing a set of instance checking calls.

In recent years, considerable efforts have been dedicated to the optimization of algorithms for ontology reasoning and query answering [2]-[4]. However, due to the enormous amount of ABox data in realistic applications, existing DL systems, such as Hermit [4] [5], Pellet [6], Racer [7] and FaCT++ [8], still have difficulties in handling the large ABoxes, as they are all based on the (*hyper*) *tableau* algorithm that is computationally expensive for expressive DLs (e.g. up to EXPTIME for instance checking in DL *SHIQ*), where the complexity is usually measured in the size of the TBox, the ABox and the query [9]-[13]. In practice, since the TBox and the query are usually much smaller compared with the ABox, the reasoning efficiency could be mostly affected by the size of the ABox.

One of the solutions to this reasoning scalability problem is to develop a much more efficient algorithm that can easily handle large amount of ABox data. While another one is to reduce size of the data by either partitioning the ABox into small and independent fragments that can be easily handled in parallel by existing systems [14]-[16], or converting the ABox reasoning into a TBox reasoning task (*i.e.* ontology reasoning without an ABox), which could be “*somewhat*” independent of the data size, the TBox is static and relatively simple, as demonstrated in this paper.

A common intuition about converting instance checking into a TBox reasoning task is the so-called most specific concept (MSC) method [10] [17] [18] that computes the MSC of a given individual and reduces any instance checking of this individual into a *subsumption test* (*i.e.* test if one concept is more general than the other). More precisely, for a given individual, its most specific concept should summarize all information of the individual in a given ontology ABox, and should be specific enough to be subsumed by any concept that the individual belongs to. Therefore, once the most specific concept C of an individual a is known, in order to check if a is an instance of any given concept D , it is sufficient to test if C is subsumed by D . With the MSC of every individual in the ABox, the efficiency of online object queries can then be boosted by performing an offline classification of all MSCs that can pre-compute many instance checks [10]. Moreover, if a large ontology ABox consists of data with great diversity and isolation, using the MSC method for instance checking could be more efficient than the original ABox reasoning, since the MSC would have the tableau algorithm to explore only the related information of the given individual, potentially restricted to a small subset of the ABox. Also, this method allows the reasoning to be parallelized and distributed, since MSCs are independent of each other and each preserves complete information of the corresponding individual.

Despite these appealing properties possessed by the MSC method, the computation of a MSC could be difficult even for a very simple description logic such as $\mathcal{AL}\mathcal{E}$. The difficulty arises mainly from the support of qualified existential restrictions (e.g. $\exists R.C$) in DLs, such that when converting a role assertion (e.g. $R(a,b)$) of some individuals into an existential restriction, so that information of that given individual may not be preserved completely. For a simple example, consider converting assertions

$$R(a,b) \text{ and } A(a)$$

into a concept for individual a . In this case, we can always find a more specific concept for a in the form of

$$A \sqcap \underbrace{\exists R.\exists R.\dots\exists R}_n.A$$

by increasing n , and none of them would capture the complete information of individual a . Such information loss is due to the occurrence of cycles in the role assertions, and none of the existential restrictions in DL could

impose a circular interpretation (model) unless *nominals* (e.g. $\{a\}$) are involved or (local) reflexivity is presented [5].

Most importantly, due to the support of existential restrictions, computation of the MSC for a given individual may involve assertions of other entities that are connected to it through role assertions. This implies not only the complexity of the computation for MSCs but also the potential that the resulting MSCs may have larger than desired sizes. In fact, in many of the practical ontology ABoxes (e.g. a social network or semantic webs), most of the individuals could be connected to each other through role assertions, forming a huge connected component in the ABox graph. Under this situation, the resulting MSC could be extremely large and reasoning with it may completely degenerate into an expensive reasoning procedure.

In this paper, we propose a revised MSC method for DL \mathcal{SHI} that attempts to tackle the above mentioned problems, by applying a *call-by-need* strategy together with optimizations. That is, instead of computing the most specific concepts that could be used to answer any queries in the future, the revised method takes into consideration only the related ABox information with *current* query and computes a concept for each individual that is only *specific enough* to answer it w.r.t. the TBox. Based on this strategy, the revision allows the method to generate much simpler and smaller concepts than the original MSCs by ignoring irrelevant ABox assertions. On the other hand, the complexity reduction comes with the price of re-computation (*i.e.* online computation of MSCs) for every new coming query, if no optimization is applied. Nevertheless, as shown in our experimental evaluation, the simplicity achieved could be significant in many practical ontologies, and the overhead is thus negligible compared with the reasoning efficiency gained for each instance checking and query answering. Moreover, due to the re-computations, *we do not assume a static ontology or query*, and the ABox data is amenable to frequent modifications, such as insertion or deletion, which is in contrast to the original MSC method where a relatively static ABox is assumed. A procedure for instance retrieval based on our revised MSC method is shown in [Figure 1](#).

The revised MSC method could be very useful for efficient instance checking in many practical ontologies, where the TBox is usually small and manageable while the ABox is in large scale as a database and tends to change frequently. Particularly, this method would be appealing to large ontologies in *non-Horn* DLs, where current optimization techniques such as rule-based reasoning or pre-computation may fall short. Moreover, the capability to parallelize the computation is another compelling reason to use this technique, in cases where answering object queries may demand thousands or even millions of instance checking tasks.

Our contributions in this paper are summarized as follows:

- 1) We propose a call-by-need strategy for the original MSC method, instead of computing the most specific concepts offline to handle any given query, which allows us to focus on the current queries and to generate online much smaller concepts that are sufficient to compute the answers. This strategy makes our MSC method suitable for query answering in ontologies, where frequent modifications to the ontology data are not uncommon;
- 2) We propose optimizations that can be used to further reduce sizes of computed concepts in practical ontologies for more efficient instance checking;
- 3) Finally, we evaluate our approach on a range of test ontologies with large ABoxes, including ones generated by existing benchmark tools and realistic ones used in biomedical research. The evaluation shows the efficacy of our proposed approach that can generate significantly smaller concepts than the original MSC. It also

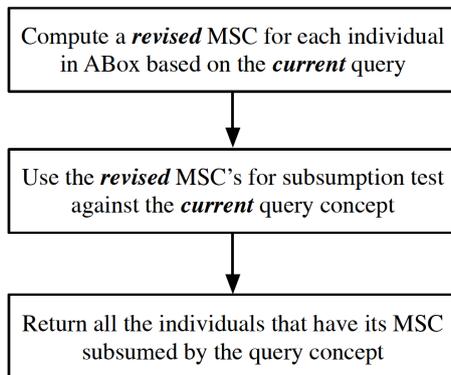


Figure 1. A procedure for instance retrieval for a given query based on our revised MSC method.

shows the great reasoning efficiency that can be achieved when using the revised MSC method for instance checking and query answering.

The rest of the paper is organized as follows: in Section 2, we introduce the background knowledge of a description logic and DL ontology; in Section 3, we give more detailed discussion about the MSC method and our call-by-need strategy; Section 4 presents the technical details of the revised MSC method; Section 5 discusses the related work; Section 6 presents an empirical evaluation on our proposed method; and finally, Section 7 concludes our work.

2. Preliminaries

The technique proposed in this paper is for description logic \mathcal{SHI} . For technical reasons, we need a *constrained* use of nominals on certain conditions (*i.e.* assertion cycles), which requires logic \mathcal{SHIO} . Thus, in this section, we give a brief introduction to formal syntax and semantics of logic \mathcal{SHIO} , DL ontologies, and basic reasoning tasks for derivation of logical entailments from DL ontologies.

2.1. Description Logic \mathcal{SHIO}

The vocabulary of description logic \mathcal{SHIO} includes a set \mathbf{R} of named roles with a subset $\mathbf{R}_+ \subseteq \mathbf{R}$ of transitive roles, a set \mathbf{C} of named (atomic) concepts, and a set \mathbf{I} of named individuals.

Definition 2.1 (\mathcal{SHIO} Role) A role in \mathcal{SHIO} is either a named (atomic) one $R \in \mathbf{R}$ or an inverse one R^- with $R \in \mathbf{R}$, and the complete role set of \mathcal{SHIO} can be denoted $\mathbf{R}^* = \mathbf{R} \cup \{R^- \mid R \in \mathbf{R}\}$. To avoid role representation such as R^- , a function $\text{Inv}(\cdot)$ is defined, such that $\text{Inv}(R) = R^-$ if R is a role name, and $\text{Inv}(R) = P$ if $R = P^-$ for some role name P . A role R is transitive, denoted $\text{Trans}(R)$, if either R or $\text{Inv}(R)$ belongs to \mathbf{R}_+ .

Definition 2.2 (\mathcal{SHIO} Concept) A \mathcal{SHIO} -concept is either an atomic (named) concept or a complex one that can be defined using the following constructs recursively

$$C, D ::= A \mid \{o\} \mid \top \mid \perp \mid \neg C \mid C \sqcap D \mid C \sqcup D \mid \forall R.C \mid \exists R.C$$

where A is an atomic concept in \mathbf{C} , o is a named individual, and $R \in \mathbf{R}^*$.

Description logic \mathcal{SHI} is then defined as a fragment of \mathcal{SHIO} , which disallows the use of nominal (*i.e.* $\{o\}$) as a construct for building complex concepts.

Definition 2.3 (\mathcal{SHIO} Semantics) The meaning of an entity in \mathcal{SHIO} is defined by a model-theoretical semantics using an interpretation denoted $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta^{\mathcal{I}}$ is referred to as a non-empty domain and $\cdot^{\mathcal{I}}$ is an interpretation function. The function $\cdot^{\mathcal{I}}$ maps every atomic concept in \mathbf{C} to a subset of $\Delta^{\mathcal{I}}$, every ABox individual to an element of $\Delta^{\mathcal{I}}$, and every role to a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$. Interpretation for other concepts and role are given below:

$$\begin{aligned} \top^{\mathcal{I}} &= \Delta^{\mathcal{I}} \\ \perp^{\mathcal{I}} &= \emptyset \\ \{o\}^{\mathcal{I}} &= \{o^{\mathcal{I}}\} \\ \neg C^{\mathcal{I}} &= \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (R^-)^{\mathcal{I}} &= \{(y, x) \mid (x, y) \in R^{\mathcal{I}}\} \\ (C \sqcap D)^{\mathcal{I}} &= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (C \sqcup D)^{\mathcal{I}} &= C^{\mathcal{I}} \cup D^{\mathcal{I}} \\ (\exists R.C)^{\mathcal{I}} &= \{x \mid \exists y. (x, y) \in R^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \\ (\forall R.C)^{\mathcal{I}} &= \{x \mid \forall y. (x, y) \in R^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \end{aligned}$$

Definition 2.4 (Simple-Form Concept) A concept is said to be in simple form, if the maximum level of nested quantifiers in this concept is less than 2.

For example, given an atomic concept A , both A and $\exists R.A$ are simple-form concepts, while $\exists R_1.(A \sqcap \exists R_2.A)$ is not, since its maximum level of nested quantifiers is two. Notice however, an arbitrary concept can be *linearly* reduced to the simple form by assigning new concept names for fillers of the quantifiers. For example, $\exists R_1.\exists R_2.C$ can be converted to $\exists R_1.D$ by letting $D \equiv \exists R_2.C$ where D is a new concept name.

Assumption: For accuracy of the technique presented in this paper, without loss of generality, we assume all ontology concepts are in simple form as defined previously, and the concept in any concept assertion is atomic.

2.2. DL Ontologies and Reasoning

Definition 2.5 (SHIO Ontology) A SHIO ontology is a tuple, denoted $\mathcal{K} = (\mathcal{T}, \mathcal{A})$, where \mathcal{T} is called a TBox and \mathcal{A} is called an ABox.

The TBox \mathcal{T} is constituted by a finite set of role inclusion axioms (i.e. $R_1 \sqsubseteq R_2$ with $R_1, R_2 \in R^*$) and a finite set of concept inclusion axioms in the form of $C \sqsubseteq D$ and $C \equiv D$, where C, D are SHI concepts. The former is called a *general concept inclusion axiom (GCI)*, and the latter can be simply converted into two GCIs as $C \sqsubseteq D$ and $D \sqsubseteq C$.

The ABox \mathcal{A} consists of a finite set of assertions, in the form of $A(a)$ (concept assertion) and $R(a, b)$ (role assertion), where A is a concept, R is a role, and a, b are named individuals in I . In a role assertion $R(a, b)$, individual a is referred to as a *R-predecessor* of b , and b is a *R-successor* (or *R⁻-predecessor*) of a . If b is a *R-successor* of a , b is also called a *R-neighbor* of a .

An interpretation \mathcal{I} satisfies an axiom $C \sqsubseteq D$ (written $\mathcal{I} \models C \sqsubseteq D$), iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, and \mathcal{I} satisfies an axiom or assertion:

$$\begin{aligned} R_1 \sqsubseteq R_2 &\text{ iff } R_1^{\mathcal{I}} \subseteq R_2^{\mathcal{I}} \\ C(a) &\text{ iff } a^{\mathcal{I}} \in C^{\mathcal{I}} \\ R(a, b) &\text{ iff } (a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}} \end{aligned}$$

If \mathcal{I} satisfies every axiom and assertion of an ontology \mathcal{K} , \mathcal{I} is called a *model* of \mathcal{K} , written $\mathcal{I} \models \mathcal{K}$. In turn, \mathcal{K} is said *satisfiable* iff it has at least one model; otherwise, it is *unsatisfiable* or *inconsistent*.

Definition 2.6 (Logical Entailment) Given an ontology \mathcal{K} and an axiom α , α is called a *logical entailment* of \mathcal{K} , denoted $\mathcal{K} \models \alpha$, if α is satisfied in every model of \mathcal{K} .

Definition 2.7 (Instance Checking) Given an ontology \mathcal{K} , a DL concept C and an individual $a \in I$, *instance checking* is defined to test if $\mathcal{K} \models C(a)$ holds.

Notice that, instance checking is considered the central reasoning service for information retrieval from ontology ABoxes [19], and more complex reasoning services, such as instance retrieval, can be realized based on this basic service. Instance checking can also be viewed as a procedure of individual “*classification*” that verifies if an individual can be classified into some defined DL concepts. An intuition to implement this instance checking service is to convert it into a concept subsumption test by using the so-called *most specific concept (MSC)* method.

Definition 2.8 (Most Specific Concept [20]) Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be an ontology, and a be an individual in I . A concept C is called the *most specific concept* for a w.r.t. \mathcal{A} , written $\text{MSC}(\mathcal{A}, a)$, if for every concept D that $\mathcal{K} \models D(a)$, $\mathcal{I} \models C \sqsubseteq D$.

The MSC method turns the instance checking into a TBox reasoning problem. That is, once the most specific concept $\text{MSC}(\mathcal{A}, a)$ of an individual a is known, to decide if $\mathcal{K} \models D(a)$ holds for an arbitrary concept D , it suffices to test if $\mathcal{T} \models \text{MSC}(\mathcal{A}, a) \sqsubseteq D$ [10].

Ontology reasoning algorithm in current systems (e.g. Pellet, and HermiT, etc.) are based on (hyper) tableau algorithms [4] [6] [7] [21]. For details of a standard tableau algorithm for SHIO, we refer readers to the work in [22].

3. Classification of Individuals

The MSC method for individual checking is based on the idea that, an individual can be classified into a given

concept D , if and only if there exists a concept behind its ABox assertions subsumed by D [17] [18] [20]. Computation of the MSC for a given individual then demands converting its ABox assertions into a concept. This task can be easily accomplished if the individual possesses only concept assertions, by simply collapsing the involved concepts into a single term using the concept conjunction. When role assertions are involved, however, a more complex procedure is demanded, and the method we used here is called *rolling-up* [23], which is elaborated in the next section.

Using the MSC method for instance checking might eliminate the memory limitation for reasoning with large ABoxes, especially when the ABox \mathcal{A} consists of data in great diversity and isolation. This is simply because each computed $\text{MSC}(\mathcal{A}, a)$ should comprise of only related information of the given individual, and makes the subsumption test (*i.e.* $\mathcal{T} \models \text{MSC}(\mathcal{A}, a) \sqsubseteq D$) as efficient as an ontology reasoning that explores only a (small) portion of \mathcal{A} .

However, as discussed in Section 1, due to the support of existential restrictions in DLs, great complexity for computation of MSC's may arise when role assertions are involved. Besides, due to the *completeness* that should be guaranteed by each $\text{MSC}(\mathcal{A}, a)$ (*i.e.* the MSC should be subsumed by any concept that the individual a belongs to.), the resulting MSC's may turn out to be a very large concept, whenever there is a great number of individuals in \mathcal{A} connected to each other by role assertions. In the *worst* case, reasoning with a MSC may degenerate into a complete ABox reasoning that could be prohibitively expensive. For example, when $\text{MSC}(\mathcal{A}, a)$ preserves complete information of \mathcal{A} , its interpretation will form a tableau, the size of which can be in the same scale of \mathcal{A} .

3.1. The Call-by-Need Strategy

Since the larger than desired sizes of MSCs are usually caused by its completeness as discussed above, a possible optimization to the MSC method is thus to abandon the completeness that is required to deal with any query concepts, and to apply a “*call-by-need*” strategy. That is, for an arbitrary query concept D , instead of computing the MSC for each individual a , we compute a concept that is only *specific-enough* to determine if a can be classified into D . As suggested by its name, this revision to the original MSC method, instead of taking the complete information of individual a when computing the “MSC”, will consider only the ABox assertions that are relevant to the current query concept.

A simple way to realize this strategy is to assign a fresh name A every time to a given (complex) query concept D by adding the axiom $A \equiv D$ to \mathcal{T} ¹, and to concentrate only on ABox assertions that would (probably) classify an individual a into A w.r.t. \mathcal{T} . Consequently, this implementation requires an analysis of the ontology axioms/assertions, such that the possibility of each role assertion to affect individual classification (w.r.t. named concept in \mathcal{T}) can be figured out. Computation of a specific-enough concept should then concentrate on role assertions that are *not impossible*. We abuse the notation here to denote this specific-enough concept for individual a w.r.t. ABox \mathcal{A} , current query concept \mathcal{Q} , and named concepts in \mathcal{T} as $\text{MSC}_{\mathcal{T}}(\mathcal{A}, \mathcal{Q}, a)$, and we call the method that uses $\text{MSC}_{\mathcal{T}}$ for instance checking the $\text{MSC}_{\mathcal{T}}$ method.

Definition 3.1 Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be an ontology, a be an individual in \mathcal{A} , and \mathcal{Q} a current query concept for individuals. A concept C is called a *specific-enough concept* for a w.r.t. named concepts in \mathcal{T} , \mathcal{Q} and \mathcal{A} , written $\text{MSC}_{\mathcal{T}}(\mathcal{A}, \mathcal{Q}, a)$, if $\mathcal{K} \models \mathcal{Q}(a)$, $\mathcal{T} \models C \sqsubseteq \mathcal{Q}$.

Since in our procedure we will add the query concept \mathcal{Q} into \mathcal{T} as a named concept, we can simplify the notation $\text{MSC}_{\mathcal{T}}(\mathcal{A}, \mathcal{Q}, a)$ as $\text{MSC}_{\mathcal{T}}(\mathcal{A}, a)$.

3.2. A Syntactic Premise

To decide whether a role assertion could affect classification of a given individual, a sufficient and necessary condition as stated previously is that, the concept behind this assertion conjuncted with other essential information of the individual should be subsumed by the given concept w.r.t. \mathcal{T} [17] [18] [20]. Formally, for a role assertion $R(a, b)$ that makes individual a classified into a concept A , the above sufficient and necessary condition in *SHI* can be expressed as:

$$\mathcal{K} \models \exists R.B \sqcap A_0 \sqsubseteq A \quad (1)$$

where $b \in B$ is entailed by \mathcal{K} , and concept A_0 summarizes the rest of the information of a that is also

¹Note that, to follow the simple-form concept restriction, multiple axioms may be added.

essential to this classification, with $A_0 \sqsubseteq A$.

As shown in [16], for subsumption (1) to hold when A is a named concept, there must exist some role restriction $\exists R'.C$ with $R \sqsubseteq R'$ in left-hand side of TBox axioms (see (2) and the following axiom equivalency) for concept definition; otherwise $\exists R.B$ is not comparable (w.r.t. subsumption) with other concepts (except \top and its equivalents). This syntactic condition for the deduction of (1) is formally expressed in the following proposition.

Proposition 3.1 ([16]) *Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a \mathcal{SHI} ontology with simple-form concepts only, $\exists R.B$, A_0 and A be \mathcal{SHI} concepts, where A is named. If*

$$\mathcal{K} \models \exists R.B \sqcap A_0 \sqsubseteq A$$

with $A_0 \not\sqsubseteq A$, there must exist some GCIs in \mathcal{T} in the form of:

$$\exists R'.C_1 \triangleright \triangleleft C_2 \sqsubseteq C_3 \quad (2)$$

where $R \sqsubseteq R'$ and $\triangleright \triangleleft$ is a place holder for \sqcup and \sqcap , C_i 's are \mathcal{SHI} concepts. Also note the following equivalence:

$$\begin{aligned} \exists R.C \sqsubseteq D &\equiv \neg D \sqsubseteq \forall R.\neg C \\ \exists R.C \sqsubseteq D &\equiv C \sqsubseteq \forall R^{\neg}.D \\ C_1 \sqcap C_2 \sqsubseteq D &\equiv C_1 \sqsubseteq D \sqcup \neg C_2 \end{aligned}$$

This proposition is proven in [16]. It states in fact a syntactic premise in \mathcal{SHI} for a role assertion to be essential for some individual classification. That is, if a role assertion $R(a,b)$ is essential for derivation of $A(a)$ for some named concept A , there must exist a related axiom in \mathcal{T} in the form of (2) for $R \sqsubseteq R'$. We denote this syntactic premise for $R(a,b)$ to affect a 's classification as SYN_COND . Using this condition, we can easily rule out role assertions that are definitely irrelevant to the query concept and will not be considered during the computation of a $\text{MSC}_{\mathcal{T}}$.

4. Computation of $\text{MSC}_{\mathcal{T}}$

In this section, we present the technique that computes a $\text{MSC}_{\mathcal{T}}$ for a given individual w.r.t. a given query. We assume the ABox considered here is consistent, since for any inconsistent ABox, the $\text{MSC}_{\mathcal{T}}$ is always the bottom concept \perp [24]. Essentially, the task is to convert ABox assertions into a single concept for a given individual, using the concept conjunction and the so-called *rolling-up* technique. This rolling-up technique was introduced in [23] to convert conjunctive queries into concept terms, and was also used by [25] to transform datalog rules into DL axioms. We adapt this technique here to roll up ABox assertions into DL concepts.

4.1. The Rolling-Up Procedure

Converting concept assertions into a concept is straightforward by simply taking conjunction of the involved concepts. When role assertions are involved, the rolling-up technique can be used to transform assertions into a concept by eliminating individuals in them. For example, given the following assertions

$$\text{Male}(\text{Tom}), \text{has Parent}(\text{Tom}, \text{Mary}), \text{Lawyer}(\text{Mary}), \quad (3)$$

transforming them for individual Tom using the rolling up and concept conjunction can generate a single concept assertion

$$(\text{Male} \sqcap \exists \text{has Parent.Lawyer})(\text{Tom})$$

Generalize the Information: The transformation here is for individual Tom, and if individual Mary is not explicitly indicated in the query, it should be sufficient to rewrite $\text{has Parent}(\text{Tom}, \text{Mary}), \text{Lawyer}(\text{Mary})$ into $\exists \text{has Parent.Lawyer}(\text{Tom})$, without loss of any information that is essential for query answering. Even if Mary is explicitly indicated in the query, we can still eliminate it by using a *representative* concept that stands for this particular individual in the given ABox [26]. For example, we can add an assertion $A_{\text{mary}}(\text{Mary})$ to the ABox, where A_{mary} is a new concept name and a representative concept for individual Mary. The above role assertions for Tom then can be transformed into concept $\exists \text{has Parent.}(\text{Lawyer} \sqcap A_{\text{mary}})(\text{Tom})(\text{Mary})$; and if the query is

also rewritten using concept Mary, the *completeness* of the query answering can be guaranteed, as indicated by the following theorem [26].

Theorem 4.1 ([26]) *Let $\mathcal{K} = (\mathcal{T}, \mathcal{A})$ be a DL ontology, a, b be two individuals in \mathcal{A} , R a role, and C_1, \dots, C_n DL concepts. Given a representative concept name A_b for b not occurring in \mathcal{K} :*

$$\mathcal{K} \models R(a, b) \wedge C_1(b) \wedge \dots \wedge C_n(b)$$

if and only if

$$\mathcal{K} \cup \{A_b(b)\} \models \exists R.(A_b \sqcap C_1 \sqcap \dots \sqcap C_n)(a)$$

The rolling-up procedure here can be better understood by considering a *graph* induced by the role assertions to be rolled up, which is defined as follows:

Definition 4.1 *A set of ABox role assertions in \mathcal{SHI} can be represented by a graph \mathcal{G} , in which there is a node x for each individual x in the assertions, and an edge between node x and y for each role assertion $R(x, y)$.*

Notice that, due to the support of inverse roles in \mathcal{SHI} , edges in \mathcal{G} are not directed. A *role path* in the graph \mathcal{G} is then defined as a set of roles corresponding to the set of edges (no duplicate allowed) leading from one node to another. For example, given assertions $R_1(x, y)$ and $R_2(x, y)$, the role path from x to z is $\{R_1, R_2^-\}$, and its reverse from z to x is $\{R_2, R_1^-\}$.

The rolling-up for a given individual x is then able to generate concepts by eliminating individuals in branches of the *tree-shaped* graph \mathcal{G} , starting from the leaf nodes and rolling up towards the root node indicated by x . Moreover, all the information of each individual being rolled up should be absorbed into a single concept by conjunction during the procedure. For example, if we have additional assertions

$$\text{has Sister}(\text{Mary}, \text{Ana}) \text{ and } \text{Professor}(\text{Ana})$$

for Mary in (3), the rolling-up for Tom should then generate concept

$$\text{Male} \sqcap \text{has Parent} . (\text{Lawyer} \sqcap \text{has Sister} . \text{Professor})$$

Inverse Role: The support of inverse roles in \mathcal{SHI} makes this rolling-up procedure bidirectional, thus, making it applicable to computing $\text{MSC}_{\mathcal{T}}$ for any individual in the ABox. For example, to compute a $\text{MSC}_{\mathcal{T}}$ for individual Mary in example (3), we simply treat this individual as the root, and roll up assertions from leaves to root to generate the concept

$$\text{Lawyer} \sqcap \text{has Parent}^- . \text{Male}$$

Transitive Role: In the rolling-up procedure, no particular care needs to be taken to deal with transitive roles, since any role assertions derived from transitive roles will be automatically preserved [26]. For example, given R a transitive role, $R(a, b)$, $R(b, c)$ two role assertions, and $B(b)$, $C(c)$ two concept assertions in the ABox, rolling-up these four assertions for individual a can generate assertion $\exists R.(B \sqcap \exists R.C)(a)$, from which together with the TBox axioms, we can still derive the fact that

$$(\exists R.(B \sqcap \exists R.C) \sqcap \exists R.C)(a)$$

Assertion Cycles: This rolling-up technique, however, may suffer information loss if the graph \mathcal{G} contains cycles (*i.e.* a role path leading from one node to itself without duplicate graph edges). For example, given the following two assertions:

$$R_1(x, y) \text{ and } R_2(x, y) \tag{4}$$

individuals x and y are related by two roles, and a cycle is thus induced in the corresponding graph. Rolling-up assertions for individual x using the method described above might generate concept $\exists R_1. \top \sqcap \exists R_2. \top$, and the fact that x is connected to the same individual y through different roles is lost. Consequently, this may compromise the resulting concept as a specific-enough concept for x to answer the current query. For example, let C be a query concept defined as:

$$\exists R_1. \exists R_2^- . \exists R_1. \top$$

It can be found out through ABox reasoning that individual x is an instance of C ; while on the other hand, it is also not difficult to figure out that $\exists R_1. \top \sqcap \exists R_2. \top$ is not subsumed by C .

Multiple solutions to this problem have been proposed, such as an approximation developed by [27], and the use of cyclic concept definition with greatest fixpoint semantics [24] [28]. In this paper, we choose to use the *nominal* (e.g. $\{x\}$) to handle circles as suggested by [19] [20], which allows explicit indication of named individuals in a concept, hence, being able to indicate the joint node of a cycle. The above two assertions in (4) then can be transformed into a concept for individual x as either $\exists R_1. \{y\} \sqcap \exists R_2. \{y\}$ or $\{x\} \sqcap \exists R_1. \exists R_2. \{x\}$, each with the nominal used for a chosen joint node, and both preserve complete information of the cycle. In our approach, when rolling up a cycle in \mathcal{G} , we always treat the cycle as a single branch and generate concepts of the second style. That is, our procedure will treat a chosen joint node x as both the tail (*i.e.* leaf) and the head (*i.e.* root) of the branch. For clarity of the following presentation, we denote the tail as x^t and the head as x^h .

Based on the discussion so far, the transformation of assertions for a given individual now can be formalized as follows. Let x be a named individual, and γ be an ABox assertion for x . γ can be transformed into a concept C_γ for x :

$$C_\gamma = \begin{cases} C & \text{if } \gamma = C(x) \mid C(x^h) \\ \exists R.D & \text{if } \gamma = R(x, y), \text{ and } D(y) \\ \exists R^-.D & \text{if } \gamma = R(y, x), \text{ and } D(y) \\ \{x^h\} \sqcap \exists R.\{x^t\} & \text{if } \gamma = R(x^h, x^t) \\ \{x^h\} \sqcap \exists R.D & \text{if } \gamma = R(x^h, y), \text{ and } D(y) \\ \{x^h\} \sqcap \exists R^-.D & \text{if } \gamma = R(y, x^h), \text{ and } D(y) \\ \{x^t\} & \text{if } \gamma = R(x^t, y) \mid R(y, x^t) \mid C(x^t) \end{cases}$$

Notice that, concept D here is a obtained concept when rolling up branch(es) in \mathcal{G} up to node y , and transforming any assertion of a cycle tail x^t always generates $\{x^t\}$, as complete information of x will be preserved when rolling up to the head x^h . Thereafter, given a set S of all assertions of individual x , $\text{MSC}_T(\mathcal{A}, x)$ can be obtained by rolling up all branches induced by role assertions in S and taking the conjunction of all obtained concepts. When S is empty, however, individual x in the ABox can only be interpreted as an element of the entire domain, and thus, the resulting concept is simply the top entity \top . Computation of a $\text{MSC}_T(\mathcal{A}, x)$ then can be formalized using the following equation:

$$\text{MSC}_T(\mathcal{A}, x) = \begin{cases} \top & \text{if } S = \emptyset \\ \bigcap_{\gamma \in S} C_\gamma & \text{if otherwise} \end{cases}$$

4.2. Branch Pruning

To apply the call-by-need strategy, the previously defined syntactic premise `SYN_COND` is employed, and a branch to be rolled up in graph \mathcal{G} will be truncated at the point where the edge does not have `SYN_COND` satisfied. More precisely, if an assertion $R(x, y)$ in a branch does not have the corresponding `SYN_COND` satisfied, it will not affect any classification of individual x w.r.t. T . Moreover, any effects of the following assertions down the branch will not be able to propagate through $R(x, y)$ to x , and thus should not be considered during the rolling-up of this branch.

This branch pruning technique could be a simple yet an efficient way to reduce complexity of a MSC_T , especially for those practical ontologies, where many of the ABox individuals may have a huge number of role assertions and only a (small) portion of them have `SYN_COND` satisfied. For a simple example, consider an individual x in an ontology ABox with the following assertions:

$$R_1(x, y_1), R_2(x, y_2), \dots, R_n(x, y_n)$$

where n could be a very large number and only R_1 has `SYN_COND` satisfied. Rolling up these assertions for individual x without the pruning will generate the concept

$$\exists R_1.C_1 \sqcap \exists R_2.C_2 \sqcap \dots \sqcap \exists R_n.C_n$$

where $y_i \in C_i$. Using this concept for any instance checking of x could be expensive, as its interpretation might completely restore the tableau structure that is induced by these assertions. However, when the pruning is applied, the new $\text{MSC}_{\mathcal{T}}$ should be $\exists R_1.C_1$, the only role restriction that is possible to affect individual classification of x w.r.t. named concepts in \mathcal{T} .

Going beyond such simple ontologies, this optimization technique may also work in complex ontologies, where most of the role assertions in ABox could have SYN_COND satisfied. For example, consider the following assertions

$$R_0(x_0, x_1), R_1(x_1, x_2), \dots, R_n(x_n, x_{n+1})$$

with all roles except R_2 having SYN_COND satisfied. Rolling up these assertions for individual x_0 will start from the leaf x_{n+1} up towards the root x_0 , and generate the concept

$$\exists R_0.\exists R_1.\dots.\exists R_n.C$$

where $x_{n+1} \in C$. However, with pruning applied, the rolling-up in this branch will start from x_2 instead of x_{n+1} , since $R_2(x_2, x_3)$ will not affect classification of individual x_2 w.r.t. \mathcal{T} and the branch is truncated at this point.

Furthermore, with branch pruning, cycles should only be considered in the truncated graph, which may further simplify the computation of $\text{MSC}_{\mathcal{T}}$'s.

4.3. Further Optimization and Implementation

The branch pruning here is based on SYN_COND to rule out irrelevant assertions, which in fact can be further improved by developing a more rigorous premise for a role assertion to affect individual classification. For exposition, consider the following ontology \mathcal{K} :

$$\{\{\exists R.C \sqsubseteq D\}, \{\neg D(a), R(a, b), \neg C(b)\}\} \quad (5)$$

When computing $\text{MSC}_{\mathcal{T}}(\mathcal{A}, a)_{\mathcal{T}}(\mathcal{A}, a)$ using the proposed method, assertion $R(a, b)$ will be rolled up as the corresponding SYN_COND is satisfied. However, it is not difficult to see that, $R(a, b)$ here actually makes no contribution to a 's classification, since individual b is in the complement of concept C , making a an instance of $\exists R.\neg C$. Besides, individual a has already been asserted as an instance of concept $\neg D$, and hence cannot be classified into D unless the ABox is inconsistent.

With these observations, a more rigorous premise based on SYN_COND can be derived. That is, to determine the possibility for $R(a, b)$ to affect classification of individual a , beyond checking in \mathcal{T} the existence of any axiom in the form of

$$\exists R'.C_1 \triangleright \triangleleft C_2 \sqsubseteq C_3$$

with $R \sqsubseteq R'$ and $\triangleright \triangleleft$ a place holder for \sqcup and \sqcap , we also check the following cases for any found axiom:

Case 1 If there is any concept B_0 in explicit concept assertions of individual b , such that $\mathcal{K} \models B_0 \sqsubseteq \neg C_1$,

Case 2 If there is any concept A_0 in explicit concept assertions of individual a , such that $\mathcal{K} \models A_0 \sqsubseteq \neg(C_3 \sqcup \neg C_2)^2$ or $\mathcal{K} \models A_0 \sqsubseteq \neg C_3$, respectively for $\triangleright \triangleleft$ standing for \sqcap or \sqcup .

If either one of the above cases happens, that particular $\exists R'.C_1$ in the left hand side of the axiom in fact makes no contribution to the inference of a 's classification, unless the ABox is *inconsistent* where MSCs are always \perp [24]. Thus, a revised condition requires not only the existence of a related axiom in the form of (2) but also with none of the above cases happening. We denote this condition as SYN_COND^* , and use it to rule out assertions that are irrelevant to the current query.

This optimization is useful to prevent rolling-up of role assertions in an arbitrary direction on existence of related axioms in \mathcal{T} . Instead, it limits the procedure to the direction that is desired by the original intention underlying the design of the given ontology. For example, in (5), the axiom $\exists R.C \sqsubseteq D$ specifies that, any individual having a R -neighbor in C is an instance of D and any individual having a R^- -neighbor in $\neg D$

²Note the axiom equivalence $C_1 \sqcap C_2 \sqsubseteq D \equiv C_1 \sqsubseteq D \sqcup \neg C_2$.

is an instance of $\neg C$ ³. However, if individual x is asserted to have a R -neighbor in $\neg C$ or a R^- -neighbor in D , that role assertion should not be rolled-up for x just on existence of this axiom.

With all the insights discussed so far, an algorithm for computation of $\text{MSC}_{\mathcal{T}}(\mathcal{A}, a)$ is presented here as a *recursive* procedure, the steps of which are summarized in **Figure 2**.

Proposition 4.1 (Algorithm Correctness) *The algorithm presented in Figure 2 computes a $\text{MSC}_{\mathcal{T}}(\mathcal{A}, a)$ for a given SHI ontology $(\mathcal{T}, \mathcal{A})$ and an individual x in \mathcal{A} .*

Proof. We prove by induction.

Basis: For a leaf node x in \mathcal{G} , which has no other role assertions except those up the branches, rolling it up yields the conjunction of concepts in its concept assertions, which preserves sufficient information of the part of the branch being rolled so far. If x is the tail of a cycle, returning $\{x^t\}$ is sufficient, as other information of x will be gathered when the rolling-up comes to the head.

Inductive Step: Let x be a node in the middle of some branch(es) in \mathcal{G} . For every role assertion $R_i(x, y_i)$ of x down the branch, assume the procedure generates a concept D_i for rolling up to each node y_i , which preserves sufficient information (w.r.t. current query) of the part of branches being rolled up so far. Then, rolling up each $R_i(x, y_i)$ generates $\exists R_i.D_i$, and together with concept assertions of x , the concept conjunction preserves sufficient information of all branches being rolled up to x . If x is marked as a joint node of a cycle, $\{x^h\}$ is also in the conjunction, so that the circular path property can be preserved. If x is the root node, the conjunction is thus a $\text{MSC}_{\mathcal{T}}(\mathcal{A}, x)$ that preserves sufficient information of x w.r.t. current query.

This algorithm visit every relevant ABox assertion at most once, and it terminates after all related assertions are visited.

5. Related Work

The idea of most specific concept for instance checking was first discussed in [18], and later extensively studied by [17] [20] for the algorithms and the computational complexity. To deal with existential restrictions when computing the most specific concept, [24] [28] [29] discussed the use of cyclic concepts with greatest fixpoint semantics for preservation of information induced by the role assertions, and [27] also proposed an approximation for most specific concept in DLs with existential restrictions.

On the other hand, for efficient ABox reasoning and instance checking, various optimization techniques have been developed, including lazy unfolding, absorption, heuristic guided search, exploration of Horn clauses of DLs [4] [5] [22], model merging [2] and extended pseudo model merging technique [3] [30].

A common direction of these optimization techniques is to reduce the high degree of nondeterminism that is mainly introduced by GCIs in the TBox: given an GCI $C \sqsubseteq D$, it can be converted to a disjunction $C \sqcup \neg D$, for which a tableau algorithm will have to nondeterministically choose one of the disjuncts for tableau expansion, resulting in an exponential-time behavior of the tableau algorithm w.r.t. the data size. Absorption optimizations [22] [31] [32] were developed to reduce such nondeterminism by combining GCIs for unfoldable concepts, such that the effectiveness of lazy unfolding can be maximized. For example, axioms $A \sqsubseteq C$ and

1. In this recursive procedure, if x has already been visited before (cycle detected), mark x as the joint node and return $\{x^t\}$.
2. Obtain a set S of all explicit assertions in \mathcal{A} of the given individual x , which have not been visited before.
3. For every role assertion $\gamma : R_i(x, y_i) \in S$ (respectively $R_i(y_i, x)$) that has SYN_COND^* satisfied and has not been visited yet, invoke this procedure *recursively* to compute concept D_i for y_i . The rolling-up in this branch for x then yields $C_\gamma = \exists R_i.D_i$ (respectively $\exists R_i^-.D_i$).
4. For every concept assertion $\gamma : C_i(x) \in S$, $C_\gamma = C_i$.
5. Return $\text{MSC}_{\mathcal{T}}(\mathcal{A}, x)$ that equals to:

$$\begin{cases} \top & \text{if } S = \emptyset \\ \bigcap_{\gamma \in S} C_\gamma \sqcap \{x^h\} & \text{if } S \neq \emptyset, \text{ and } x \text{ is marked} \\ \bigcap_{\gamma \in S} C_\gamma & \text{if otherwise} \end{cases}$$

Figure 2. A recursive procedure for computation of $\text{MSC}_{\mathcal{T}}(\mathcal{A}, x)$.

³Note that $\exists R.C \sqsubseteq D$ is equivalent with $\exists R.\neg D \sqsubseteq \neg C$.

$A \sqsubseteq D$ can be combined into $A \sqsubseteq C \sqcap D$, where A is a named concept; then the inference engine can deduce $C \sqcap D(a)$ if the ABox contains $A(a)$. Notice however, this technique may allow only parts of TBox axioms to be absorbed, thus, may not be able to eliminate all sources of nondeterminism especially when ontologies are complex. Based on the absorption optimization, [33] proposed an approach for efficient ABox reasoning for \mathcal{ALCIQ} that will convert ABox assertions into TBox axioms, apply a absorption technique on the TBox, and covert instance retrieval into concept satisfaction problems.

Another way to reduce nondeterminism is the exploration of Horn clauses in DLs, since there exist reasoning techniques for Horn clauses that can be deterministic [5] [34]. [5] takes advantage of this in their HerMiT reasoner by preprocessing a DL ontology into DL-clauses and invoking the hyperresolution for the Horn clauses, avoiding unnecessary nondeterministic handling of Horn problems in existing DL tableau calculi.

For non-Horn DL, techniques such as model merging [2] and pseudo model merging [30] can be used to capture some deterministic information of named individuals. These techniques are based on the assumption of a consistent ABox and the observation that usually individuals are members of a small number of concepts. The (pseudo) model merging technique merges clash-free tableau models that are constructed by disjunction rules for a consistent ABox, and can figure out individuals that are obviously non-instance of a given concept. For example, if in one tableau model individual a belongs to concept C while in another a belongs to $\neg C$, it is then obvious that individual a cannot be deterministically inferred to be an instance of concept C , thus, eliminating the unnecessary instance checking for $C(a)$.

Another option for scalable ABox reasoning is the use of tractable DL languages. For example, the description logic EL and its extension EL^{++} , which allow existential restrictions and conjunction as introduced by [35] [36], possess intriguing algorithmic properties such that the satisfiability problem and implication in this DL language can be determined in polynomial time. Another notable example of lightweight DLs is the so-called $DL-LITE$ family identified by [37], which is specifically tailored to capture basic DL properties and expressivity while still be able to achieve low computational complexity for both TBox and ABox reasoning. In [9] [38] they further identified that, for conjunctive queries that are FOL-reducible, answering them in ontologies of any DL-LITE logic enjoys a LOGSPACE data complexity.

Based on the above lightweight DLs, efficient DL reasoners are developed, such as OWLIM [39], ELK reasoner [40], and Oracle's native inference engine for RDF data sets [41].

[42] proposed an approximation technique for instance retrieval, which computes both lower bound and upper bound of an answer set of individuals for a given query concept. Their approach invokes an axiom rewriting procedure that converts an ontology in Horn DL into a datalog program, and then uses Oracle's native inference engine to derive the bounds for query answering.

Recently, techniques for partitioning or modularizing ABoxes into logically-independent fragments have been developed [15] [16]. These techniques partition ABoxes into logically-independent modules, such that each will preserve complete information of a given set of individuals, and thus can be reasoned independently w.r.t. the TBox and be able to take advantage of existing parallel-processing techniques.

6. Empirical Evaluation

We implemented the rolling-up procedures for computation of $MSC_{\mathcal{T}}$'s based on the OWL API⁴, and evaluated the MSC method for instance checking and retrieving on a lab PC with Intel(R) Xeon(R) 3.07 GHz CPU, Windows 7, and 1.5 GB Java Heap. For the test suite, we have collected a set of well-known ontologies with large ABoxes:

- 1) LUBM(s) (LM) are benchmark ontologies generated using the tool provided by [43];
- 2) Arabidopsis thaliana (AT) and Caenorhabditis elegans (CE) are two biomedical ontologies⁵, sharing a common TBox called *Biopax* that models biological pathways;
- 3) DBpedia* (DP) ontologies are *extended* from the original DBpedia ontology [44]: expressivity of their TBox is extended from \mathcal{ALF} to \mathcal{SHI} by adding complex roles and concepts defined on role restrictions; their ABoxes are obtained by random sampling on the original triple store.

Details of these ontologies can be found in **Table 1**, in terms of DL expressivity, number of atomic concepts (# Cpts), TBox axioms (# Axms), named individuals (# Ind.), and ABox assertions (# Ast.). Notice that, DL

⁴<http://sourceforge.net/projects/owlapi>

⁵<http://www.reactome.org/download>

Table 1. Information of tested ontologies.

Ontology	Expressivity	# Cpts.	# Axms.	# Ind.	# Ast.
LM1	\mathcal{SHI}	43	42	17,175	67,465
LM2	\mathcal{SHI}	43	42	78,579	319,714
AT	\mathcal{SHI}	59	344	42,695	117,405
CE	\mathcal{SHI}	59	344	37,162	105,238
DP1	\mathcal{SHI}	449	465	273,663	402,062
DP2	\mathcal{SHI}	449	465	298,103	419,505

expressivity of AT and CE is originally \mathcal{SHIN} , but in our experiments, number restrictions (*i.e.* \mathcal{N}) in their ontology TBox are removed.

6.1. Complexity of $\text{MSC}_{\mathcal{T}}$

Using the MSC (or $\text{MSC}_{\mathcal{T}}$) method, the original instance checking problem is converted to a subsumption test, the complexity of which could be computationally high w.r.t. both size of a TBox and size of the testing concepts [10]. Therefore, when evaluating the rolling-up procedure for computation of $\text{MSC}_{\mathcal{T}}$'s, one of the most important criteria is the size of each resultant $\text{MSC}_{\mathcal{T}}$, as it is the major factor to the time efficiency of a subsumption test, given a relatively static ontology TBox.

As we already know, one of the major source of complexity in ontology reasoning is the so-called “*and-branching*”, which introduces new individuals in the tableau expansion through the \exists -rule, and affects the searching space of the reasoning algorithm as discussed in [10]. Thus, when evaluating sizes of computed $\text{MSC}_{\mathcal{T}}$'s, we measure both the level of nested quantifiers (*i.e.* quantification depth) and the number of conjuncts of each $\text{MSC}_{\mathcal{T}}$. For example, the concept

$$\exists R_1.C_1 \sqcap \exists R_2.(C_2 \sqcap \exists R_3.C_3)$$

has quantification depth 2 and number of conjuncts 2 (*i.e.* $\exists R_1.C_1$ and $\exists R_2.(C_2 \sqcap \exists R_3.C_3)$).

6.1.1. Experiment Setup

To evaluate and show efficacy of the proposed strategy and optimization, we have implemented the following three versions of the rolling-up method for comparison:

V1 The original rolling-up procedure adapted to ABox assertions without applying the call-by-need strategy, which computes the most specific concept w.r.t. \mathcal{A} for a given individual;

V2 The rolling-up procedure with the proposed *call-by-need* strategy based on SYN_COND, which features the branch pruning as fully discussed in Section 4.2;

V3 The rolling-up procedure with the *call-by-need* strategy based on SYN_COND* discussed in Section 4.3.

We compute the $\text{MSC}_{\mathcal{T}}$ for each individual in every ontology using the three methods respectively, and report in **Table 2** and **Table 3** the maximum and the average of quantification depth and number of conjuncts of the concepts, respectively. We also demonstrate the running-time efficiency of the optimized rolling-up procedure by showing the average time spent on computation of a $\text{MSC}_{\mathcal{T}}$ for each individual in **Figure 3**.

6.1.2. Result Analysis

As we can see from **Table 2** and **Table 3**, the sizes of $\text{MSC}_{\mathcal{T}}$'s generated by V2 and V3 are significantly smaller than those generated by V1 (the original method), which are almost in the same scale of size of the corresponding ontology ABox. The large size of $\text{MSC}_{\mathcal{T}}$'s from V1 is caused by the fact that most individuals (greater than 99%) in each of these ontologies are connected together by role paths in the graph. The bulk of each $\text{MSC}_{\mathcal{T}}$ makes the original MSC method completely inefficient and unscalable for answering object queries, as a subsumption test based on these concepts would be prohibitively expensive as a complete ABox reasoning. Thus, the comparison here reflects the potential and the importance of our proposed optimizations in

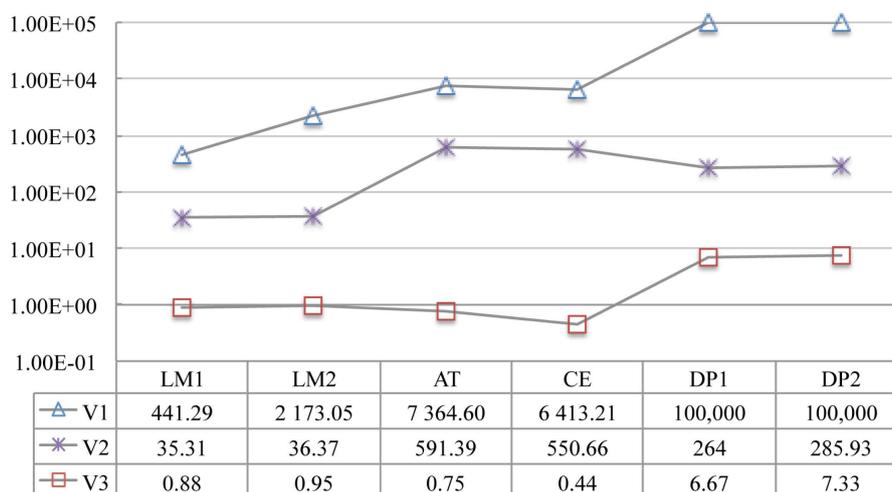


Figure 3. Average time (ms) on computation of a MSC_T . Timeout is set to be 100,000 ms.

Table 2. Quantification depth of MSC_T 's from different rolling-up procedures.

	V1	V2	V3
Ontology	Max./Avg.	Max./Avg.	Max./Avg.
LM1	5103/4964.68	215/98.8	2/1.48
LM2	23,015/22654.01	239/103.59	2/1.51
AT	2605/2505.50	1008/407.97	8/3.02
CE	3653/3553.4	1106/437.18	8/2.76
DP1	3906/3070.80	50/2.98	4/1.13
DP2	3968/3865.60	58/2.94	5/1.12

Table 3. Number of conjuncts of MSC_T 's from different rolling-up procedures.

	V1	V2	V3
LM1	104/31.34	8/3.58	4/1.56
LM2	203/64.71	8/3.78	4/1.59
AT	88/87.99	19/5.92	13/1.97
CE	52/50.70	16/6.25	17/1.94
DP1	33,591/26864.50	71/5.19	12/1.83
DP2	60,047/60011.60	64/3.95	12/1.79

this paper, which revive the MSC (*i.e.* MSC_T method as an efficient way for instance checking and object query answering.

The comparison between V2 and V3 demonstrates the efficacy of the optimization technique discussed in Section 4.3, which could prevent the rolling-up in arbitrary directions by providing a more rigorous precondition based on SYN_COND . This optimization could be useful in many practical ontologies, especially when their ABoxes contain “hot-spots” individuals that connect (tens of) thousands of individuals together and could cause the rolling-up to generate concepts with a prohibitive quantification depth.

In particular, in our previous study of modularization for ontology ABoxes [16], the biomedical ontologies

(i.e. AT and CE) are found to be complex with many of their ontology roles (33 out of 55) used for concept definitions, and their ABoxes are hard to be modularized even with various optimization techniques applied [16]. However, in this paper, we found much simpler $MSC_{\mathcal{T}}$'s can also be achieved in these complex ontologies when the optimization (i.e. SYN_COND*) is applied. For example, the maximum quantification depths of computed $MSC_{\mathcal{T}}$'s in both AT and CE are decreased significantly from more than 1000 to less than 10. Nevertheless, it should also be noted that, effectiveness of this optimization may vary on different ontologies, depending on their different levels of complexity and different amount of explicit information in their ABoxes that can be explored for optimization.

6.2. Reasoning with $MSC_{\mathcal{T}}$

In this section, we will show the efficiency that can be achieved when using the computed $MSC_{\mathcal{T}}$ for instance checking and retrieving. We conduct the experiments on the collected ontologies, and measure the average reasoning time that is required when performing instance checking (for every ABox individual) and instance retrieval using the $MSC_{\mathcal{T}}$ method, respectively.

6.2.1. Experiment Setup

We will not compare our method with a particular optimization technique for ABox reasoning, such as lazy unfolding, absorption, or model merging, etc., since they have already been built into existing reasoners and it is usually hard to control reasoners to switch on or off a particular optimization technique. Additionally, the $MSC_{\mathcal{T}}$ method still relies on the reasoning services provided by the state-of-art reasoners. Nevertheless, we do compare the reasoning efficiency between the $MSC_{\mathcal{T}}$ method and a regular complete ABox reasoning using existing reasoners, but only to show the effectiveness of the proposed $MSC_{\mathcal{T}}$ method for efficient instance checking and data retrieving. Moreover, we also compare the $MSC_{\mathcal{T}}$ method with the ABox partitioning method (*modular reasoning*) developed in [16], as they are developed based on the similar principles and both allow parallel or distributed reasoning.

The $MSC_{\mathcal{T}}$'s here are computed using algorithm V3, and the ABox partitioning technique used is the most optimized one presented in [16]. For a regular complete ABox reasoning, the reasoners used are OWL DL reasoners, Hermit [5] and Pellet [6], each of which has its particular optimization techniques implemented for the reasoning algorithm. Both the $MSC_{\mathcal{T}}$ method and the modular reasoning are based on reasoner Hermit, and they are not parallelized but instead running in an arbitrary *sequential* order of $MSC_{\mathcal{T}}$'s or ABox partitions.

Queries: LUBM comes with 14 standard queries. For biomedical and DBpedia* ontologies respectively, queries listed in Figure 4 are used.

```

Query on biomedical ontologies :
Q 1 = Control  $\sqcap$   $\forall$ controlled.Catalysis  $\sqcap$   $\forall$ controller.PhysicalEntity
Q 2 = Interaction  $\sqcap$   $\forall$ participant.PhysicalEntity
Q 3 = Interaction  $\sqcap$   $\exists$ interactionType.T  $\sqcap$   $\exists$ participant.T
     $\sqcap$   $\forall$ participant.Gene  $\sqcap$   $\exists$ phenotype.T
Q 4 = PhysicalEntity  $\sqcap$   $\forall$ entityReference.DnaReference
     $\sqcap$   $\forall$ memberPhysicalEntity.Dna
Q 5 = PhysicalEntity  $\sqcap$   $\exists$ entityReference.SmallMoleculeReference
     $\sqcap$   $\exists$ feature.BindingFeature  $\sqcap$   $\exists$ notfeature.BindingFeature
     $\sqcap$   $\exists$ memberPhysicalEntity.SmallMolecule

Query on DBpedia* ontologies :
Q 1 = Person  $\sqcap$   $\exists$ nationality.(Country  $\sqcap$   $\exists$ officialLanguage.English)
Q 2 = Music  $\sqcap$   $\exists$ composer.MusicalArtist
Q 3 = Person  $\sqcap$   $\exists$ child.Human  $\sqcap$   $\exists$ spouse.Person
Q 4 = Event  $\sqcap$   $\exists$ commander.Person  $\sqcap$   $\exists$ place.City
Q 5 = Produce  $\sqcap$   $\exists$ manufacturer.(Organization  $\sqcap$   $\exists$ place.Country)

```

Figure 4. Queries for biomedical and DBpedia* ontologies.

For each test ontology, we run the reasoning for each of the given queries. We report the average reasoning time spent on instance checking (Figure 5) and instance retrieval (Figure 6), respectively. The reasoning time reported here does not include the time spent for resource initialization (*i.e.* ontology loading and reasoner initialization), since the initialization stage can be done offline for query answering. However, it is obvious that the MSC_T method should be more efficient, since it only requires to load an ontology TBox while a regular ABox reasoning requires to load an entire ontology (including large ABoxes). For reasoning with MSC_T 's and ABox partitions, any updates during the query answering procedure (e.g. update the reasoner for different ABox partitions or different MSC_T 's) is counted into the reasoning time.

Another point worth noting here is that, for answering object queries using either modular reasoning or the MSC_T method, the overhead (time for ABox modularization or computation of MSC_T 's) should be taken into account. However, as shown in previous section and in [16], this overhead is negligible comparing with the efficiency gained on the reasoning, not to mention when these two methods get parallelized using existing frameworks such as MapReduce [45].

6.2.2. Result Analysis

As can be seen from the above two figures, using the MSC_T method, reasoning efficiency for both instance checking and instance retrieval in the testing ontologies has been improved significantly: 1) by more than three

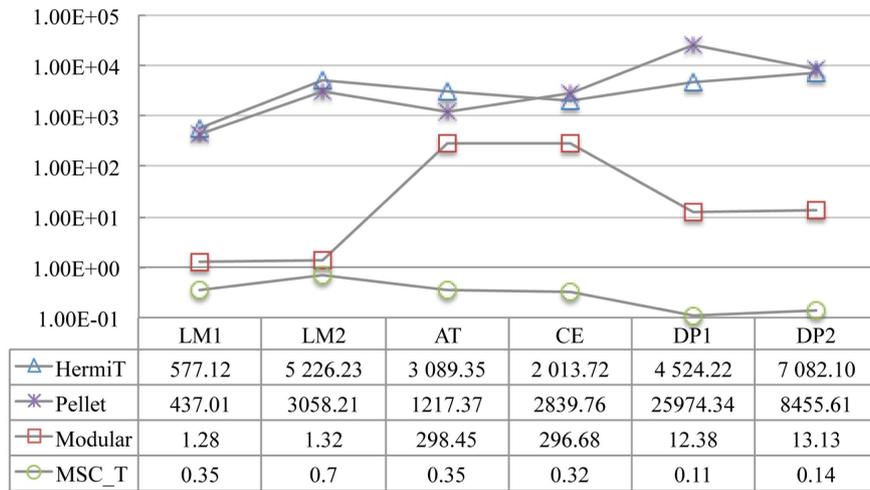


Figure 5. Average time (ms) on instance checking.

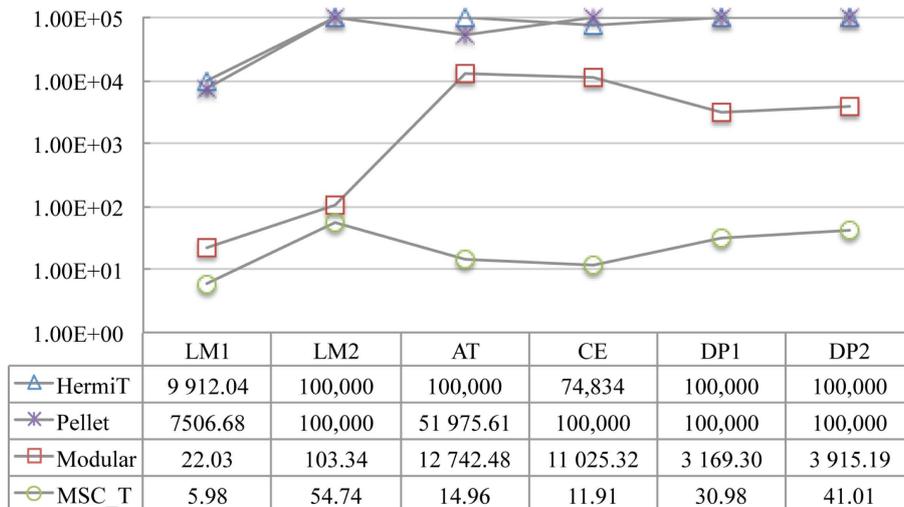


Figure 6. Average time(s) spent on instance retrieval. Timeout is set to be 100,000 s.

orders of magnitude when comparing with a complete reasoning; 2) and by about two orders of magnitude (except in LUBM1 and LUBM2) when comparing with the modular reasoning. For the latter, the improvement in LUBM1 and LUBM2 are not as significant as in others, which is because of the simplicity of these two ontologies that allows fine granularity of ABox partitions to be achieved [45].

On the other hand, using the $MSC_{\mathcal{T}}$ method in complex ontologies, such as AT and CE, the great improvement in reasoning efficiency comes from the reduction of searching space for reasoning algorithms, by branch pruning and also *concept absorption* during the computation of $MSC_{\mathcal{T}}$'s. For example, consider an individual x having the following n role assertions:

$$R(x, y_1), R(x, y_2), \dots, R(x, y_n)$$

where $y_i \in D$ and n tends to be large in these practical ontologies. Rolling up these assertions may generate a set of $\exists R.D$'s, the conjunction of which is still $\exists R.D$. Thus, when using this concept for instance checking, the interpretation may generate only one R -neighbor of individual x instead of n .

6.3. Scalability Evaluation

Using the $MSC_{\mathcal{T}}$ method for query answering over large ontologies is intended for distributed (parallel) computing. However, even if it is executed sequentially in a single machine, linear scalability may still be achieved on large ontologies that are not extremely complex; and there are mainly two reasons for that: first, the computation of $MSC_{\mathcal{T}}$'s focuses on only the query-relevant assertions instead of the entire ABox; second, the obtained $MSC_{\mathcal{T}}$'s could be very simple, sizes of which could be significantly smaller than that of the ABoxes. We test the scalability of this method for query answering (sequentially executed) using the benchmark ontology LUBM, which models organization of universities with each university constituted about 17,000 related individuals. The result is show in Figure 7.

7. Conclusions and Outlook

In this paper, we proposed a revised MSC method for efficient instance checking. This method allows the ontology reasoning to explore only a much smaller subset of ABox data that is relevant to a given instance checking problem, thus being able to achieve great efficiency and to solve the limitation of current memory-based reasoning techniques. It can be particularly useful for answering object queries over those large *non-Horn* DL ontologies, where existing optimization techniques may fall short and answering object queries may demand thousands or even millions of instance checking tasks. Most importantly, due to the independence between $MSC_{\mathcal{T}}$'s, scalability for query answering over huge ontologies (e.g. in the setting of semantic webs) could also be achieved by parallelizing the computations.

Our technique currently works for logic *SHI*, which is semi-expressive and is sufficient for many of the practical ontologies. However, the use of more expressive logic in modeling application domains requires more advanced technique for efficient data retrieving from ontology ABoxes. For the future work, we will investigate on how to extend the current technique to support *SHIN* or *SHIQ* that are featured with (qualified) number restrictions. We will concentrate on extending the rolling-up procedure to generate number restrictions, such as

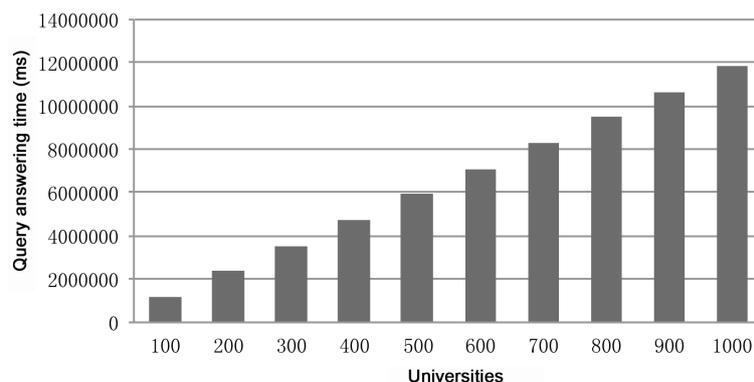


Figure 7. Scalability evaluation.

$\geq nR.T$ or $\geq nR.C$, whenever there is a need. We will also have to take a particular care of the identical individual problem, where concepts and role assertions of an individual can be derived via individual equivalence.

Acknowledgements

This work is partly supported by grant # R44GM097851 from the National Institute of General Medical Sciences, part of the U.S. National Institutes of Health (NIH).

References

- [1] Horrocks, I. (2008) Ontologies and the Semantic Web. *Communications of the ACM*, **51**, 58-67. <http://dx.doi.org/10.1145/1409360.1409377>
- [2] Horrocks, I.R. (1997) Optimising Tableaux Decision Procedures for Description Logics. Ph.D. Dissertation, University of Manchester, Manchester.
- [3] Haarslev, V. and Möller, R. (2008) On the Scalability of Description Logic Instance Retrieval. *Journal of Automated Reasoning*, **41**, 99-142. <http://dx.doi.org/10.1007/s10817-008-9104-7>
- [4] Motik, B., Shearer, R. and Horrocks, I. (2007) Optimized Reasoning in Description Logics using Hypertableaux. *Proceedings of Conference on Automated Deduction (CADE)*, **4603**, 67-83.
- [5] Motik, B., Shearer, R. and Horrocks, I. (2009) Hypertableau Reasoning for Description Logics. *Journal of Artificial Intelligence Research*, **36**, 165-228.
- [6] Sirin, E., Parsia, B., Grau, B.C., Kalyanpur, A. and Katz, Y. (2007) Pellet: A Practical Owl-DI Reasoner. *Journal of Web Semantics*, **5**, 51-53. <http://dx.doi.org/10.1016/j.websem.2007.03.004>
- [7] Haarslev, V. and Möller, R. (2001) RACER System Description. *Proceedings of the First International Joint Conference on Automated Reasoning*, Siena, June 2001, 701-705.
- [8] Horrocks, I. (1998) Using an Expressive Description Logic: Fact or Fiction? *Proceedings of Knowledge Representation and Reasoning*, **98**, 636-645.
- [9] Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. (2007) Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family. *Journal of Automated Reasoning*, **39**, 385-429. <http://dx.doi.org/10.1007/s10817-007-9078-x>
- [10] Donini, F. (2007) The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, Cambridge.
- [11] Glimm, B., Horrocks, I., Lutz, C. and Sattler, U. (2008) Conjunctive Query Answering for the Description Logic. *Journal of Artificial Intelligence Research*, **31**, 157-204.
- [12] Ortiz, M., Calvanese, D. and Eiter, T. (2008) Data Complexity of Query Answering in Expressive Description Logics via Tableaux. *Journal of Automated Reasoning*, **41**, 61-98. <http://dx.doi.org/10.1007/s10817-008-9102-9>
- [13] Tobies, S. (2001) Complexity Results and Practical Algorithms for Logics in Knowledge Representation. Ph.D. Dissertation, RWTH Aachen, Aachen.
- [14] Guo, Y. and Heflin, J. (2006) A Scalable Approach for Partitioning OWL Knowledge Bases. *International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS)*, Georgia, November 2006, 636-641.
- [15] Wandelt, S. and Möller, R. (2012) Towards ABox Modularization of Semi-Expressive Description Logics. *Applied Ontology*, **7**, 133-167.
- [16] Xu, J., Shironoshita, P., Visser, U., John, N. and Kabuka, M. (2013) Extract ABox Modules for Efficient Ontology Querying. ArXiv e-Prints. arXiv: 1305.4859
- [17] Donini, F., Lenzerini, M., Nardi, D. and Schaerf, A. (1994) Deduction in Concept Languages: From Subsumption to Instance Checking. *Journal of Logic and Computation*, **4**, 423-452. <http://dx.doi.org/10.1093/logcom/4.4.423>
- [18] Nebel, B. (1990) Reasoning and Revision in Hybrid Representation Systems. Vol. 422, Springer-Verlag, Germany.
- [19] Schaerf, A. (1994) Reasoning with Individuals in Concept Languages. *Data and Knowledge Engineering*, **13**, 141-176. [http://dx.doi.org/10.1016/0169-023X\(94\)90002-7](http://dx.doi.org/10.1016/0169-023X(94)90002-7)
- [20] Donini, F. and Era, A. (1992) Most Specific Concepts for Knowledge Bases with Incomplete Information. *Proceedings of CIKM*, Baltimore, November 1992, 545-551.
- [21] Tsarkov, D. and Horrocks, I. (2006) FaCT++ Description Logic Reasoner: System Description. *Proceedings of 3rd International Joint Conference on Automated Reasoning*, Seattle, 17-20 August 2006.
- [22] Horrocks, I. and Sattler, U. (2007) A Tableau Decision Procedure for Mathcal SHOIQ. *Journal of Automated Reasoning*, **39**, 249-276. <http://dx.doi.org/10.1007/s10817-007-9079-9>

- [23] Horrocks, I. and Tessaris, S. (2000) A Conjunctive Query Language for Description Logic ABoxes. *Proceedings of AAAI*, Austin, August 2000, 399-404.
- [24] Baader, F. and Küsters, R. (1998) Computing the Least Common Subsumer and the Most Specific Concept in the Presence of Cyclic ALN-Concept Descriptions. In: Herzog, O. and Gunter, A., Eds., *KI-98: Advances in Artificial Intelligence*, Springer, Bremen, 129-140.
- [25] Krötzsch, M., Rudolph, S. and Hitzler, P. (2008) Description Logic Rules. *European Conference on AI*, **178**, 80-84.
- [26] Horrocks, I., Sattler, U. and Tobies, S. (2000) Reasoning with Individuals for the Description Logic SHIQ. *Proceedings of Automated Deduction (CADE)*, Pittsburgh, June 2000, 482-496.
- [27] Küsters, R. and Molitor, R. (2001) Approximating Most Specific Concepts in Description Logics with Existential Restrictions. In: Baader, Franz, Brewka, Gerhard, Eiter and Thomas, Eds., *KI 2001: Advances in Artificial Intelligence*, Springer, Vienna, 33-47.
- [28] Baader, F. (2003) Least Common Subsumers and Most Specific Concepts in a Description Logic with Existential Restrictions and Terminological Cycles. *International Joint Conference on Artificial Intelligence*, **3**, 319-324.
- [29] Baader, F., Küsters, R. and Molitor, R. (1999) Computing Least Common Subsumers in Description Logics with Existential Restrictions. *International Joint Conference on Artificial Intelligence*, **99**, 96-101.
- [30] Haarslev, V., Möller, R. and Turhan, A.Y. (2001) Exploiting Pseudo Models for TBox and ABox Reasoning in Expressive Description Logics. *International Joint Conference, IJCAR 2001*, Siena.
- [31] Hudek, A.K. and Weddell, G. (2006) Binary Absorption in Tableaux-Based Reasoning for Description Logics. *Proceedings of the International Workshop on Description Logics (DL 2006)*, **189**, 86-96.
- [32] Tsarkov, D. and Horrocks, I. (2004) Efficient Reasoning with Range and Domain Constraints. *Proceedings of the 2004 Description Logic Workshop (DL 2004)*, **104**, 41-50.
- [33] Wu, J., Hudek, A.K., Toman, D. and Weddell, G.E. (2012) Assertion Absorption in Object Queries over Knowledge Bases. *International Conference on the Principles of Knowledge Representation and Reasoning*, Rome, June 2012.
- [34] Groszof, B., Horrocks, I., Volz, R. and Decker, S. (2003) Description Logic Programs: Combining Logic Programs with Description Logic. *Proceedings of WWW*, Budapest, May 2003, 48-57.
- [35] Baader, F., Brand, S. and Lutz, C. (2005) Pushing the EL Envelope. *Proceedings of IJCAI*, Edinburgh, August 2005, 364-369.
- [36] Baader, F., Brandt, S. and Lutz, C. (2008) Pushing the EL Envelope Further. *Proceedings of the OWLED 2008 DC Workshop on OWL: Experiences and Directions*, Karlsruhe, October 2008.
- [37] Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. (2005) DL-Lite: Tractable Description Logics for Ontologies. *Proceedings of AAAI*, **5**, 602-607.
- [38] Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. (2006) Data Complexity of Query Answering in Description Logics. *Proceedings of Knowledge Representation and Reasoning (KR)*, **6**, 260-270.
- [39] Bishop, B., Kiryakov, A., Ognyanoff, D., Peikov, I., Tashev, Z. and Velkov, R. (2011) OWLIM: A Family of Scalable Semantic Repositories. *Journal of Semantic Web*, **2**, 33-42.
- [40] Kazakov, Y., Krötzsch, M. and Simančík, F. (2011) Concurrent Classification of EL Ontologies. *International Semantic Web Conference*, Bonn, October 2011, 305-320.
- [41] Wu, Z., Eadon, G., Das, S., Chong, E.I., Kolovski, V., Annamalai, M. and Srinivasan, J. (2008) Implementing an Inference Engine for RDFS/OWL Constructs and User-Defined Rules in Oracle. *Proceedings of IEEE 24th International Conference on Data Engineering (ICDE)*, Cancun, April 2008, 1239-1248.
- [42] Zhou, Y., Cuenca Grau, B., Horrocks, I., Wu, Z. and Banerjee, J. (2013) Making the Most of Your Triple Store: Query Answering in OWL 2 Using an RL Reasoner. *Proceedings of the 22nd International Conference on World Wide Web*, Rio, May 2013, 1569-1580.
- [43] Guo, Y., Pan, Z. and Heflin, J. (2005) LUBM: A Benchmark for OWL Knowledge Base Systems. *Journal of Web Semantics*, **3**, 158-182. <http://dx.doi.org/10.1016/j.websem.2005.06.005>
- [44] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R. and Ives, Z. (2007) DBpedia: A Nucleus for a Web of Open Data. *Proceedings of ISWC*, Busan, November 2007, 722-735.
- [45] Dean, J. and Ghemawat, S. (2008) MapReduce: Simplified Data Processing on Large Clusters. *Communications of the ACM*, **51**, 107-113. <http://dx.doi.org/10.1145/1327452.1327492>

Scientific Research Publishing (SCIRP) is one of the largest Open Access journal publishers. It is currently publishing more than 200 open access, online, peer-reviewed journals covering a wide range of academic disciplines. SCIRP serves the worldwide academic communities and contributes to the progress and application of science with its publication.

Other selected journals from SCIRP are listed as below. Submit your manuscript to us via either submit@scirp.org or [Online Submission Portal](#).

