




Contents lists available at ScienceDirect

Data & Knowledge Engineering

journal homepage: www.elsevier.com/locate/datak

Editorial

Revising the constraints of lightweight mediated schemas

Marco A. Casanova^{a,*}, Tanara Lauschner^b, Luiz André P. Paes Leme^c, Karin K. Breitman^a, Antonio L. Furtado^a, Vânia M.P. Vidal^d

^a Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro-Rio de Janeiro, RJ, Brazil

^b Departamento de Ciência da Computação-Universidade Federal do Amazonas-Manaus, AM, Brazil

^c Instituto de Computação, Universidade Federal Fluminense, Niterói, RJ, Brazil

^d Departamento de Computação, Universidade Federal do Ceará, Fortaleza, CE, Brazil

ARTICLE INFO

Article history:

Received 12 March 2010

Received in revised form 18 September 2010

Accepted 4 October 2010

Available online xxx

Keywords:

Constraint revision

Mediated schema

Description logics

ABSTRACT

In this article, we address the problem of changing the constraints of a mediated schema to accommodate the set of constraints of a new export schema. The relevance of this problem lies in that the constraints of a mediated schema capture the common semantics of the data sources and, as such, they must be maintained and made available to the users of the mediation environment. We first argue that such problem can be solved by computing the greatest lower bound of two theories induced by sets of constraints, defined as the intersection of the theories. Then, for an expressive family of conceptual schemas, we show how to efficiently decide logical implication and how to compute the greatest lower bound of two theories induced by sets of constraints. The family of conceptual schemas we work with partly corresponds to OWL Lite and supports the equivalent of named classes, datatype and object properties, minCardinalities and maxCardinalities, InverseFunctionalProperties, subset constraints, and disjointness constraints. Such schemas are also sufficiently expressive to encode commonly used UML constructs, such as classes, attributes, binary associations without association classes, cardinality of binary associations, multiplicity of attributes, and ISA hierarchies with disjointness, but not with complete generalizations.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

A *mediation environment* contains *export schemas* E_1, \dots, E_n , which describe data sources, *import schemas* I_1, \dots, I_n , such that I_i is a view over E_i , and a *mediated schema* M , which intuitively combines I_1, \dots, I_n . For each import schema I_i , the environment features and a *local mapping* γ_i that defines the concepts of I_i in terms of the concepts of E_i . The environment also has a *mediated mapping* γ that defines the concepts of M in terms of those of I_1, \dots, I_n . Fig. 1 depicts these notions.

The constraints of the mediated schema are relevant for a correct understanding of what the semantics of the external schemas have in common. For example, consider a virtual store mediating access to online booksellers. The class hierarchy of the mediated schema indicates what the booksellers' book classifications have in common; if the mediated schema enforces that all books must have ISBNs, then it means that all booksellers must abide by the same requirement; if it allows books with no (known) authors, then at least one bookseller must so allow; and so on.

In this article, we focus on the process of adding to the mediation environment a new export schema E_0 , with import schema I_0 and local mapping γ_0 . We may break this process into three steps. The *concept revision step* adjusts the vocabulary of M to perhaps include classes and properties originally defined in I_0 . The *mapping revision step* may modify the mediated mapping. Finally, the *constraint revision step* applies a minimum set of changes to the set of constraints of M to account for the set of constraints of I_0 .

* Corresponding author.

E-mail addresses: casanova@inf.puc-rio.br (M.A. Casanova), tanara@dcc.ufam.edu.br (T. Lauschner), lpaesleme@ic.uff.br (L.A.P. Paes Leme), karin@inf.puc-rio.br (K.K. Breitman), furtado@inf.puc-rio.br (A.L. Furtado), vvidal@lia.ufc.br (V.M.P. Vidal).




Fig. 1. Components of a mediation environment.

One may have to iterate through these three steps since, in particular, revising the constraints of the mediated schema may interfere with the definition of the schema mappings. For example, the local mapping γ_0 may have to be adjusted to preserve the class hierarchy of the mediated schema, or the class hierarchy of the mediated schema may have to be changed to reflect the class hierarchy of I_0 .

In this article, we are primarily concerned with the constraint revision step, with a bias to mediation environments in the context of the Web. Maintaining mediation environments in such context becomes a challenge because the number of data sources may be very large and, moreover, the mediator does not have much control over the data sources, which may join or leave the mediation environment at will.

We break the constraint revision step in two sub-steps. The *constraint translation step* translates the set EC_0 of constraints of E_0 to I_0 , creating a set of constraints IC_0 in such a way that γ_0 maps states of E_0 that satisfy EC_0 into states of I_0 that satisfy IC_0 . Intuitively, as a result of this step, we express the semantics of E_0 in terms of I_0 .

The *least constraint change step* applies a minimum set of changes to the constraints of M to accommodate IC_0 in such a way that all schema mappings remain correct. This step intuitively means to harmonize the semantics of E_0 with the semantics of all export schemas previously added to the mediation environment, captured in the constraints of M . The key questions here are how to precisely define what it means to apply a minimum set of changes to a set of constraints, and how to guarantee that the mappings remain correct.

The contribution of this article is twofold. First, we formulate the problem of changing the constraints of the mediated schema as the problem of computing the greatest lower bound (g.l.b.) of the theories induced by two sets of constraints (the g.l.b. of two theories is defined as their intersection). Second, for an expressive family of conceptual schemas, we show how to efficiently decide logical implication and how to compute a representation of the g.l.b. of two theories induced by two sets of constraints.

In more detail, we work with schemas that partly correspond to OWL Lite [1] and support the equivalent of named classes, datatype and object properties, minCardinalities and maxCardinalities, InverseFunctionalProperties, subset constraints, and disjointness constraints. The schemas we work with are also sufficiently expressive to encode commonly used UML constructs, such as classes, attributes, binary associations without association classes, cardinality of binary associations, multiplicity of attributes, and ISA hierarchies with disjointness, but not with complete generalizations.

The decision procedure for logical implication and the procedure to compute a representation of the g.l.b. of two theories induced by two sets of constraints are based on the satisfiability algorithm for Boolean formulas in conjunctive normal form with at most two literals per clause, described in [2]. The intuition is that the constraints we consider can be treated much in the same way as Boolean implications. However, cardinality constraints pose considerable technical problems to the proof of the theorems. The decision procedure essentially explores the structure of a set of constraints, captured as a graph. The procedure to compute the greatest lower bound of two theories induced by two sets of constraints is a direct consequence of the decision procedure. These results are new, and cover an expressive and useful family of constraints, outlined in the previous paragraph.

This article is organized as follows. Section 2 surveys related work. Section 3 reviews concepts of Description Logics and introduces the notion of mediation environment. Section 4 shows how to generate the revised set of constraints of the mediated schema and presents the proofs for the main results. Section 5 contains the conclusions.

2. Related work

Research on the construction of mediated schemas concentrates on vocabulary matching techniques, on the definition of schema mappings, and on query processing, mostly ignoring the question of constraint revision.

Matching techniques are useful for the process of revising the vocabulary of the mediated schema, a topic we do not directly address, but mention on Section 4.1. Euzenat and Shvaiko [3] present a comprehensive survey of ontology matching. Rahm and Bernstein [4] survey schema matching, and Bernstein and Melnik [5] list the requirements for model management systems that support the matching process. Köpcke and Rahm [6] comparatively analyze eleven frameworks for entity matching.

Schema matching techniques may be classified as syntactic, semantic, or hybrid. For example, Melnik et al. [7] and Madhavan et al. [8] describe syntactic techniques based on modeling the schemas as graphs. Bilke and Naumann [9] propose a semantic technique based on an analysis of duplicated instances. Brauner et al. [10] adopt this strategy to align thesauri. Wang et al. [11] describe a semantic technique based on probing the databases.

Departing from this classification, Qi and Ling [12] present algorithms to resolve schematic discrepancies by transforming metadata into the attribute values of entity types, keeping the information and constraints of original schemas. Zhao and Ram [13] propose an iterative procedure for detecting both schema-level and instance-level matchings from heterogeneous data sources.

Schema and ontology reuse, as proposed in Lonsdale et al. [14] and in Simperla [15], is a fruitful strategy to overcome interoperability issues. The use of templates to help exchange schemas, as proposed in Papott and Torlone [16], is a similar strategy that may also be used to circumvent interoperability problems.

As for the mappings between the external schema and the mediated schema, two basic approaches have been used [17]. The first approach, called *global-as-view* (GAV), requires that the mediated schema be expressed in terms of the data sources. More precisely, a view over the data sources is associated with each element of the global schema, so that the meaning of the element is specified in terms of the data stored at the data sources. This means that adding a new data source may impact the previously defined mappings, which may need to be updated. Several projects, such as TSIMMIS [18], IBIS [19] and INFOMIX [20] adopt the GAV approach.

The second approach, called *local-as-view* (LAV), requires that the mediated schema be specified independently from the data sources. The data sources are in turn defined as views over the mediated schema [21]. This means that adding a new data source only requires adding a new assertion to the mediated mapping. This approach improves maintainability and extensibility of the systems [9]. Pícsel [22] is an example of a LAV system.

Mappings may also be classified according to their accuracy into sound, complete and exact [17,23]. Let V be a view associated with an element E of the mediated schema. In the GAV approach, V is *sound* when all data provided by V satisfies E , but there may be additional data satisfying E that V does not provide. View V is *complete* when not all data provided by V satisfies E , but all data satisfying E is provided by V . Finally, V is *exact*, when all data provided by V satisfies E , and all data satisfying E is provided by V [23].

Rull et al. [24] present an approach for validating schema mappings that allows the mapping designer to ask whether they have certain desirable properties.

The approach we take in Section 3.2 to define the mediation environment is akin to the idea of sound views. Yet, we consider that constraints should be included in the mediated schema to capture the common semantics of the data sources, unlike most proposals based on the concept of exact views, which assume that the mediated schema has no constraints, as observed in [17].

Calì et al. [23] argue that the constraints of a mediated schema should be taken into account during query processing and that the schema definition language should incorporate flexible and powerful representation mechanisms for integrity constraints. The authors also argue that, when the mediated schema contains constraints, the semantics of the data integration system is best described in terms of a set of databases, and that query processing should be based on the notion of querying incomplete databases.

Calvanese et al. [25] introduce a Description Logics framework, similar to that in Section 3.1, to address schema integration and query answering. Atzeni et al. [26] cover the problem of rewriting a schema from one model to another, but they do not touch on the more complex problem of generating a new set of constraints that generalizes a pair of sets of constraints from different schemas, which we address in Section 4. Hick and Hainaut [27] show how requirements changes are propagated to database schemas, to data and to programs through a general strategy. Hartmann et al. [28] apply techniques from Propositional Logic to offer decision support for specifying Boolean and multivalued dependencies.

When compared with the DL-Lite family, as aptly condensed in [1], the schemas we consider treat *maxCardinality* as a negated form of *minCardinality* and formulate concept inclusions and disjunctions in such a way that negated descriptions occur only on the right-hand side of inclusions. This and other limitations are formulated as restrictions on the schema constraints themselves, and not on the descriptions that occur in the constraints, as in [1]. These limitations restrict the interaction between the schema constraints, and are easily grasped by the database designer. While retaining expressiveness, they permit using a novel approach that leads to polynomial time decision procedures that do not destroy the constraint structure of a schema.

In more detail, the subsumption problem in Description Logics (DL) refers to the question of deciding if a concept description always denotes a subset of the set denoted by another concept description. The subsumption problem is decidable for expressive dialects of DL, but typically belongs to hard complexity classes [29], especially in the presence of axioms (or constraints) [30]. For certain dialects of DL, there are polynomial time decision procedures for the subsumption problem that explore the structure of the concept descriptions and that are, for this reason, called *structural subsumption procedures* [31,32]. However, such procedures do not take axioms into account. Furthermore, the reductions suggested to encode the axioms lead us back to dialects for which the subsumption problem is hard [30].

From the point of view of deciding logical implication and computing a representation of the g.l.b. of the theories induced by two sets of constraints, we depart from the tradition of Description Logics deduction services, which are mostly based on tableaux techniques [29]. The decision procedure described in Section 4.2 is based on the satisfiability algorithm for Boolean formulas in conjunctive normal form with at most two literals per clause, described in [2]. The procedure to compute a representation of the g. l.b. of the theories induced by two sets of constraints is a direct consequence of the decision procedure. These results also depend on the notion of Herbrand interpretation for DL, defined in Section 4.3.

3. Mediation environment

In this section, we introduce the notation used to describe schema mappings and constraints, and define what we mean by extralite schemas and mediation environments.

3.1. A brief review of concepts from description logics

We adopt a family of *attributive languages* [33] defined as follows. A *language* \mathcal{L} in the family is characterized by an *alphabet* \mathcal{A} , consisting of a set of *atomic concepts*, a set of *atomic roles*, the *universal concept* and the *bottom concept*, denoted by \top and \perp , respectively, the *universal role* and the *bottom role*, also denoted by \top and \perp , respectively, and a set of *constants*.

The set of *role descriptions* of \mathcal{L} is inductively defined as

- An atomic role, and the universal and bottom roles are role descriptions
- If p and q are role descriptions, then the following expressions are role descriptions

p^- (the inverse of p)
 $p \circ q$ (the composition of p and q)
 $p \sqcup q$ (the union of p and q)

The set of *concept descriptions* of \mathcal{L} is inductively defined as

- An atomic concept, and the universal and bottom concepts are concept descriptions
- If a_1, \dots, a_n are constants, then $\{a_1, \dots, a_n\}$ is a concept description
- If e and f are concept descriptions and p is a role description, then the following expressions are concept descriptions

$\neg e$ (negation)
 $e \sqcap f$ (intersection)
 $e \sqcup f$ (union)
 $\exists p$ (existential quantification)
 $\exists p.e$ (full existential quantification)
 $\forall p.e$ (value restriction)
 $(\leq n p)$ (at most restriction)
 $(\geq n p)$ (at least restriction)

Role inverse, concept negation, existential quantification, at most restriction and at least restriction are required to express constraints, and will be extensively used in Section 4. The other types of expressions will be exclusively used to describe schema mappings, and are required only in the examples. Also, the universal and the bottom roles will be used exclusively to express schema mappings. Albeit not standard, the overloading of the symbols \top and \perp will not create notational ambiguities and their use will be kept to a minimum.

An *interpretation* \mathbf{s} for \mathcal{A} consists of a nonempty set $\Delta^{\mathbf{s}}$, the *domain* of \mathbf{s} , whose elements are called *individuals*, and an *interpretation function*, also denoted \mathbf{s} , where:

- $\mathbf{s}(\perp) = \emptyset$, when \perp denotes the bottom concept or the bottom role
- $\mathbf{s}(\top) = \Delta^{\mathbf{s}}$, when \top denotes the universal concept
- $\mathbf{s}(\top) = \Delta^{\mathbf{s}} \times \Delta^{\mathbf{s}}$, when \top denotes the universal role
- $\mathbf{s}(A) \subseteq \Delta^{\mathbf{s}}$, for each atomic concept A of \mathcal{L}
- $\mathbf{s}(P) \subseteq \Delta^{\mathbf{s}} \times \Delta^{\mathbf{s}}$, for each atomic role P of \mathcal{L}
- $\mathbf{s}(a) \in \Delta^{\mathbf{s}}$, for each constant a of \mathcal{L} , such that distinct constants denote distinct individuals (the *uniqueness assumption*)

The function \mathbf{s} is extended to role and concept descriptions of \mathcal{L} as follows:

- $\mathbf{s}(p^-) = \mathbf{s}(p)^-$ (the inverse of $\mathbf{s}(p)$)
- $\mathbf{s}(p \circ q) = \mathbf{s}(p) \circ \mathbf{s}(q)$ (the composition of $\mathbf{s}(p)$ with $\mathbf{s}(q)$)
- $\mathbf{s}(p \sqcup q) = \mathbf{s}(p) \cup \mathbf{s}(q)$ (the union of $\mathbf{s}(p)$ with $\mathbf{s}(q)$)
- $\mathbf{s}(\{a_1, \dots, a_n\}) = \{\mathbf{s}(a_1), \dots, \mathbf{s}(a_n)\}$ (the set consisting of the individuals $\mathbf{s}(a_1), \dots, \mathbf{s}(a_n)$)
- $\mathbf{s}(\neg e) = \Delta^{\mathbf{s}} - \mathbf{s}(e)$ (the complement of $\mathbf{s}(e)$ w.r.t. $\Delta^{\mathbf{s}}$)
- $\mathbf{s}(e \sqcap f) = \mathbf{s}(e) \cap \mathbf{s}(f)$ (the intersection of $\mathbf{s}(e)$ and $\mathbf{s}(f)$)
- $\mathbf{s}(e \sqcup f) = \mathbf{s}(e) \cup \mathbf{s}(f)$ (the union of $\mathbf{s}(e)$ and $\mathbf{s}(f)$)
- $\mathbf{s}(\exists p) = \{I \in \Delta^{\mathbf{s}} / (\exists J \in \Delta^{\mathbf{s}})((I, J) \in \mathbf{s}(p))\}$ (the set of individuals that $\mathbf{s}(p)$ relates to some individual)
- $\mathbf{s}(\exists p.e) = \{I \in \Delta^{\mathbf{s}} / (\exists J \in \Delta^{\mathbf{s}})((I, J) \in \mathbf{s}(p) \wedge J \in \mathbf{s}(e))\}$ (the set of individuals that $\mathbf{s}(p)$ relates to some individual in $\mathbf{s}(e)$)
- $\mathbf{s}(\forall p.e) = \{I \in \Delta^{\mathbf{s}} / (\forall J \in \Delta^{\mathbf{s}})((I, J) \in \mathbf{s}(p) \Rightarrow J \in \mathbf{s}(e))\}$
 (the set of individuals I such that, if $\mathbf{s}(p)$ relates I to an individual J , then J is in $\mathbf{s}(e)$)
- $\mathbf{s}(\geq n p) = \{I \in \Delta^{\mathbf{s}} / |\{J \in \Delta^{\mathbf{s}} / (I, J) \in \mathbf{s}(p)\}| \geq n\}$ (the set of individuals that $\mathbf{s}(p)$ relates to at least n distinct individuals)
- $\mathbf{s}(\leq n p) = \{I \in \Delta^{\mathbf{s}} / |\{J \in \Delta^{\mathbf{s}} / (I, J) \in \mathbf{s}(p)\}| \leq n\}$ (the set of individuals that $\mathbf{s}(p)$ relates to at most n distinct individuals)

A *formula* of \mathcal{L} is an expression of the form $u \sqsubseteq v$, called an *inclusion*, or of the form $u \sqcup v$, called a *disjunction*, or of the form $u \equiv v$, called an *equivalence*, where u and v are both concept descriptions or they are both role descriptions of \mathcal{L} . A *definition* is an equivalence of the form $T \equiv u$, where T is an atomic concept and u is a concept description, or T is an atomic role and u is a role description.

An interpretation \mathbf{s} for \mathcal{L} satisfies $u \sqsubseteq v$ iff $\mathbf{s}(u) \subseteq \mathbf{s}(v)$, \mathbf{s} satisfies $u \mid v$ iff $\mathbf{s}(u) \cap \mathbf{s}(v) = \emptyset$, and \mathbf{s} satisfies $u \equiv v$ iff $\mathbf{s}(u) = \mathbf{s}(v)$. We adopt the following familiar notation, where σ is a formula and Σ and Γ are sets of formulas:

- $\mathbf{s} \models \sigma$ indicates that \mathbf{s} satisfies σ
- $\mathbf{s} \models \Sigma$ indicates that \mathbf{s} satisfies all formulas in Σ ; in this case, we say that \mathbf{s} is a *model* of Σ
- Σ is *satisfiable* iff there is a model of Σ
- $\Sigma \models \sigma$ indicates that any model of Σ satisfies σ ; in this case, we say that Σ *logically implies* σ
- $\Sigma \models \Gamma$ indicates that any model of Σ is also a model of Γ ; in this case, we say that Σ *logically implies* Γ
- $Th(\Sigma)$ denotes the *theory induced* by Σ , which is the smallest set of formulas that contains Σ and is closed under logical implication.

Also, in Sections 3 and 4, we will use concept and role descriptions over an alphabet \mathcal{A} which is the union of disjoint alphabets $\mathcal{A}_1, \dots, \mathcal{A}_n$. The syntax of concept and role descriptions remains the same. An interpretation \mathbf{s} for \mathcal{A} is constructed from interpretations $\mathbf{s}_1, \dots, \mathbf{s}_n$ for $\mathcal{A}_1, \dots, \mathcal{A}_n$ in the obvious way, except that we assume that

- (*Domain Disjointness Assumption*) Any pair of interpretations for \mathcal{A}_i and \mathcal{A}_j have disjoint domains, for each $i, j \in [1, n]$, with $i \neq j$

3.2. Extralite schemas

We will work with *extralite schemas* that partly correspond to OWL Lite [1]. Extralite schemas support the equivalent of named classes, datatype and object properties, minCardinalities and maxCardinalities, InverseFunctionalProperties, which capture simple keys, subset constraints, and disjointness constraints. Extralite schemas are also sufficiently expressive to encode commonly used UML constructs, such as classes, attributes, binary associations without association classes, cardinality of binary associations, multiplicity of attributes, and ISA hierarchies with disjointness, but not with complete generalizations.

Formally, an *extralite schema* is a pair $S = (\mathcal{A}, C)$ such that

- \mathcal{A} is an alphabet, called the *vocabulary* of S , whose atomic concepts and atomic roles are called the *classes* and *properties* of S , respectively
- C is a set of formulas, called the *constraints* of S , which must be of one of the forms
 - *Domain Constraint*: $\exists P \sqsubseteq D$ (property P has class D as domain)
 - *Range Constraint*: $\exists P^- \sqsubseteq R$ (property P has class R as range)
 - *minCardinality constraint*: $C \sqsubseteq (\geq k P)$ or $C \sqsubseteq (\geq k P^-)$
(property P or its inverse P^- maps each individual in class C to at least k distinct individuals)
 - *maxCardinality constraint*: $C \sqsubseteq (\leq k P)$ or $C \sqsubseteq (\leq k P^-)$
(property P or its inverse P^- maps each individual in class C to at most k distinct individuals)
 - *Subset Constraint*: $E \sqsubseteq F$ (class E is a subclass of class F)
 - *Disjointness Constraint*: $E \perp F$ (classes E and F are disjoint).

We also admit constraints of one of the forms:

- $C \sqsubseteq \perp$ (class C is always empty)
- $\exists P \sqsubseteq \perp$ or $\exists P^- \sqsubseteq \perp$ (property P is always empty, i.e., P has an empty domain or an empty range)

We will use the terms *class*, *property*, *vocabulary* and *state* interchangeably with *atomic concept*, *atomic role*, *alphabet* and *interpretation*, respectively. In the examples that follow, we note that the data types, such as String, Decimal, etc. should also be treated as classes.

Example 1. Fig. 2 contains schemas for fragments of the Amazon and the eBay databases, using the namespace prefixes “a:” and “e:” to refer to their vocabularies, respectively.

Fig. 2(a) and (c) show the schemas using an informal notation. Fig. 2(b) and (d) formalize the constraints: the first column shows the domain and range constraints; the second column depicts the cardinality constraint; and the third column contains the subset and disjointness constraints.

For example, the first column of Fig. 2(b) indicates that:

- a:title is a property with domain a:Product and range string (the set of XML Schema strings)
- a:pub is a property with domain a:Book and range a:Publ

The second column of Fig. 2(b) shows the cardinalities of the Amazon schema:

- all properties have maxCardinality equal to 1, except a:author, a:pub and a:city
- a:author has unbounded maxCardinality, consistently with the fact that a book may have multiple authors (hence, a:author has no maxCardinality constraint)
- a:pub has minCardinality equal to 2
- a:city has minCardinality equal to 3

The third column of Fig. 2(b) indicates that a:Book and a:Music are subclasses of a:Product, and that a:Book and a:Music are disjoint classes.

<pre>a:Product a:title range string a:price range decimal a:currency range string a:Book a:isbn range string a:author range string a:pub range a:Publ a:Publ a:name range string a:city range string</pre>	<pre>a:Book is-a a:Product a:Music is-a a:Product a:Book disjoint-from a:Music</pre>	
(a)		
<pre>∃ a:title ⊆ a:Product ∃ a:title ⊆ string ... ∃ a:pub ⊆ a:Book ∃ a:pub ⊆ a:Publ</pre>	<pre>a:Product ⊆ (≤ 1 a:title) a:Product ⊆ (≤ 1 a:price) a:Product ⊆ (≤ 1 a:currency) a:Book ⊆ (≤ 1 a:isbn) a:Book ⊆ (≥ 2 a:pub) a:Publ ⊆ (≤ 1 a:name) a:Publ ⊆ (≥ 3 a:city)</pre>	<pre>a:Book ⊆ a:Product a:Music ⊆ a:Product a:Book a:Music</pre>
(b)		
<pre>string integer string string integer integer string</pre>	<pre>e:Seller e:name range string e:Offer e:qty range integer e:price range double e:currency range string e:seller range e:Seller e:product range e:Product</pre>	<pre>e:Product e:type range e:ean range e:title range e:author range e:edition range e:year range e:place range</pre>
(c)		
<pre>disjointness constraints)</pre>	<pre>∃ e:name ⊆ e:Seller ∃ e:name ⊆ string ... ∃ e:seller ⊆ e:Offer ∃ e:seller ⊆ e:Seller ∃ e:product ⊆ e:Offer ∃ e:product ⊆ e:Product ...</pre>	<pre>e:Seller ⊆ (≤ 1 e:name) e:Offer ⊆ (≤ 1 e:qty) e:Offer ⊆ (≤ 1 e:price) ... e:Product ⊆ (≤ 1 e:type) e:Product ⊆ (≤ 1 e:ean) e:Product ⊆ (≤ 1 e:title) ... e:Product ⊆ (≥ 1 e:place)</pre>
(d)		

Fig. 2. (a). Informal definition of the Amazon schema. (b). Formal definition of (some of) the constraints of the Amazon schema. (c). Informal definition of the eBay schema. (d). Formal definition of (some of) the constraints of the eBay schema.

Fig. 2(d) likewise describes the constraints of the eBay schema. In particular, the second column indicates that all properties have maxCardinality equal to 1, except `e:place`. □

3.3. Components of a mediation environment

A mediation environment contains a mediated schema M , a mediated mapping γ and, for each $k = 1, \dots, n$, an export schema E_k , an import schema I_k and a local mapping γ_k .

As mentioned in the introduction, we stress that import schemas are a notational convenience to divide the definition of the mappings into two stages: the definition of the local mappings and the definition of the mediated mapping. We restrict the import schemas as follows:

- (1) for $k = 1, \dots, n$, the vocabulary of I_k is equal to the vocabulary of M , in the sense that the two vocabularies have the same classes and properties, but different namespaces.

Assume that the classes and properties in M are C_1, \dots, C_u and P_1, \dots, P_v . We adopt namespace prefixes, as in the examples, to distinguish the occurrence of a symbol in the vocabulary of M from the occurrence of the same symbol in the vocabulary of I_k . However, in the formal development, we follow a more abstract notation. For each class C_i (or property P_j) in the vocabulary of M , we denote the occurrence of C_i (or P_j) in the vocabulary of I_k by C_i^k (or P_j^k), and say that C_i^k (or P_j^k) matches C_i (or P_j).

For each $k = 1, \dots, n$, the local mapping γ_k defines the classes and properties of I_k in the terms of the vocabulary of the export schema E_k . We restrict γ_k as follows:

- for each class C_i^k of I_k , the local mapping γ_k contains a definition of the form

$$(2) C_i^k \equiv \rho_i^k$$

where ρ_i^k is a concept description over the vocabulary of E_k

- for each property P_j^k of I_k , the local mapping γ_k contains a definition of the form

$$(3) P_j^k \equiv \pi_j^k$$

where π_j^k is a role description over the vocabulary of E_k

Note that ρ_i^k may be the bottom concept \perp to indicate that E_k does not contribute with any individual to class C_i^k . In other words, the interpretation of C_i^k is always an empty set. Combined with the requirement that the vocabulary of I_k be equal to the vocabulary of M , this might seem an unnecessary complication. However, these technical details simplify the computation of the revised set of constraints of a mediated schema. Likewise, π_j^k may be the bottom role \perp , when E_k does not contribute with any individual to P_j^k .

We introduce $\bar{\gamma}_k$ as the function induced by γ_k , defined as the function from states of E_k into states of I_k such that, for each state s of E_k , $\bar{\gamma}_k(s) = r$ iff

- $r(C_i^k) = s(\rho_i^k)$, if $C_i^k \equiv \rho_i^k$ is the definition for class C_i^k in γ_k
- $r(P_j^k) = s(\pi_j^k)$, if $P_j^k \equiv \pi_j^k$ is the definition for property P_j^k in γ_k

For each $k = 1, \dots, n$, let EC_k be the set of constraints of E_k . The set IC_k of constraints of the import schema I_k should be defined so that $\bar{\gamma}_k$ maps consistent states of E_k into consistent states of I_k . We refer the reader to Lauschner et al. [34] for efficient strategies to generate IC_k , when EC_k is the family of schema constraints considered in Section 3.2 and the local mapping γ_k uses an expressive family of concept and role expressions.

We illustrate the concepts just introduced with the help of an example.

Example 2. Consider the Sales mediated schema with the vocabulary shown in Fig. 4(a), distinguished by the namespace prefix “s:”.

Fig. 3(a) defines the vocabulary of the Amazon import schema, which is equal to that of the Sales mediated schema, but is identified by the namespace prefix “ai:”. Fig. 3(b) contains the translation of the constraints of the Amazon export schema, shown in Fig. 2(b), to the Amazon import schema. Fig. 3(c) contains the local mapping that defines the concepts of the vocabulary of the Amazon import schema in terms of the concepts of the vocabulary of the Amazon export schema of Fig. 2(a).

For example, the definitions $ai:city \equiv a:pub \circ a:city$ and $ai:Book \equiv a:Book$ have several consequences. First, the domain and range of $ai:city$ are $ai:Book$ and $string$. Second, $ai:city$ has $minCardinality$ 3 with respect to $ai:Book$ since, observing Fig. 2(b), $a:pub$ has $minCardinality$ 2 with respect to $a:Book$, $a:city$ has $minCardinality$ 3 with respect to $a:Publ$, and $a:Publ$ is both the range of $a:pub$ and the domain of $a:city$.

Intuitively, in Fig. 2(b), we assumed that each book is associated with at least 2 publishers and that each publisher is located in at least 3 cities, which are not necessarily distinct from the cities associated with other publishers. Hence, all we can assert is that each book is associated with at least 3 publishers' cities, which is expressed in Fig. 3(b) by the $minCardinality$ constraint for cities with respect to books. As a concrete example, suppose that: (1) the book “Semantic Web” is associated with two publishers, “Springer Verlag” and “Ed. Campus”; (2) “Springer Verlag” is located in three cities “London”, “Berlin” and “Sidney”; “Ed. Campus” is also located in “London”, “Berlin” and “Sidney”. Note that these individuals do not violate the cardinality constraints of the Amazon export schema. Then, the book “Semantic Web” is associated with three cities, “London”, “Berlin” and “Sidney”.

The other constraints of the Amazon import schema follow directly from those of the Amazon export schema, since each of the other classes and properties of the import schema is defined in terms of a single class or property of the export schema.

Fig. 3(d) defines the vocabulary of the eBay import schema, which is again equal to that of the Sales mediated schema, but is identified by the namespace prefix “ei:”. Fig. 3(e) contains the translation of the constraints of the eBay export schema, shown in Fig. 2(c), to the eBay import schema. Fig. 3(f) contains the local mapping for the eBay export schema of Fig. 2(c).

In particular, observe that, in Fig. 3(f), $ei:Music$ and $ei:Book$ are defined as restrictions of $e:Product$ (given an atomic concept A , a restriction of A is an intersection of the form $A \sqcap e$). As a consequence, we have the two subset constraints and the disjointness constraint shown on the third column of Fig. 3(e), albeit the original eBay schema has no such constraints (see Fig. 2(d)). Note that the disjointness constraint requires assuming that distinct constants denote distinct individuals. \square

We now complete the description of a mediation environment with the definition of the mediated mapping. We restrict a mediated mapping γ as follows:

- for each $i = 1, \dots, u$, the mapping γ contains a definition of the form

$$(4) C_i \equiv C_i^1 \sqcup \dots \sqcup C_i^u$$

where C_i^k is the class of I_k that matches C_i (which always exists by (1)), for each $k = 1, \dots, n$

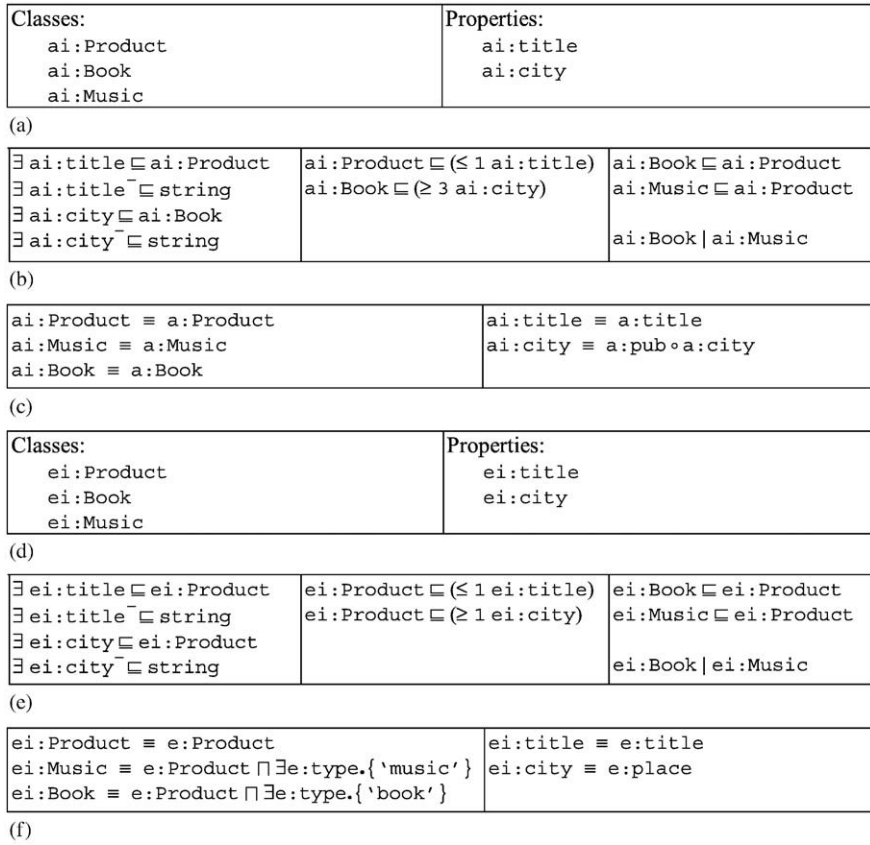


Fig. 3. (a). Vocabulary of the Amazon import schema. (b). Constraints of the Amazon import schema. (c). Local mapping from the Amazon export schema to Amazon import schema. (d). Vocabulary of the eBay import schema. (e). Constraints of the eBay import schema. (f). Local mapping from the eBay export schema to the eBay import schema.

• for each $j = 1, \dots, v$, the mapping γ contains a definition of the form

$$(5) P_j \equiv P_j^1 \sqcup \dots \sqcup P_j^v$$

where P_j^k is the property of I_k that matches P_j (which always exists by (1)), for each $k = 1, \dots, v$

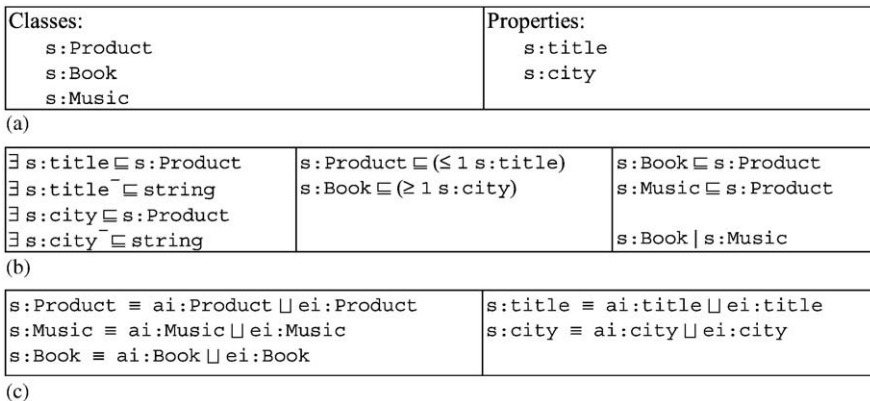


Fig. 4. (a). Vocabulary of the Sales mediated schema. (b). Constraints of the Sales mediated schema. (c). Mediated mapping.

We introduce $\bar{\gamma}$ as the *function induced* by the mediated mapping γ and the local mapping $\gamma_1, \dots, \gamma_n$ as the mapping from states of E_1, \dots, E_n into states of M such that, for states $\mathbf{s}_1, \dots, \mathbf{s}_n$ of E_1, \dots, E_n , $\bar{\gamma}(s_1, \dots, s_n) = \mathbf{r}$ iff, for $i = 1, \dots, u$ and $j = 1, \dots, v$

- $\mathbf{r}(C_i) = \mathbf{s}_i(C_i^1) \cup \dots \cup \mathbf{s}_n(C_i^n)$, if $C_i \equiv C_i^1 \sqcup \dots \sqcup C_i^n$ is the definition of C_i in γ
- $\mathbf{r}(P_j) = \mathbf{s}_i(P_j^1) \cup \dots \cup \mathbf{s}_n(P_j^n)$, if $P_j \equiv P_j^1 \sqcup \dots \sqcup P_j^n$ is the definition of P_j in γ

Example 3. A complete description of a mediation environment would be as follows:

- for the mediated schema sales
 - the vocabulary listed in Fig. 4(a)
 - the constraints shown in Fig. 4(b), whose construction is discussed in Example 4 in Section 4.1
 - the mediated mapping shown in Fig. 4(c)
- for the Amazon database fragment
 - the export schema shown in Fig. 2(a) and (b)
 - the import schema with the vocabulary listed in Fig. 3(a) and the constraints shown in Fig. 3(b)
 - the local mapping shown in Fig. 3(c)
- for the eBay database fragment:
 - the export schema shown in Fig. 2(c) and (d)
 - the import schema with the vocabulary listed in Fig. 3(d) and the constraints shown in Fig. 3(e)
 - the local mapping shown in Fig. 3(f).

□

4. Construction of the mediated schema constraints

4.1. Basic steps of the constraint revision process

Consider a mediation environment with mediated schema M and mediated mapping γ . Assume that MV is the vocabulary and MC is the set of constraints of M . Let E_0 be a new export schema, with vocabulary EV_0 and set of constraints EC_0 .

To create a revised mediation environment that includes E_0 , we treat M much in the same way as a data source, as follows:

1. (Concept revision step)

1.1. Define the vocabulary MV_r of the revised mediated schema M_r with the same classes and properties as MV and perhaps new classes and properties that reflect classes and properties in EV_0 that convey new information not captured in the current vocabulary.

1.2. Define a new vocabulary MV^+ by adding to MV these new classes and properties.

1.3. Define the vocabulary IV_0 of the import schema I_0 for E_0 with the same classes and properties as MV_r .

2. (Mapping revision step)

2.1. Define the local mapping γ_0 between I_0 and E_0 .

2.2. Define a new mediated mapping γ^+ by adding to γ definitions for the new classes and properties in MV^+ .

2.3. Define the mediated mapping γ_r as in Eq (4) and (5) on pages 7 and 8.

3. (Constraint revision step)

3.1. Define the set IC_0 of constraints of I_0 by inspecting EC_0 and γ_0 .

3.2. Define a new set of constraints MC^+ by adding to MC constraints for the new classes and properties in MV^+ .

3.3. Define the set of constraints MC_r of M_r by applying a minimum set of changes to MC^+ to account for IC_0 .

Step 3.3 is the main thrust of this article and is discussed in detail in this and the next sections. Steps 1.1, 1.2, 1.3 and 2.1 may be carried out by the automated matching process we discussed in [35–37]. Step 3.1 was discussed in [34]. Steps 2.2, 2.3 and 3.2 are quite simple, but raise a few points that we address in what follows.

As in Section 3.3, assume that the classes and properties in MV are C_i and P_j , for $i = 1, \dots, u$ and $j = 1, \dots, v$. Suppose that the classes and properties in MV_r are C_i^r and P_j^r , for $i = 1, \dots, u + p$ and $j = 1, \dots, v + q$. Then, for $i = u + 1, \dots, u + p$ and $j = v + 1, \dots, v + q$

- the new classes and properties in MV^+ are C_i and P_j , which *match* C_i^r and P_j^r
- the new definitions in γ^+ are $C_i \equiv \perp$ and $P_j \equiv \perp$
- the new constraints in MC^+ are $C_i \sqsubseteq \perp$, $\exists P_j \sqsubseteq \perp$ and $\exists (P_j)^- \sqsubseteq \perp$

Observe that the new constraints in MC^+ are a trivial consequence of the fact that, for $i = u + 1, \dots, u + p$ and $j = v + 1, \dots, v + q$, the new definitions in γ^+ force C_i and P_j to always have empty interpretations. In particular, the constraints for P_j capture that P_j is an empty property by saying that the domain and range of P_j are always empty. This strategy is necessary since the constraints we consider do not allow expressions of the form $P_j \sqsubseteq \perp$. Furthermore, note that it is redundant (though not wrong) to add constraints saying that both the domain and the range of P_j are always empty.

Also observe that IC_0 will likewise have a constraint of the form $C_i^0 \sqsubseteq \perp$, whenever γ_0 contains a definition of the form $C_i^0 \sqsubseteq \perp$, and constraints of the forms $\exists P_j^0 \sqsubseteq \perp$ and $\exists (P_j^0)^- \sqsubseteq \perp$, whenever γ_0 contains a definition of the form $P_j^0 \sqsubseteq \perp$. The revised mapping can then be written as follows:

- for each $i = 1, \dots, u + p$, the revised mediated mapping γ_r contains a definition of the form
 - (6) $C_i^r \sqsubseteq C_i^0 \sqcup C_i$ where C_i^0 is the class of I_0 that matches C_i^0 and C_i is the class of M that matches C_i^r
- for each $j = 1, \dots, v + q$, the revised mediated mapping γ_r contains a definition of the form
 - (7) $P_j^r \sqsubseteq P_j^0 \sqcup P_j$ where P_j^0 is the property of I_0 that matches P_j^0 and P_j is the property of M that matches P_j^r

We focus on how to create the revised set of constraints MC_r . The reader should bear in mind the notation just introduced, which will be used in what follows.

There are two questions here: (1) what it means to apply a minimum set of changes to a set of constraints; (2) how to maintain the correctness of the schema mappings. To address the first question, we introduce a lattice of sets of constraints.

Recall from Section 3.1 that $Th(\Phi)$ denotes the theory induced by a set of formulas Φ . Let \mathcal{T} be the set of all theories induced by sets of constraints. Then, $(\mathcal{T}, \sqsupseteq)$ is a lattice where, given any two sets of constraints, Φ_1 and Φ_2 , the least upper bound (l.u.b.) of their induced theories is $\Phi_1 \sqsupseteq \Phi_2 = Th(\Phi_1) \cup Th(\Phi_2)$ and the greatest lower bound (g.l.b.) of their induced theories is $\Phi_1 \sqcap \Phi_2 = Th(\Phi_1) \cap Th(\Phi_2)$.

In what follows, we will sometimes refer to the g.l.b. of two sets of constraints, rather than the g.l.b. of the theories induced by the sets of constraints.

We argue that MC_r can be taken as the g.l.b. of the translation of MC^+ to MV_r , and the translation of IC_0 to MV_r . Note that a translation step is necessary since, technically, no two constraints respectively in $Th(MC^+)$ and in $Th(IC_0)$ would be equal since they are written in different vocabularies. Intuitively, the translation would be just a matter of changing namespaces.

Let \mathcal{L}_1 and \mathcal{L}_2 be two languages with alphabets \mathcal{A}_1 and \mathcal{A}_2 , respectively.

- An injective mapping $\lambda: \mathcal{A}_1 \rightarrow \mathcal{L}_2$ is called a *substitution function* from \mathcal{A}_1 into \mathcal{L}_2 iff
 - $\lambda(\perp) = \perp$ and $\lambda(\top) = \top$
 - if s is an atomic concept of \mathcal{A}_1 and $\lambda(s) = e$ then e is a concept expression of \mathcal{L}_2
 - if s is an atomic role of \mathcal{A}_1 and $\lambda(s) = e$ then e is a role expression of \mathcal{L}_2
- The *translation* of a formula φ of \mathcal{L}_1 to \mathcal{L}_2 via λ is the formula of \mathcal{L}_2 , denoted by $\varphi[\lambda]$, obtained by replacing in φ each symbol A of \mathcal{A}_1 by $\lambda(A)$.
- The *translation* of a set of formulas Θ of \mathcal{L}_1 to \mathcal{L}_2 via λ is the set of formulas of \mathcal{L}_2 , denoted $\Theta[\lambda]$, obtained by translating each formula in Θ to \mathcal{L}_2 via λ .

In particular, the mediated mapping γ_r induces three *canonical substitution functions*:

- $\hat{\gamma}_r^0$ from IV_0 into MV_r such that $\hat{\gamma}_r^0(A) = B$ iff A is an atomic concept or an atomic role of IV_0 that occurs in the body of the definition for B in γ_r
- $\hat{\gamma}_r^+$ from MV^+ into MV_r such that $\hat{\gamma}_r^+(A) = B$ iff A is an atomic concept or an atomic role of MV^+ that occurs in the body of the definition for B in γ_r
- $\hat{\gamma}_r$ from MV_r into $IV_0 \cup MV^+$ such that $\hat{\gamma}_r(B) = e$ iff the definition of B in γ_r is $B \sqsubseteq e$

To improve the notation, we write the translation of a constraint φ of IC_0 from IV_0 to MV_r using $\hat{\gamma}_r^0$ as $\varphi[IV_0 \rightarrow MV_r]$, the translation of a constraint φ of MV^+ from MV^+ to MV_r using $\hat{\gamma}_r^+$ as $\varphi[MV^+ \rightarrow MV_r]$, and the translation of a constraint φ of M_r from MV_r to $IV_0 \cup MV^+$ using $\hat{\gamma}_r$ as $\varphi[MV_r \rightarrow IV_0 \cup MV^+]$.

Therefore, the translation of IC_0 to MV_r is the set of constraints $IC_0[IV_0 \rightarrow MV_r]$ and the translation of MC^+ to MV_r is the set of constraints $MC^+[MV^+ \rightarrow MV_r]$.

We are now ready to state that MC_r can be taken as the g.l.b. of $IC_0[IV_0 \rightarrow MV_r]$ and $MC^+[MV^+ \rightarrow MV_r]$ without impairing consistency preservation.

Theorem 1. Let $MC_r = IC_0[IV_0 \rightarrow MV_r] \sqcap MC^+[MV^+ \rightarrow MV_r]$. Suppose that:

- (i) (*Domain Disjointness Assumption*) Any pair of interpretations for EV_i and EV_j have disjoint domains.
- (ii) The mediated mapping γ and the local mapping $\gamma_1, \dots, \gamma_n$ induce a mapping from consistent states of E_1, \dots, E_n into consistent states of M .
- (iii) The local mapping γ_0 induces a mapping from consistent states of E_0 into consistent states of I_0 . Then, the revised mediated mapping γ_r and the local mappings $\gamma_0, \gamma_1, \dots, \gamma_n$ induce a mapping from consistent states of EC_0, EC_1, \dots, EC_n into states of the revised mediated schema that satisfy MC_r .

Proof. The proof depends on the definition of the mediated mapping with the help of union expressions, as in Eqs. (4) and (5), and on the Domain Disjointness Assumption, introduced at the end of Section 3.1. In detail, let $\sigma \in Th(MC_r)$. Then, by definition of g.l.b., we have:

- (1) $\sigma \in Th(IC_0[IV_0 \rightarrow MV_r])$
- (2) $\sigma \in Th(MC^+[MV^+ \rightarrow MV_r])$

But, by definition of the canonical translation functions, we have:

- (3) $\sigma \in Th(IC_0[IV_0 \rightarrow MV_r])$ iff $\sigma[MV_r \rightarrow IV_0] \in Th(IC_0)$
 (4) $\sigma \in Th(MC^+[MV^+ \rightarrow MV_r])$ iff $\sigma[MV_r \rightarrow MV^+] \in Th(MC^+)$

Let $k = 0, \dots, n$. Let \mathbf{s}_k be a consistent state of E_k . Since $\bar{\gamma}_k$ preserves consistency, $\bar{\gamma}_k(\mathbf{s}_k) = \mathbf{t}_k$ is a consistent state of I_k . Furthermore, since $\bar{\gamma}$ preserves consistency, $\bar{\gamma}(s_1, \dots, s_n) = \mathbf{s}$ is a consistent state of M . Note that, by definition of MC^+ , \mathbf{s} is also consistent with respect to MC^+ .

Therefore, we have

- (5) $\mathbf{s}_0 \models \sigma[MV_r \rightarrow IV_0]$
 (6) $\mathbf{s} \models \sigma[MV_r \rightarrow MV^+]$

Recall that $\sigma[MV_r \rightarrow MV^+ \cup IV_0]$ denotes the constraint obtained from σ by replacing each class C_i^r of MV_r by the union expression $C_i^0 \sqcup C_i$, where C_i^0 and C_i respectively are the classes of I_0 and M that match C_i^r , and likewise for the properties of MV_r . Note that $\sigma[MV_r \rightarrow MV^+ \cup IV_0]$ is a constraint written in $MV^+ \cup IV_0$, the union of the vocabularies MV^+ and IV_0 .

Let $\mathbf{s} \cup \mathbf{t}_0$ denote the interpretation for $MV^+ \cup IV_0$ induced by \mathbf{s} and \mathbf{t}_0 in the obvious way. Then, using the domain disjointness assumption, we can prove that:

- (7) $\mathbf{s} \cup \mathbf{t}_0 \models \sigma[MV_r \rightarrow MV^+ \sqcup IV_0]$

Now, by definition by definition of $\bar{\gamma}_r$, from (7), we finally have:

- (8) $\bar{\gamma}_r(\mathbf{s}, \mathbf{t}_0) \models \sigma$

Therefore, recalling that $\bar{\gamma}(s_1, \dots, s_n) = \mathbf{s}$, we finally have that $\bar{\gamma}_r(\bar{\gamma}(s_1, \dots, s_n), \mathbf{t}_0)$ is a consistent state of M_r , as desired. \square
 Since MC_r is defined as the g.l.b. of $IC_0[IV_0 \rightarrow MV_r]$ and $MC^+[MV^+ \rightarrow MV_r]$ with respect to (\mathcal{T}, \models) , we consider that MC_r is the best way to revise MC and to retain correctness of the mappings, in view of [Theorem 1](#). We now give an example that illustrates how the constraints of a mediated schema can be defined.

Example 4. We illustrate how the constraints of the Sales mediated schema, listed in [Fig. 4\(b\)](#), can be gradually constructed from the constraints of the Amazon and the eBay import schemas, shown in [Fig. 3\(b\)](#) and (e). Then, we discuss how to include a third import schema.

(A) Assume that the Sales mediation environment contains just the definition of the vocabulary listed in [Fig. 4\(a\)](#). Suppose that one wishes to add to the mediation environment the Amazon fragment described in [Fig. 2\(a\)](#) and (b), with the import schema defined in [Fig. 3\(a\)](#) and (b), and the local mapping introduced in [Fig. 3\(c\)](#).

Then, after this initial step, the Amazon import schema is treated as the mediated schema, and the mediated mapping is simply empty. Furthermore, the initial vocabulary of the mediated schema is in fact that of the Amazon import schema, identified by the namespace prefix “ai:”, with classes ai:Book, ai:Music and ai:Product, and properties ai:title and ai:city.

(B) Consider adding to the mediation environment the eBay fragment described in [Fig. 2\(c\)](#) and (d), with the import schema defined in [Fig. 3\(d\)](#) and (e), and the local mapping introduced in [Fig. 3\(f\)](#).

We perform three steps:

(Concept revision step) Assume for the sake of argument that no new classes or properties are added. Thus, the Sales vocabulary, now identified by the namespace prefix “s:”, has classes s:Book, s:Music and s:Product, and properties s:title and s:city.

(Mapping revision step) [Fig. 5\(a\)](#) shows the revised mediated mapping of the Sales mediation environment.

(Constraint revision step) Consider the following sets of constraints:

- Ψ_A, Ψ_E – the sets of constraints of the Amazon and eBay import schemas, shown in [Fig. 3\(b\)](#) and (e).
- Φ_A, Φ_E – the sets of constraints obtained by translating, respectively, Ψ_A and Ψ_E to the vocabulary of the mediated schema. The translation is simply a process that replaces ai:Product by s:Product, etc.

We stress that it does not make sense to compute the g.l.b. of Ψ_A and Ψ_E , since these constraints are written in different vocabularies. Therefore, we compute the g.l.b. of Φ_A and Φ_E , which are constraints in the same vocabulary (that of the mediated schema). Since

$$\Phi_A \Delta \Phi_E = Th(\Phi_A) \cap Th(\Phi_E)$$

$s:\text{Product} \equiv ai:\text{Product} \sqcup bi:\text{Product}$ $s:\text{Music} \equiv ai:\text{Music} \sqcup bi:\text{Music}$ $s:\text{Book} \equiv ai:\text{Book} \sqcup bi:\text{Book}$	$s:\text{title} \equiv ai:\text{title} \sqcup bi:\text{title}$ $s:\text{city} \equiv ai:\text{city} \sqcup bi:\text{city}$
---	---

(a)

$\exists s:\text{title} \sqsubseteq s:\text{Product}$ $\exists s:\text{title}^- \sqsubseteq \text{string}$ $\exists s:\text{city} \sqsubseteq s:\text{Product}$ $\exists s:\text{city}^- \sqsubseteq \text{string}$	$s:\text{Product} \sqsubseteq (\leq 1 s:\text{title})$ $s:\text{Book} \sqsubseteq (\geq 1 s:\text{city})$	$s:\text{Book} \sqsubseteq s:\text{Product}$ $s:\text{Music} \sqsubseteq s:\text{Product}$ $s:\text{Book} s:\text{Music}$
--	--	---

(b)

Classes: $b:\text{Product}$ $b:\text{CultProd}$ $b:\text{Music}$ $b:\text{Book}$	Properties: $b:\text{title}$
---	--

(c)

$\exists b:\text{title} \sqsubseteq b:\text{Product}$ $\exists b:\text{title}^- \sqsubseteq \text{string}$	(no cardinality constraints)	$b:\text{CultProd} \sqsubseteq b:\text{Product}$ $b:\text{Book} \sqsubseteq b:\text{CultProd}$ $b:\text{Music} \sqsubseteq b:\text{CultProd}$
---	------------------------------	---

(d)

Classes: $sr:\text{Product}$ $sr:\text{Music}$ $sr:\text{Book}$	Properties: $sr:\text{title}$ $sr:\text{city}$
---	---

(e)

Classes: $s:\text{Product}$ $s:\text{Music}$ $s:\text{Book}$	Properties: $s:\text{title}$ $s:\text{city}$
--	---

(f)

Classes: $bi:\text{Product}$ $bi:\text{Music}$ $bi:\text{Book}$	Properties: $bi:\text{title}$ $bi:\text{city}$
---	---

(g)

$bi:\text{Product} \equiv b:\text{Product}$ $bi:\text{Music} \equiv b:\text{Music}$ $bi:\text{Book} \equiv b:\text{Book}$	$bi:\text{title} \equiv b:\text{title}$ $bi:\text{city} \equiv \perp$
---	--

(h)

$sr:\text{Product} \equiv bi:\text{Product} \sqcup s:\text{Product}$ $sr:\text{Music} \equiv bi:\text{Music} \sqcup s:\text{Music}$ $sr:\text{Book} \equiv bi:\text{Book} \sqcup s:\text{Book}$	$sr:\text{title} \equiv bi:\text{title} \sqcup s:\text{title}$ $sr:\text{city} \equiv bi:\text{city} \sqcup s:\text{city}$
---	---

(i)

$\exists bi:\text{title} \sqsubseteq bi:\text{Product}$ $\exists bi:\text{title}^- \sqsubseteq \text{string}$ $\exists bi:\text{city} \sqsubseteq \perp$ $\exists bi:\text{city}^- \sqsubseteq \perp$		$bi:\text{Book} \sqsubseteq bi:\text{Product}$ $bi:\text{Music} \sqsubseteq bi:\text{Product}$
--	--	---

(j)

$\exists sr:\text{title} \sqsubseteq sr:\text{Product}$ $\exists sr:\text{title}^- \sqsubseteq \text{string}$ $\exists sr:\text{city} \sqsubseteq sr:\text{Product}$ $\exists sr:\text{city}^- \sqsubseteq \text{string}$	$\exists sr:\text{city} \sqsubseteq (\leq 1 s:\text{title})$	$sr:\text{Book} \sqsubseteq sr:\text{Product}$ $sr:\text{Music} \sqsubseteq sr:\text{Product}$
--	--	---

(k)

Fig. 5. (a). Revised mediated mapping of the Sales mediation environment. (b). Constraints of the Sales mediated schema. (c). Vocabulary of the BN export schema. (d). Constraints of the BN export schema. (e). Vocabulary of the Sales/BN mediated schema. (f). Vocabulary of the Sales mediated schema. (g). Vocabulary of BN import schema. (h). Local mapping from the BN export schema to the BN import schema. (i). Mediated mapping of the Sales/BN mediation environment. (j). Constraints of the BN import schema. (k). Constraints of the revised Sales/BN mediated schema.

we have to find the constraints that are simultaneously derivable from Φ_A and from Φ_E . For ease of reference, Fig. 5(b) repeats the constraints of the Sales mediated schema.

We first analyze in detail what minCardinality constraints for property $s:city$ are in $\Phi_A \Delta \Phi_E$. From Fig. 3(b) and (e), we have the following minCardinality constraints for $city$ in Ψ_A and Ψ_E :

- (1) $ai:Book \sqsubseteq (\geq 3 \text{ } ai:city)$ (in Ψ_A)
- (2) $ei:Product \sqsubseteq (\geq 1 \text{ } ei:city)$ (in Ψ_E)

We also have the following subset constraint in Ψ_E :

- (3) $ei:Book \sqsubseteq ei:Product$ (in Ψ_E)

When translated to the vocabulary of the mediated schema, identified by the prefix “s:”, the constraints in (1) to (3) become:

- (4) $s:Book \sqsubseteq (\geq 3 \text{ } s:city)$ (in Φ_A)
- (5) $s:Product \sqsubseteq (\geq 1 \text{ } s:city)$ (in Φ_E)
- (6) $s:Book \sqsubseteq s:Product$ (in Φ_E)

Hence, the only minCardinality constraint for property $s:city$ that is simultaneously derivable from Φ_A and Φ_E is

- (7) $s:Book \sqsubseteq (\geq 1 \text{ } s:city)$ (in $\Phi_A \Delta \Phi_E$)

Indeed, we have that:

- (4) implies (7), if we observe that a minCardinality of n implies a minCardinality of m , if $m \leq n$
- (5) and (6) imply (7)

By a simpler argument, we also have:

- (8) $s:Product \sqsubseteq (\leq 1 \text{ } s:title)$ (in $\Phi_A \Delta \Phi_E$)

The subset and disjointness constraints in $\Phi_A \Delta \Phi_E$ are those shown in the third column of Fig. 5(b); in fact, they are in the intersection of Φ_A and Φ_E .

The domain and range constraints in $\Phi_A \Delta \Phi_E$ are those shown in the first column of Fig. 5(b); in fact, they are in the intersection of Φ_A and Φ_E , except for the domain constraint $\exists s:city \sqsubseteq s:Product$, which is derived as follows. From Fig. 3(b) and (e), we have the following domain constraints in Ψ_A and Ψ_E :

- (9) $\exists ai:city \sqsubseteq ai:Book$ (in Ψ_A)
- (10) $\exists ei:city \sqsubseteq ei:Product$ (in Ψ_E)

We also have the following subset constraints in Ψ_A :

- (11) $ai:Book \sqsubseteq ai:Product$ (in Ψ_A)

When translated to the vocabulary of the mediated schema, once again, identified by the prefix “s:”, the constraints in (9) to (11) become:

- (12) $\exists s:city \sqsubseteq s:Book$ (in Φ_A)
- (13) $\exists s:city \sqsubseteq s:Product$ (in Φ_E)
- (14) $s:Book \sqsubseteq s:Product$ (in Φ_A)

Hence, the domain constraint for property $s:city$ that is simultaneously derivable from Φ_A and Φ_E is

- (15) $\exists s:city \sqsubseteq s:Product$ (in $\Phi_A \Delta \Phi_E$)

This illustrates the computation of the constraints of a mediated schema as the g.l.b. of the sets of constraints of the import schemas, after proper translation.

(C) Let BN be a new export schema (say, a fragment of the Barnes&Noble database), shown in Fig. 5(c) and (d).

To include BN in the Sales mediation environment, creating the Sales/BN mediation environment, we again perform three steps:

(Concept revision step) Assume for the sake of argument that the vocabulary of the Sales/BN mediated schema, with namespace “sr:”, as in Fig. 5(e), is equal to that of the Sales mediated schema. The vocabulary of the Sales mediated schema is still identified with namespace prefix “s:”, as in Fig. 5(f). The BN import schema has the vocabulary shown in Fig. 5(g).

(Mapping revision step) Fig. 5(h) shows the local mapping from the BN export schema to the BN import schema. Note that the definition $bi:city \sqsubseteq \perp$ indicates that property $bi:city$ will always be empty in the BN import schema. Fig. 5(i) depicts the mediated mapping of the Sales/BN mediation environment.

(Constraint revision step) Fig. 5(j) contains the constraints of the BN import schema. Note that the constraints $\exists bi:city \sqsubseteq \perp$ and $\exists bi:city^- \sqsubseteq \perp$ in Fig. 5(j) follow from the definition $bi:city \sqsubseteq \perp$ in Fig. 5(h). Indeed, these constraints capture that $bi:city$ is an empty property by saying that its domain and range are always empty. This strategy is necessary since the constraints we consider do not allow expressions of the form $bi:city \sqsubseteq \perp$. Furthermore, note that it is redundant (but not wrong) to add a constraint saying that the domain of $bi:city$ is always empty, as well as a constraint saying that the range of $bi:city$ is always empty.

Since the BN external schema has no explicit cardinality constraints, the BN import schema has no non-trivial cardinality constraints. However, $\exists bi:city \sqsubseteq \perp$ logically implies that $\top \sqsubseteq (\leq k bi:city)$, where k is any positive integer. Hence, $\exists bi:city \sqsubseteq \perp$ trivially implies \maxCardinality constraints of the form $e \sqsubseteq (\leq k bi:city)$, where e is any concept expression and k is any positive integer. Likewise, $\exists bi:city \sqsubseteq \perp$ trivially implies disjointness constraints of the form $\exists sr:city \mid C$, where C is any expression. Any of these constraints need not be made explicit since they will be in the theory of the constraints of the BN import schema. Similar observations apply to $\exists bi:city^- \sqsubseteq \perp$.

We translate the set of constraints of the BN import schema to the vocabulary of the Sales/BN mediated schema simply by replacing $bi:Book$ by $sr:Book$, etc. This results in the set of constraints Φ_B , where

- (16) $\exists sr:title \sqsubseteq sr:Product$ (in Φ_B)
- (17) $\exists sr:title - \sqsubseteq string$ (in Φ_B)
- (18) $\exists sr:city \sqsubseteq \perp$ (in Φ_B)
- (19) $\exists sr:city^- \sqsubseteq \perp$ (in Φ_B)
- (20) $sr:Book \sqsubseteq sr:Product$ (in Φ_B)
- (21) $sr:Music \sqsubseteq sr:Product$ (in Φ_B)

Now, recalling that $sr:city$ is an empty property in the BN import schema, $Th(\Phi_B)$ also contains

- (22) $\exists sr:city \sqsubseteq sr:Book$ (in $Th(\Phi_B)$)
- (23) $\exists sr:city^- \sqsubseteq string$ (in $Th(\Phi_B)$)
- (24) $\exists sr:city \sqsubseteq (\leq 1 sr:title)$ (in $Th(\Phi_B)$)

We also translate the set of constraints of the old Sales mediate schema, shown in Fig. 5(b), to the vocabulary of the Sales/BN mediated schema, obtaining the set of constraints Φ_S , where

- (25) $\exists sr:title \sqsubseteq sr:Product$ (in Φ_S)
- (26) $\exists sr:title - \sqsubseteq string$ (in Φ_S)
- (27) $\exists sr:city \sqsubseteq sr:Product$ (in Φ_S)
- (28) $\exists sr:city^- \sqsubseteq string$ (in Φ_S)
- (29) $sr:Product \sqsubseteq (\leq 1 sr:title)$ (in Φ_S)
- (30) $sr:Book \sqsubseteq (\geq 1 sr:city)$ (in Φ_S)
- (31) $sr:Book \sqsubseteq sr:Product$ (in Φ_S)
- (32) $sr:Music \sqsubseteq sr:Product$ (in Φ_S)
- (33) $sr:Book \mid sr:Music$ (in Φ_S)

Observe that, by (27) and (29), $Th(\Phi_S)$ contains the following constraint:

- (34) $\exists sr:city \sqsubseteq (\leq 1 sr:title)$

The constraints of the (revised) Sales/BN mediated schema are computed as $SC_r = \Phi_B \Delta \Phi_S = Th(\Phi_B) \cap Th(\Phi_S)$. Fig. 5(k) lists the constraints in SC_r . By inspection, observe that $SC_r = Th(\Phi_B) \cap Th(\Phi_S)$ contains:

- the domain and range constraints for $sr:title$, by (16), (17), (25) and (26)
- the domain and range constraints for $sr:city$, by (22), (23), (27) and (28)

- the subset constraints for $sr:Product$, by (20), (21), (31) and (32)
- a single cardinality constraint, of a rather unanticipated nature, by (24) and (34)
- no disjointness constraints since $Th(\Phi_S)$ does not contain any of the trivial disjointness constraints in $Th(\Phi_B)$ of the form $\exists sr:city \mid C$ or of the form $\exists sr:city^- \mid C$, where C is any expression. \square

4.2. Testing logical implication and computing the greatest lower bound of two sets of constraints

In this section, we introduce a procedure to test if a constraint is a logical consequence of a set of constraints, and a procedure to compute a representation of the g.l.b. of two theories induced by sets of constraints. In Section 4.3, we prove the correctness of the procedures introduced.

We stress that computing a representation of the g.l.b. of two theories, $Th(\Sigma_1)$ and $Th(\Sigma_2)$, induced by two sets of constraints, Σ_1 and Σ_2 , is not a trivial problem. Recall that the g.l.b. of $Th(\Sigma_1)$ and $Th(\Sigma_2)$ is defined as $Th(\Sigma_1) \cap Th(\Sigma_2)$. The major points to be addressed are: (1) we have to deal with all the constraint types that extralite schemas allow, including disjointness and cardinality constraints; (2) we have to construct a representation for $Th(\Sigma_1) \cap Th(\Sigma_2)$. The first point takes us beyond the early DL classifiers, whereas the second point leads us to be conservative and adopt a procedure that retains enough of the constraint structure of Σ_1 and Σ_2 to allow for the construction of a set of constraints Γ such that $Th(\Gamma) = Th(\Sigma_1) \cap Th(\Sigma_2)$. Recall from Section 3.2 that the constraints of a schema are of one of the following forms:

- $\exists P \sqsubseteq D$ (property P has domain D)
- $\exists P^- \sqsubseteq R$ (property P has range R)
- $C \sqsubseteq (\geq k P)$ or $C \sqsubseteq (\geq k P^-)$ (P or P^- maps each individual in C to at least k distinct individuals)
- $C \sqsubseteq (\leq k P)$ or $C \sqsubseteq (\leq k P^-)$ (P or P^- maps each individual in C to at most k distinct individuals)
- $C \sqsubseteq D$ (class C is a subclass of class D)
- $C \mid D$ (classes C and D are disjoint)

We also admit constraints of one of the forms:

- $C \sqsubseteq \perp$ (class C is always empty)
- $\exists P \sqsubseteq \perp$ or $\exists P^- \sqsubseteq \perp$ (property P is always empty, i.e., P has an empty domain or an empty range)

We normalize a set of constraints by rewriting:

- $\exists P \sqsubseteq D$ as $(\geq 1 P) \sqsubseteq D$
- $\exists P^- \sqsubseteq R$ as $(\geq 1 P^-) \sqsubseteq R$
- $C \sqsubseteq (\leq k P)$ as $C \sqsubseteq \neg(\geq k+1 P)$
- $C \sqsubseteq (\leq k P^-)$ as $C \sqsubseteq \neg(\geq k+1 P^-)$
- $C \mid D$ as $C \sqsubseteq \neg D$ (or, equivalently, $D \sqsubseteq \neg C$)
- $\exists P \sqsubseteq \perp$ as $(\geq 1 P) \sqsubseteq \perp$
- $\exists P^- \sqsubseteq \perp$ as $(\geq 1 P^-) \sqsubseteq \perp$

We observe that, after normalization, negated expressions (including the bottom concept \perp) occur only on the right-hand side of the constraints.

The question of computing the greatest lower bound of two sets of constraints is not straightforward since constraints may interact in unanticipated ways, as the following simple example illustrates.

Example 5. Suppose that $\Sigma = \{A \sqsubseteq B, A \sqsubseteq C, B \mid C\}$. Since B and C are disjoint and A is a subset of both B and C , the set of constraints Σ implies that A will always be empty, that is, $\Sigma \models A \sqsubseteq \perp$.

As a second example, assume that $\Sigma = \{A \sqsubseteq (\leq m P), A \sqsubseteq (\geq n P)\}$. Suppose that $m < n$. Then, since $(\leq m P)$ and $(\geq n P)$ denote disjoint sets, and A is a subset of both constraints, we again have that $\Sigma \models A \sqsubseteq \perp$.

Finally, note that $A \sqsubseteq \perp$ logically implies $A \sqsubseteq e$, for any expression e , which affects how we compute $Th(\Sigma)$ and, consequently, how we compute $\Sigma \Delta \Gamma$, where Γ is a second set of constraints. \square

The following sequence of definitions indicates how to construct a graph that captures the structure of a set of constraints. We say that the *complement* of a non-negated expression e is $\neg e$, and vice-versa; furthermore, the *complement* of \perp is \top , and vice-versa. If c is an expression, we denote its complement by \bar{c} . A *constraint expression* is an expression that may occur on the right- or left-hand sides of a normalized constraint.

Let Σ be a set of normalized constraints and Ω be a set of constraint expressions.

Definition 1. The labeled graph $g(\Sigma, \Omega) = (\gamma, \delta, \kappa)$ that captures Σ and Ω , where κ labels each node with an expression, is defined as follows:

- (i) For each concept expression e that occurs on the right- or left-hand side of an inclusion in Σ , or that occurs in Ω , there is exactly one node in γ labeled with e . If necessary, the set of nodes is augmented with new nodes until the following conditions are met:
 - (a) For each atomic concept C , there must be exactly one node in γ labeled with C .
 - (b) For each atomic role P , there must be exactly one node in γ labeled with $(\geq 1 P)$ and one node labeled with $(\geq 1 P^-)$.
- (ii) For each atomic role P , there must be exactly one node in γ labeled with P (this is just a theoretical convenience, explored in Definitions 6, 7 and 8).
- (iii) If there is a node in γ labeled with a concept expression e , then there must be exactly one node in γ labeled with \bar{e} .
- (iv) For each inclusion $e \sqsubseteq f$ in Σ , there is an arc (M, N) in δ , where M and N are the nodes labeled with e and f , respectively.
- (v) If there are nodes M and N in γ labeled with $(\geq m p)$ and $(\geq n p)$, where p is either P or P^- and $m < n$, then there is an arc (N, M) in δ .
- (vi) If there is an arc (M, N) in δ , where M and N are the nodes labeled with e and f respectively, then there is an arc (K, L) in δ , where K and L are the nodes labeled with \bar{f} and \bar{e} , respectively.
- (vii) These are the only nodes and arcs of $g(\Sigma)$. □

Definition 2. The labeled graph $G(\Sigma, \Omega) = (\eta, \varepsilon, \lambda)$ that represents Σ and Ω , where λ labels each node with a set of expressions, is defined from $g(\Sigma, \Omega)$ by collapsing each clique of $g(\Sigma, \Omega)$ into a single node labeled with the expressions that previously labeled the nodes in the clique. When Ω is the empty set, we simply write $G(\Sigma)$ and say that the graph represents Σ . □

If a node K of $G(\Sigma, \Omega)$ is labeled with an expression e , then \bar{K} denotes the node labeled with \bar{e} (which may be K itself). We say that K and \bar{K} are *dual nodes* of $G(\Sigma, \Omega)$. We use $K \rightarrow M$ to indicate that there is a path in $G(\Sigma, \Omega)$ (or in $g(\Sigma, \Omega)$) from K to M , and $K \nrightarrow M$ to indicate that no such path exists. Also, to simplify the notation, we use $e \rightarrow f$ to denote that there is a path in $G(\Sigma, \Omega)$ (or in $g(\Sigma, \Omega)$) from the node labeled with e to the node labeled with f , and $e \nrightarrow f$ to indicate that no such path exists.

Definition 3. Let $G(\Sigma, \Omega) = (\eta, \varepsilon, \lambda)$ be the labeled graph that represents Σ and Ω . We say that a node K of $G(\Sigma, \Omega)$ is a \perp -node with level n , for a non-negative integer n , iff one of the following conditions holds:

- (i) K is a \perp -node with level 0 iff
 - a. K is labeled with \perp , or
 - b. There are nodes M and N , not necessarily distinct from K , and a non-negated concept expression h such that M and N are respectively labeled with h and $\neg h$, and $K \rightarrow M$ and $K \rightarrow N$.
- (ii) K is a \perp -node with level $n + 1$ iff
 - a. There is a \perp -node M of level n , distinct from K , such that $K \rightarrow M$, and M is the \perp -node with the smallest level such that $K \rightarrow M$, or
 - b. K is labeled with a minCardinality constraint of the form $(\geq 1 P)$ (or of the form $(\geq 1 P^-)$) and there is a \perp -node M of level n , distinct from K , such that M is labeled with $(\geq 1 P^-)$ (or with $(\geq 1 P)$), and M is the \perp -node with the smallest level labeled with $(\geq 1 P^-)$ or $(\geq 1 P)$. □

In view of Case (ii-b), the notion of level is necessary to avoid a circular definition. In Case (i-b), note that, if $K = M = N$, then K is labeled with both h and $\neg h$; other special cases occur when $K = M$, and when $K = N$. Also note that K may be labeled with both $(\geq 1 P)$ and $(\geq 1 P^-)$, and yet be a \perp -node by virtue of Cases (i) and (ii-a), but not because of Case (ii-b).

Definition 4. Let $G(\Sigma, \Omega) = (\eta, \varepsilon, \lambda)$ be the labeled graph that represents Σ and Ω . Let K be a node of $G(\Sigma, \Omega)$. We say that

- (i) K is a \perp -node iff K is a \perp -node with level n , for some non-negative integer n .
- (ii) K is a role \perp -node iff K is labeled with an atomic role P and the node labeled with $(\geq 1 P)$ is a \perp -node.
- (iii) K is a \top -node iff \bar{K} is a \perp -node.
- (iv) K satisfies the consistency check iff K is not a \perp -node.
- (v) K satisfies the dual of the consistency check iff K is not a \top -node.
- (vi) $G(\Sigma, \Omega)$ satisfies the consistency check iff all nodes labeled with an atomic concept or with a minCardinality of the form $(\geq 1 P)$ satisfy the consistency check. □

$G(\Sigma, \Omega)$ satisfies the following properties (see Proposition 1 in Section 4.3):

- There is a path in $G(\Sigma, \Omega)$ from a node labeled with e to a node labeled with f iff there is a path in $G(\Sigma, \Omega)$ from a node labeled with \bar{f} to a node labeled with \bar{e} .
- If two concept expressions, e and f , label the same node of $G(\Sigma, \Omega)$, then $\Sigma \models e \sqsubseteq f$, that is, Σ forces e and f to denote the same set of individuals.
- If a concept expression e labels a \perp -node of $G(\Sigma, \Omega)$, then $\Sigma \models e \sqsubseteq \perp$, that is, Σ forces e to denote an empty set of individuals.

- If a concept expression f labels a \top -node of $G(\Sigma, \Omega)$, then $\Sigma \models \top \sqsubseteq f$, that is, Σ forces f to denote the set of all individuals.
 - If there is a path in $G(\Sigma, \Omega)$ from a node labeled with e to a node labeled with f , then $\Sigma \models e \sqsubseteq f$.
- Based on the previous definitions, we introduce a procedure to test logical implication:

IMPLIES($\Sigma, e \sqsubseteq f$)

input: a set Σ of normalized constraints and a normalized constraint $e \sqsubseteq f$

output: “YES – Σ logically implies $e \sqsubseteq f$ ”

“NO – Σ does not logically imply $e \sqsubseteq f$ ”

begin Construct $G(\Sigma, \{e, f\})$;

if the node of $G(\Sigma, \{e, f\})$ labeled with e is a \perp -node, or
the node of $G(\Sigma, \{e, f\})$ labeled with f is a \top -node, or
there is a path in $G(\Sigma, \{e, f\})$ from the node labeled with e to the node labeled with f ,
then return “YES – Σ logically implies $e \sqsubseteq f$ ”;
else return “NO – Σ does not logically imply $e \sqsubseteq f$ ”;

end

Note that **IMPLIES** has polynomial time complexity on the size of $\Sigma \cup \{e \sqsubseteq f\}$. **Theorem 2** in **Section 4.3** establishes the soundness and completeness of **IMPLIES**.

Example 6. Consider the constraints of the Sales mediated schema, listed in **Fig. 5(b)**. Abbreviate the names of the classes and properties by just their first letter, ignoring the namespace prefix. Let Σ be the set obtained by normalizing such constraints:

- (1) $\exists t \sqsubseteq P$ normalized as: $(\geq 1 t) \sqsubseteq P$
- (2) $\exists t^- \sqsubseteq S$ normalized as: $(\geq 1 t^-) \sqsubseteq S$
- (3) $\exists c \sqsubseteq P$ normalized as: $(\geq 1 c) \sqsubseteq P$
- (4) $\exists c^- \sqsubseteq S$ normalized as: $(\geq 1 c^-) \sqsubseteq S$
- (5) $P \sqsubseteq (\leq 1 t)$ normalized as: $P \sqsubseteq \neg(\geq 2 t)$
- (6) $B \sqsubseteq (\geq 1 c)$
- (7) $B \sqsubseteq P$
- (8) $M \sqsubseteq P$
- (9) $B \mid M$ normalized as: $B \sqsubseteq \neg M$

Fig. 6 depicts $g(\Sigma)$, the graph capturing Σ , using the normalized form of the constraints. In this case, $g(\Sigma)$ is equal to $G(\Sigma)$, the graph representing Σ . By inspecting $G(\Sigma)$, note that:

- There is a path from the node labeled with $(\geq 1 c)$ to the node labeled with $\neg(\geq 2 t)$, which implies that

$$(10) \Sigma \models (\geq 1 c) \sqsubseteq \neg(\geq 2 t)$$

- There are paths from the node K labeled with $(\geq 2 t)$ to the node labeled with $\neg P$ and the node labeled with P . Hence, K is a \perp -node with level 0, which implies that

$$(11) \Sigma \models (\geq 2 t) \sqsubseteq \perp$$

Intuitively, t never maps an individual to two or more individuals, in the presence of the constraints in Σ . □

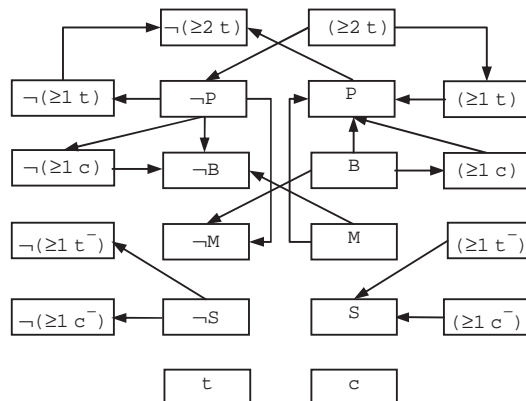


Fig. 6. The graph $G(\Sigma)$ that represents Σ .

Example 7. Consider the constraints of the BN import schema, listed in Fig. 5(j), and again abbreviate the names of the classes and properties by just their first letter, ignoring the namespace prefix for the moment. Let Φ be the set obtained by normalizing such constraints:

- (12) $\exists t \sqsubseteq P$ normalized as: $(\geq 1 t) \sqsubseteq P$
- (13) $\exists t^- \sqsubseteq S$ normalized as: $(\geq 1 t^-) \sqsubseteq S$
- (14) $\exists c \sqsubseteq \perp$ normalized as: $(\geq 1 c) \sqsubseteq \perp$
- (15) $\exists c^- \sqsubseteq \perp$ normalized as: $(\geq 1 c^-) \sqsubseteq \perp$
- (16) $B \sqsubseteq P$
- (17) $M \sqsubseteq P$

Fig. 7 depicts the graph $G(\Phi)$ representing Φ (using the normalized form of constraints). Note that the structure of $G(\Phi)$ is quite simple in this case. □

Let G^* denote the transitive closure of a graph G . Based on **IMPLIES**, we define a second procedure that creates a representation of the g.l.b. of two theories induced by sets of normalized constraints as follows:

GLB($\Sigma_1, \Sigma_2; \Gamma$)

input: two sets Σ_1 and Σ_2 of normalized constraints

output: a set Γ of normalized constraints such that $Th(\Gamma) = \Sigma_1 \Delta \Sigma_2 = Th(\Sigma_1) \cap Th(\Sigma_2)$

begin A constraint $e \sqsubseteq f$ is in Γ iff there are $i, j \in \{1, 2\}$, with $i \neq j$, such that one of the following conditions holds:

- (a) There is a \perp -node M of $G(\Sigma_i)$ and a \perp -node P of $G(\Sigma_j)$ and
 - e is a non-negated constraint expression that labels both M and P
 - f is the bottom concept \perp
- (b) There is a \perp -node M of $G(\Sigma_i)$ and an arc (P, Q) of $G^*(\Sigma_j)$ such that P is not a \perp -node of $G(\Sigma_j)$ and
 - e is a non-negated constraint expression that labels both M and P
 - f is a constraint expression that labels Q
- (c) There is a \top -node N of $G(\Sigma_i)$ and an arc (P, Q) of $G^*(\Sigma_j)$ such that Q is not a \top -node of $G(\Sigma_j)$ and
 - e is a non-negated constraint expression that labels P
 - f is a constraint expression that labels both N and Q
- (d) There is an arc (M, N) of $G^*(\Sigma_i)$ and an arc (P, Q) of $G^*(\Sigma_j)$ such that none of the nodes M, N, P and Q is a \perp -node or a \top -node, and
 - e is a non-negated constraint expression that labels both M and P
 - f is a constraint expression that labels both N and Q

end

Note that Γ is a normalized set of constraints since, by construction, e is always a non-negated constraint expression and f is a constraint expression. Furthermore, note that Γ can be constructed in $O(n^2)$, where $n = \max(n_1, n_2)$ and n_i is the number of nodes of $G(\Sigma_i)$. However, we do not claim that Γ is the best set of constraints that generates $\Sigma_1 \Delta \Sigma_2$, in the sense of having the smallest number of constraints. But we shall show in Section 4.3 that Γ is correctly constructed, in the sense that

$$Th(\Gamma) = \Sigma_1 \Delta \Sigma_2 = Th(\Sigma_1) \cap Th(\Sigma_2)$$

We close the section with a final example that illustrates how to systematically obtain the set of constraints of the (revised) Sales/BN mediated schema informally derived in Step (C) of Example 4.

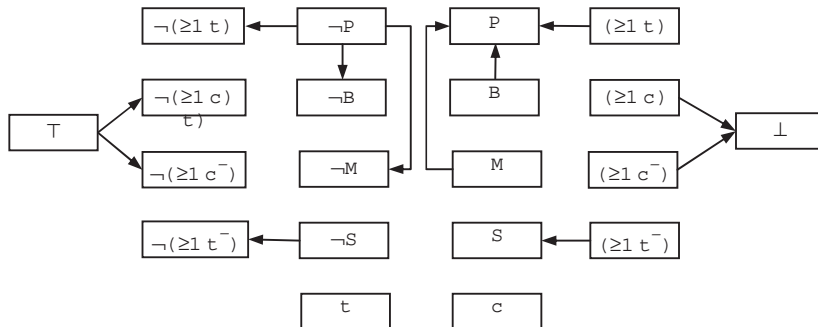


Fig. 7. The graph $G(\Phi)$ that represents Φ .

Example 8. Let Σ be the set of normalized constraints of Example 6, and Φ be the set of normalized constraints of Example 7. Fig. 5 (k) shows a set of (unnormalized) constraints whose theory is the g.l.b. of the theories induced by Σ and Φ . Let Γ be the set of constraints obtained by normalizing those in Fig. 5(k). Again abbreviating the names of the classes and properties by their first letter, and ignoring the namespace prefix, the constraints and their normalized forms are:

- (1) $\exists t \sqsubseteq P$ normalized as: $(\geq 1 t) \sqsubseteq P$
- (2) $\exists t^- \sqsubseteq S$ normalized as: $(\geq 1 t^-) \sqsubseteq S$
- (3) $\exists c \sqsubseteq P$ normalized as: $(\geq 1 c) \sqsubseteq P$
- (4) $\exists c^- \sqsubseteq S$ normalized as: $(\geq 1 c^-) \sqsubseteq S$
- (5) $\exists c \sqsubseteq (\leq 1 t)$ normalized as: $(\geq 1 c) \sqsubseteq \neg(\geq 2 t)$
- (6) $B \sqsubseteq P$
- (7) $M \sqsubseteq P$

Consider the graph $G(\Sigma)$ of Fig. 6 and the graph $G(\Phi)$ of Fig. 7. We systematically construct Γ as follows.

Table 1(a) and (b) show the arcs of $G^*(\Sigma)$ and $G^*(\Phi)$. Note that a tabular presentation of the arcs, as opposed to a graphical representation, is much more convenient since we are working with transitive closures. For example, line 3 of Table 1(a) indicates that $G^*(\Sigma)$ has arcs from the node labeled with B to the three nodes respectively labeled with $(\geq 1 c)$, P and $\neg M$.

In this specific example, Table 1(c) induces Γ as follows:

- (8) Lines 1, 7, 9 and 12 are discarded since they correspond to arcs in just one of the graphs, $G^*(\Sigma)$ or $G^*(\Phi)$.
- (9) Lines 2 and 5 are discarded since they have a negated expression on the left-hand side cell.
- (10) Lines 8 and 11 correspond to Case (b) of the **GLB** procedure.
- (11) Lines 3, 4, 6 and 10 correspond to Case (d) of the **GLB** procedure.

The case corresponding to lines 2 and 5 deserves an additional comment. Consider line 5, for example. Note that the pair $(\neg S, \neg(\geq 1 t^-))$ occurs in line 5 of Table 1(a) and (b). However, we need not add $\neg S \sqsubseteq \neg(\geq 1 t^-)$ to Γ since line 10 forces the addition of the equivalent constraint $(\geq 1 t^-) \sqsubseteq S$.

Finally, we warn the reader that the example does not illustrate all cases of the **GLB** procedure. □

4.3. Correctness of the procedures

In this section, we prove the correctness of the procedures to test logical implication and to construct a representation of the greatest lower bound of two theories induced by sets of constraints. To avoid repetitions, in what follows, let Σ be a set of normalized constraints and Ω be a set of constraint expressions. Let $G(\Sigma, \Omega)$ be the graph that represents Σ and Ω .

Table 1
Construction of the set of constraints Γ that generates $\Sigma \Delta \Phi$.

	(a) $G^*(\Sigma)$		(b) $G^*(\Phi)$		(c) Γ	
1	P	$\neg(\geq 2 t)$				
2	$\neg P$	$\neg(\geq 1 t)$	$\neg P$	$\neg(\geq 1 t)$		
		$\neg(\geq 2 t)$		$\neg B$		
		$\neg(\geq 1 c)$		$\neg M$		
		$\neg B$				
		$\neg M$				
3	B	$(\geq 1 c)$	B	P	B	P
		P				
		$\neg M$				
4	M	$\neg B$	M	P	M	P
		P				
		$\neg(\geq 2 t)$				
5	$\neg S$	$\neg(\geq 1 t^-)$	$\neg S$	$\neg(\geq 1 t^-)$		
		$\neg(\geq 1 c^-)$				
6	$(\geq 1 t)$	P	$(\geq 1 t)$	P	$(\geq 1 t)$	P
		$\neg(\geq 2 t)$				
7	$\neg(\geq 1 t)$	$\neg(\geq 2 t)$				
8	$(\geq 1 c)$	P	$(\geq 1 c)$	\perp	$(\geq 1 c)$	P
		$\neg(\geq 2 t)$				$\neg(\geq 2 t)$
9	$\neg(\geq 1 c)$	$\neg B$				
10	$(\geq 1 t^-)$	S	$(\geq 1 t^-)$	S	$(\geq 1 t^-)$	S
11	$(\geq 1 c^-)$	S	$(\geq 1 c^-)$	\perp	$(\geq 1 c^-)$	S
12			\top	$(\geq 1 c)$		
				$\neg(\geq 1 c)$		

Proposition 1.

- (i) $G(\Sigma, \Omega)$ is acyclic.
- (ii) For any pair of nodes M and N , we have that $M \rightarrow N$ iff $\overline{N} \rightarrow \overline{M}$.
- (iii) For any node K of $G(\Sigma, \Omega)$, for any expression e , we have that e labels K iff \bar{e} labels \bar{K} .
- (iv) For any node K of $G(\Sigma, \Omega)$,
 - (a) K is labeled only with \perp , or
 - (b) K is labeled only with \top , or
 - (c) K is labeled only with a single atomic role, or
 - (d) K is labeled only with non-negated concept expressions, which must be atomic concepts or minCardinality constraints of the form $(\geq m p)$, where p is either P or P^- and $m \geq 1$, or
 - (e) K is labeled only with negated concept expressions, which must be negated atomic concepts or minCardinality constraints of the form $\neg(\geq m p)$, where p is either P or P^- and $m \geq 1$.
- (v) For any pair of nodes M and N of $G(\Sigma, \Omega)$, for any pair of expressions e and f that label M and N , respectively, if $M \rightarrow N$ then $\Sigma \models e \sqsubseteq f$.
- (vi) For any node K of $G(\Sigma, \Omega)$, for any pair of expressions e and f that label K , $\Sigma \models e \sqsubseteq f$.
- (vii) For any node K of $G(\Sigma, \Omega)$, for any expression e that labels K , if K is a \perp -node, then $\Sigma \models e \sqsubseteq \perp$.
- (viii) For any node K of $G(\Sigma, \Omega)$, for any expression e that labels K , if K is a \top -node, then $\Sigma \models \top \sqsubseteq e$.
- (ix) For any node K of $G(\Sigma, \Omega)$ labeled with an atomic role P , if K is a role \perp -node, then any model \mathbf{s} of Σ is such that $\mathbf{s}(P) = \emptyset$.
- (x) Let K be a node of $G(\Sigma, \Omega)$. Assume that K is a \perp -node and K is not labeled with \perp . Then, K is labeled only with atomic concepts or minCardinality constraints of the form $(\geq m p)$, where p is either P or P^- and $m \geq 1$.
- (xi) Let L be a node of $G(\Sigma, \Omega)$. Assume that L is a \top -node and L is not labeled with \top . Then, L is labeled only with negated atomic concepts or negated minCardinality constraints of the form $\neg(\geq m p)$, where p is either P or P^- and $m \geq 1$. □

The proof of Proposition 1 follows from the definition of $G(\Sigma, \Omega)$ and the details can be found in [38].

Definition 5. A set Φ of distinct function symbols is called a set of *Skolem function symbols* for $G(\Sigma, \Omega)$ iff:

- (i) For any node N of $G(\Sigma, \Omega)$ labeled with $(\geq n P)$, Φ has n distinct unary function symbols, denoted $f_1[N, P], \dots, f_n[N, P]$.
- (ii) For any node N of $G(\Sigma, \Omega)$ labeled with $(\geq n P^-)$, Φ has n distinct unary function symbols, denoted $g_1[N, P], \dots, g_n[N, P]$.
- (iii) For any node N of $G(\Sigma, \Omega)$ labeled with an atomic concept or with $(\geq 1 P)$, Φ has a distinct constant, denoted $c[N]$ (a constant is a 0-ary function symbol).

The *Herbrand Universe* $\Delta[\Phi]$ for Φ is the set of first-order terms constructed using the function symbols in Φ . The terms in $\Delta[\Phi]$ are called *individuals*. □

Again, to avoid repetitions, let Φ be a set of distinct Skolem function symbols for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$ be the Herbrand Universe for Φ .

Definition 6.

- (i) An *instance labeling function* for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$ is a function s' that associates a set of individuals in $\Delta[\Phi]$ to each node of $G(\Sigma, \Omega)$ not labeled with an atomic role, and a set of pairs of individuals in $\Delta[\Phi]$ to each node of $G(\Sigma, \Omega)$ labeled with an atomic role.
- (ii) Let N be a node of $G(\Sigma, \Omega)$ labeled with an atomic concept or with $(\geq 1 P)$. Assume that N is not a \perp -node. Then, the Skolem constant $c[N]$ is a *seed term* of N , and N is the *seed node* of $c[N]$.
- (iii) Let N_p be the node of $G(\Sigma, \Omega)$ labeled with the atomic role P . Assume that N_p is not a role \perp -node. For each term a , for each node M labeled with $(\geq m P)$, if $a \in s'(M)$ and there is no node K labeled with $(\geq k P)$ such that $m \leq k$ and $a \in s'(K)$, then the pair $(a, f_r[M, P](a))$ is called a *seed pair* of N_p triggered by $a \in s'(M)$. We also say that the term $f_r[M, P](a)$ is a *seed term* of the node L labeled with $(\geq 1 P^-)$, and L is called the *seed node* of $f_r[M, P](a)$, for $r \in [2, m]$, if a is of the form $g_i[J, P](b)$, for some node J and some term b , and for $r \in [1, m]$, otherwise.
- (iv) Let N_p be the node of $G(\Sigma, \Omega)$ labeled with the atomic role P . Assume that N_p is not a role \perp -node. For each term b , for each node N labeled with $(\geq n P^-)$, if $b \in s'(N)$ and there is no node K labeled with $(\geq k P^-)$ such that $n \leq k$ and $b \in s'(K)$, then the pair $(g_r[N, P](b), b)$ is called a *seed pair* of N_p triggered by $b \in s'(N)$. We also say that the term $g_r[N, P](b)$ is a *seed term* of the node L labeled with $(\geq 1 P)$, and L is called the *seed node* of $g_r[N, P](b)$, for $r \in [2, n]$, if b is of the form $f_i[J, P](a)$, for some node J and some term a , and for $r \in [1, n]$, otherwise. □

Definition 7. A *canonical instance labeling function* for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$ is an instance labeling function that satisfies the following restrictions, for each node K of $G(\Sigma, \Omega)$:

- (i) Assume that K is not labeled with an atomic role, and that K is neither a \perp -node nor a \top -node. Then, $t \in s'(K)$ iff t is a seed term of a node J and there is a path from J to K (nodes J and K may be equal, in which case the path is trivial).
- (ii) Assume that K is labeled with an atomic role P , and that K is not a role \perp -node. Then, $(t, u) \in s'(K)$ iff (t, u) is a seed pair of K .

- (iii) Assume that K is not labeled with an atomic role, and that K is a \perp -node. Then, $s'(K) = \emptyset$.
- (iv) Assume that K is not labeled with an atomic role, and that K is a \top -node. Then, $s'(K) = \Delta[\Phi]$.
- (v) Assume that K is labeled with an atomic role P , and that K is a role \perp -node. Then, $s'(K) = \emptyset$. □

Proposition 2. Let s' be canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$. Then

- (i) For any pair of nodes M and N of $G(\Sigma, \Omega)$ that are not labeled with an atomic role, if $M \rightarrow N$ then $s'(M) \subseteq s'(N)$.
- (ii) For any pair of nodes M and N of $G(\Sigma, \Omega)$ that are not labeled with an atomic role, and that are not a \perp -node, $s'(M) \cap s'(N) \neq \emptyset$ iff
 - a. either M or N is a \top -node, or
 - b. both M and N are not a \top -node, and there is a seed node K such that $K \rightarrow M$ and $K \rightarrow N$ (nodes K and M , and K and N may be equal, in which case the respective path is trivial).
- (iii) For any node N_p of $G(\Sigma, \Omega)$ labeled with an atomic role P , for any node M of $G(\Sigma, \Omega)$ labeled with $(\geq m P)$, for any term $t \in s'(M)$, either $s'(N_p)$ contains all seed pairs triggered by $t \in s'(M)$, or there are no seed pairs triggered by $t \in s'(M)$.
- (iv) For any node N_p of $G(\Sigma, \Omega)$ labeled with an atomic role P , for any node N of $G(\Sigma, \Omega)$ labeled with $(\geq n P^-)$, for any term $t \in s'(N)$, either $s'(N_p)$ contains all seed pairs triggered by $t \in s'(N)$, or there are no seed pairs triggered by $t \in s'(N)$. □

The proof of [Proposition 2](#) follows from the definition of canonical instance labeling function and the details can be found in [\[38\]](#).

Definition 8. Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$. The interpretation s for Σ induced by s' is defined as follows:

- (i) $\Delta[\Phi]$ is the domain of s .
- (ii) $s(C) = s'(M)$, for each atomic concept C , where M is the node of $G(\Sigma, \Omega)$ labeled with C (there is just one such node).
- (iii) $s(P) = s'(N)$, for each atomic role P , where N is the node of $G(\Sigma, \Omega)$ labeled with P (again, there is just one such node). □

Lemma 1. Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$. Let s be the interpretation induced by s' . Then, we have:

- (i) For each node N of $G(\Sigma, \Omega)$, for each non-negated concept expression e that labels N , $s'(N) = s(e)$.
- (ii) For each node N of $G(\Sigma, \Omega)$, for each negated concept expression $\neg e$ that labels N , $s'(N) \subseteq s(\neg e)$.

Proof. Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$. Let s be the interpretation induced by s' .

- (i) Let N be a node of $G(\Sigma, \Omega)$.
Let e be a non-negated concept expression that labels N . We have to prove that $s'(N) = s(e)$.

Case 1. N is not a \perp -node or a \top -node.

By the restrictions on constraints and constraint expressions – and this is important – there are 3 cases to consider:

Case 1.1. e is an atomic concept C .

By [Definition 8\(ii\)](#), $s'(N) = s(C)$.

Case 1.2. e is of the form $(\geq n P)$.

Let N_p be the node labeled with P . Then, N_p is not a role \perp -node. Indeed, assume otherwise. Then, the node L labeled with $(\geq 1 P)$ would be a \perp -node, by definition of role \perp -node. But, by construction of $G(\Sigma, \Omega)$, there would be an arc from N (the node labeled with $(\geq n P)$) to L . Hence, N would be a \perp -node, contradicting the assumption of Case 1.

Then, since N_p is not a role \perp -node, [Definition 7\(ii\)](#) applies to $s'(N_p)$.

Recall that N is the node labeled with $(\geq n P)$ and that N is not a \perp -node or a \top -node. We first prove that

- (1) $a \in s'(N)$ implies that $a \in s((\geq n P))$.

Let $a \in s'(N)$. Let K be the node labeled with $(\geq k P)$ such that $a \in s'(K)$ and k is the largest possible. Since $a \in s'(K)$ and k is the largest possible, there are k pairs in $s'(N_p)$ whose first element is a , by [Proposition 2\(iii\)](#). By [Definition 8\(iii\)](#), $s(P) = s'(N_p)$. Hence, by definition of minCardinality , $a \in s((\geq k P))$. But again by definition of minCardinality , $s((\geq k P)) \subseteq s((\geq n P))$, since $n \leq k$, by the choice of k . Therefore, $a \in s((\geq n P))$.

We now prove that

- (2) $a \in s((\geq n P))$ implies that $a \in s'(N)$.

Let $a \in s(\geq n P)$. By definition of minCardinality, there must be n distinct pairs $(a, b_1), \dots, (a, b_n)$ in $s(P)$ and, consequently, in $s'(N_p)$, since $s(P) = s'(N_p)$, by Definition 8(iii).

Recall that N_p is not a role \perp -node. Then, by Definition 7(ii) and Definition 6(iii), possibly by reordering b_1, \dots, b_n , we then have that there are nodes L_0, L_1, \dots, L_v such that

- (3) (a, b_1) is a seed pair of N_p of the form $(g_{i_0}[L_0, P](u), u)$, triggered by $u \in s'(L_0)$, where L_0 is labeled with $(\geq l_0 P^-)$, for some $i_0 \in [1, l_0]$ or
- (4) (a, b_1) is a seed pair of N_p of the form $(a, f_{l_1}[L_1, P](a))$, triggered by $a \in s'(L_1)$, where L_1 is labeled with $(\geq l_1 P)$ and
- (5) (a, b_j) is a seed pair of N_p of the form $(a, f_{w_j}[L_i, P](a))$, triggered by $a \in s'(L_i)$, where L_i is labeled with $(\geq l_i P)$, $j \in [(\sum_{r=1}^{i-1} l_r) + 1, \sum_{r=1}^i l_r]$, with $w_j \in [1, l_i]$ and $i \in [2, v]$

Furthermore, $l_i \neq l_j$, for $i, j \in [2, v]$, with $i \neq j$, since only one node is labeled with $(\geq l_i P)$. We may therefore assume without loss of generality that $l_1 > l_2 > \dots > l_v$. But note that we then have that $a \in s'(L_i)$ and $a \in s'(L_j)$ and $l_i > l_j$, for each $ij \in [1, v]$, with $i < j$. But this contradicts the fact that $(a, f_{w_j}[L_j, P](a))$ is a seed pair of N_p triggered by $a \in s'(L_j)$ since, by Definition 6(iii), there could be no node L_i labeled with $(\geq l_i P)$ with $l_i > l_j$ and $a \in s'(L_i)$. This means that in fact there is just one node, L_1 , that satisfies (5).

We are now ready to show that $a \in s'(N)$.

Case 1.2.1. $n = 1$.

Case 1.2.1.1. a is of the form $g_{i_0}[L_0, P](u)$.

Recall that N_p is not a role \perp -node. Then, by Definition 6(iv), $g_{i_0}[L_0, P](u)$ is a seed term of the node labeled with $(\geq 1 P)$, which must be N , since $n = 1$ and there is just one node labeled with $(\geq 1 P)$. Therefore, since N is not a \perp -node or a \top -node, by Definition 7(i), $a \in s'(N)$.

Case 1.2.1.2. a is not of the form $g_{i_0}[L_0, P](u)$.

Then, by (4) and assumptions of the case, $a \in s'(L_1)$. Since, L_1 is labeled with $(\geq l_1 P)$ and N with $(\geq 1 P)$, either $n = l_1 = 1$ and $N = L_1$, or $l_1 > n = 1$ and (L_1, N) is an arc of $G(\Sigma, \Omega)$, by definition of $G(\Sigma, \Omega)$. Then, $s'(L_1) \subseteq s'(N)$, using Proposition 2(i), for the second alternative. Therefore, $a \in s'(N)$ as desired, since $a \in s'(L_1)$.

Case 1.2.2. $n > 1$.

We first show that $n \leq l_1$. First observe that, by (5) and $n > 1$, $s'(N_p)$ contains a seed pair $(a, f_{w_j}[L_1, P](a))$ triggered by $a \in s'(L_1)$. Then, by Proposition 2(iii), $s'(N_p)$ contains all seed pairs triggered by $a \in s'(L_1)$. In other words, we have that $a \in s(\geq n P)$ and $(a, b_1), \dots, (a, b_n) \in s'(N_p)$ and $(a, b_1), \dots, (a, b_n)$ are triggered by $a \in s'(L_1)$. Therefore, either $(a, b_1), \dots, (a, b_n)$ are all pairs triggered by $a \in s'(L_1)$, in which case $n = l_1$, or $(a, b_1), \dots, (a, b_n), (a, b_{n+1}), \dots, (a, b_{l_1})$, in which case $n < l_1$. Hence, we have that $n \leq l_1$.

Since L_1 is labeled with $(\geq l_1 P)$ and N with $(\geq n P)$, with $n \leq l_1$, either $n = l_1$ and $N = L_1$, or $l_1 > n$ and (L_1, N) is an arc of $G(\Sigma, \Omega)$, by definition of $G(\Sigma, \Omega)$. Then, $s'(L_1) \subseteq s'(N)$, using Proposition 2(i), for the second alternative. Therefore, $a \in s'(N)$ as desired, since $a \in s'(L_1)$.

Therefore, we established that (2) holds. Hence, from (1) and (2), $s'(N) = s(\geq n P)$, as desired.

Case 1.3. e is of the form $(\geq n P^-)$.

The proof of this case is entirely similar to that of Case 1.2.

Case 2. N is a \perp -node.

By Definition 7(iii), we then have $s'(N) = \emptyset$. Let e be a non-negated concept expression that labels N . We show that $s'(N) = s(e) = \emptyset$.

We begin by observing that, by Proposition 1(x), either N is labeled with \perp , or N is labeled only with non-negated atomic concepts or minCardinality constraints of the form $(\geq n p)$, where p is either P or P^- and $1 \leq n$.

Then, there are two cases to consider.

Case 2.1. e is a non-negated atomic concept C .

Then, we trivially have, by Definition 8 (ii), that $s(C) = \emptyset$.

Case 2.2. e is a minCardinality constraint of the form $(\geq n p)$, where p is either P or P^- and $1 \leq n$.

We prove that $s(\geq n p) = \emptyset$, using an argument similar to that in Case 1.2.

Let N_p be the node labeled with P .

Case 2.2.1. N_p is a role \perp -node.

Then, by [Definition 7\(v\)](#) and [Definition 8\(iii\)](#), $s(P) = s'(N_p) = \emptyset$. Hence, $s(\geq n p) = \emptyset$.

Case 2.2.2. N_p is not a role \perp -node.

Then, [Definition 7\(ii\)](#) applies to $s'(N_p)$.

Assume that $s(\geq n p) \neq \emptyset$ and let $a \in s(\geq n p)$. By definition of minCardinality and since $s(P) = s'(N_p)$, there must be n distinct pairs $(a, b_1), \dots, (a, b_n)$ in $s'(N_p)$. Using an argument similar to that in [Case 1.2](#), there are nodes L_0 and L_1 such that

- (6) (a, b_1) is a seed pair of N_p of the form $(g_{i_0}[L_0, P](u), u)$, triggered by $u \in s'(L_0)$, where L_0 is labeled with $(\geq l_0 P^-)$, for some $i_0 \in [1, l_0]$ or
- (7) (a, b_1) is a seed pair of N_p of the form $(a, f_1[L_1, P](a))$, triggered by $a \in s'(L_1)$, where L_1 is labeled with $(\geq l_1 P)$ and
- (8) (a, b_j) is a seed pair of N_p of the form $(a, f_{w_j}[L_1, P](a))$, triggered by $a \in s'(L_1)$, where L_1 is labeled with $(\geq l_1 P)$, with $j \in [2, l_1]$

We are now ready to show that no such $a \in s(\geq n p)$ exists. Recall that $n > 1$. We first show that $n \leq l_1$. First observe that, by (8) and $n > 1$, $s'(N_p)$ contains a seed pair $(a, f_{w_j}[L_1, P](a))$ triggered by $a \in s'(L_1)$. Then, by [Proposition 2\(iii\)](#), $s'(N_p)$ contains all seed pairs triggered by $a \in s'(L_1)$. In other words, we have that $a \in s(\geq n P)$ and $(a, b_1), \dots, (a, b_n) \in s'(N_p)$ and $(a, b_1), \dots, (a, b_n)$ are triggered by $a \in s'(L_1)$. Therefore, either $(a, b_1), \dots, (a, b_n)$ are all pairs triggered by $a \in s'(L_1)$, in which case $n = l_1$, or $(a, b_1), \dots, (a, b_n), (a, b_{n+1}), \dots, (a, b_{l_1})$, in which case $n < l_1$. Hence, we have that $n \leq l_1$. Since L_1 is labeled with $(\geq l_1 P)$ and N with $(\geq n P)$, with $n \leq l_1$, either $n = l_1$ and $N = L_1$, or $l_1 > n$ and (L_1, N) is an arc of $G(\Sigma, \Omega)$, by definition of $G(\Sigma, \Omega)$. Then, $s'(L_1) \subseteq s'(N)$, using [Proposition 2\(i\)](#), for the second alternative. Therefore, $a \in s'(N)$, since $a \in s'(L_1)$. But this is impossible, since $s'(N) = \emptyset$.

Hence, we conclude that $s(\geq n p) = \emptyset$.

Therefore, we have that, if N is a \perp -node, then $s'(N) = s(e) = \emptyset$, for any non-negated concept expression e that labels N .

Case 3. N is a \top -node.

By [Definition 7\(iv\)](#), we then have $s'(N) = \Delta[\Phi]$. Let e be a non-negated expression that labels N . We show that $s'(N) = s(e) = \Delta[\Phi]$.

By [Proposition 1\(xi\)](#), N is either labeled only with \top , or labeled only with negated expressions. Therefore, e can only be the top concept \top . Therefore, trivially, $s(e) = \Delta[\Phi]$.

Therefore, we established in all three cases that [Lemma 1\(i\)](#) holds.

(ii) Let N be a node of $G(\Sigma, \Omega)$.

Let $\neg e$ be a negated expression that labels N . We have to prove that $s'(N) \subseteq s(\neg e)$.

Case 1. N is not a \perp -node or a \top -node.

Suppose, by contradiction, that there is a term t such that $t \in s'(N)$ and $t \notin s(\neg e)$.

Since $t \notin s(\neg e)$, we have that $t \in s(e)$, by definition. Let M be the node labeled with e . Hence, by [Lemma 1\(i\)](#), $t \in s'(M)$. That is, $t \in s'(M) \cap s'(N)$.

Note that M and N are in fact dual nodes. Therefore, since N is not a \perp -node or a \top -node, M is also not a \top -node or a \perp -node, by definition of \top -node. Hence, by [Proposition 2\(ii\)](#) and [Definition 7\(i\)](#), since both M and N are not a \perp -node or a \top -node, there is a seed node K such that $K \rightarrow M$ and $K \rightarrow N$ and $t \in s'(K)$. But this is impossible. Indeed, we would have that $K \rightarrow M$ and $K \rightarrow N$, M is labeled with e , and N is labeled with $\neg e$, which implies that K is a \perp -node. Hence, by [Definition 7\(iii\)](#), $s'(K) = \emptyset$, which implies that $t \notin s'(K)$.

Therefore, we established that, for all terms t , if $t \in s'(N)$ then $t \in s(\neg e)$. That is, $s'(N) \subseteq s(\neg e)$, as desired.

Case 2. N is a \perp -node.

By [Definition 7\(iii\)](#), we then have $s'(N) = \emptyset$, which trivially implies that $s'(N) \subseteq s(\neg e)$.

Case 3. N is a \top -node.

By [Definition 7\(iv\)](#), we then have $s'(N) = \Delta[\Phi]$. We show that $s(\neg e) = \Delta[\Phi]$. Let \bar{N} be the dual node of N . Since N is a \top -node, we have that \bar{N} is a \perp -node. Furthermore, since $\neg e$ labels N , e labels \bar{N} . Since e is a positive expression, by [Lemma 1\(i\)](#), $s'(\bar{N}) = s(e) = \emptyset$. Thus, $s(\neg e) = \Delta[\Phi]$, which trivially implies that $s'(N) \subseteq s(\neg e)$.

Therefore, we established that, in all three cases, [Lemma 1\(ii\)](#) holds. □

Lemma 2. Let s be the interpretation for Σ induced by a canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$. Then, we have

- (i) s is a model of Σ .
- (ii) Let N be a node of $G(\Sigma, \Omega)$. Let e be an atomic concept or a minCardinality of the form $(\geq 1 P)$ that labels N . Assume that N is not a \perp -node. Then $s(e) \neq \emptyset$.
- (iii) Let N be a node of $G(\Sigma, \Omega)$. Let P be an atomic role that labels N . Assume that N is not a role \perp -node. Then, $s(P) \neq \emptyset$.

Proof. Let Σ be a set of normalized constraints and Ω be a set of constraint expressions. Let $G(\Sigma, \Omega)$ be the graph that represents Σ and Ω . Let Φ be a set of distinct function symbols and $\Delta[\Phi]$ be the Herbrand Universe for Φ . Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and $\Delta[\Phi]$ and s be the interpretation for Σ induced by s' .

(i) We prove that s satisfies all constraints in Σ .

Let $e \sqsubseteq f$ be a constraint in Σ . By the restrictions on the constraints in Σ , e must be non-negated and f can be negated or not. Therefore, there are two cases to consider.

Case 1. e and f are both non-negated.

Then, by Lemma 1(i), $s'(M) = s(e)$ and $s'(N) = s(f)$, where M and N are the nodes labeled with e and f , respectively. If $M = N$, then we trivially have that $s'(M) = s'(N)$. So assume that $M \neq N$. Since $e \sqsubseteq f$ is in Σ and $M \neq N$, there must be an arc (M, N) of $G(\Sigma, \Omega)$. By Proposition 2(i), we then have $s'(M) \subseteq s'(N)$. Hence, $s(e) = s'(M) \subseteq s'(N) = s(f)$.

Case 2. e is non-negated and f is negated.

Then, by Lemma 1(i), $s'(M) = s(e)$ and, by Lemma 1(ii), $s'(N) \subseteq s(f)$, where M and N are the nodes labeled with e and f , respectively. Since negated expressions do not occur on the left-hand side of constraints in Σ , e and f cannot label nodes that belong to the same clique in the original graph. Therefore, we have that $M \neq N$. Since $e \sqsubseteq f$ is in Σ and $M \neq N$, there must be an arc (M, N) of $G(\Sigma, \Omega)$. By Proposition 2(i), we then have $s'(M) \subseteq s'(N)$. Hence, $s(e) = s'(M) \subseteq s'(N) \subseteq s(f)$.

Thus, in both cases, $s(e) \subseteq s(f)$. Therefore, for any constraint $e \sqsubseteq f$ in Σ , we have that $s \models e \sqsubseteq f$, which implies that s is a model of Σ .

- (ii) Let N be a node of $G(\Sigma, \Omega)$. Let e be an atomic concept or a minCardinality of the form $(\geq 1 P)$ that labels N . Assume that N is not a \perp -node. Then, by Lemma 1(i), $s(e) = s'(N)$.

Case 1. N is a \top -node.

Then, we trivially have that $s(e) = s'(N) = \Delta[\Phi] \neq \emptyset$.

Case 2. N is not a \top -node.

Then, N is neither a \perp -node nor a \top -node. By Definition 6(ii) and Definition 7(i), the seed term $c[N]$ of N is such that $c[N] \in s(e)$. Hence, trivially, $s(e) \neq \emptyset$.

- (iii) Let N be a node of $G(\Sigma, \Omega)$ and P be an atomic role that labels N . Assume that N is not a role \perp -node. Then, the node labeled with $(\geq 1 P)$ is not a \perp -node. Then, by (ii), $s((\geq 1 P)) \neq \emptyset$. Hence, $s(P) \neq \emptyset$. □

We are now ready to prove that the **IMPLIES** procedure is sound and complete.

Theorem 2. Let Σ be a set of normalized constraints. Let $e \sqsubseteq f$ be a constraint and $\Omega = \{e, f\}$. Let $G(\Sigma, \Omega)$ be the graph that represents Σ and Ω . Then, $\Sigma \models e \sqsubseteq f$ iff one of the following conditions holds:

- (a) The node labeled with e is a \perp -node; or
- (b) The node labeled with f is a \top -node; or
- (c) There is a path in $G(\Sigma, \Omega)$ from the node labeled with e to the node labeled with f .

Proof. Let Σ be a set of normalized constraints. Let $e \sqsubseteq f$ be a constraint and $\Omega = \{e, f\}$. Let $G(\Sigma, \Omega)$ be the graph that represents Σ and Ω . Observe that, by construction, $G(\Sigma, \Omega)$ has a node labeled with e and a node labeled with f . Let M and N be such nodes, respectively.

(\Leftarrow) We show that $\Sigma \models e \sqsubseteq f$. There are three cases to consider:

Case 1. M is a \perp -node.

Then, by Proposition 1 (vii), $\Sigma \models e \sqsubseteq \perp$, which trivially implies that $\Sigma \models e \sqsubseteq f$.

Case 2. N is a \top -node.

Then, by [Proposition 1](#) (viii), $\Sigma \models \top \sqsubseteq f$, which trivially implies that $\Sigma \models e \sqsubseteq f$.

Case 3. There is a path in $G(\Sigma, \Omega)$ from M to N .

Then, by [Proposition 1](#) (v) and (vi), we have that $\Sigma \models e \sqsubseteq f$.

(\Rightarrow) We prove that, if the conditions of the theorem do not hold, then $\Sigma \not\models e \sqsubseteq f$.

Since $e \sqsubseteq f$ is a constraint, we have:

- (1) e is either an atomic concept C or minCardinalities of the form of the form $(\geq k p)$, where p is either P or P^- , and
- (2) f is either the bottom concept \perp , an atomic concept C , a negated atomic concept $\neg D$, minCardinality constraints of the form $(\geq k p)$, or negated minCardinality constraints of the form $\neg(\geq k p)$, where p is either P or P^-

Assume that the conditions of the theorem do not hold, that is:

- (3) The node M labeled with e is not a \perp -node; and
- (4) The node N labeled with f is not a \top -node; and
- (5) There is no path in $G(\Sigma, \Omega)$ from M to N .

To prove that $\Sigma \not\models e \sqsubseteq f$, it suffices to exhibit a model r of Σ such that $r \not\models e \sqsubseteq f$. Recall that $r \not\models e \sqsubseteq f$ iff there is an individual t such that $t \in r(e)$ and $t \notin r(f)$ or, equivalently, $t \in r(\neg f)$.

Recall that, to simplify the notation, $e \rightarrow f$ denotes that there is a path in $G(\Sigma, \Omega)$ from the node labeled with e to the node labeled with f , and $e \not\rightarrow f$ to indicate that no such path exists.

Since $e \sqsubseteq f$ is a constraint, e must be non-negated and f can be negated or not. Hence, there are 2 cases to consider.

Case 1. e and f are both non-negated.

Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and s be the model induced by s' . By [Lemma 2](#), s is a model of Σ . We show that $s \not\models e \sqsubseteq f$.

Case 1.1. N is a \perp -node.

Since N is a \perp -node, by [Proposition 1](#) (vii), we have that $\Sigma \models f \sqsubseteq \perp$, which implies that $s(f) = \emptyset$, since s is a model of Σ .

By (1), e is either an atomic concept C or minCardinalities of the form $(\geq k p)$, where p is either P or P^- . By (3), M is not a \perp -node. Hence, by [Definition 7](#) (i), if M is not a \top -node, or by [Definition 7](#) (iv), if M is a \top -node, and by [Definition 8](#) (ii), $s(e) \neq \emptyset$. Hence, we trivially have that $s \not\models e \sqsubseteq f$.

Case 1.2. N is not a \perp -node.

Observe that M and N are neither a \perp -node nor a \top -node. Indeed, by assumption of the case and by (4), N is neither a \perp -node nor a \top -node. Now, by (3), M is not a \perp -node. Furthermore, since $e \sqsubseteq f$ is a constraint, either M and N are the same node or there is an arc (M, N) in $G(\Sigma, \Omega)$. Therefore, M cannot be a \top -node as otherwise N would be a \top -node, contradicting (4).

By [Lemma 1](#) (i), since e is non-negated by assumption, and [Definition 6](#) (ii) and [Definition 7](#) (i), since M is neither a \perp -node nor a \top -node, we have

- (6) $s'(M) = s(e)$ and there is a seed term $c[M] \in s'(M)$

By definition of canonical instance labeling function, we have:

- (7) For each node K of $G(\Sigma, \Omega)$ that is neither a \perp -node nor a \top -node or labeled with an atomic role, $c[M] \in s'(K)$ iff there is a path from M to K

By (5), we have $e \not\rightarrow f$. Furthermore, N is neither a \perp -node nor a \top -node. Hence, by (7), we have:

- (8) $c[M] \notin s'(N)$

Since f is positive, by [Lemma 1](#) (i), $s'(N) = s(f)$. Hence, we have

- (9) $c[M] \notin s(f)$

Therefore, by (6) and (9), $s(e) \not\subseteq s(f)$, that is, $s \not\models e \sqsubseteq f$, as desired.

Case 2. e is non-negated and f is negated.

Assume that f is a negated expression of the form $\neg g$, where g is non-negated (if f is \perp then g is \top).

Case 2.1. $e \rightarrow g$.

Let s' be a canonical instance labeling function for $G(\Sigma, \Omega)$ and s be the model induced by s' . By Lemma 2, s is a model of Σ . We show that $s \not\models e \sqsubseteq f$.

By Proposition 1(v) and (vi), and since s is a model of Σ , we have that $s \models e \sqsubseteq g$, if e and g label the same node, and $s \models e \sqsubseteq g$, otherwise. Hence, we have that $s \not\models e \sqsubseteq \neg g$. Now, since f is $\neg g$, we have $s \not\models e \sqsubseteq f$, as desired.

Case 2.2. $e \nrightarrow g$.

Construct Φ as follows:

(10) Φ is Σ with two new constraints, $H \sqsubseteq e$ and $H \sqsubseteq g$, where H is a new atomic concept

Let r' be a canonical instance labeling function for $G(\Phi, \Omega)$ and r be the model induced by r' . By Lemma 2, r is a model of Φ . We show that $r \not\models e \sqsubseteq f$.

We first observe that

(11) There is no expression h such that $e \rightarrow h$ and $g \rightarrow \neg h$ are paths in $G(\Sigma, \Omega)$

Indeed, by construction of $G(\Sigma, \Omega)$, $g \rightarrow \neg h$ iff $h \rightarrow \neg g$. But $e \rightarrow h$ and $h \rightarrow \neg g$ implies $e \rightarrow \neg g$, contradicting (5), since f is $\neg g$. Hence, (11) follows.

We now prove that

(12) There is no non-negated expression h such that $H \rightarrow h$ and $H \rightarrow \neg h$ are paths in $G(\Phi, \Omega)$

Assume otherwise. Let h be a non-negated expression such that $H \rightarrow h$ and $H \rightarrow \neg h$ are paths in $G(\Phi, \Omega)$.

Case 2.2.1. $H \rightarrow e \rightarrow h$ and $H \rightarrow g \rightarrow \neg h$ are paths in $G(\Phi, \Omega)$.

Then, $e \rightarrow h$ and $g \rightarrow \neg h$ must be paths in $G(\Sigma, \Omega)$, which contradicts (11).

Case 2.2.2. $H \rightarrow e \rightarrow \neg h$ and $H \rightarrow g \rightarrow h$ are paths in $G(\Phi, \Omega)$.

Then, $e \rightarrow \neg h$ and $g \rightarrow h$ must be paths in $G(\Sigma, \Omega)$. But, since $g \rightarrow h$ iff $\neg h \rightarrow \neg g$, we have $e \rightarrow \neg h \rightarrow \neg g$ is a path in $G(\Sigma, \Omega)$, which contradicts (5), recalling that f is $\neg g$.

Case 2.2.3. $H \rightarrow e \rightarrow h$ and $H \rightarrow e \rightarrow \neg h$ are paths in $G(\Phi, \Omega)$.

Then, $e \rightarrow h$ and $e \rightarrow \neg h$ must be paths in $G(\Sigma, \Omega)$, which contradicts (3), by definition of \perp -node.

Case 2.2.4. $H \rightarrow g \rightarrow h$ and $H \rightarrow g \rightarrow \neg h$ are paths in $G(\Phi, \Omega)$.

Then, $g \rightarrow h$ and $g \rightarrow \neg h$ must be paths in $G(\Sigma, \Omega)$. Now, observe that, since $\neg g$ is f , that is, f and g are complementary expressions, g labels \bar{N} , the dual node of N in $G(\Sigma, \Omega)$. Then, $g \rightarrow h$ and $g \rightarrow \neg h$ implies that \bar{N} is a \perp -node of $G(\Sigma, \Omega)$, that is, N is a \top -node, which contradicts (4). Hence, we established (12).

Let K be the node of $G(\Phi, \Omega)$ labeled with H . Note that, by construction of Φ , K is labeled only with H . Then, by (12), K is not a \perp -node. By Lemma 2(i), r is a model of Φ . Furthermore, by Lemma 2(ii) and since K is not a \perp -node, we have

(13) $r(H) \neq \emptyset$

Since $H \sqsubseteq e$ and $H \sqsubseteq g$ are in Φ , and since r is a model of Φ , we also have:

(14) $r(H) \subseteq r(e)$ and $r(H) \subseteq r(g)$

Therefore, by (13) and (14) and since $f = -g$

(15) $r(e) \cap r(g) \neq \emptyset$ or, equivalently, $r(e) \not\subseteq r(-g)$ or, equivalently, $r(e) \not\subseteq r(f)$ or, equivalently, $r \neq e \sqsubseteq f$

But since $\Sigma \subseteq \Phi$, r is also a model of Σ . Therefore, for [Case 2.2](#), we also exhibited a model r of Σ such that $r \neq e \sqsubseteq f$, as desired. Therefore, in all cases, we exhibited a model of Σ that does not satisfy $e \sqsubseteq f$, as desired. \square

Corollary 1. Let Σ be a set of normalized constraints. Let $e \sqsubseteq f$ be a constraint and $\Omega = \{e, f\}$. Let $G(\Sigma, \Omega)$ be the graph that represents Σ and Ω , and $G(\Sigma)$ be the graph that represents Σ . Suppose that $\Sigma \models e \sqsubseteq f$. Then:

- (a) Either e labels a node of $G(\Sigma)$ or e is of the form $(\geq k P)$ and there is a node of $G(\Sigma)$ labeled with $(\geq j P)$, where $j < k$.
- (b) Either f labels a node of $G(\Sigma)$ or f is of the form $\neg(\geq n P)$ and there is a node of $G(\Sigma)$ labeled with $\neg(\geq m P)$, where $m < n$. \square

To conclude, we state [Corollary 2](#) which establishes that the GLB procedure is correct.

Corollary 2. Let Σ_1 and Σ_2 be two sets of normalized constraints. Let Γ be the set of constraints that the GLB procedure outputs when called with Σ_1 and Σ_2 . Then, $Th(\Gamma) = \Sigma_1 \Delta \Sigma_2 = Th(\Sigma_1) \cap Th(\Sigma_2)$. \square

The proof of [Corollary 2](#) follows from [Theorem 2](#) and [Corollary 1](#) and the details can be found in [\[38\]](#).

5. Conclusions

In this article, we addressed the problem of changing the constraints of a mediated schema to accommodate the set of constraints of a new export schema. We argued that such problem can be solved by computing the greatest lower bound of two sets of constraints. The approach we took to define the mediation environment is akin to the idea of exact views [\[17,23\]](#). Yet, we considered that constraints should be included in the mediated schema to capture the common semantics of the data sources.

For the family of extralite schema, we described efficient procedures to decide logical implication and to compute the greatest lower bound of two sets of constraints. The procedures essentially explore the structure of a set of constraints, captured as a graph. However, cardinality constraints posed considerable technical problems to the proof of the theorems, which we overcame with the help of the notion of canonical Herbrand interpretation introduced in [Section 4.3](#). These developments are new, and cover an expressive and useful family of constraints, which justifies the detailed proofs included in [Section 4.3](#).

As for future work, we plan to extend the schema matching framework described in [\[39\]](#) to a full-fledged tool that helps create mediation environments, by including the strategy described in this article. Additional work should also be devoted to minimize the set of constraints that generates $\Sigma \Delta \Phi$, which will require a careful analysis of the graphs that represent Σ and Φ .

Acknowledgements

This work was partly supported by CNPq under grants 301497/2006-0, 473110/2008-3, 557128/2009-9, FAPERJ E-26/170028/2008, and CAPES/PROCAD NF 21/2009.

References

- [1] A. Artale, D. Calvanese, R. Kontchakov, M. Zakharyashev, The DL-Lite family and relations, *Journal of Artificial Intelligence Research* 36 (2009) 1–69.
- [2] B. Aspvall, M.F. Plass, R.E. Tarjan, A linear-time algorithm for testing the truth of certain quantified boolean formulas, *Information Processing Letters* 8 (3) (March 1979) 121–123.
- [3] J. Euzenat, P. Shvaiko, *Ontology matching*, Springer-Verlag, 2007.
- [4] E. Rahm, P.A. Bernstein, A survey of approaches to automatic schema matching, *The VLDB Journal* 10 (4) (2001) 334–350.
- [5] P.A. Bernstein, S. Melnik, *Model Management 2.0: Manipulating Richer Mappings*, Proc. of the 27th ACM SIGMOD Int'l. Conf. Management of Data, Beijing, China, 2007, pp. 1–12.
- [6] H. Köpcke, E. Rahm, Frameworks for entity matching: a comparison, *Data & Knowledge Engineering* 69 (2) (Feb. 2010) 197–210.
- [7] S. Melnik, H. Garcia-Molina, E. Rahm, Similarity Flooding: A Versatile Graph Matching Algorithm and its Application to Schema Matching, Proc. of the 18th Int'l. Conf. on Data Engineering, 2002, pp. 117–128.
- [8] J. Madhavan, P.A. Bernstein, E. Rahm, Generic Schema Matching with Cupid, Proc. of the 27th Int'l. Conf. on Very Large Data Bases, 2001, pp. 49–58.
- [9] A. Bilke, F. Naumann, Schema Matching Using Duplicates, Proc. of the 21st Int'l. Conf. on Data Engineering, 2005, pp. 69–80.
- [10] D.F. Brauner, M.A. Casanova, R. Milidiu, Towards Gazetteer Integration Through an Instance-based Thesauri Mapping Approach, in: C.A. Davis Jr., M.A.V. Monteiro (Eds.), *Advances in Geoinformatics*, Springer, Heidelberg, 2007, pp. 235–245.
- [11] J. Wang, J. Wen, F. Lochovsky, W. Ma, Instance-based Schema Matching for Web Databases by Domain-Specific Query Probing, Proc. of the 30th Int'l. Conf. on Very Large Data Bases, 2004, pp. 408–419.
- [12] Q. Qi He, T.W. Ling, An ontology based approach to the integration of entity–relationship schemas, *Data & Knowledge Engineering* 58 (3) (Sept. 2006) 299–326.
- [13] H. Zhao, S. Ram, Combining schema and instance information for integrating heterogeneous data sources, *Data & Knowledge Engineering* 61 (2) (May 2007) 281–303.
- [14] D. Lonsdale, D.W. Embley, Y. Ding, L. Xu, M. Hepp, Reusing ontologies and language components for ontology generation, *Data & Knowledge Engineering* 69 (4) (Apr. 2010) 318–330.
- [15] E. Simperla, Reusing ontologies on the Semantic Web: a feasibility study, *Data & Knowledge Engineering* 68 (10) (Oct. 2009) 905–925.

- [16] P. Papotti, R. Torlone, Schema exchange: generic mappings for transforming data and metadata, *Data & Knowledge Engineering* 68 (7) (Jul. 2009) 665–682.
- [17] M. Lenzerini, Data Integration: A Theoretical Perspective, *Proc. of the 21st ACM Symp. on Principles of Database Systems, 2002*, pp. 233–246.
- [18] H. Garcia-Molina, Y. Papakonstantinou, D. Quass, A. Rajaraman, Y. Sagiv, J.D. Ullman, V. Vassalos, J. Widom, The TSIMMIS Approach to Mediation: Data Models and Languages, *Journal of Intelligent Information Systems* 8 (2) (1997) 117–132.
- [19] A. Cali, D. Calvanese, G. De Giacomo, M. Lenzerini, P. Naggari, F. Vernacotola, IBIS: Semantic Data Integration at Work, *Proc. of the 15th Int'l. Conf. on Advanced Information Systems Engineering, 2003*, pp. 79–94.
- [20] N. Leone, T. Eiter, W. Faber, M. Fink, G. Gottlob, G. Greco, E. Kalka, G. Ianni, D. Lembo, M. Lenzerini, V. Lio, B. Nowicki, R. Rosati, M. Ruzzi, W. Staniszkis, G. Terracina, The INFOMIX System for Advanced Integration of Incomplete and Inconsistent Data, *Proc. of the 25th ACM SIGMOD Int'l. Conf. Management of Data, 2005*, pp. 915–917.
- [21] A. Halevy, Answering queries using views: a survey, *The VLDB Journal* 10 (4) (2001) 270–294.
- [22] F. Goasdoué, V. Lattes, M.-C. Rousset, The use of CARIN language and algorithms for information integration: the Pictel system, *International Journal of Cooperative Information Systems* 9 (4) (2000) 383–401.
- [23] A. Cali, D. Calvanese, G. De Giacomo, M. Lenzerini, Data integration under integrity constraints, *Information Systems* 2 (29) (2004) 147–163.
- [24] G. Rull, C. Farré, E. Teniente, T. Urpi, Validation of mappings between schemas, *Data & Knowledge Engineering* 66 (3) (Sept. 2008) 414–437.
- [25] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, A. Poggi, R. Rosati, M. Ruzzi, Data Integration through DL-Lite-A Ontologies, *Proc. of the Workshop on Semantics in Data and Knowledge Bases, 2008*, pp. 26–47.
- [26] P. Atzeni, P. Cappellari, R. Torlone, P.A. Bernstein, G. Gianforme, Model-independent schema translation, *The VLDB Journal* 17 (6) (2008) 1347–1370.
- [27] J.-M. Hick, J.-H. Hainaut, Database application evolution: a transformational approach, *Data & Knowledge Engineering* 59 (3) (Dec. 2006) 534–558.
- [28] S. Hartmann, S. Link, T. Trinh, Constraint acquisition for entity-relationship models, *Data & Knowledge Engineering* 68 (10) (Oct. 2009) 1128–1155.
- [29] F. Baader, W. Nutt, Basic Description Logics, in: F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, P.F. Patel-Schneider (Eds.), *The Description Logic Handbook: Theory, Implementation and Applications*, Cambridge University Press, Cambridge, UK, 2003, pp. 43–95.
- [30] F.M. Donini, Complexity of Reasoning, in: F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, P.F. Patel-Schneider (Eds.), *The Description Logic Handbook: Theory, Implementation and Applications*, Cambridge University Press, Cambridge, UK, 2003, pp. 96–133.
- [31] E. Franconi, Structural Description Logics: FL-, *Description Logics Course* Available at: <http://www.inf.unibz.it/~franconi/dl/course/slides/struct-DL/flminus.pdf>, 2003.
- [32] H.J. Levesque, R.J. Brachman, Expressiveness and tractability in knowledge representation and reasoning, *Computational Intelligence* 3 (1987) 78–93.
- [33] D. Nardi, R.J. Brachman, An Introduction to Description Logics, in: F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, P.F. Patel-Schneider (Eds.), *The Description Logic Handbook: Theory, Implementation and Applications*, Cambridge University Press, Cambridge, UK, 2003, pp. 1–40.
- [34] T. Lauschner, M.A. Casanova, V.M.P. Vidal, J.A. Macedo, Efficient Decision Procedures for Query Containment and Related Problems, *Proc. of the XXIV Brazilian Symposium on Databases, 2009*.
- [35] D.F. Brauner, A. Gazola, M.A. Casanova, K.K. Breitman, Adaptive Matching of Database Web Services Export Schemas, *Proc. of the 10th Int'l. Conf. on Enterprise Information Systems, 2008*, pp. 49–56.
- [36] L.A.P.P. Leme, D.F. Brauner, K.K. Breitman, M.A. Casanova, A. Gazola, Matching object catalogues, *Journal of Innovations in Systems and Software Engineering* 4 (4) (2008) 315–328.
- [37] L.A.P.P. Leme, M.A. Casanova, K.K. Breitman, A.L. Furtado, Instance-based OWL Schema Matching, *Proc. of the 11th Int'l. Conf. on Enterprise Information Systems, 2009*, pp. 14–26.
- [38] M.A. Casanova, T. Lauschner, L.A.P.P. Leme, K.K. Breitman, A.L. Furtado, V.M.P. Vidal, Revising the Constraints of the Mediated Schema, Technical Report MCC06/10, Department of Informatics, PUC-Rio, 2010.
- [39] R.V. Gomes, L.A.P.P. Leme, M.A. Casanova, MatchMaking – a tool to match OWL schemas, *Journal of Theoretical and Applied Informatics* (2010) 71–76.