Inverting Schema Mappings

RONALD FAGIN
IBM Almaden Research Center

A schema mapping is a specification that describes how data structured under one schema (the source schema) is to be transformed into data structured under a different schema (the target schema). Although the notion of an inverse of a schema mapping is important, the exact definition of an inverse mapping is somewhat elusive. This is because a schema mapping may associate many target instances with each source instance, and many source instances with each target instance. Based on the notion that the composition of a mapping and its inverse is the identity, we give a formal definition for what it means for a schema mapping $\mathcal{M}'$ to be an inverse of a schema mapping $\mathcal{M}$ for a class $\mathcal{S}$ of source instances. We call such an inverse an $\mathcal{S}$-*inverse*. A particular case of interest arises when $\mathcal{S}$ is the class of all source instances, in which case an $\mathcal{S}$-inverse is a global inverse. We focus on the important and practical case of schema mappings specified by source-to-target tuple-generating dependencies, and uncover a rich theory. When $\mathcal{S}$ is specified by a set of dependencies with a finite chase, we show how to construct an $\mathcal{S}$-inverse when one exists. In particular, we show how to construct a global inverse when one exists. Given $\mathcal{M}$ and $\mathcal{M}'$, we show how to define the largest class $\mathcal{S}$ such that $\mathcal{M}'$ is an $\mathcal{S}$-inverse of $\mathcal{M}$.

Categories and Subject Descriptors: H.2.5 [**Database Management**]: Heterogeneous Databases—*Data translation*; H.2.4 [**Database Management**]: Systems—*Relational data bases*

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Data exchange, inverse, schema mapping, data integration, chase, computational complexity, dependencies, metadata model management, second-order logic

ACM Reference Format:

Fagin, R. 2007. Inverting schema mappings. ACM Trans. Datab. Syst. 32, 4, Article 25 (November 2007), 53 pages. DOI = 10.1145/1292609.1292615 http://doi.acm.org/10.1145/1292609.1292615

1. INTRODUCTION

Data exchange is the problem of materializing an instance that adheres to a target schema, given an instance of a source schema and a schema mapping that specifies the relationship between the source and the target. This is a very old problem [Shu et al. 1977] that arises in many tasks where data must be

This is an expanded version of Fagin [2006].

Author's address: IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120; email: fagin@almaden.ibm.com.

transferred between independent applications that do not have the same data format.

Because of the extensive use of schema mappings, it has become important to develop a framework for managing schema mappings and other metadata, and operators for manipulating them. Bernstein [2003] has introduced such a framework, called *model management*. Melnik et al. [2005] have developed a semantics for model-management operators that allows applying the operators to executable mappings. One important schema mapping operator, at least in principle, is the inverse operator. What do we mean by an inverse of a schema mapping? This is a delicate question, since in spite of the traditional use of the name *mapping*, a schema mapping is not simply a function that maps an instance of the source schema to an instance of the target schema. Instead, for each source instance, the schema mapping may associate many target instances. Furthermore, for each target instance, there may be many corresponding source instances.

As in Fagin et al. [2005a, 2005b, 2005c], we study the relational case where a schema is a sequence of distinct relational symbols. A *schema mapping* is a triple $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where $\mathbf{S}$ (the *source schema*) and $\mathbf{T}$ (the *target schema*) are sequences of distinct relation symbols with no relation symbols in common and Σ is a set of formulas of some logical formalism over $\langle \mathbf{S}, \mathbf{T} \rangle$. We say that Σ *specifies* the schema $\mathcal{M}$. As in Fagin et al. [2005a, 2005b, 2005c], our main focus is on the important and practical case of schema mappings where Σ is a finite set of *source-to-target tuple-generating dependencies* (which we shall call *s-t tgds* or simply *tgds*). These are formulas of the form $\forall \mathbf{x}(\varphi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$, where $\varphi(\mathbf{x})$ is a conjunction of atoms[1] over $\mathbf{S}$, and where $\psi(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms over $\mathbf{T}$.[2] They have been used to formalize data exchange [Fagin et al. 2005a]. They have also been used in data integration scenarios under the name of GLAV (global-and-local-as-view) assertions [Lenzerini 2002]. Note that tgds do not contain equality, or any other "built-in relation symbols." When we consider egds (*equality-generating dependencies*), we shall of course treat equality as a built-in relation symbol that appears in the conclusion. Later (in Section 15), we shall extend the language of tgds so that the premise may include inequalities, and also a relation symbol *Constant* that represents constants.

Intuitively, we would expect invertibility of a schema mapping to correspond to "no loss of information." As an example, assume that the source schema has only the binary relation symbol P, and the target schema has only the unary relation symbol Q. Consider the projection schema mapping that is specified by the s-t tgd $\mathtt{P}(x, y) \rightarrow \mathtt{Q}(x)$.[3] It is clear that information is lost by this mapping, and, indeed, the projection schema mapping turns out not to have an inverse. Now assume that the source schema has only the binary relation symbol P, and the target schema has only the ternary relation symbol R. Consider the schema

[1]An *atom over* $\mathbf{S}$ is a formula of the form $\mathtt{P}(v_1, \ldots, v_m)$, where P is a relation symbol of $\mathbf{S}$, and $v_1, \ldots, v_m$ are variables; similarly, we define an *atom over* $\mathbf{T}$.
[2]There is also a safety condition, which says that every variable in $\mathbf{x}$ appears in φ. However, not all of the variables in $\mathbf{x}$ need to appear in ψ.
[3]We will often drop the universal quantifiers in front of a tgd, and implicitly assume such quantification. However, we will write down all existential quantifiers.

mapping that is specified by the s-t tgd $\mathtt{P}(x, y) \rightarrow \exists z \mathtt{R}(x, y, z)$. It is clear that no information is lost by this mapping, and indeed, this schema mapping turns out to have an inverse. One such inverse is specified by the tgd that results by "reversing the arrows," namely, $\mathtt{R}(x, y, z) \rightarrow \mathtt{P}(x, y)$. However, it turns out that "reversing the arrows" does not always produce an inverse, even when one exists.

There are other flavors of "schema mappings" that have been studied in the literature, such as view definitions, where there is a unique target instance associated with each source instance. In such cases, a schema mapping is a function in the classical sense, and so it is quite clear and unambiguous as to what an inverse mapping is. An example of such work is Hull's [1986] seminal research on information capacity of relational database schemas. Although our schema mappings are not actually functions, they have the advantage of being simpler and more flexible. LAV (local-as-view) mappings, which have been widely used in data integration, are special cases of schema mappings specified by s-t tgds, where we simply add the restriction that the premise of each tgd must be a single atom rather than a conjunction of atoms.

Let us now consider how to define the inverse in our context, where schema mappings are not actually functions. Let us associate with the schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ the set S_{12} of ordered pairs $\langle I, J \rangle$ such that I is a source instance, J is a target instance, and the pair $\langle I, J \rangle$ satisfies Σ_{12} (written $\langle I, J \rangle \models \Sigma_{12}$). Perhaps the most natural definition of the inverse of the schema mapping $\mathcal{M}_{12}$ would be a schema mapping $\mathcal{M}_{21}$ that is associated with the set $S_{21} = \{\langle J, I \rangle : \langle I, J \rangle \in S_{12}\}$. This reflects the standard algebraic definition of an inverse, and is the definition that Melnik [2004] and Melnik et al. [2005] gave for the inverse. In those articles, this definition was intended for a generic model management context, where mappings can be defined in a variety of ways, including as view definitions, relational algebra expressions, etc. However, this definition does not make sense in our context. This is because S_{12}, by being associated with a schema mapping specified by s-t tgds, is automatically "closed down on the left and closed up on the right." This means that if $\langle I, J \rangle \in S_{12}$ and if $I' \subseteq I$ (that is, I' is a subinstance of I) and $J \subseteq J'$, then $\langle I', J' \rangle \in S_{12}$. However, instead of being closed down on the left and closed up on the right, S_{21} is closed up on the left and closed down on the right. This is inconsistent with a schema mapping that is specified by a set of s-t tgds, which is the case we focus on in this article. In fact, the "language of inverse" (that is, the language needed to specify inverses for schema mappings specified by s-t tgds) turns out, as we shall discuss in Section 15, to be given by a generalization of s-t tgds that are also closed down on the left and closed up on the right.

Our notion of an inverse of a schema mapping is based on another algebraic property of inverses, that the composition of a function with its inverse is the identity mapping. In our context, the identity mapping is specified by tgds that "copy" the source instance to the target instance. Our definition of inverse says that the schema mapping $\mathcal{M}_{21}$ is an inverse of the schema mapping $\mathcal{M}_{12}$ for the class $\mathcal{S}$ of source instances if the schema mapping specified by their composition is equivalent on $\mathcal{S}$ to the identity mapping. We refer then to $\mathcal{M}_{21}$ as an $\mathcal{S}$-*inverse* of $\mathcal{M}_{12}$. When $\mathcal{S}$ is the class of all source instances, then $\mathcal{M}_{21}$ is said to be a *global inverse* of $\mathcal{M}_{12}$. When $\mathcal{S}$ is a singleton set containing only the source instance I, then $\mathcal{M}_{21}$ is said to be a *local inverse*, or simply an *inverse*, of $\mathcal{M}_{12}$ for I. Note that our definition of what it means for $\mathcal{M}_{21}$ to be an inverse of $\mathcal{M}_{12}$ corresponds exactly to what we would like an inverse mapping to do in data exchange: if after applying $\mathcal{M}_{12}$, we then apply $\mathcal{M}_{21}$, the resulting effect of $\mathcal{M}_{21}$ is to "undo" the effect of $\mathcal{M}_{12}$. Fortunately, because of work by Fagin et al. [2005c], we now understand very well the composition of schema mappings, and so we are in a good position to study our notion of inverse. This article is the first step in exploring the very rich theory that arises.

This article is an expanded version of a conference article [Fagin 2006]. The most significant difference between the two is that this article contains the proofs that were missing in Fagin [2006].

2. OVERVIEW OF RESULTS

If $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ is a schema mapping, I is a source instance, and J is a target instance, then J is a *solution* for I if $\langle I, J \rangle \models \Sigma_{12}$. A simple necessary condition for $\mathcal{M}_{12}$ to have a global inverse is the *unique-solutions property*, which says that no two distinct source instances have the same set of solutions. For a fixed choice of $\mathcal{M}_{12}$, let f be the set-valued function where $f(I)$ is the set of solutions for the source instance I. The unique-solutions property is equivalent to the condition that f be one-to-one. The fact that this condition is necessary for there to be a global inverse is analogous to the standard algebraic condition that an invertible function be one-to-one. We show that, surprisingly and pleasingly, in the important special case of LAV schema mappings, the unique-solutions property is not only necessary for $\mathcal{M}_{12}$ to have a global inverse but also sufficient.

Assume that $\mathcal{M}$ is a schema mapping specified by a finite set of s-t tgds, and I is a source instance. We derive a *canonical candidate tgd local inverse*, which is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}$ for I if there is any such inverse. We also derive a *canonical candidate tgd global inverse*, which is a schema mapping specified by a finite set of s-t tgds that is a global inverse of $\mathcal{M}$ if there is any such global inverse. In fact, we do something a little more general than this. Let $\mathcal{S}$ be a class of source instances specified (as source integrity constraints) by a set Γ of tgds and egds that always have a finite chase.[4] We generalize the notion of a canonical candidate tgd global inverse to that of a *canonical candidate tgd $\mathcal{S}$-inverse*, which is a schema mapping specified by a finite set of s-t tgds that is an $\mathcal{S}$-inverse of $\mathcal{M}$ if there is any such $\mathcal{S}$-inverse. (When Γ is the empty set, then $\mathcal{S}$ is the class of all source instances, and we obtain the canonical candidate tgd global inverse.) It might seem that the canonical candidate tgd local inverse is of theoretical interest only: after all, we typically care only about an inverse that "works" for a large class, not for a single instance. However, it turns out that the canonical candidate tgd local inverse plays a key role in the proof of correctness of the canonical candidate tgd global inverse (and the canonical candidate tgd $\mathcal{S}$-inverse).

[4]As usual, an egd, or *equality-generating dependency*, has the same form as a tgd, except that the conclusion must be an equality of variables.

Our canonical candidate tgd inverses are each specified by finite sets of *full* tgds (those with no existential quantifiers). This is not an accident: we show that if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are schema mappings that are each specified by a finite set of s-t tgds, $\mathcal{S}$ is a class of source instances, and $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, then there is a schema mapping specified by a finite set of *full* s-t tgds and that is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$.

It is folk wisdom that an inverse can be obtained by simply "reversing the arrows" in a tgd. We show that even a weak form of this folk wisdom is false. Instead, our canonical candidate tgd inverses are obtained by a slightly more complicated but still very natural procedure.

Since a local inverse may be quite tailored to a particular instance, it is natural to ask whether it is possible for a schema mapping specified by a finite set of s-t tgds to have an inverse for every source instance yet not have a global inverse. We show that this can indeed happen.

Given schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each specified by a finite set of s-t tgds, an analyst might want to investigate when (that is, for which source instances) $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$. (We give an example later, where $\mathcal{M}_{12}$ does projections and $\mathcal{M}_{21}$ joins the projections.) If we hold $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ fixed, then we show that the problem of deciding whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I is in the complexity class NP. It therefore follows from Fagin's Theorem [Fagin 1974] that the class $\mathcal{S}$ of source instances such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely the class $\mathcal{S}$ can be defined by a formula Γ in existential second-order logic. Remarkably, we are able to obtain such a formula Γ by a purely syntactical transformation of the formula that specifies the composition of the schema mappings. Furthermore, when $\mathcal{M}_{12}$ is specified by *full* s-t tgds, this formula is first-order.

Finally, we obtain other complexity results about deciding local or global invertibility.

3. APPLICATIONS OF INVERSE MAPPINGS

There are potentially a number of applications for inverse mappings, especially in schema evolution. For example, assume that data has been migrated from one schema to another with a schema mapping $\mathcal{M}$. At some point, we might decide to "roll back" to the original schema, and so we might want to apply an inverse schema mapping $\mathcal{M}^{-1}$. In fact, if we think this scenario is probable, we might deliberately choose a schema mapping $\mathcal{M}$ that has an inverse $\mathcal{M}^{-1}$.

As a more intricate example, assume that there are two different schema mappings from schema $\mathbf{S_1}$: the schema mapping $\mathcal{M}_1$ from schema $\mathbf{S_1}$ to schema $\mathbf{T_1}$, and the schema mapping $\mathcal{M}'_1$ from $\mathbf{S_1}$ to $\mathbf{S'_1}$. Assume that there is also a schema mapping $\mathcal{M}_2$ from $\mathbf{T_1}$ to $\mathbf{T'_1}$. If there is an "inverse schema mapping" ${\mathcal{M}'_1}^{-1}$ of $\mathcal{M}'_1$, then these schema mappings can be composed to give a schema mapping directly from $\mathbf{S'_1}$ to $\mathbf{T'_1}$, by taking the composition of the schema mapping ${\mathcal{M}'_1}^{-1}$ (from $\mathbf{S'_1}$ to $\mathbf{S_1}$) with the schema mapping $\mathcal{M}_1$ (from $\mathbf{S_1}$ to $\mathbf{T_1}$) and composing the result with the schema mapping $\mathcal{M}_2$ (from $\mathbf{T_1}$ to $\mathbf{T'_1}$).

There are several obstacles to practicality for inverse schema mappings. One obstacle to practicality is that the notion of an inverse of a schema mapping is rather restrictive, since it is rare that a schema mapping possesses an inverse. There are several ways to possibly mitigate this problem. One way, as noted earlier, is to strive to design a schema mapping to be invertible. Another way is to relax the notion of an inverse. One such relaxed notion is the "quasi-inverse" [Fagin et al. 2007]; practical applications of the quasi-inverse are discussed in Fagin et al. [2007].

Perhaps an even more serious obstacle to practicality is the complexity of even deciding if a schema mapping has a global inverse. If $\mathcal{M}_{12}$ is specified by a finite set of s-t tgds, we show (Corollary 14.10) that the problem of deciding whether $\mathcal{M}_{12}$ has a global inverse is coNP-hard. In fact, for all we know, this problem is undecidable! Possibly there is a natural class of schema mappings that arise in practice where this complexity is greatly reduced, but that remains to be seen.

4. BACKGROUND

We now review basic concepts from data exchange.

A *schema* is a finite sequence $\mathbf{R} = \langle \mathtt{R}_1, \ldots, \mathtt{R}_k \rangle$ of distinct relation symbols, each of a fixed arity. An *instance I (over the schema $\mathbf{R}$)* is a sequence $\langle \mathtt{R}_1^I, \ldots, \mathtt{R}_k^I \rangle$ such that each $\mathtt{R}_i^I$ is a finite relation of the same arity as $\mathtt{R}_i$. We call $\mathtt{R}_i^I$ the $\mathtt{R}_i$*-relation of I*. We shall often abuse the notation and use $\mathtt{R}_i$ to denote both the relation symbol and the relation $\mathtt{R}_i^I$ that interprets it.

Let $\mathbf{S} = \langle \mathtt{S}_1, \ldots, \mathtt{S}_n \rangle$ and $\mathbf{T} = \langle \mathtt{T}_1, \ldots, \mathtt{T}_m \rangle$ be two schemas with no relation symbols in common. We write $\langle \mathbf{S}, \mathbf{T} \rangle$ to denote the schema that is the result of concatenating the members of $\mathbf{S}$ with the members of $\mathbf{T}$. If I is an instance over $\mathbf{S}$ and J is an instance over $\mathbf{T}$, then we write $\langle I, J \rangle$ for the instance K over the schema $\langle \mathbf{S}, \mathbf{T} \rangle$ such that $\mathtt{S}_i^K = \mathtt{S}_i^I$ and $\mathtt{T}_j^K = \mathtt{T}_j^J$, for $1 \le i \le n$ and $1 \le j \le m$.

If K is an instance and σ is a formula in some logical formalism, then we write $K \models \sigma$ to mean that K satisfies σ. If Σ is a set of formulas, then we write $K \models \Sigma$ to mean that $K \models \sigma$ for every formula $\sigma \in \Sigma$.

Given a tuple $(t_1, \ldots, t_r)$ occurring in a relation R, we denote by $\mathtt{R}(t_1, \ldots, t_r)$ the association between $(t_1, \ldots, t_r)$ and R, and call it a *fact*. We will identify an instance with its set of facts. We call each t_i in the tuple $(t_1, \ldots, t_r)$ a *value*. We refer to the set of values that appear in an instance as its *active domain*. We assume that we have a fixed infinite set $\underline{\mathsf{Const}}$ of constants and an infinite set $\underline{\mathsf{Var}}$ of nulls that is disjoint from $\underline{\mathsf{Const}}$. We shall assume that the active domain of every source instance consists only of constants. We will sometimes emphasize this by referring to source instances as *ground instances*. Of course, it is possible for a "source instance" to arise that contains nulls, by chasing a source instance with s-t tgds from $\mathbf{S}$ to $\mathbf{T}$, and then chasing again with s-t tgds from $\mathcal{T}$ to $\mathbf{S}$. But we shall think of these "source instances" that contain nulls as simply artifacts, or "canonical instances."

If K is an instance with values in $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$, then $\underline{\mathsf{Var}}(K)$ denotes the set of nulls appearing in relations in K. Let K_1 and K_2 be two instances over the same schema with values in $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$. A *homomorphism* $h : K_1 \to K_2$ is a mapping from $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}(K_1)$ to $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}(K_2)$ such that (1) $h(c) = c$, for every

$c \in \underline{\text{Const}}$; and (2) for every fact $\mathtt{R}(\mathbf{t})$ of K_1, we have that $\mathtt{R}(h(\mathbf{t}))$ is a fact of K_2 (where, if $\mathbf{t} = (t_1, \ldots, t_s)$, then $h(\mathbf{t}) = (h(t_1), \ldots, h(t_s))$).

Consider a schema mapping $(\mathbf{S}, \mathbf{T}, \Sigma)$, as defined in the Introduction. Recall that if I is a ground instance, and J is a target instance, then J is a *solution* for I if $\langle I, J\rangle \models \Sigma$. If I is a ground instance, then a *universal solution for* I is a solution J for I such that for every solution J' for I, there exists a homomorphism $h : J \to J'$. When Σ is a finite set of s-t tgds, and I is a ground instance, then the result of chasing I with Σ is a universal solution for I [Fagin et al. 2005a].

Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{23} = (\mathbf{S_2}, \mathbf{S_3}, \Sigma_{23})$ be two schema mappings such that the schemas $\mathbf{S_1}, \mathbf{S_2}, \mathbf{S_3}$ have no relation symbol in common pairwise. The *composition formula* [Fagin et al. 2005c], denoted by $\Sigma_{12} \circ \Sigma_{23}$, has the semantics that, if I is an instance of $\mathbf{S_1}$ and J is an instance of $\mathbf{S_3}$, then $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{23}$ precisely if there is an instance J^* of $\mathbf{S_2}$ such that $\langle I, J^*\rangle \models \Sigma_{12}$ and $\langle J^*, J\rangle \models \Sigma_{23}$. It was proven in [Fagin et al. 2005c] that, when Σ_{12} and Σ_{23} are finite sets of s-t tgds, then the composition formula is given by a *second-order tgd (SO tgd)*. We give the definition of SO tgds later (Definition 12.1). We now give an example (from Fagin et al. [2005c]) of an SO tgd that defines the composition formula.

Example 4.1. Consider the following three schemas $\mathbf{S}_1$, $\mathbf{S}_2$ and $\mathbf{S}_3$. Schema $\mathbf{S}_1$ consists of a single unary relation symbol $\mathtt{Emp}$ of employees. Schema $\mathbf{S}_2$ consists of a single binary relation symbol $\mathtt{Mgr}_1$ that associates each employee with a manager. Schema $\mathbf{S}_3$ consists of a similar binary relation symbol $\mathtt{Mgr}$ that is intended to provide a copy of $\mathtt{Mgr}_1$, and an additional unary relation symbol $\mathtt{SelfMgr}$ that is intended to store employees who are their own manager. Consider now the schema mappings $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$ and $\mathcal{M}_{23} = (\mathbf{S}_2, \mathbf{S}_3, \Sigma_{23})$, where Σ_{12} consists of the tgd $\mathtt{Emp}(e) \to \exists m \mathtt{Mgr}_1(e, m)$, and Σ_{23} consists of the two tgds $\mathtt{Mgr}_1(e, m) \to \mathtt{Mgr}(e, m)$ and $\mathtt{Mgr}_1(e, e) \to \mathtt{SelfMgr}(e)$. Then the composition formula $\Sigma_{12} \circ \Sigma_{23}$ is defined by the following second-order tgd:

$$\exists f (\forall e(\mathtt{Emp}(e) \to \mathtt{Mgr}(e, f(e))) \land \forall e(\mathtt{Emp}(e) \land (e = f(e)) \to \mathtt{SelfMgr}(e))). \quad (1)$$

Intuitively, $f(e)$ is the manager of employee e.

We now give a lemma that corresponds to our statement in the Introduction that the set $\langle I, J\rangle$ of pairs that satisfy an s-t tgd is “closed down on the left and closed up on the right.”

Lemma 4.2. *Let Σ be a finite set of s-t tgds. Assume that $\langle I, J\rangle \models \Sigma$, and $I' \subseteq I$ and $J \subseteq J'$. Then $\langle I', J'\rangle \models \Sigma$.*

Proof. This follows easily from the definitions. □

Corollary 4.3. *Let Σ_{12} and Σ_{23} be finite sets of s-t tgds. Assume that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{23}$ and $I' \subseteq I$ and $J \subseteq J'$. Then $\langle I', J'\rangle \models \Sigma_{12} \circ \Sigma_{23}$.*

Proof. Since $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{23}$, there is J^* such that $\langle I, J^*\rangle \models \Sigma_{12}$ and $\langle J^*, J\rangle \models \Sigma_{23}$. Since $\langle I, J^*\rangle \models \Sigma_{12}$ and $I' \subseteq I$, it follows from Lemma 4.2 that $\langle I', J^*\rangle \models \Sigma_{12}$. Since $\langle J^*, J\rangle \models \Sigma_{23}$ and $J \subseteq J'$, it follows from Lemma 4.2 that $\langle J^*, J'\rangle \models \Sigma_{23}$. Since $\langle I', J^*\rangle \models \Sigma_{12}$ and $\langle J^*, J'\rangle \models \Sigma_{23}$, it follows that $\langle I', J'\rangle \models \Sigma_{12} \circ \Sigma_{23}$, as desired. □

5. WHAT IS AN INVERSE MAPPING?

Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ is a schema mapping. For each relation symbol $\mathtt{R}$ of $\mathbf{S_1}$, let $\widehat{\mathtt{R}}$ be a new relation symbol (different from any relation symbol of $\mathbf{S_1}$ or $\mathbf{S_2}$) of the same arity as $\mathtt{R}$. Define $\widehat{\mathbf{S_1}}$ to be $\{\widehat{\mathtt{R}} : \mathtt{R} \in \mathbf{S_1}\}$. Thus, $\widehat{\mathbf{S_1}}$ is a schema disjoint from $\mathbf{S_1}$ and $\mathbf{S_2}$ that can be thought of as a copy of $\mathbf{S_1}$. If I is an instance of $\mathbf{S_1}$, define $\widehat{I}$ to be the corresponding instance of $\widehat{\mathbf{S_1}}$. Thus, $\widehat{\mathtt{R}}^{\widehat{I}} = \mathtt{R}^I$ for every $\mathtt{R}$ in $\mathbf{S_1}$.

Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be a schema mapping, where the source schema $\mathbf{S_2}$ is the target schema of $\mathcal{M}_{12}$, and where the target schema is $\widehat{\mathbf{S_1}}$. The issue we are concerned with is: what does it mean for $\mathcal{M}_{21}$ to be an inverse of $\mathcal{M}_{12}$, and what can we say about such inverse mappings? We are most interested in the case where Σ_{12} and Σ_{21} are finite sets of s-t tgds. We now introduce an example that we shall use as a running example to demonstrate some of the issues that arise.

Example 5.1. Let $\mathbf{S_1}$ consist of the ternary relation symbol EDL (*Employee-Department-Location*). Let $\mathbf{S_2}$ consist of the binary relation symbol ED (*Employee-Department*) and the binary relation symbol DL (*Department-Location*). Let Σ_{12} consist of the s-t tgd $\mathtt{EDL}(x, y, z) \rightarrow \mathtt{ED}(x, y) \wedge \mathtt{DL}(y, z)$ that corresponds to projecting EDL onto ED and DL. Let Σ_{21} consist of the s-t tgd $(\mathtt{ED}(x, y) \wedge \mathtt{DL}(y, z)) \rightarrow \widehat{\mathtt{EDL}}(x, y, z)$, where the source schema is $\mathbf{S_2}$ and the target schema is $\widehat{\mathbf{S_1}}$, that corresponds to taking the join of the projections. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$.

Let Γ be the multivalued dependency[5]

$$\mathtt{EDL}(x, y, z') \wedge \mathtt{EDL}(x', y, z)) \rightarrow \mathtt{EDL}(x, y, z). \tag{2}$$

It is known [Fagin 1977] that, if we project the EDL relation onto ED and DL and then join the resulting projections, we obtain the original EDL relation precisely if the multivalued dependency Γ holds for the EDL relation. We want our definition of inverse to have the property that the schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely those ground instances I that satisfy Γ.

Let us now define some preliminary notions that will allow us to define what it means for the mapping $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ to be an $\mathcal{S}$-inverse of the mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$. (In Example 5.1, the class $\mathcal{S}$ would consist of those ground instances that satisfy Γ.) Define Σ_{Id} (where Id stands for *identity*) to consist of the tgds $\mathtt{R}(x_1, \ldots, x_k) \rightarrow \widehat{\mathtt{R}}(x_1, \ldots, x_k)$, where $x_1, \ldots, x_k$ are distinct variables, when $\mathtt{R}$ is a k-ary relation symbol of $\mathbf{S_1}$. Define the *identity mapping* to be $\mathcal{M}_{Id} = (\mathbf{S_1}, \widehat{\mathbf{S_1}}, \Sigma_{Id})$. Note that J is a solution for I under the identity mapping if and only if $\widehat{I} \subseteq J$. The reason we have $\widehat{I} \subseteq J$ rather than simply $\widehat{I} = J$ is that Σ_{Id} is a set of s-t tgds, and hence whenever J is a solution, then so is every J' with $J \subseteq J'$. Let us say that two schema mappings with source schema $\mathbf{S_1}$ and target schema $\widehat{\mathbf{S_1}}$ are *equivalent on* I if they have the same ground instances as solutions for I.[6]

[5] Note that Γ is not an s-t tgd, since the premise and conclusion use the same relation symbol EDL. Of course, Γ is a tgd in the classical sense of Beeri and Vardi [1984].

[6] Technically, a ground instance is an instance of $\mathbf{S_1}$ whose active domain consists only of constants.

We are now ready to define the notion of inverse. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be schema mappings. Let σ be the composition formula $\Sigma_{12} \circ \Sigma_{21}$ of $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$, and let $\mathcal{M}_{11} = (\mathbf{S_1}, \widehat{\mathbf{S_1}}, \sigma)$. Let I be a ground instance of $\mathbf{S_1}$. Let us say that $\mathcal{M}_{21}$ *is an inverse of* $\mathcal{M}_{12}$ *for* I if $\mathcal{M}_{11}$ and the identity mapping $\mathcal{M}_{Id}$ are equivalent on I. Thus, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for the ground instance I precisely if for every ground instance J,

$$\langle I, J \rangle \models \sigma \text{ if and only if } \widehat{I} \subseteq J. \tag{3}$$

We note that in the case where Σ_{12} and Σ_{21} are each a finite set of s-t tgds (the case we shall mainly consider in this article), (3) holds for every ground instance J if and only if (3) holds for every instance J (ground or not) of $\widehat{\mathbf{S_1}}$ (this can be shown by replacing each null by a new constant; if a null appears several times, then it is replaced by the same constant each time). However, this is not necessarily true when either Σ_{12} or Σ_{21} is not a finite set of s-t tgds.

If $\mathcal{S}$ is a class of ground instances, then we say that $\mathcal{M}_{21}$ *is an* $\mathcal{S}$*-inverse of* $\mathcal{M}_{12}$ if $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I, for each I in $\mathcal{S}$. A particularly important case arises when $\mathcal{S}$ is the class of all ground instances. In that case, we say that $\mathcal{M}_{21}$ is a *global inverse* of $\mathcal{M}_{12}$.

We now discuss further the fact that there is an inclusion rather than an equality in the right-hand side of (3). On the face of it, this seems to be an artifact that arises because the identity schema mapping Σ_{Id}, as we have given it, is specified by s-t tgds, and so the results are "closed up on the right," as in Lemma 4.2. We now show that the reason for the inclusion rather than an equality in the right-hand side of (3) is more fundamental. Assume that we are considering a global inverse of a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where Σ_{12} is a finite set of s-t tgds. The next proposition implies that no matter how rich a language we are able to use to define Σ_{Id}, there still must be an inclusion rather than an equality in the right-hand side of (3).

PROPOSITION 5.2. *Let* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *be a schema mapping, where* Σ_{12} *is a finite set of s-t tgds. Assume that* Σ_{21} *is an arbitrary formula (not necessarily a finite set of s-t tgds), such that* $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$*, for every ground instance* I*. Then* $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$ *whenever* J *is a ground instance with* $\widehat{I} \subseteq J$*.*

PROOF. Assume that $\widehat{I} \subseteq J$. Find I' such that $\widehat{I'} = J$. So $I \subseteq I'$. By assumption, $\langle I', \widehat{I'} \rangle \models \Sigma_{12} \circ \Sigma_{21}$, that is, $\langle I', J \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Therefore, by definition of $\Sigma_{12} \circ \Sigma_{21}$, there is J' such that $\langle I', J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{21}$. Since $I \subseteq I'$ and $\langle I', J' \rangle \models \Sigma_{12}$, and since Σ_{12} is a finite set of s-t tgds, it follows from Lemma 4.2 that $\langle I, J' \rangle \models \Sigma_{12}$. Since $\langle I, J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{21}$, it follows that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$. This was to be shown. □

We now return to our running example of Example 5.1 to further investigate the question of when (that is, for which ground instances) one schema mapping is an inverse of another.

But we may also refer to an instance of $\widehat{\mathbf{S_1}}$ whose active domain consists only of constants as a ground instance.

Example 5.3. We said in Example 5.1 that we want our definition of inverse to have the property that the schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely those ground instances I that satisfy Γ. We now show that satisfying Γ is a sufficient condition for $\mathcal{M}_{21}$ to be an inverse of $\mathcal{M}_{12}$. In Example 12.8, we shall show that Γ is also a necessary condition.

If we apply the composition algorithm of Fagin et al. [2005c], we find that the composition formula $\Sigma_{12} \circ \Sigma_{21}$, which we denote by σ, is

$$\mathtt{EDL}(x, y, z') \wedge \mathtt{EDL}(x', y, z)) \rightarrow \widehat{\mathtt{EDL}}(x, y, z). \tag{4}$$

Let I be a ground instance of $\mathbf{S_1}$ satisfying Γ. We must show that (3) holds for every ground instance J. Assume first that $\langle I, J \rangle \models \sigma$; we must show that $\widehat{I} \subseteq J$. Now Σ_{Id} consists of the tgd $\mathtt{EDL}(x, y, z) \rightarrow \widehat{\mathtt{EDL}}(x, y, z)$. It is clear that σ logically implies Σ_{Id} (we let the roles of x' and z' be played by x and z, respectively). Therefore, since $\langle I, J \rangle \models \sigma$, it follows that $\langle I, J \rangle \models \Sigma_{Id}$. So $\widehat{I} \subseteq J$, as desired.

Assume now that $\widehat{I} \subseteq J$; we must show that $\langle I, J \rangle \models \sigma$. Since σ is given by (4), this means that we must show that if $\mathtt{EDL}(x, y, z')$ and $\mathtt{EDL}(x', y, z)$ hold in I, then $\widehat{\mathtt{EDL}}(x, y, z)$ holds in J. So assume that $\mathtt{EDL}(x, y, z')$ and $\mathtt{EDL}(x', y, z)$ hold in I. Since $I \models \Gamma$, it follows that $\mathtt{EDL}(x, y, z)$ holds in I. Since $\widehat{I} \subseteq J$, it follows that $\widehat{\mathtt{EDL}}(x, y, z)$ holds in J, as desired.

Note the unexpected similarity of the composition formula (4) and Γ (the multivalued dependency (2)). We shall explain this surprising connection between the composition formula and Γ later (in Example 12.8).

The next example shows that there need not be a unique inverse. Therefore, we refer to "*an* inverse mapping" rather than "*the* inverse mapping."

Example 5.4. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where $\mathbf{S_1}$ consists of the unary relation symbol $\mathtt{R}$, where $\mathbf{S_2}$ consists of the binary relation symbol $\mathtt{S}$, and where Σ_{12} consists of the tgd $\mathtt{R}(x) \rightarrow \mathtt{S}(x, x)$. Let Σ_{21} consist of the tgd $\mathtt{S}(x, x) \rightarrow \widehat{\mathtt{R}}(x)$, and let Σ'_{21} consist of the tgd $\mathtt{S}(x, y) \rightarrow \widehat{\mathtt{R}}(x)$. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, and let $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$. In both cases (for $\mathcal{M}_{21}$ and for $\mathcal{M}'_{21}$), the composition formula is $\mathtt{R}(x) \rightarrow \widehat{\mathtt{R}}(x)$, which specifies the identity mapping. So both $\mathcal{M}_{21}$ and $\mathcal{M}'_{21}$ are global inverses of $\mathcal{M}_{12}$, and so there is not a unique global inverse of $\mathcal{M}_{12}$.

Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where Σ_{12} is a finite set of s-t tgds. Assume that $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of s-t tgds. Later, we show (Corollary 9.4) that there is then a schema mapping ("the canonical candidate tgd global inverse") $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{21} is a finite set of s-t tgds, that is, the "most general" global inverse of $\mathcal{M}_{12}$ that is specified by s-t tgds, in the sense that Σ'_{21} logically implies Σ_{21} for every global inverse $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ of $\mathcal{M}_{12}$ where Σ'_{21} is a finite set of s-t tgds. In the case of Example 5.4, the schema mapping $\mathcal{M}_{21}$ is the canonical candidate tgd global inverse of $\mathcal{M}_{12}$ (note in particular that Σ'_{21} logically implies Σ_{21}).

The notion *global inverse* is not symmetric. Thus, in Example 5.4, although $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, it is not true that $\mathcal{M}_{12}$ is a global inverse of $\mathcal{M}_{21}$.[7] In fact, it is straightforward to show that $\mathcal{M}_{21}$ has no inverse for the instance consisting of the single fact $\texttt{S}(0, 1)$, and hence $\mathcal{M}_{21}$ has no global inverse.

The final example of this section shows that, in some cases, we may get uniqueness.

Example 5.5. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where $\mathbf{S_1}$ consists of the unary relation symbol R, where $\mathbf{S_2}$ consists of the unary relation symbol S, and where Σ_{12} consists of the tgd $\texttt{R}(x) \rightarrow \texttt{S}(x)$. Let Σ_{21} consist of the tgd $\texttt{S}(x) \rightarrow \widehat{\texttt{R}}(x)$, and let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. It is not hard to see that $\mathcal{M}_{21}$ is the unique global inverse of $\mathcal{M}_{12}$ (up to logical equivalence of Σ_{21}) specified by s-t tgds.

6. THE UNIQUE-SOLUTIONS PROPERTY

Unlike the rest of this article, in this section we do not restrict our attention to schema mappings $(\mathbf{S}, \mathbf{T}, \Sigma)$ where Σ is a finite set of s-t tgds. Instead, we may allow Σ to be an arbitrary constraint between source and target instances. Our only requirement is that the satisfaction relation between formulas and instances be preserved under isomorphism. This means that, if $\langle I, J\rangle \models \Sigma$, and if $\langle I', J'\rangle$ is isomorphic to $\langle I, J\rangle$, then $\langle I', J'\rangle \models \Sigma$. This is a mild condition that is true of all standard logical formalisms, such as first-order logic, second-order logic, fixed-point logics, and infinitary logics.

Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ be a schema mapping, and let I be a ground instance. Intuitively, as far as $\mathbf{S_2}$ is concerned, the only information about I is the set of solutions for I, that is, the set of target instances J such that $\langle I, J\rangle \models \Sigma_{12}$. Therefore, we would expect that, if $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for two distinct instances I_1 and I_2, then I_1 and I_2 would have different sets of solutions. Otherwise, intuitively, there would not be enough information to allow $\mathcal{M}_{21}$ to reconstruct I_1 after applying $\mathcal{M}_{12}$. We now give a theorem (Theorem 6.1), which says that this intuition is correct. We then give three corollaries of Theorem 6.1, each of which provides a useful necessary condition for the existence of local inverses, and each of which we shall utilize later. As an amusing application of Theorem 6.1, we show that there is a schema mapping specified by a finite set of s-t tgds that has an inverse for every ground instance yet does not have a global inverse. We conclude this section by considering global rather than local properties. Specifically, we say that a schema mapping has the *unique-solutions property* if no two distinct instances have the same set of solutions. We show that the unique-solutions property is a necessary condition for a schema mapping to have a global inverse. Furthermore, in the case of LAV schema mappings, we show that the unique-solutions property is not only necessary but also sufficient for the existence of a global inverse.

THEOREM 6.1. *Let $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ be schema mappings. Assume that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for distinct ground instances I_1 and I_2. Then the set of solutions for I_1 under $\mathcal{M}_{12}$ is different from the set of solutions for I_2 under $\mathcal{M}_{12}$.*

[7] To even make sense syntactically of the question as to whether $\mathcal{M}_{12}$ is a global inverse of $\mathcal{M}_{21}$, we must first rename $\widehat{\mathbf{S_1}}$ in $\mathcal{M}_{21}$ to be $\mathbf{S_1}$, and rename $\mathbf{S_2}$ in $\mathcal{M}_{12}$ to be $\widehat{\mathbf{S_2}}$.

Proof. Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. Let σ be the composition formula of $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$. Assume that I_1 and I_2 are ground instances such that the set of solutions for I_1 under $\mathcal{M}_{12}$ equals the set of solutions for I_2 under $\mathcal{M}_{12}$, that is,

$$\left\{J' : \langle I_1, J'\rangle \models \Sigma_{12}\right\} = \left\{J' : \langle I_2, J'\rangle \models \Sigma_{12}\right\}. \quad (5)$$

We shall show that $I_1 = I_2$. We have (where we always take J to be ground) the following:

$\left\{J : \widehat{I_1} \subseteq J\right\}$
$= \{J : \langle I_1, J\rangle \models \sigma\}$ (by (3) with I_1 for I)
$= \{J :$ There exists J' such that $\langle I_1, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma_{21}\}$
(by definition of the composition formula)
$= \{J :$ There exists J' such that $\langle I_2, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma_{21}\}$ (by (5))
$= \{J : \langle I_2, J\rangle \models \sigma\}$ (by definition of the composition formula)
$= \left\{J : \widehat{I_2} \subseteq J\right\}$ (by (3) with I_2 for I).

We just showed that $\left\{J : \widehat{I_1} \subseteq J\right\} = \left\{J : \widehat{I_2} \subseteq J\right\}$. Since $\widehat{I_1} \in \left\{J : \widehat{I_1} \subseteq J\right\}$, it follows that $\widehat{I_1} \in \left\{J : \widehat{I_2} \subseteq J\right\}$, that is, $\widehat{I_2} \subseteq \widehat{I_1}$. Identically, we have $\widehat{I_1} \subseteq \widehat{I_2}$, and so $\widehat{I_1} = \widehat{I_2}$. Therefore, $I_1 = I_2$, as desired. □

The reader might wonder why we are considering this property, which says that distinct instances have distinct solution sets, rather than considering a stronger property that says that distinct instances should not even have a solution in common. Let $\mathcal{M}_{12}$ be a schema mapping specified by a finite set of s-t tgds, and let I_1 and I_2 be distinct ground instances. We now show that I_1 and I_2 always have a solution in common, whether or not $\mathcal{M}_{12}$ is invertible. This is because if J_1 is an arbitrary solution of I_1, and J_2 is an arbitrary solution of I_2, then Lemma 4.2 implies that $J_1 \cup J_2$ is a solution of both I_1 and I_2 (where $J_1 \cup J_2$ is the target instance whose set of facts consists of the union of the facts of J_1 and J_2).

As a corollary of Theorem 6.1, we obtain a necessary condition for $\mathcal{M}_{12}$ to have an inverse for a fixed ground instance. (The proof depends on our assumption of preservation under isomorphism.)

Corollary 6.2. *Let $\mathcal{M}_{12}$ be a schema mapping, and let I_1 and I_2 be distinct but isomorphic ground instances. Assume that there is an inverse of $\mathcal{M}_{12}$ for I_1. Then the set of solutions for I_1 under $\mathcal{M}_{12}$ is different from the set of solutions for I_2 under $\mathcal{M}_{12}$.*

Proof. Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, and that $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ is an inverse of $\mathcal{M}_{12}$ for I_1. Let σ be the composition formula $\Sigma_{12} \circ \Sigma_{21}$, and let $\mathcal{M}_{11} = (\mathbf{S_1}, \widehat{\mathbf{S_1}}, \sigma)$. Since $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I_1, it follows by definition of inverse that $\mathcal{M}_{11}$ and the identity mapping are equivalent on I_1. Since I_1 and I_2 are isomorphic, it follows in a straightforward way from our assumption of preservation under isomorphism that $\mathcal{M}_{11}$ and the identity mapping are equivalent on I_2. Hence, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I_2. So by Theorem 6.1, the set of solutions for I_1 and I_2 are different. □

We now give a simple example of the use of Corollary 6.2.

Example 6.3. Let $\mathbf{S_1}$ consist of the unary relation symbols R and R′, let $\mathbf{S_2}$ consist of the unary relation symbol S, and let $\Sigma_{12} = \{\mathtt{R}(x) \rightarrow \mathtt{S}(x), \mathtt{R}'(x) \rightarrow \mathtt{S}(x)\}$. Assume that the facts of I_1 are precisely R(0) and R′(1); we now show that $\mathcal{M}_{12}$ does not have an inverse for I_1. Let I_2 be the ground instance whose facts are precisely R(1) and R′(0). Let J be the target instance whose facts are precisely S(0) and S(1). Then the solutions under Σ_{12} for I_1 are exactly those J' where $J \subseteq J'$. But these are also exactly the solutions for I_2. Since I_1 and I_2 are distinct isomorphic ground instances with the same set of solutions, it follows from Corollary 6.2 that $\mathcal{M}_{12}$ does not have an inverse for I_1. □

We now give two more corollaries of Theorem 6.1, both of which give necessary conditions for the existence of an inverse of schema mappings specified by s-t tgds. Both corollaries make use of the fundamental notion of the chase. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where Σ_{12} is a finite set of s-t tgds. Assume that I is an instance of $\mathbf{S_1}$. If the result of chasing $\langle I, \emptyset\rangle$ with Σ_{12} is $\langle I, J\rangle$, then we define $chase_{12}(I)$ to be J.[8] We may say loosely that J is the result of chasing I with Σ_{12}.

Corollary 6.4. *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ be a schema mapping, where Σ_{12} is a finite set of s-t tgds. Assume that some schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for distinct ground instances I_1 and I_2. Then $chase_{12}(I_1) \neq chase_{12}(I_2)$.*

Proof. It follows from results in Fagin et al. [2005a] that the solutions for a ground instance I are exactly the homomorphic images of $chase_{12}(I)$. Therefore, if $chase_{12}(I_1) = chase_{12}(I_2)$, then I_1 and I_2 have the same solutions. But this is a contradiction, by Theorem 6.1. □

Let I be a ground instance. Let us say that the schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ has the *constant-propagation property for I* if every member of the active domain of I is a member of the active domain of $chase_{12}(I)$. If $\mathcal{M}_{12}$ has the constant-propagation property for every ground instance I, then we say simply that $\mathcal{M}_{12}$ has the *constant-propagation property*. The next corollary says that the constant-propagation property for I is a necessary condition for invertibility for I.

Corollary 6.5. *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ be a schema mapping, where Σ_{12} is a finite set of s-t tgds. If $\mathcal{M}_{12}$ has an inverse for I (not necessarily specified by s-t tgds), then $\mathcal{M}_{12}$ has the constant-propagation property for I.*

Proof. Assume that the constant a in the active domain of I is not in the active domain of $chase_{12}(I)$; we shall derive a contradiction. Let a' be a new constant that does not appear in I, and let I' be the result of replacing every occurrence of a in I by a'. Then I and I' are isomorphic, and

[8]For definiteness, we use the version of the chase as defined in Fagin et al. [2005c], although it does not really matter.

$chase_{12}(I) = chase_{12}(I')$. Hence, as in the proof of Corollary 6.4, it follows that I and I' have the same solutions. But this is a contradiction of Corollary 6.2. □

The next proposition is an amusing application of Theorem 6.1.

PROPOSITION 6.6. *There is a schema mapping $\mathcal{M}_{12}$ specified by a finite set of full s-t tgds that has no global inverse, but where, for every ground instance I, there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I.*

PROOF. Let $\mathbf{S_1}$ consist of the unary relation symbols P and Q, and let $\mathbf{S_2}$ consist of the binary relation symbol R and the unary relation symbol S. Let $\Sigma_{12} = \{\mathtt{P}(x) \wedge \mathtt{Q}(y) \rightarrow \mathtt{R}(x, y), \mathtt{P}(x) \rightarrow \mathtt{S}(x), \mathtt{Q}(x) \rightarrow \mathtt{S}(x)\}$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$.

We now show that for every ground instance I, the schema mapping $\mathcal{M}_{12}$ has an inverse that is specified by a finite set of s-t tgds. There are three cases:

—$\mathtt{P}^I$ is empty. Then an inverse is $\mathtt{S}(x) \rightarrow \widehat{\mathtt{Q}}(x)$.
—$\mathtt{Q}^I$ is empty. Then an inverse is $\mathtt{S}(x) \rightarrow \widehat{\mathtt{P}}(x)$.
—Neither $\mathtt{P}^I$ nor $\mathtt{Q}^I$ is empty. Then an inverse is $\mathtt{R}(x, y) \rightarrow \widehat{\mathtt{P}}(x) \wedge \widehat{\mathtt{Q}}(y)$.

Now we will show that $\mathcal{M}_{12}$ does not have a global inverse. Let $I_1 = \{P(0)\}$, and let $I_2 = \{Q(0)\}$. Then the set of solutions for I_1 under $\mathcal{M}_{12}$ equals the set of solutions for I_2 under $\mathcal{M}_{12}$ (both equal the set of target instances J that contain $\{S(0)\}$). It then follows from Theorem 6.1 that there is no schema mapping $\mathcal{M}_{21}$ that is an inverse of $\mathcal{M}_{12}$ for both I_1 and I_2. Therefore $\mathcal{M}_{12}$ does not have a global inverse. □

Let us say that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ has the *unique-solutions property* if whenever I_1 and I_2 are distinct ground instances, then the set of solutions for I_1 is distinct from the set of solutions for I_2. In the case where Σ_{12} is a finite set of s-t tgds, it follows from results of Fagin et al. [2005a] that I_1 and I_2 have the same set of solutions if and only if they share a universal solution. Therefore, when Σ_{12} is a finite set of tgds, the unique-solutions property is equivalent to the *unique-universal-solution property*, which says that whenever I_1 and I_2 are distinct ground instances, then no universal solution for I_1 is a universal solution for I_2.

In this section, the result of greatest interest is the following important special case of Theorem 6.1.

THEOREM 6.7. *Every schema mapping with a global inverse satisfies the unique-solutions property.*

Recall that a *LAV* (local-as-view) schema mapping is a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ where Σ_{12} is a finite set of s-t tgds all with a singleton premise. The next theorem says that, for LAV schema mappings, the unique-solutions property is not only necessary for global invertibility but also sufficient. This shows robustness of our notion of inverse, since (at least in the case of LAV mappings) our notion of global invertibility is equivalent to the unique-solutions property, which is another natural notion. Before we state and prove this theorem, we need two simple lemmas. The first lemma, which is standard, gives the key property of the chase.

LEMMA 6.8. *Let* $(\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *be a schema mapping, where* Σ_{12} *is a finite set of s-t tgds. Let* I *be a ground instance and* J *a target instance. Then* $\langle I, J\rangle \models \Sigma_{12}$ *if* $chase_{12}(I) \subseteq J$. *If the s-t tgds in* Σ_{12} *are all full, then* $\langle I, J\rangle \models \Sigma_{12}$ *if and only if* $chase_{12}(I) \subseteq J$.

The next lemma gives a sufficient condition for an instance to be a universal solution.

LEMMA 6.9. *Let* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *be a schema mapping, where* Σ_{12} *is a finite set of s-t tgds. Let* J *be a universal solution for* I, *and let* J' *be a subinstance of* J *that is also a homomorphic image of* J. *Then* J' *is also a universal solution for* I.

PROOF. Since J' is a homomorphic image of the universal solution J, and since Σ_{12} is a finite set of s-t tgds, it follows that J' is a solution for I. Let J'' be an arbitrary solution for I. Since J is universal, there is a homomorphism from J into J''. But this homomorphism, when restricted to the subinstance J' of J, is a homomorphism from J' into J''. So J' is indeed a universal solution for I. □

THEOREM 6.10. *A LAV schema mapping has a global inverse if and only if it has the unique-solutions property.*

PROOF.[9] By Theorem 6.7, we know that when a schema mapping (LAV or otherwise) has a global inverse, then it has the unique-solutions property. Assume now that the LAV schema $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ has the unique-solutions property; we must show that it has a global inverse.

Define Σ_{21} to have the meaning that it defines all pairs $(chase_{12}(I'), \widehat{I'})$ where I' is a ground instance. That is, $\langle J_2, J_1\rangle \models \Sigma_{21}$ precisely if there is a ground instance I' such that $J_1 = \widehat{I'}$ and $J_2 = chase_{12}(I')$. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. We must show that, for every ground instance J,

$$\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21} \text{ if and only if } \widehat{I} \subseteq J. \tag{6}$$

We first show that if $\widehat{I} \subseteq J$, then $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$. By Proposition 5.2, we need only show that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ for every ground instance I. Let $J^* = chase_{12}(I)$. Then $\langle I, J^*\rangle \models \Sigma_{12}$ by Lemma 6.8. By definition of Σ_{21}, it follows that $\langle J^*, \widehat{I}\rangle \models \Sigma_{21}$. Since $\langle I, J^*\rangle \models \Sigma_{12}$ and $\langle J^*, \widehat{I}\rangle \models \Sigma_{21}$, it follows that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired.

Assume now that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$; we must show that $\widehat{I} \subseteq J$. Since $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$, there is J^* such that $\langle I, J^*\rangle \models \Sigma_{12}$ and $\langle J^*, J\rangle \models \Sigma_{21}$. Since $\langle J^*, J\rangle \models \Sigma_{21}$, it follows by definition of Σ_{21} that there is I' such that $J^* = chase_{12}(I')$ and $J = \widehat{I'}$. Let $I'' = I \cup I'$ (that is, the set of facts of I'' consists of the union of the facts of I and I'). Let $I''' = I \setminus I'$ (that is, the set of facts of I''' consists of

[9]Catriel Beeri simplified the proof.

those facts in I that are not in I'). So $I'' = I''' \cup I'$. Since Σ_{12} is LAV, it follows that

$$chase_{12}(I'') = chase_{12}(I''') \cup chase_{12}(I'), \tag{7}$$

that is, the chase of the union is the union of the chases. This follows immediately from the fact that the premise of every member of Σ_{12} is a singleton.

In chasing I'' with Σ_{12}, we can first chase I''' with Σ_{12} and then chase I' with Σ_{12}. Let X''' be the set of nulls introduced in chasing I''' with Σ_{12}, and let X' be the set of nulls introduced in chasing I' with Σ_{12}. Then X''' and X' are disjoint. It was shown in Fagin et al. [2005a] that $chase_{12}(I)$ is a universal solution for I under $\mathcal{M}_{12}$. Therefore, since J^* is a solution for I under $\mathcal{M}_{12}$, there is a homomorphism h such that $h(chase_{12}(I)) \subseteq J^*$. Combining this with the fact that $J^* = chase_{12}(I')$, we obtain

$$h(chase_{12}(I)) \subseteq chase_{12}(I'). \tag{8}$$

Since $I''' \subseteq I$, we have $chase_{12}(I''') \subseteq chase_{12}(I)$. So

$$h(chase_{12}(I''')) \subseteq h(chase_{12}(I)). \tag{9}$$

By (8) and (9), it follows that

$$h(chase_{12}(I''') \subseteq chase_{12}(I'). \tag{10}$$

Define h'' on $chase_{12}(I'')$ by letting $h''(x) = h(x)$ if $x \in X'''$, and $h''(x) = x$ if $x \in X'$. By (7), we have

$$h''(chase_{12}(I'')) = h''(chase_{12}(I''')) \cup h''(chase_{12}(I')). \tag{11}$$

But $h''(chase_{12}(I''')) = h(chase_{12}(I'''))$, and $h''(chase_{12}(I')) = chase_{12}(I')$. So from (11), we see that $h''(chase_{12}(I'')) = h(chase_{12}(I''')) \cup chase_{12}(I')$. Hence, by (10), it follows that $h''(chase_{12}(I'')) = chase_{12}(I')$. Since $I' \subseteq I''$, we know $chase_{12}(I') \subseteq chase_{12}(I'')$. We have shown that $chase_{12}(I')$ is both a subinstance of $chase_{12}(I'')$ and (under h'') a homomorphic image of $chase_{12}(I'')$. But $chase_{12}(I'')$ is a universal solution for I''. So by Lemma 6.9, $chase_{12}(I')$ is a universal solution for I''.

Hence, $chase_{12}(I')$ is a universal solution for both I' and I''. So I' and I'' share a universal solution. We noted earlier that, for schema mappings specified by a finite set of s-t tgds, the unique-solutions property is equivalent to the unique-universal-solution property. By assumption, $\mathcal{M}_{12}$ has the unique-solutions property, and hence it has the unique-universal-solution property. Therefore, since I' and I'' share a universal solution, it follows that $I' = I''$, that is, $I' = I \cup I'$. But this implies that $I \subseteq I'$, and so $\widehat{I} \subseteq \widehat{I'}$. Since $J = \widehat{I'}$, this implies that $\widehat{I} \subseteq J$. This was to be shown. □

It was shown in Fagin et al. [2007] that there is a schema mapping specified by a finite set of s-t tgds that has the unique-solutions property, but does not have a global inverse. Hence, the LAV assumption is needed in the statement of Theorem 6.10.

The schema mapping that is a global inverse in our proof of Theorem 6.10 is (at least on the face of it) not defined in terms of s-t tgds, and might require an infinitary logic to specify. For the rest of this article, we shall focus only on the important and practical case of schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each specified by a finite set of s-t tgds.

7. CHARACTERIZING INVERTIBILITY

In this section, we give a useful characterization of invertibility in terms of the chase (Theorem 7.3). As a corollary (Corollary 7.4), we give an especially simple characterization when the schema mappings are full. We then show (Proposition 7.5) that, in the nonfull case, a weakened version of this condition provides a necessary condition for invertibility. We also show (Proposition 7.6) that this necessary condition is not sufficient. As a tool to help prove these results, we give a proposition (Proposition 7.2) that gives an explicit universal solution for the composition of schema mappings. Some of the results of this section (in particular, Corollary 7.4) are interesting in their own right. But this section is mainly here to provide a useful set of tools for inverses. We shall make use of these tools a number of times.

We begin with a lemma. This lemma gives us a different viewpoint of composition. It will be used to prove the subsequent proposition.[10] Note that in this lemma, when we write $\langle I_1, \langle I_2, I_3\rangle\rangle \models \Sigma_{12} \cup \Sigma_{23}$, we are thinking of Σ_{12} as consisting of s-t tgds and Σ_{23} as consisting of target tgds. Although in this article, we do not allow target tgds in our schema mappings, they are allowed in Fagin et al. [2005a].

LEMMA 7.1. *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{23} = (\mathbf{S_2}, \mathbf{S_3}, \Sigma_{23})$ be schema mappings, with $\mathbf{S_1}$, $\mathbf{S_2}$, and $\mathbf{S_3}$ pairwise disjoint. Let I_1 be an instance of schema $\mathbf{S_1}$ and I_3 an instance of schema $\mathbf{S_3}$. Then $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$ if and only if there is an instance I_2 of schema $\mathbf{S_2}$ such that $\langle I_1, \langle I_2, I_3\rangle\rangle \models \Sigma_{12} \cup \Sigma_{23}$.*

PROOF. Assume first that $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$. This means that there is I_2 such that $\langle I_1, I_2\rangle \models \Sigma_{12}$ and $\langle I_2, I_3\rangle \models \Sigma_{23}$. Since $\mathbf{S_1}$, $\mathbf{S_2}$, and $\mathbf{S_3}$ are pairwise disjoint, it follows easily that $\langle I_1, \langle I_2, I_3\rangle\rangle \models \Sigma_{12} \cup \Sigma_{23}$. Conversely, assume that $\langle I_1, \langle I_2, I_3\rangle\rangle \models \Sigma_{12} \cup \Sigma_{23}$. So $\langle I_1, I_2\rangle \models \Sigma_{12}$ and $\langle I_2, I_3\rangle \models \Sigma_{23}$. Therefore, $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$. □

For the next proposition, we define $chase_{23}$ based on $\mathcal{M}_{23} = (\mathbf{S_2}, \mathbf{S_3}, \Sigma_{23})$ just as we defined $chase_{12}$ based on $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$.

PROPOSITION 7.2. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{23} = (\mathbf{S_2}, \mathbf{S_3}, \Sigma_{23})$ are schema mappings, where $\mathbf{S_1}$, $\mathbf{S_2}$, and $\mathbf{S_3}$ are pairwise disjoint, and where Σ_{12} and Σ_{23} are finite sets of s-t tgds. Let $\mathcal{M}_{13} = (\mathbf{S_1}, \mathbf{S_3}, \Sigma_{12} \circ \Sigma_{23})$. Then $chase_{23}(chase_{12}(I))$ is a universal solution for I under $\mathcal{M}_{13}$.*

PROOF. We prove this by making use of the change in viewpoint given by Lemma 7.1. Let $\mathcal{M} = (\mathbf{S_1}, \langle \mathbf{S_2}, \mathbf{S_3}\rangle, \Sigma_{12} \cup \Sigma_{23})$, where we are now thinking of $\mathbf{S_1}$ as the source schema, $\langle \mathbf{S_2}, \mathbf{S_3}\rangle$ as the target schema, Σ_{12} as s-t tgds, and Σ_{23} as target tgds. Let $I_1 = I$, let $I_2 = chase_{12}(I_1)$, and let $I_3 = chase_{23}(I_2)$. Then $I_3 = chase_{23}(chase_{12}(I_1))$. It is easy to see that $\langle I_1, \langle I_2, I_3\rangle\rangle$ is a result of chasing

[10]Lemma 7.1 and Proposition 7.2 are due to Lucian Popa and Wang-Chiew Tan.

$\langle I_1, \emptyset\rangle$ with $\Sigma_{12} \cup \Sigma_{23}$. It then follows from results in Fagin et al. [2005a] that $\langle I_2, I_3\rangle$ is a universal solution for I_1 under $\mathcal{M}$.

We must show that I_3 is a universal solution for I_1 under $\mathcal{M}_{13}$. First, we show that I_3 is a solution. By Lemma 6.8, it follows that $\langle I_1, I_2\rangle \models \Sigma_{12}$ and $\langle I_2, I_3\rangle \models \Sigma_{23}$. Therefore, $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$. So I_3 is a solution for I_1 under $\mathcal{M}_{13}$.

We now show that I_3 is a universal solution for I_1 under $\mathcal{M}_{13}$. Let I_3' be an arbitrary solution for I_1 under $\mathcal{M}_{13}$. By Lemma 7.1, there is an instance I_2' of schema $\mathbf{S_2}$ such that $\langle I_1, \langle I_2', I_3'\rangle\rangle \models \Sigma_{12} \cup \Sigma_{23}$. Since, as we showed, $\langle I_2, I_3\rangle$ is a universal solution for I_1 under $\mathcal{M}$, there is a homomorphism h from $\langle I_2, I_3\rangle$ to $\langle I_2', I_3'\rangle$. Since also $\mathbf{S_2}$ and $\mathbf{S_3}$ are disjoint, it follows that h gives a homomorphism from I_3 to I_3'. So I_3 is a universal solution for I_1 under $\mathcal{M}_{13}$, as desired. □

Our next theorem gives a useful characterization of when one schema mapping is the inverse of another schema mapping for a given ground instance. We define $chase_{21}$ based on $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ just as we defined $chase_{12}$ based on $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$.

THEOREM 7.3. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$.*

PROOF. Assume first that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. Thus, (6) holds for every ground instance J. It follows immediately that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. Let $J_1 = chase_{12}(I)$ and $J_2 = chase_{21}(J_1)$. Let J_1', J_2' be obtained from J_1, J_2 by replacing each null by a new constant (as before, if a null appears several times, then it is replaced by the same constant each time). By Lemma 6.8, we know that $\langle I, J_1\rangle \models \Sigma_{12}$ and $\langle J_1, J_2\rangle \models \Sigma_{21}$. Therefore, $\langle I, J_1'\rangle \models \Sigma_{12}$ and $\langle J_1', J_2'\rangle \models \Sigma_{21}$. Hence, $\langle I, J_2'\rangle \models \Sigma_{12} \circ \Sigma_{21}$. So by (6), where the role of J is played by J_2', it follows that $\widehat{I} \subseteq J_2'$. Therefore, $\widehat{I} \subseteq J_2$. That is, $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$, as desired.

Conversely, assume that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and that $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$; we must show that (6) holds for each J. Assume first that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$; we must show that $\widehat{I} \subseteq J$. Let $\mathcal{M}_{11} = (\mathbf{S_1}, \widehat{\mathbf{S_1}}, \Sigma_{12} \circ \Sigma_{21})$. By Proposition 7.2, we know that $chase_{21}(chase_{12}(I))$ is a universal solution for I under $\mathcal{M}_{11}$, and so there is a homomorphism from $chase_{21}(chase_{12}(I))$ to J. Since $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$, and since homomorphisms map values in I onto themselves, it follows that $\widehat{I} \subseteq J$, as desired.

Assume now that $\widehat{I} \subseteq J$; we must show that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Since $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq J$, it follows from Corollary 4.3 that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. □

As a corollary, we obtain a particularly simple characterization when Σ_{12} and Σ_{21} consist of *full* tgds.

COROLLARY 7.4. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of full s-t tgds. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\widehat{I} = chase_{21}(chase_{12}(I))$.*

PROOF. Assume first that $\widehat{I} = chase_{21}(chase_{12}(I))$. By Lemma 6.8, we know that $\langle I, chase_{12}(I)\rangle \models \Sigma_{12}$ and $\langle chase_{12}(I), chase_{21}(chase_{12}(I))\rangle \models \Sigma_{21}$. But because $chase_{21}(chase_{12}(I)) = \widehat{I}$, we know that $\langle chase_{12}(I), \widehat{I}\rangle \models \Sigma_{21}$. From the facts that $\langle I, chase_{12}(I)\rangle \models \Sigma_{12}$ and $\langle chase_{12}(I), \widehat{I}\rangle \models \Sigma_{21}$, it follows that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Since also $\widehat{I} = chase_{21}(chase_{12}(I))$, it follows from Theorem 7.3 that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I.

Conversely, assume that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. By Theorem 7.3, $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. Since $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$, we know that there is J such that $\langle I, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I}\rangle \models \Sigma_{21}$. Since $\langle I, J\rangle \models \Sigma_{12}$, it follows from Lemma 6.8 that $chase_{12}(I) \subseteq J$. Since $chase_{12}(I) \subseteq J$ and $\langle J, \widehat{I}\rangle \models \Sigma_{21}$, we know from Lemma 4.2 that $\langle chase_{12}(I), \widehat{I}\rangle \models \Sigma_{21}$. It follows from Lemma 6.8 that $chase_{21}(chase_{12}(I)) \subseteq \widehat{I}$. Since also $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$, we have $\widehat{I} = chase_{21}(chase_{12}(I))$, as desired. □

Proposition 7.2, Theorem 7.3, and Corollary 7.4 are all useful tools that we shall make use of in later proofs. The next result[11] (Proposition 7.5) is not used later, but is interesting in that it gives a necessary condition for invertibility, in the spirit of Corollary 7.4, but that holds even when the tgds are not full. We will then conclude the section by proving (Proposition 7.6) that this necessary condition is not sufficient.

Two instances I_1 and I_2 are *homomorphically equivalent* if there is a homomorphism from I_1 into I_2 and a homomorphism from I_2 into I_1.

PROPOSITION 7.5. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. If $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I, then $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent.*

PROOF. Assume that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. Let σ be the composition formula $\Sigma_{12} \circ \Sigma_{21}$, and let $\mathcal{M}_{11} = (\mathbf{S_1}, \widehat{\mathbf{S_1}}, \sigma)$. By Theorem 7.3, $\langle I, \widehat{I}\rangle \models \sigma$ and $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. Since $\langle I, \widehat{I}\rangle \models \sigma$ and since $chase_{21}(chase_{12}(I))$ is a universal solution for I under $\mathcal{M}_{11}$ (by Proposition 7.2), it follows that there is a homomorphism from $chase_{21}(chase_{12}(I))$ to $\widehat{I}$. Since $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$, we know that the identity function is a homomorphism from $\widehat{I}$ to $chase_{21}(chase_{12}(I))$. Hence, $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent, as desired. □

The next theorem implies that the necessary condition for invertibility in Proposition 7.5 is not sufficient. Thus, the next theorem implies that, even if $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent, then $\mathcal{M}_{21}$ is not necessarily an inverse of $\mathcal{M}_{12}$ for I. In fact, this theorem says even more: it says that, even if $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent for *every* I, there can be an I such that $\mathcal{M}_{21}$ is not inverse of $\mathcal{M}_{12}$ for I.

PROPOSITION 7.6. *There are schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{12} and Σ_{21} are finite sets of s-t tgds, such that $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent for every instance I of $\mathbf{S_1}$, but $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for some instance I of $\mathbf{S_1}$.*

[11]This result is due to Lucian Popa.

PROOF. Let $\mathbf{S_1}$ consist of the binary relation symbol R, let $\mathbf{S_2}$ consist of the binary relation symbol S, let Σ_{12} consist of the s-t tgd $\mathtt{R}(a, b) \rightarrow \exists x(\mathtt{S}(a, x) \wedge \mathtt{S}(x, b))$, and let Σ_{21} consist of the s-t tgd $\mathtt{S}(a, x) \wedge \mathtt{S}(x, b) \rightarrow \widehat{\mathtt{R}}(a, b)$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$.

We first show that $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent for every instance I of $\mathbf{S_1}$. Let I be an arbitrary instance of $\mathbf{S_1}$. It is clear that we have $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. Therefore, the identity function is a homomorphism from $\widehat{I}$ to $chase_{21}(chase_{12}(I))$. We now show that there is a homomorphism from $chase_{21}(chase_{12}(I))$ to $\widehat{I}$. Each null x in $chase_{21}(chase_{12}(I))$ is a null in $chase_{12}(I)$, which is obtained by applying Σ_{12} to a fact $\mathtt{R}(a, b)$ to obtain the facts $\mathtt{S}(a, x)$ and $\mathtt{S}(x, b)$. Let h be the function that is the identity on I, and which maps such a null x to a.

What are the facts that appear in $chase_{21}(chase_{12}(I))$? First, there are the facts $\widehat{\mathtt{R}}(a, b)$ such that $\mathtt{R}(a, b)$ is a fact of I. For such facts, $\widehat{\mathtt{R}}(h(a), h(b))$ is simply $\widehat{\mathtt{R}}(a, b)$, which is consistent with h being a homomorphism. The other facts of $chase_{21}(chase_{12}(I))$ are facts $\widehat{\mathtt{R}}(x, y)$ such that $\mathtt{S}(x, a)$ and $\mathtt{S}(a, y)$ appear in $chase_{12}(I)$, where x and y are nulls, and where a is a constant (a value in I). We must show that $\widehat{\mathtt{R}}(h(x), h(y))$ is a fact of $\widehat{I}$.

Since $\mathtt{S}(x, a)$ appears in $chase_{12}(I)$, there is a constant c such that $\mathtt{R}(c, a)$ is a fact of I, and $\mathtt{S}(c, x)$ is a fact of $chase_{12}(I)$. Then $h(x) = c$, and $h(y) = a$. Since $\mathtt{R}(c, a)$ is a fact of I, we know that $\widehat{\mathtt{R}}(c, a)$ is a fact of $\widehat{I}$, that is, $\widehat{\mathtt{R}}(h(x), h(y))$ is a fact of $\widehat{I}$, as desired. This completes the proof that $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ are homomorphically equivalent.

We now show that there is an instance I of $\mathbf{S_1}$ such that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I. Let I consist of the facts $\mathtt{R}(0, 1)$ and $\mathtt{R}(1, 0)$. We need only show that $\langle I, \widehat{I} \rangle \not\models \Sigma_{12} \circ \Sigma_{21}$. Assume that $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$; we shall derive a contradiction. Then there is J such that $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \widehat{I} \rangle \models \Sigma_{21}$. Since $\langle I, J \rangle \models \Sigma_{12}$, there is some X (either a null or a constant) such that $\mathtt{S}(0, X)$ and $\mathtt{S}(X, 1)$ are facts of J. Similarly, there is some Y (either a null or a constant) such that $\mathtt{S}(1, Y)$ and $\mathtt{S}(Y, 0)$ are facts of J. Since $\mathtt{S}(X, 1)$ and $\mathtt{S}(1, Y)$ are facts of J, and since $\langle J, \widehat{I} \rangle \models \Sigma_{21}$, it follows that $\widehat{\mathtt{R}}(X, Y)$ is a fact of $\widehat{I}$, that is, $\mathtt{R}(X, Y)$ is a fact of I. But I contains only the facts $\mathtt{R}(0, 1)$ and $\mathtt{R}(1, 0)$, so either $X = 0$ and $Y = 1$, or $X = 1$ and $Y = 0$. Assume that $X = 0$ and $Y = 1$ (a symmetric proof works if $X = 1$ and $Y = 0$). Now $\mathtt{S}(0, X)$ is a fact of J, that is, $\mathtt{S}(0, 0)$ is a fact of J. Since $\langle J, \widehat{I} \rangle \models \Sigma_{21}$, it follows that $\widehat{\mathtt{R}}(0, 0)$ is a fact of $\widehat{I}$ (this is because we apply the tgd $(\mathtt{S}(a, x) \wedge \mathtt{S}(x, b)) \rightarrow \widehat{\mathtt{R}}(a, b)$ to J where the roles of a, x, and b are all played by 0). But this is a contradiction, since $\widehat{\mathtt{R}}(0, 0)$ is not a fact of $\widehat{I}$. □

8. THE CANONICAL CANDIDATE TGD LOCAL INVERSE

Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds, and let I be a ground instance. In this section, we give a schema mapping (the *canonical candidate tgd local inverse*) that is guaranteed to be an inverse of $\mathcal{M}$ for I if there is any inverse at all that is specified by a finite set of s-t tgds. There are two reasons why this result is useful for us. First, given a schema mapping $\mathcal{M}$ and an instance I, we can answer the question about whether $\mathcal{M}$ has an inverse for I that is specified by a finite set of s-t tgds by simply checking whether the canonical candidate tgd local inverse is an inverse of $\mathcal{M}$ for I. Second, the canonical candidate tgd local inverse will help us in our real interest, which is in finding a global inverse. In particular, in the next section, we shall develop a schema mapping that is guaranteed to be a global inverse of $\mathcal{M}$ if there is any global inverse at all that is specified by a finite set of s-t tgds. A key tool in the proof of correctness of this global inverse is the canonical candidate tgd local inverse.

We begin with a definition. Assume that I and J are instances (of different schemas) where every member of the active domain of I is in the active domain of J. Define $\beta_{J,I}$ to be the full tgd where the premise is the conjunction of the facts of J, and the conclusion is the conjunction of the facts of I (we are treating the values in J as universally quantified variables in the tgd).

Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ is a schema mapping where Σ_{12} is a finite set of s-t tgds. Assume that there is a schema mapping $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ that is an inverse of $\mathcal{M}_{12}$ for I, where Σ_{21} is a finite set of s-t tgds. Let $J^* = \mathit{chase}_{12}(I)$. It follows from Corollary 6.5 that every member of the active domain of I (and hence of the active domain of $\widehat{I}$) is in the active domain of J^*, and so $\beta_{J^*,\widehat{I}}$ is a full tgd. Define the *canonical candidate tgd local inverse of* $\mathcal{M}_{12}$ *for* I to be $(\mathbf{S_2}, \widehat{\mathbf{S_1}}, \{\beta_{J^*,\widehat{I}}\})$. For example, assume that I consists of the facts $\mathtt{P}(c_1, c_2)$ and $\mathtt{P}(c_2, c_3)$, where c_1, c_2, c_3 are constants. Assume that Σ consists of the s-t tgd

$$\mathtt{P}(x_1, x_2) \wedge \mathtt{P}(x_2, x_3) \rightarrow \exists y (\mathtt{Q}(x_1, x_2, x_3) \wedge \mathtt{R}(x_3, y)).$$

Then J^* consists of the facts $\mathtt{Q}(c_1, c_2, c_3)$ and $\mathtt{R}(c_3, y)$, where y is being treated as a null , and $\beta_{J^*,\widehat{I}}$ is the full tgd

$$\mathtt{Q}(c_1, c_2, c_3) \wedge \mathtt{R}(c_3, y) \rightarrow (\widehat{\mathtt{P}}(c_1, c_2) \wedge \widehat{\mathtt{P}}(c_2, c_3)), \quad (12)$$

where c_1, c_2, c_3, y are treated in (12) as universally quantified variables.

We call $\mathcal{M}_{21}$ the *most general tgd inverse of* $\mathcal{M}_{12}$ *for* I if Σ'_{21} logically implies Σ_{21} for every inverse $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ of $\mathcal{M}_{12}$ for I where Σ'_{21} is a finite set of s-t tgds.

Before we show that the canonical candidate tgd local inverse has its desired properties, we must prove a simple lemma.

LEMMA 8.1. *Assume that* Σ'_{23} *logically implies* Σ_{23}. *Then* $\Sigma_{12} \circ \Sigma'_{23}$ *logically implies* $\Sigma_{12} \circ \Sigma_{23}$.

PROOF. Assume that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma'_{23}$. We must show that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{23}$. Since $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma'_{23}$, there is J' such that $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma'_{23}$. Since Σ'_{23} logically implies Σ_{23}. and $\langle J', J\rangle \models \Sigma'_{23}$, it follows that $\langle J', J\rangle \models \Sigma_{23}$. Since $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma_{23}$, it follows that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{23}$, as desired. □

THEOREM 8.2. *Let* $\mathcal{M}$ *be a schema mapping specified by a finite set of s-t tgds, and let* I *be a ground instance. Assume that* $\mathcal{M}$ *has an inverse for* I *that is specified by a finite set of s-t tgds. Then the canonical candidate tgd local inverse of* $\mathcal{M}$ *for* I *is an inverse of* $\mathcal{M}$ *for* I, *and in fact the most general tgd inverse of* $\mathcal{M}$ *for* I.

PROOF. Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, and that $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ is an inverse of $\mathcal{M}_{12}$ for I, where Σ_{12} and Σ'_{21} are finite sets of s-t tgds. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be the canonical candidate tgd local inverse of $\mathcal{M}_{12}$ for I. We first prove that Σ'_{21} logically implies Σ_{21}. Let $J^* = \mathit{chase}_{12}(I)$. By Lemma 6.8, we know that $\langle I, J^* \rangle \models \Sigma_{12}$. Let J' be the result of chasing J^* with Σ'_{21}, and let J'' be the result of replacing each null in J' by a new, distinct constant. So J'' is a ground instance isomorphic to J', where the isomorphism maps constants into themselves. By Lemma 6.8, we know that $\langle J^*, J' \rangle \models \Sigma'_{21}$, and so $\langle J^*, J'' \rangle \models \Sigma'_{21}$. Since $\langle I, J^* \rangle \models \Sigma_{12}$ and $\langle J^*, J'' \rangle \models \Sigma'_{21}$, it follows that $\langle I, J'' \rangle \models \Sigma_{12} \circ \Sigma'_{21}$. Since $\mathcal{M}'_{21}$ is an inverse of $\mathcal{M}_{12}$, it follows that $\widehat{I} \subseteq J''$. Therefore, $\widehat{I} \subseteq J'$, since the constants introduced into J'' when constructing J'' from J' were new. Thus, the result of chasing J^* with Σ'_{21} necessarily contains $\widehat{I}$. It follows from standard results in dependency theory that Σ'_{21} logically implies $\beta_{J^*,\widehat{I}}$ and hence Σ_{21}, as desired.

We now show that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. We know that $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma'_{21}$, since $\mathcal{M}'_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. Since Σ'_{21} logically implies Σ_{21}, it follows from Lemma 8.1 that $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$. When we chase J^* with $\beta_{J^*,\widehat{I}}$ we clearly obtain at least $\widehat{I}$. That is, $\widehat{I} \subseteq \mathit{chase}_{21}(\mathit{chase}_{12}(I))$. So by Theorem 7.3, it follows that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I, as desired. □

In the next section, we make use of the canonical candidate tgd local inverse (and the fact that it is most general).

9. THE CANONICAL CANDIDATE TGD GLOBAL INVERSE

Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ be a schema mapping, where Σ_{12} is a finite set of s-t tgds. In this section, we shall define the *canonical candidate tgd global inverse of* $\mathcal{M}_{12}$, which is a schema mapping specified by a finite set of s-t tgds, and show that it is a global inverse of $\mathcal{M}_{12}$ if there is any such global inverse. In fact, we shall do something a little more general. We shall consider certain classes $\mathcal{S}$ of ground instances, and show how to define a *canonical candidate tgd $\mathcal{S}$-inverse of* $\mathcal{M}_{12}$, which is a schema mapping specified by a finite set of s-t tgds that is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$ if there is any such $\mathcal{S}$-inverse. When $\mathcal{S}$ is the class of all ground instances, we obtain the canonical candidate tgd global inverse.

Let us say that a set Γ of tgds and egds (all on the source schema) is *finitely chasable* if for every (finite) ground instance I, some result of chasing I with Γ is a (finite) instance, or else some chase of I with Γ fails (by trying to equate two distinct values in I). It follows from results in Fagin et al. [2005a] that, when Γ is the union of a weakly acyclic set of tgds (as defined in Fagin et al. [2005a]) with a set of egds, then Γ is finitely chasable. We now give a simple example where the converse fails.

Example 9.1. Let Γ' consist of the single tgd $\mathtt{R}(x, y) \rightarrow \exists z \mathtt{R}(y, z)$. It follows easily from the definition of weak acyclicity that Γ' is not weakly acyclic, and in fact not finitely chasable. Let Γ'' consist of the single egd $\mathtt{R}(x, y) \rightarrow (x = y)$. Now let Γ be $\Gamma' \cup \Gamma''$. Then Γ is finitely chasable (since in this case, we need only chase with Γ'' alone).

Let Γ be a finitely chasable set of tgds and egds, and let $\mathcal{S}$ be the class of all ground instances that satisfy Γ. Assume that $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, and Σ_{21} is a finite set of s-t tgds. For each relational symbol R of $\mathbf{S_1}$, let $I_{\mathtt{R}}$ be a one-tuple instance that contains only the fact $\mathtt{R}(\mathbf{x})$, where the variables in $\mathbf{x}$ are distinct. Note that $I_{\mathtt{R}}$ is not an instance in the usual sense, because the active domain consists of variables, not constants or nulls. Instead, it is a type of canonical instance. Let $I_{\mathtt{R}}^{\Gamma}$ be a finite instance that is a result of chasing $I_{\mathtt{R}}$ with Γ, where it is all right to allow distinct variables in $\mathbf{x}$ to be equated by the chase. In our case of greatest interest, where Γ is the empty set, we have $I_{\mathtt{R}}^{\Gamma} = I_{\mathtt{R}}$. Let $J_{\mathtt{R}}^{\Gamma}$ be $chase_{12}(I_{\mathtt{R}}^{\Gamma})$, a result of chasing $I_{\mathtt{R}}^{\Gamma}$ with Σ_{12}.[12]

Since $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, in particular $\mathcal{M}_{21}$ is a local inverse of $\mathcal{M}_{12}$ for $I_{\mathtt{R}}^{\Gamma}$ (this is because $I_{\mathtt{R}}^{\Gamma}$ is a member of $\mathcal{S}$). It follows from Corollary 6.5 that every member of the active domain of $I_{\mathtt{R}}^{\Gamma}$ (and hence of the active domain of $\widehat{I_{\mathtt{R}}^{\Gamma}}$) is in the active domain of $J_{\mathtt{R}}^{\Gamma}$. Therefore, $\beta_{J,\widehat{I}}$ is a full tgd, where I is $I_{\mathtt{R}}^{\Gamma}$, and J is $J_{\mathtt{R}}^{\Gamma}$, with $\beta_{\cdot,\cdot}$ as defined in Section 8. Let us denote this full tgd by $\delta_{\mathtt{R}}^{\Gamma}$, and let $\Sigma_{21}^{\mathcal{S}}$ consist of all of the tgds $\delta_{\mathtt{R}}^{\Gamma}$, one for every relation symbol R of $\mathbf{S_1}$. Define the *canonical candidate tgd $\mathcal{S}$-inverse of* $\mathcal{M}_{12}$ to be $\mathcal{M}_{21}^{\mathcal{S}} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^{\mathcal{S}})$. In the case where $\mathcal{S}$ is the class of all ground instances, we may write Σ_{21}^{c} for $\Sigma_{21}^{\mathcal{S}}$, and $\mathcal{M}_{21}^{c}$ for $\mathcal{M}_{21}^{\mathcal{S}}$, where c stands for "canonical." We call $\mathcal{M}_{21}$ the *most general tgd $\mathcal{S}$-inverse of* $\mathcal{M}_{12}$ if Σ'_{21} logically implies Σ_{21} for every $\mathcal{S}$-inverse $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ of $\mathcal{M}_{12}$ where Σ'_{21} is a finite set of s-t tgds.

Example 9.2. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$. Assume that $\mathbf{S_1}$ consists of the binary relation symbol R and the unary relation symbol S, and that $\mathbf{S_2}$ consists of the binary relation symbols T and U. Let Σ_{12} consist of the s-t tgds $\mathtt{R}(x_1, x_2) \rightarrow \exists y(\mathtt{T}(x_1, y) \wedge \mathtt{U}(y, x_2))$, $\mathtt{R}(x, x) \rightarrow \mathtt{U}(x, x)$, and $\mathtt{S}(x) \rightarrow \exists y \mathtt{U}(x, y)$. Let Γ consist of the egd $\mathtt{R}(x_1, x_2) \rightarrow (x_1 = x_2)$.

Now $I_{\mathtt{R}}$ consists of the fact $\mathtt{R}(x_1, x_2)$, and so $I_{\mathtt{R}}^{\Gamma}$ consists of the fact $\mathtt{R}(x_1, x_1)$. Then $J_{\mathtt{R}}^{\Gamma}$ consists of the facts $\mathtt{T}(x_1, y)$, $\mathtt{U}(y, x_1)$, and $\mathtt{U}(x_1, x_1)$. So $\delta_{\mathtt{R}}^{\Gamma}$ is the tgd

$$\mathtt{T}(x_1, y) \wedge \mathtt{U}(y, x_1) \wedge \mathtt{U}(x_1, x_1) \rightarrow \widehat{\mathtt{R}}(x_1, x_1).$$

Also, $I_{\mathtt{S}}$ and $I_{\mathtt{S}}^{\Gamma}$ each consist of the fact $\mathtt{S}(x_1)$, and $J_{\mathtt{S}}^{\Gamma}$ consists of the fact $\mathtt{U}(x_1, y)$. So $\delta_{\mathtt{S}}^{\Gamma}$ is the tgd $\mathtt{U}(x_1, y) \rightarrow \widehat{\mathtt{S}}(x_1)$. Finally, $\mathcal{M}_{21}^{\mathcal{S}} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^{\mathcal{S}})$, where $\Sigma_{21}^{\mathcal{S}}$ consists of the tgds $\delta_{\mathtt{R}}^{\Gamma}$ and $\delta_{\mathtt{S}}^{\Gamma}$. □

THEOREM 9.3. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. Let Γ be a finitely chasable set of tgds and egds, and let $\mathcal{S}$ be the class of ground instances that satisfy Γ. Assume that $\mathcal{M}$ has an $\mathcal{S}$-inverse that is specified by a finite set of s-t tgds. Then the canonical candidate tgd $\mathcal{S}$-inverse of $\mathcal{M}$ is an $\mathcal{S}$-inverse of $\mathcal{M}$, and in fact the most general tgd $\mathcal{S}$-inverse of $\mathcal{M}$.*

PROOF. Let $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ be an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, where Σ'_{21} is s finite set of s-t tgds. Let $\mathcal{M}_{21}^{\mathcal{S}} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^{\mathcal{S}})$ be the canonical candidate tgd $\mathcal{S}$-inverse of $\mathcal{M}_{12}$.

[12]Even though $J_{\mathtt{R}}^{\Gamma}$ depends not just on R and Γ, but also on Σ_{12}, for simplicity we do not reflect the dependency on Σ_{12} in the notation $J_{\mathtt{R}}^{\Gamma}$.

Let $\mathtt{R}$ be a relational symbol of $\mathbf{S_1}$. Since $\mathcal{M}'_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, and since $I^\Gamma_{\mathtt{R}}$ is in $\mathcal{S}$, certainly $\mathcal{M}'_{21}$ is an inverse of $\mathcal{M}_{12}$ for $I^\Gamma_{\mathtt{R}}$. Now $(\mathbf{S_2}, \widehat{\mathbf{S_1}}, \{\delta^\Gamma_{\mathtt{R}}\})$ is the canonical candidate tgd local inverse of $\mathcal{M}_{12}$ for $I^\Gamma_{\mathtt{R}}$, so, by Theorem 8.2, we know that Σ'_{21} logically implies $\delta^\Gamma_{\mathtt{R}}$. Since $\mathtt{R}$ is an arbitrary relational symbol of $\mathbf{S_1}$, it follows that Σ'_{21} logically implies $\Sigma^{\mathcal{S}}_{21}$, which was one thing to be shown.

Since $\mathcal{M}'_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, we know that, for every I that satisfies Γ, we have

$$\langle I, J\rangle \models \Sigma_{12} \circ \Sigma'_{21} \text{ if and only if } \widehat{I} \subseteq J. \tag{13}$$

We wish to show that $\mathcal{M}^{\mathcal{S}}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$. Thus, we must show that, for every I that satisfies Γ, we have

$$\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^{\mathcal{S}}_{21} \text{ if and only if } \widehat{I} \subseteq J. \tag{14}$$

Assume first that I satisfies Γ, and $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^{\mathcal{S}}_{21}$. Then there is J' such that $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma^{\mathcal{S}}_{21}$. Let $\mathtt{R}(\mathbf{c})$ be a fact of I. Then each of the equalities among members of $\mathbf{c}$ that are "forced" by chasing the database whose only fact is $\mathtt{R}(\mathbf{c})$ with the set Γ necessarily already holds in $\mathbf{c}$, since I satisfies Γ. Let J^* be an instance that is obtained by replacing $\mathbf{x}$ by $\mathbf{c}$ in $J^\Gamma_{\mathtt{R}}$. Then a homomorphic image of J^* appears in J' (under a homomorphism that maps each member of $\mathbf{c}$ onto itself). Therefore, there is a homomorphism from the premise of $\delta^\Gamma_{\mathtt{R}}$ into J' that maps $\mathbf{x}$ onto $\mathbf{c}$. Hence, since $\langle J', J\rangle \models \delta^\Gamma_{\mathtt{R}}$ (because $\delta^\Gamma_{\mathtt{R}}$ is in $\Sigma^{\mathcal{S}}_{21}$), it follows that $\widehat{\mathtt{R}}(\mathbf{c})$ is in J. Since $\mathtt{R}(\mathbf{c})$ is an arbitrary fact of I, this implies that $\widehat{I} \subseteq J$, as desired.

Assume now that $\widehat{I} \subseteq J$. By (13), we know that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma'_{21}$. Since $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma'_{21}$ and Σ'_{21} logically implies $\Sigma^{\mathcal{S}}_{21}$, it follows from Lemma 8.1 that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^{\mathcal{S}}_{21}$. This was to be shown. □

COROLLARY 9.4. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. Assume that $\mathcal{M}$ has a global inverse that is specified by a finite set of s-t tgds. Then the canonical candidate tgd global inverse of $\mathcal{M}$ is a global inverse of $\mathcal{M}$, and in fact the most general tgd global inverse of $\mathcal{M}$.*

PROOF. This follows from Theorem 9.3 by letting Γ be the empty set. □

10. FULL TGDS SUFFICE

The canonical candidate tgd local inverse and the canonical candidate tgd global inverse (and more generally the canonical candidate tgd $\mathcal{S}$-inverse) are each specified by a finite set of *full* tgds. In this section, we show that this is no accident: if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are schema mappings that are each specified by a finite set of s-t tgds, $\mathcal{S}$ is an arbitrary class of ground instances (not necessarily defined by a set Γ, or even closed under isomorphism), and $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, then there is a schema mapping $\mathcal{M}^f_{21}$ specified by a finite set of *full* s-t tgds and that is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$. While the canonical candidate tgd local inverse is tailored to a particular instance I, the mapping M^f_{21} is, as we shall see, constructed only from $\mathcal{M}_{21}$. From a technical point of view, this contrasts also with the canonical candidate tgd global inverse, which is constructed only from $\mathcal{M}_{12}$.

Let $\mathcal{M}_{12}$ be a schema mapping that is specified by a finite set of s-t tgds. It is important to note that (in the case of a global inverse) we are *not* claiming that, if $\mathcal{M}_{12}$ has a global inverse, then $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of full s-t tgds. Instead, in this section we show that, if $\mathcal{M}_{12}$ has a global inverse *that is specified by a finite set of s-t tgds*, then $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of *full* s-t tgds. The question of "the language of inverses," which tells how rich a language needs to be to specify a global inverse of $\mathcal{M}_{12}$, was resolved in Fagin et al. [2007]. The results are summarized in Section 15 of this article.

We begin with some definitions. Let γ be an s-t tgd. Assume that γ is $\forall\mathbf{x}(\varphi_{\mathbf{S}}(\mathbf{x}) \rightarrow \exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y}))$, where $\varphi_{\mathbf{S}}(\mathbf{x})$ is a conjunction of atoms over $\mathbf{S}$ and $\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y})$ is a conjunction of atoms over $\mathbf{T}$. Assume also that all of the variables in $\mathbf{x}$ appear in $\varphi_{\mathbf{S}}(\mathbf{x})$ (but not necessarily in $\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y})$). Let $\psi^f_{\mathbf{T}}(\mathbf{x})$ be the conjunction of all atoms in $\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y})$ that do not contain any variables in $\mathbf{y}$ (the f stands for "full"). Define γ^f (the *full part of* γ) to be the full tgd $\forall\mathbf{x}(\varphi_{\mathbf{S}}(\mathbf{x}) \rightarrow \psi^f_{\mathbf{T}}(\mathbf{x}))$. As an example (where we do not bother to write the universal quantifiers), if γ is $\mathtt{P}(x,y) \rightarrow \exists z(\mathtt{Q}(x,x) \wedge \mathtt{Q}(y,z))$, then γ^f is $\mathtt{P}(x,y) \rightarrow \mathtt{Q}(x,x)$. If $\psi^f_{\mathbf{T}}(\mathbf{x})$ is an empty conjunction, then γ^f is a *dummy tgd* where the conclusion is "Truth" (and so the dummy tgd itself is "Truth").

Let $\psi^n_{\mathbf{T}}(\mathbf{x})$ be the conjunction of all atoms in $\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y})$ that contain some variable in $\mathbf{y}$ (the n stands for "nonfull"). Define γ^n (the *nonfull part of* γ) to be the tgd $\forall\mathbf{x}(\varphi_{\mathbf{S}}(\mathbf{x}) \rightarrow \exists\mathbf{y}\psi^n_{\mathbf{T}}(\mathbf{x},\mathbf{y}))$. If we again take γ to be $\mathtt{P}(x,y) \rightarrow \exists z(\mathtt{Q}(x,x) \wedge \mathtt{Q}(y,z))$, then γ^n is $\mathtt{P}(x,y) \rightarrow \exists z\mathtt{Q}(y,z)$. As before, if $\psi^n_{\mathbf{T}}(\mathbf{x})$ is an empty conjunction, then γ^n is a *dummy tgd* where the conclusion is "Truth"(and so the dummy tgd itself is "Truth"). If Σ is a set of tgds, let Σ^f be the set of γ^f where $\gamma \in \Sigma$ and where γ^f is not a dummy tgd. Similarly, let Σ^n be the set of γ^n where $\gamma \in \Sigma$ and where γ^n is not a dummy tgd. It is easy to see that Σ is logically equivalent to $\Sigma^f \cup \Sigma^n$. The next theorem tells us that only full tgds play a role in the inverse.

THEOREM 10.1. *Assume that* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *and* $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ *are schema mappings where* Σ_{12} *and* Σ_{21} *are finite sets of s-t tgds. Let* $\mathcal{M}^f_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma^f_{21})$. *Let* $\mathcal{S}$ *be an arbitrary class of ground instances (not necessarily closed under isomorphism). If* $\mathcal{M}_{21}$ *is an* $\mathcal{S}$*-inverse of* $\mathcal{M}_{12}$*, then so is* $\mathcal{M}^f_{21}$.

PROOF. Assume that $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$, and that I is in $\mathcal{S}$. Since $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I, we know that (6) holds. We must show that

$$\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^f_{21} \text{ if and only if } \widehat{I} \subseteq J. \tag{15}$$

Assume that $\widehat{I} \subseteq J$. From (6), we know that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Therefore, there is J' such that $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma_{21}$. Since $\langle J', J\rangle \models \Sigma_{21}$, and also Σ_{21} logically implies Σ^f_{21}, it follows that $\langle J', J\rangle \models \Sigma^f_{21}$. Since $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma^f_{21}$, it follows that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^f_{21}$. This proves the "if" direction of (15).

Conversely, assume that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^f_{21}$; we must prove that $\widehat{I} \subseteq J$. Since $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma^f_{21}$, there is J' such that $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma^f_{21}$. Let J'' be the result of chasing J' with Σ^n_{21}, and let J''' be the result of replacing each null in J'' by a new constant. Then $\langle J', J''\rangle \models \Sigma_{21}^n$ by Lemma 6.8, and so $\langle J', J'''\rangle \models \Sigma_{21}^n$. Since $\langle J', J\rangle \models \Sigma_{21}^f$ and $\langle J', J'''\rangle \models \Sigma_{21}^n$, it follows that $\langle J', J \cup J'''\rangle \models \Sigma_{21}^f \cup \Sigma_{21}^n$. Since Σ_{21} is logically equivalent to $\Sigma_{21}^f \cup \Sigma_{21}^n$, it follows that $\langle J', J \cup J'''\rangle \models \Sigma_{21}$. Since also $\langle I, J'\rangle \models \Sigma_{12}$, it follows that $\langle I, J \cup J'''\rangle \models \Sigma_{12} \circ \Sigma_{21}$. So by (6), where the role of J is played by $J \cup J'''$, we know that $\widehat{I} \subseteq J \cup J'''$.

Every tuple introduced by doing the chase of J' with Σ_{21}^n necessarily contains nulls, by construction of Σ_{21}^n. That is, every tuple in J'' contains nulls. So every tuple in J''' contains some new constant. Since also $\widehat{I} \subseteq J \cup J'''$, it follows that $\widehat{I} \subseteq J$, which was to be shown. □

The following corollary is immediate (by letting Σ'_{21} in the corollary be Σ_{21}^f).

Corollary 10.2. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. Let $\mathcal{S}$ be an arbitrary class of ground instances (not necessarily closed under isomorphism). Assume that $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$. Then there is a finite set Σ'_{21} of full tgds such that $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$.*

11. REVERSING THE ARROWS (NOT!)

It is folk wisdom that simply "reversing the arrows" gives an inverse. In this section, we show even a weak form of this folk wisdom is not true.

What does "reversing the arrows" mean in our context? Let us call a full tgd *reversible* if the same variables appear in the premise as the conclusion. If γ is a reversible tgd $\varphi \to \psi$, define $rev(\gamma)$ to be the full tgd $\psi \to \widehat{\varphi}$, where $\widehat{\varphi}$ is the result of replacing every relational symbol $\mathtt{R}$ by $\widehat{\mathtt{R}}$. Since γ is reversible, $rev(\gamma)$ is indeed a full tgd. We think of $rev(\gamma)$ as the result of "reversing the arrow" of γ.

Example 11.1. We now give a simple example that shows that the schema mapping $(\mathbf{S_2}, \widehat{\mathbf{S_1}}, \{rev(\gamma) : \gamma \in \Sigma_{12}\})$ is not necessarily a global inverse of $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, even when Σ_{12} consists of a finite set of reversible tgds and $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of s-t tgds. Let $\mathbf{S_1}$ consist of the unary relation symbols $\mathtt{R}_1$ and $\mathtt{R}_2$. Let $\mathbf{S_2}$ consist of the unary relation symbols $\mathtt{S}_1$, $\mathtt{S}_2$, and $\mathtt{S}_3$. Let $\Sigma_{12} = \{\mathtt{R}_1(x) \to \mathtt{S}_1(x), \mathtt{R}_2(x) \to \mathtt{S}_2(x), \mathtt{R}_1(x) \to \mathtt{S}_3(x), \mathtt{R}_2(x) \to \mathtt{S}_3(x)\}$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$. Let $\Sigma_{21} = \{\mathtt{S}_1(x) \to \widehat{\mathtt{R}_1}(x), \mathtt{S}_2(x) \to \widehat{\mathtt{R}_2}(x)\}$. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. It is easy to see that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$.

Now let $\Sigma'_{21} = \{rev(\gamma) : \gamma \in \Sigma_{12}\}$. Thus $\Sigma'_{21} = \{\mathtt{S}_1(x) \to \widehat{\mathtt{R}_1}(x), \mathtt{S}_2(x) \to \widehat{\mathtt{R}_2}(x), \mathtt{S}_3(x) \to \widehat{\mathtt{R}_1}(x), \mathtt{S}_3(x) \to \widehat{\mathtt{R}_2}(x)\}$. Let $\mathcal{M}'_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma'_{21})$. It is easy to verify that $\mathcal{M}'_{21}$ is not a global inverse of $\mathcal{M}_{12}$. So simply "reversing the arrows" does not necessarily give a global inverse, even when there is a global inverse.

Note that although $\{rev(\gamma) : \gamma \in \Sigma_{12}\}$ in Example 11.1 does not specify a global inverse, some subset of it (namely, Σ_{21}) does. The next theorem says that there is an example where there is no subset of $\{rev(\gamma) : \gamma \in \Sigma_{12}\}$ that specifies a global inverse. The example is "normalized" by making each (full) tgd that appears have a singleton conclusion.

THEOREM 11.2. *There is a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ where each member of Σ_{12} is a reversible tgd with a singleton conclusion that has a global inverse specified by a finite set of s-t tgds, but where there is no subset X of Σ_{12} such that $(\mathbf{S_2}, \widehat{\mathbf{S_1}}, \{rev(\gamma) : \gamma \in X\})$ is a global inverse of $\mathcal{M}_{12}$.*

PROOF. Let $\mathbf{S_1}$ consist of the binary relation symbol R and the unary relation symbol V. Let $\mathbf{S_2}$ consist of the binary relation symbols S and T, and the unary relation symbol U. Let Σ_{12} consist of the tgds $\mathtt{R}(x, y) \rightarrow \mathtt{S}(x, y)$, $\mathtt{R}(x, y) \rightarrow \mathtt{T}(x, y)$, $\mathtt{V}(x) \rightarrow \mathtt{U}(x)$, $\mathtt{R}(x, x) \wedge \mathtt{R}(y, y) \rightarrow \mathtt{S}(x, y)$, and $\mathtt{V}(x) \rightarrow \mathtt{T}(x, x)$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$.

Let Σ_{21} consist of the tgds $\mathtt{S}(x, y) \wedge \mathtt{T}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y)$ and $\mathtt{U}(x) \rightarrow \widehat{\mathtt{V}}(x)$. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$.[13]

It is straightforward to verify that, when we apply the composition algorithm of Fagin et al. [2005c] to compute $\Sigma_{12} \circ \Sigma_{21}$, we obtain $\mathtt{R}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y)$, $\mathtt{R}(x, x) \wedge V(x) \rightarrow \widehat{\mathtt{R}}(x, x)$, $\mathtt{R}(x, x) \wedge \mathtt{R}(y, y) \wedge \mathtt{R}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y)$, and $\mathtt{V}(x) \rightarrow \widehat{\mathtt{V}}(x)$. It is easy to see that the second and third tgds are logical consequences of the first tgd. So the composition $\Sigma_{12} \circ \Sigma_{21}$ is (logically equivalent to) $\{\mathtt{R}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y), \mathtt{V}(x) \rightarrow \widehat{\mathtt{V}}(x)\}$. But this set of tgds specifies the identity mapping. It follows that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$.

We conclude by showing that there is no subset X of Σ_{12} such that the tgds $\{rev(\gamma) : \gamma \in X\}$ specify a global inverse of Σ_{12}. Define Σ_{21}^X to be $\{rev(\gamma) : \gamma \in X\}$. Let $\mathcal{M}_{21}^X = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^X)$. Assume that $\mathcal{M}_{21}^X$ is a global inverse of $\mathcal{M}_{12}$.

We first show that the set X cannot contain $\mathtt{R}(x, y) \rightarrow \mathtt{S}(x, y)$. Assume that it does. Then Σ_{21}^X contains $\mathtt{S}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y)$. Let I contain only the facts $\mathtt{R}(0, 0)$ and $\mathtt{R}(1, 1)$. Then $chase_{12}(I)$ contains the fact $\mathtt{S}(0, 1)$. So $chase_{21}^X(chase_{12}(I))$, the result of chasing $chase_{12}(I)$ with Σ_{21}^X, contains the fact $\widehat{\mathtt{R}}(0, 1)$. Therefore, $\widehat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 7.4 that $\mathcal{M}_{21}^X$ is not an inverse of $\mathcal{M}_{12}$ for I, and hence $\mathcal{M}_{21}^X$ is not a global inverse of $\mathcal{M}_{12}$, a contradiction.

We now show that the set X cannot contain $\mathtt{R}(x, y) \rightarrow \mathtt{T}(x, y)$. Assume that it does. Then Σ_{21}^X contains $\mathtt{T}(x, y) \rightarrow \widehat{\mathtt{R}}(x, y)$. Let I contain only the fact $\mathtt{V}(0)$. Then $chase_{12}(I)$ contains the fact $\mathtt{T}(0, 0)$. So $chase_{21}^X(chase_{12}(I))$ contains the fact $\widehat{\mathtt{R}}(0, 0)$. Therefore, $\widehat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 7.4 that $\mathcal{M}_{21}^X$ is not an inverse of $\mathcal{M}_{12}$ for I, and hence $\mathcal{M}_{21}^X$ is not a global inverse of $\mathcal{M}_{12}$, a contradiction.

All that is left for X now is to consist of some subset of the three tgds $\mathtt{V}(x) \rightarrow \mathtt{U}(x)$, $\mathtt{R}(x, x) \wedge \mathtt{R}(y, y) \rightarrow \mathtt{S}(x, y)$, and $\mathtt{V}(x) \rightarrow \mathtt{T}(x, x)$. In this case, Σ_{21}^X consists of some subset of the tgds $\mathtt{U}(x) \rightarrow \widehat{\mathtt{V}}(x)$, $\mathtt{S}(x, y) \rightarrow \widehat{\mathtt{R}}(x, x) \wedge \widehat{\mathtt{R}}(y, y)$, and $\mathtt{T}(x, x) \rightarrow \widehat{\mathtt{V}}(x)$. Let I contain only the fact $\mathtt{R}(0, 1)$. Then $chase_{21}^X(chase_{12}(I))$ (and in fact, $chase_{21}^X(J)$ for an arbitrary J) does not contain the fact $\widehat{\mathtt{R}}(0, 1)$, since the only facts about $\widehat{\mathtt{R}}$ that can be generated by chasing with Σ_{21}^X are of the form $\widehat{\mathtt{R}}(x, x)$. Therefore, $\widehat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 7.4 that $\mathcal{M}_{21}^X$ is not an inverse of $\mathcal{M}_{12}$ for I, and hence $\mathcal{M}_{21}^X$ is not a global inverse of $\mathcal{M}_{12}$, a contradiction. □

[13]This mapping $\mathcal{M}_{21}$ is not the same as the canonical candidate tgd global inverse in Theorem 9.3.

12. CHARACTERIZING THE CLASS $\mathcal{S}$

Given schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, we might want to know the class $\mathcal{S}$ of ground instances I such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. For example, this is a problem posed in our running example (Example 5.1). Let $\mathcal{M}_{11}$ be the schema mapping $(\mathbf{S_1}, \widehat{\mathbf{S_1}}, \sigma)$, where σ is the composition formula $\Sigma_{12} \circ \Sigma_{21}$. The class $\mathcal{S}$ we are seeking is the class of all ground instances I such that $\mathcal{M}_{11}$ and the identity mapping are equivalent on I. Therefore, the class $\mathcal{S}$ is determined completely by the composition formula σ. In this section, we show that, remarkably, there is a syntactic transformation of σ that produces a formula Γ that actually defines $\mathcal{S}$! We now begin our development, with the formal definition of second-order tgds from Fagin et al. [2005c].

Given a collection $\mathbf{x}$ of variables and a collection $\mathbf{f}$ of function symbols, a *term (based on* $\mathbf{x}$ *and* $\mathbf{f}$*)* is defined recursively as follows: (1) every variable in $\mathbf{x}$ is a term, and (2) if f is a k-ary function symbol in $\mathbf{f}$ and $t_1, \ldots, t_k$ are terms, then $f(t_1, \ldots, t_k)$ is a term. We now define a second-order tgd.

Definition 12.1. Let $\mathbf{S}$ be a source schema and $\mathbf{T}$ a target schema. A *second-order tuple-generating dependency (SO tgd)* is a formula of the form

$$\exists \mathbf{f}((\forall \mathbf{x_1}(\varphi_1 \rightarrow \psi_1)) \wedge \cdots \wedge (\forall \mathbf{x_n}(\varphi_n \rightarrow \psi_n))),$$

where (1) each member of $\mathbf{f}$ is a function symbol; (2) each φ_i is a conjunction of (a) atoms $S(y_1, \ldots, y_k)$, where S is a k-ary relation symbol of schema $\mathbf{S}$, and $y_1, \ldots, y_k$ are variables in $\mathbf{x_i}$, not necessarily distinct, and (b) equalities of the form $t = t'$ where t and t' are terms based on $\mathbf{x_i}$ and $\mathbf{f}$; (3) each ψ_i is a conjunction of atoms $T(t_1, \ldots, t_l)$, where T is an l-ary relation symbol of schema $\mathbf{T}$ and $t_1, \ldots, t_l$ are terms based on $\mathbf{x_i}$ and $\mathbf{f}$; and (4) each variable in $\mathbf{x_i}$ appears in some atom of φ_i.

It was shown in Fagin et al. [2005c] that the composition of schema mappings, each specified by a finite set of s-t tgds, is specified by an SO tgd (but not necessarily by a finite set of s-t tgds). In particular, our composition formula $\Sigma_{12} \circ \Sigma_{21}$, which we are denoting by σ, is given by an SO tgd.

If γ is an SO tgd, or a set of (first-order) tgds, from $\mathbf{S}$ to $\widehat{\mathbf{S}}$, define $\gamma^\sharp$ to be the source constraint that is the result of replacing each relational symbol $\widehat{\mathtt{R}}$ in γ by $\mathtt{R}$. For example, if γ is the s-t tgd (4), then $\gamma^\sharp$ is (2). The next proposition follows easily from the definitions of $\widehat{I}$ and of $\gamma^\sharp$.

PROPOSITION 12.2. *Let γ be an SO tgd, or a set of s-t tgds with source schema* $\mathbf{S}$ *and target schema* $\widehat{\mathbf{S}}$*, and let I be an instance of* $\mathbf{S}$*. Then $I \models \gamma^\sharp$ if and only if $\langle I, \widehat{I} \rangle \models \gamma$.*

We need some more definitions. Let γ be an SO tgd. We now define the *equality-free reduction*[14] γ^* of γ. The intuition is that we think of the function symbols as representing Skolem functions, so that the only way that two "Skolem terms" $f(\mathbf{t})$ and $g(\mathbf{t}')$ can be equal is if the function symbols f and g are the same, and if $\mathbf{t} = \mathbf{t}'$. Beginning with γ, in order to obtain

[14]A similar notion appears in Yu and Popa [2005] under the name *mapping reduction*.

the equality-free reduction γ^* of γ, we first recursively replace each equality $f(t_1, \ldots, t_k) = f(t'_1, \ldots, t'_k)$ by $(t_1 = t'_1) \wedge \cdots \wedge (t_k = t'_k)$. We replace each equality $f(\mathbf{t}) = g(\mathbf{t}')$ where f and g are different function symbols by "False." Similarly, we replace each equality $f(\mathbf{t}) = x$, where x is a variable, by "False." We then "clean up" by deleting each "tgd" that appears as a conjunct of γ and that contains "False." The remaining equalities are all of the form $x = y$, where x and y are variables. Within each "tgd," we form equivalence classes of variables based on these equalities (where two variables are in the same equivalence class if they are forced to be equal by these equalities), replace each occurrence of each variable by a fixed representative of its equivalence class, and delete the equalities. The final result γ^* is an SO tgd that contains no equalities.

For example, the equality-free reduction of the SO tgd (1) is $\exists f(\forall e(\texttt{Emp}(e) \rightarrow \texttt{Mgr}(e, f(e)))$, the result of dropping the second clause of (1). As another example, consider the following SO tgd:

$$\exists f(\forall x \forall y(\texttt{R}(x, y) \wedge (f(x) = f(y)) \rightarrow \texttt{S}(x, f(x)) \wedge \texttt{T}(x, y))). \tag{16}$$

Its equality-free reduction is $\exists f(\forall x(\texttt{R}(x, x) \rightarrow \texttt{S}(x, f(x)) \wedge \texttt{T}(x, x)))$, which is obtained by replacing $f(x) = f(y)$ by $x = y$ and simplifying.

We now define *fulltgd*(γ), which is a set of full tgds that we associate with the SO tgd γ. To obtain *fulltgd*(γ), we first find the equality-free reduction γ^* of γ. We then rewrite γ^* so that each conclusion is a singleton. Thus, we replace $\varphi \rightarrow (\psi_1 \wedge \cdots \wedge \psi_r)$, where $\psi_1, \ldots, \psi_r$ are atoms, by $(\varphi \rightarrow \psi_1) \wedge \cdots \wedge (\varphi \rightarrow \psi_r)$. We then delete each "tgd" $\varphi \rightarrow \psi$ where the conclusion ψ contains a function symbol. Then *fulltgd*(γ) is the set of s-t tgds that remain. These are real tgds, since there are no function symbols present. By construction, *fulltgd*(γ) is a set of full tgds with singleton conclusions.

As an example, when γ is the SO tgd (1), then *fulltgd*(γ) is the empty set. As another example, when γ is the SO tgd (16), then *fulltgd*(γ) contains the single tgd $\texttt{R}(x, x) \rightarrow \texttt{T}(x, x)$.

For each SO tgd γ where the source schema is $\mathbf{S}$ and the target schema is $\widehat{\mathbf{S}}$, we now define $\gamma^\dagger$. As we shall see, if σ is the composition formula, then $\sigma^\dagger$ plays a complementary role to $\sigma^\sharp$. For each k and each k-ary relational symbol R of $\mathbf{S}$, take $x_1, \ldots, x_k$ to be k distinct variables that do not appear in *fulltgd*(γ). Let $A_{\texttt{R}}$ be the set of all tgds of *fulltgd*(γ) where the relational symbol in the conclusion is $\widehat{\texttt{R}}$. For each $\alpha \in A_{\texttt{R}}$, assume that α is $\nu(\mathbf{y}) \rightarrow \widehat{\texttt{R}}(y_1, \ldots, y_k)$, where $y_1, \ldots, y_k$ are not necessarily distinct (since α is full, every y_i appears in $\mathbf{y}$). Define μ_α to be the first-order formula $\exists \mathbf{y}(\nu(\mathbf{y}) \wedge (x_1 = y_1) \wedge \cdots \wedge (x_k = y_k))$. Define $\psi_{\texttt{R}}$ to be $\texttt{R}(x_1 \ldots, x_k) \rightarrow \vee \{\mu_\alpha : \alpha \in A_{\texttt{R}}\}$. Since the empty disjunction represents "False," it follows that, if $A_{\texttt{R}} = \emptyset$, then $\psi_{\texttt{R}}$ is equivalent to $\neg\texttt{R}(x_1 \ldots, x_k)$. Now define $\gamma^\dagger$ to be the conjunction of the formulas $\psi_{\texttt{R}}$ (over all relational symbols R of $\mathbf{S}$). Note that $\gamma^\dagger$ is a first-order formula.

Example 12.3. Assume that there are two source relation symbols R and S, and assume that *fulltgd*(γ) consists of the following tgds, which we denote by α_1, α_2:

$(\alpha_1):\quad \texttt{R}(y_2, y_1, y_3) \wedge \texttt{S}(y_2, y_3, y_3) \rightarrow \widehat{\texttt{R}}(y_1, y_1, y_2);$

$(\alpha_2):\quad \texttt{S}(y_1, y_2, y_2) \rightarrow \widehat{\texttt{R}}(y_1, y_2, y_1).$

So $\mu_{\alpha_1}, \mu_{\alpha_2}$ are as follows:

$$(\mu_{\alpha_1}): \quad \exists y_1 \exists y_2 \exists y_3 (\mathtt{R}(y_2, y_1, y_3) \wedge \mathtt{S}(y_2, y_3, y_3) \wedge (x_1 = y_1) \wedge (x_2 = y_1) \wedge (x_3 = y_2));$$
$$(\mu_{\alpha_2}): \quad \exists y_1 \exists y_2 \exists y_3 (\mathtt{S}(y_1, y_2, y_2) \wedge (x_1 = y_1) \wedge (x_2 = y_2) \wedge (x_3 = y_1)).$$

Then $\psi_{\mathtt{R}}$ is $\mathtt{R}(x_1, x_2, x_3) \rightarrow (\mu_{\alpha_1} \vee \mu_{\alpha_2})$. Further, $\psi_{\mathtt{S}}$ is $\neg\mathtt{S}(x_1, x_2, x_3)$, since $A_{\mathtt{S}} = \emptyset$. Finally, $\gamma^{\dagger}$ is $(\mathtt{R}(x_1, x_2, x_3) \rightarrow (\mu_{\alpha_1} \vee \mu_{\alpha_2})) \wedge \neg\mathtt{S}(x_1, x_2, x_3)$. (Of course, this formula is universally quantified with $\forall x_1 \forall x_2 \forall x_3$, but we suppress this as usual.)

LEMMA 12.4. *Let γ be an SO tgd. Then*

(1) each tgd in $fulltgd(\gamma)$ is logically implied by γ, and

(2) if I is a ground instance, and if $\mathtt{S}(\mathbf{a})$ is a fact obtained by chasing[15] I with γ, where each member of $\mathbf{a}$ is a constant, then $\mathtt{S}(\mathbf{a})$ is obtained by chasing I with some member of $fulltgd(\gamma)$.

PROOF. It is straightforward to verify that γ logically implies its equality-free reduction γ^*, which in turn logically implies each tgd in $fulltgd(\gamma)$. Then (1) follows.

We now prove (2). It follows from the definition of the chase with SO tgds in Fagin et al. [2005c] that the result of chasing I with γ is the same as the result of chasing I with γ^*. It is easy to see that the only chase steps in chasing with γ^* that can generate a fact $\mathtt{S}(\mathbf{a})$ where each member of $\mathbf{a}$ is a constant is by chasing with members of $fulltgd(\gamma)$. □

We now have the following proposition.

PROPOSITION 12.5. *Let γ be an SO tgd with source schema $\mathbf{S}$ and target schema $\widehat{\mathbf{S}}$, and let I be a ground instance of $\mathbf{S}$. The following are equivalent:*

(1) $I \models \gamma^{\dagger}$.

(2) $\langle I, J \rangle \models \gamma$ implies that $\widehat{I} \subseteq J$, for every ground instance J.

PROOF. We first show that (1) $\Rightarrow$ (2). Assume that $I \models \gamma^{\dagger}$ and $\langle I, J \rangle \models \gamma$; we must show that $\widehat{I} \subseteq J$. Let $\mathtt{R}(a_1, \ldots, a_k)$ be a fact of I. Since $I \models \gamma^{\dagger}$, we know that $I \models \psi_{\mathtt{R}}$. Since $I \models \psi_{\mathtt{R}}$ and $\mathtt{R}(a_1, \ldots, a_k)$ is a fact of I, we know that $\psi_{\mathtt{R}}$ is not equivalent to $\neg\mathtt{R}(x_1 \ldots, x_k)$. So $A_{\mathtt{R}} \neq \emptyset$. Since $I \models \psi_{\mathtt{R}}$, there is α in $A_{\mathtt{R}}$ such that I satisfies μ_{α} when the roles of $x_1, \ldots, x_k$ are played by $a_1, \ldots, a_k$, respectively. Assume that α is $\nu(\mathbf{y}) \rightarrow \widehat{\mathtt{R}}(y_1, \ldots, y_k)$. So I satisfies $\exists \mathbf{y}(\nu(\mathbf{y}) \wedge (x_1 = y_1) \wedge \cdots \wedge (x_k = y_k))$ when the roles of $x_1, \ldots, x_k$ are played by $a_1, \ldots, a_k$, respectively. But also $\langle I, J \rangle \models \alpha$; this is because $\langle I, J \rangle \models \gamma$ and hence (by part (1) of Lemma 12.4) $\langle I, J \rangle \models fulltgd(\gamma)$ and so $\langle I, J \rangle \models \alpha$ (since $\alpha \in fulltgd(\gamma)$). Therefore, $\widehat{\mathtt{R}}(a_1, \ldots, a_k)$ is a fact of J. So $\widehat{I} \subseteq J$, as desired.

We now show that (2) $\Rightarrow$ (1). Assume that (2) holds. We must show that $I \models \gamma^{\dagger}$. Let $\mathtt{R}$ be an arbitrary relational symbol of $\mathbf{S}$; thus, we must show that $I \models \psi_{\mathtt{R}}$. There are two possibilities, depending on whether $\mathtt{R}^I$, the $\mathtt{R}$ relation of I, is empty or not. If $\mathtt{R}^I$ is empty, then clearly $I \models \psi_{\mathtt{R}}$. So assume that $\mathtt{R}^I$ is nonempty. Let $\langle I, J^* \rangle$ be the result of chasing $\langle I, \emptyset \rangle$ with γ. Then $\langle I, J^* \rangle \models \gamma$.

[15]There is a notion of *chasing* with SO tgds [Fagin et al. 2005c] that is very similar to the notion of chasing with s-t tgds.

Let J' be the result of replacing each distinct term in J^* by a new constant. Since $\langle I, J^*\rangle \models \gamma$, it is not hard to see that $\langle I, J'\rangle \models \gamma$. Since (2) holds by assumption, it follows that $\widehat{I} \subseteq J'$. Therefore, $\widehat{I} \subseteq J^*$. Let $\mathtt{R}(a_1, \ldots, a_k)$ be a fact of I. So $\widehat{\mathtt{R}}(a_1, \ldots, a_k)$ is a fact of J^*. By part (2) of Lemma 12.4, $\widehat{\mathtt{R}}(a_1, \ldots, a_k)$ is obtained by chasing I with some member α of *fulltgd*(γ). So μ_α holds in I when the roles of $x_1, \ldots, x_k$ are played by $a_1, \ldots, a_k$, respectively. Therefore, $\psi_{\mathtt{R}}$ holds for I when the roles of $x_1, \ldots, x_k$ are played by $a_1, \ldots, a_k$, respectively. Since $\mathtt{R}(a_1, \ldots, a_k)$ is an arbitrary fact of I of the form $\mathtt{R}(x_1, \ldots, x_k)$, it follows that $I \models \psi_{\mathtt{R}}$, as desired. □

The next theorem gives a formula that defines the largest class $\mathcal{S}$ of ground instances where $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$.

THEOREM 12.6. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. Let σ be $\Sigma_{12} \circ \Sigma_{21}$. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for the ground instance I if and only if $I \models \sigma^\sharp \wedge \sigma^\dagger$.*

PROOF. Assume first that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I; we must show that $I \models \sigma^\sharp \wedge \sigma^\dagger$. By definition of inverse, (3) holds. So $\langle I, \widehat{I}\rangle \models \sigma$. By Proposition 12.2 it follows that $I \models \sigma^\sharp$. Furthermore, by (3), we know that whenever I and J are ground instances and $\langle I, J\rangle \models \sigma$, necessarily $\widehat{I} \subseteq J$. Therefore, by Proposition 12.5, it follows that $I \models \sigma^\dagger$. So $I \models \sigma^\sharp \wedge \sigma^\dagger$, as desired.

Conversely, assume that $I \models \sigma^\sharp \wedge \sigma^\dagger$; we must show that (3) holds for every ground instance J. Let J be a ground instance. Since $I \models \sigma^\sharp$, we know from Proposition 12.2 that $\langle I, \widehat{I}\rangle \models \sigma$. So by Corollary 4.3, we know that $\langle I, J\rangle \models \sigma$ whenever $\widehat{I} \subseteq J$. Since $I \models \sigma^\dagger$, it follows from Proposition 12.5 that $\widehat{I} \subseteq J$ whenever $\langle I, J\rangle \models \sigma$. So (3) holds for every ground instance J. □

From our earlier Theorem 7.3, we can prove that, if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each specified by a finite set of s-t tgds and held fixed, then the problem of deciding, given I, whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I is in NP. (Complexity results appear in Section 14.) So by Fagin's Theorem [Fagin 1974], the class $\mathcal{S}$ of ground instances I such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I can be specified by a formula Γ in existential second-order logic. What is remarkable is that, as Theorem 12.6 says, there is such a formula Γ, namely, $\sigma^\sharp \wedge \sigma^\dagger$, that can be obtained from the composition formula σ by a purely syntactical transformation.

The following corollary gives an important case where $\mathcal{S}$ is first-order definable.

COROLLARY 12.7. *Assume that $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are schema mappings that are each specified by a finite set of s-t tgds, and the s-t tgds of $\mathcal{M}_{12}$ are full. There is a first-order formula φ such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for the ground instance I if and only if $I \models \varphi$.*

PROOF. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds, and the s-t tgds in Σ_{12} are full. Then the composition formula $\Sigma_{12} \circ \Sigma_{21}$, which we are denoting by σ, is defined by a finite set of s-t tgds [Fagin et al. 2005c]. In particular, σ is first-order. So $\sigma^\sharp$

is first-order. Also, $\sigma^\dagger$ is always first-order. So $\sigma^\sharp \wedge \sigma^\dagger$ is first-order. The result now follows from Theorem 12.6. □

Example 12.8. We continue with our running example (from Examples 5.1 and 5.3). We shall fulfill our promise to show that the schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely those ground instances I that satisfy Γ. We noted that Γ (as given by (2)) looks mysteriously similar to the composition formula (as given by (4)). We shall explain this mystery.

We observed in Example 5.3 that σ logically implies Σ_{Id}. We now show that this implies that $\sigma^\dagger$ is valid. Let I be an arbitrary instance of $\mathbf{S_1}$; we must show that $I \models \sigma^\dagger$. We can assume without loss of generality that I is a ground instance. By Proposition 12.5, we need only show that $\langle I, J\rangle \models \sigma$ implies that $\widehat{I} \subseteq J$ when J is a ground instance. Let J be an arbitrary ground instance such that $\langle I, J\rangle \models \sigma$. Since σ logically implies Σ_{Id}, it follows that $\langle I, J\rangle \models \Sigma_{Id}$. Therefore, $\widehat{I} \subseteq J$, as desired. So indeed, $\sigma^\dagger$ is valid.

Therefore, by Theorem 12.6, we know that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $I \models \sigma^\sharp$. But in the case we are considering, $\sigma^\sharp$ is exactly Γ. This not only proves our claim that the schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely those ground instances I that satisfy Γ, but also explains the mystery of the resemblance of σ and Γ. In fact, this mysterious resemblance in this example is what inspired us to search for and discover Theorem 12.6.

13. SMALL MODEL THEOREM

A *small model theorem* says that, when there is an instance with certain properties, then there is a "small" (polynomial-size) instance with the same properties. In this section, we give a small model theorem, which says that for a schema mapping $\mathcal{M}_{12}$ specified by a finite set of full tgds, if there is an instance I such that $\mathcal{M}_{12}$ does not have an inverse for I, then there is a small such instance I. We use this result to prove an upper bound on complexity in Section 14. In the case of schema mappings specified by tgds that are not necessarily full, the small model theorem fails because of a recent undecidability result by Arenas [2006]. In this case, we give a concrete example that demonstrates the failure of the methodology (a "small submodel theorem") used to prove the small model theorem. We now state and prove our small submodel theorem.

THEOREM 13.1 (SMALL SUBMODEL THEOREM). *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be schema mappings where Σ_{12} and Σ_{21} are finite sets of full s-t tgds. Let I be a ground instance. Assume that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I. Then there is a subinstance I' of I, with size polynomial in the size of Σ_{12} and Σ_{21}, such that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I'.*

PROOF. Let us denote $chase_{21}(chase_{12}(I))$ by J. By Corollary 7.4, $\widehat{I} \neq J$. There are two cases.

Case 1: $\widehat{I}$ is not a subinstance of J. Then there is a fact $\mathtt{R}(\mathbf{t})$ of I such that $\widehat{\mathtt{R}}(\mathbf{t})$ is not in J. Let I' be the one-tuple subinstance that contains only the fact $\mathtt{R}(\mathbf{t})$. By monotonicity of the chase, $\widehat{\mathtt{R}}(\mathbf{t})$ is not in $chase_{21}(chase_{12}(I'))$, Hence, $\widehat{I'} \neq chase_{21}(chase_{12}(I'))$, and so, by Corollary 7.4, we know that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I'.

Case 2: $\widehat{I}$ is a subinstance of J. So $\widehat{I}$ is a proper subinstance of J. Let k_{12} be the maximum, over all members ψ of Σ_{12}, of the number of conjuncts that appear in the premise of ψ. Similarly, let k_{21} be the maximum, over all members ψ of Σ_{21}, of the number of conjuncts that appear in the premise of ψ. Since $\widehat{I}$ is a proper subset of J, there is a fact $\widehat{\mathtt{R}}(\mathbf{t})$ that is in J but not in $\widehat{I}$. Note that the members of $\mathbf{t}$ are all constants, since all of the tgds under consideration are full. Now $\widehat{\mathtt{R}}(\mathbf{t})$ was generated in one chase step when chasing $chase_{12}(I)$ with a tgd γ in Σ_{21}. There is a set W of at most k_{21} facts of $chase_{12}(I)$ such that the conjuncts in the premise of γ map to members of W in order to obtain $\widehat{\mathtt{R}}(\mathbf{t})$ with a chase step. Each member w of W was obtained from I by chasing with a member γ_w of Σ_{12}. There is a subinstance I_w of at most k_{12} facts of I that the conjuncts in the premise of γ_w each map to in order to obtain w by chasing I with γ_w. Let I' be the union over all $w \in W$ of I_w. The number of facts in I' is at most $k_{12}k_{21}$, which is polynomial in the size of Σ_{12} and Σ_{21}. By definition of I', it follows easily that $chase_{21}(chase_{12}(I'))$ contains the fact $\widehat{\mathtt{R}}(\mathbf{t})$. Since $\widehat{\mathtt{R}}(\mathbf{t})$ is not in $\widehat{I}$, it follows that $\widehat{\mathtt{R}}(\mathbf{t})$ is not in $\widehat{I'}$. Since $\widehat{\mathtt{R}}(\mathbf{t})$ is not in $\widehat{I'}$ but is in $chase_{21}(chase_{12}(I'))$, it follows that $\widehat{I'} \neq chase_{21}(chase_{12}(I'))$. By Corollary 7.4, we know that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I'. □

As an immediate corollary of the small submodel theorem, we obtain the following small model theorem.

THEOREM 13.2 (SMALL MODEL THEOREM). *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ be schema mappings where Σ_{12} and Σ_{21} are finite sets of full s-t tgds. Assume that $\mathcal{M}_{21}$ is not a global inverse of $\mathcal{M}_{12}$. Then there is an instance I, with size polynomial in the size of Σ_{12} and Σ_{21}, such that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I.*

The small model theorem (and hence the small submodel theorem) fails when the tgds are not necessarily full. This follows easily from a recently announced result of Arenas [2006] that the problem of deciding whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, when $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each specified by a finite set of s-t tgds, is undecidable.

The next theorem gives a concrete example that demonstrates a dramatic failure of the small submodel theorem when the tgds are not necessarily full.

THEOREM 13.3. *There are schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{12} is a finite set of s-t tgds, and Σ_{21} is a finite set of full s-t tgds, and where, for arbitrarily large n, there is an instance I of $\mathbf{S_1}$ consisting of n facts, such that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I, but $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for every proper subinstance I' of I.*

PROOF. The source schema $\mathbf{S_1}$ has binary relation symbols P and D. The target schema $\mathbf{S_2}$ has binary relation symbols P′, D′, and C. The set Σ_{12} consists of the tgds $\mathtt{P}(x, y) \rightarrow \mathtt{P}'(x, y)$, $\mathtt{D}(u, v) \rightarrow \mathtt{D}'(u, v)$, and $\mathtt{P}(x, y) \rightarrow \exists u \exists v(\mathtt{C}(x, u) \wedge \mathtt{C}(y, v))$.

The set Σ_{21} consists of the tgds $\mathtt{P}'(x, y) \rightarrow \widehat{\mathtt{P}}(x, y)$, $\mathtt{D}'(u, v) \rightarrow \widehat{\mathtt{D}}(u, v)$, and $\mathtt{P}'(x, y) \wedge \mathtt{P}'(y, x) \wedge \mathtt{D}'(w, z) \wedge \mathtt{D}'(z, w) \wedge \mathtt{C}(x, u) \wedge C(y, v) \rightarrow \widehat{\mathtt{D}}(u, v)$.

We now define I. We let $\mathtt{P}^I$, the $\mathtt{P}$ relation of I, be an undirected cycle of odd length $2m+1$. Specifically, define $i \oplus 1$ to be $i+1$ mod $2m+1$, and let $\mathtt{P}^I$ consist of all tuples $(i, i \oplus 1)$ and $(i \oplus 1, i)$, for $0 \leq i \leq 2m$. We let $\mathtt{D}^I$ consist of the two tuples $(g, b), (b, g)$, where g and b are distinct (and intuitively represent "green" and "blue," since we will be considering 2-colorings). Then the number n of facts of I is $4m+4$, which can be made arbitrarily large by taking m arbitrarily large.

We now show that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I. Assume that it were; we shall derive a contradiction. Since $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$, it follows that $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Hence, there is J such that $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \widehat{I} \rangle \models \Sigma_{21}$. Since $\langle I, J \rangle \models \Sigma_{12}$, it follows that, for each i with $0 \leq i \leq 2m$, there is u such that $C(i, u)$ holds in J. For each i with $0 \leq i \leq 2m$, define $c(i)$ to be some such u. Since $\mathtt{P}$ and $\mathtt{D}$ are copied into $\mathtt{P}'$ and $\mathtt{D}'$, respectively (according to the first two tgds in Σ_{12}), and since $\langle J, \widehat{I} \rangle \models \Sigma_{21}$, the third tgd in Σ_{21} tells us (when the roles of x, y, w, z, u, v are played by, respectively, $i, i \oplus 1, g, b, c(i), c(i \oplus 1)$) that $(c(i), c(i \oplus 1))$ is in $\widehat{D}^{\widehat{I}}$, which equals D^I. Hence, $c(i)$ and $c(i \oplus 1)$ are each either g or b and are distinct. Therefore, c is a 2-coloring of the odd cycle $\mathtt{P}^I$. But this is impossible, since an odd cycle does not have a 2-coloring.

Let I' be a proper subinstance of I. We now show that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I'. Because of the first two tgds of Σ_{12} and the first two tgds of Σ_{21}, it is clear that $\widehat{I'} \subseteq chase_{21}(chase_{12}(I'))$. Therefore, by Theorem 7.3, we need only show that $\langle I', \widehat{I'} \rangle \models \Sigma_{12} \circ \Sigma_{21}$. There are two cases.

Case 1: $D^{I'}$ is a proper subset of D^I. Define the $\mathbf{S_2}$ instance J by letting $\mathtt{P}'^J = \mathtt{P}^{I'}$ and $\mathtt{D}'^J = \mathtt{D}^{I'}$, and defining $\mathtt{C}^J$ to consist of the tuples (i, g), for $0 \leq i \leq 2m$. It is clear that $\langle I', J \rangle \models \Sigma_{12}$. So we need only show that $\langle J, \widehat{I'} \rangle \models \Sigma_{21}$. Clearly $\langle J, \widehat{I'} \rangle$ satisfies the first two tgds of Σ_{21}. Let γ be the third tgd of Σ_{21}. Then $\langle J, \widehat{I'} \rangle$ satisfies γ also, since there are no w, z such that (w, z) and (z, w) are each tuples of $\mathtt{D}'^J$.

Case 2: $D^{I'} = D^I$. Since I' is a proper subinstance of I, there is a tuple $(i_0, i_0 \oplus 1)$ of $\mathtt{P}^I$ that is not in $\mathtt{P}^{I'}$. Define $c : \{0, \ldots, 2m\} \rightarrow \{g, b\}$ by letting $c(i_0) = c(i_0 \oplus 1) = g$, and having $c(i) \neq c(i \oplus 1)$ if $i \neq i_0$. Intuitively, the points on the cycle are alternately colored g or b, except that i_0 and $i_0 \oplus 1$ are both colored g. This is possible, since the cycle is of odd length. Define the $\mathbf{S_2}$ instance J by letting $\mathtt{P}'^J = \mathtt{P}^{I'}$ and $\mathtt{D}'^J = \mathtt{D}^{I'}$, and defining $\mathtt{C}^J$ to consist of the tuples $(i, c(i))$, for $0 \leq i \leq 2m$. It is clear that $\langle I', J \rangle \models \Sigma_{12}$. So we need only show that $\langle J, \widehat{I'} \rangle \models \Sigma_{21}$. Clearly $\langle J, \widehat{I'} \rangle$ satisfies the first two tgds of Σ_{21}. We now show that $\langle J, \widehat{I'} \rangle$ satisfies the third tgd γ also. If (x, y) is either $(i_0, i_0 \oplus 1)$ or $(i_0 \oplus 1, i_0)$, then the premise of γ fails (since $\mathtt{P}'(x, y) \wedge \mathtt{P}'(y, x)$ fails), and hence γ holds. Otherwise (if (x, y) is not $(i_0, i_0 \oplus 1)$ and not $(i_0 \oplus 1, i_0)$), then again γ holds. So in all cases, γ holds. □

14. COMPLEXITY RESULTS

This section presents complexity results dealing with both local and global invertibility. We do not consider complexity issues for $\mathcal{S}$-invertibility except when $\mathcal{S}$ is a singleton (local invertibility) or when $\mathcal{S}$ is the class of all ground instances (global invertibility). It might be interesting to consider complexity issues for other choices of $\mathcal{S}$. The results are summarized in Tables I and II. In both tables,

Table I. Local Invertibility: Input Is Ground Instance I

Held fixed	Complexity
$\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	NP; may be NP-complete
Full $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	polytime
$\mathcal{M}_{12}$	Σ_2^P; may be coNP-hard
Full $\mathcal{M}_{12}$	coNP; may be coNP-complete

Table II. Global Invertibility

Input	Complexity
$\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	undecidable[17]
Full $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	DP-complete
$\mathcal{M}_{12}$	coNP-hard
Full $\mathcal{M}_{12}$	coNP-complete

we consider separately the cases where the tgds that define $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are full. In Table I, the input is a ground instance I. The first line of the table says that, if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are fixed schema mappings each specified by a finite set of s-t tgds, then the problem of deciding if $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I is in NP, and there is a choice of $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ where the problem is NP-complete.[16] The second line deals with the same problem as the first line, but where $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each specified by a finite set of *full* tgds. Then the complexity drops to polynomial time (in fact, by Corollary 12.7, the problem is even definable in first-order logic, which makes it logspace computable). The third line deals with the problem of whether $\mathcal{M}_{12}$ has an inverse for I specified by a finite set of s-t tgds. Since we have shown how to obtain a canonical candidate tgd local inverse that is an inverse for I if there is any inverse for I specified by a finite set of tgds, the reader may be puzzled as to why this problem does not reduce to the problem in the first line, where $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are given. The reason is that the size of Σ_{21} that defines the canonical candidate tgd local inverse grows with I, unlike the situation in the first line where Σ_{21} is given and so of fixed size. The fourth line deals with the same problem as the third line, but where $\mathcal{M}_{12}$ is specified by a finite set of full tgds. The problem is then in coNP, and there is a choice of $\mathcal{M}_{12}$ where the problem is coNP-complete. The third line inherits its lower bound from the full case (the fourth line). There is a complexity gap in the third line, where we have an upper bound of Σ_2^P in the polynomial-time hierarchy, and a lower bound of coNP-hardness.

In the first line of Table II, the input consists of schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each specified by a finite set of s-t tgds, and the problem is deciding whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. Arenas [2006] has recently announced that this problem is undecidable. The second line deals with the same problem as the first line, but where $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each defined by a finite set of full tgds, and the problem is then DP-complete.[18] The third line deals with the problem of whether $\mathcal{M}_{12}$ has a global inverse specified by a finite set of s-t tgds.

[16]The NP-hardness result was obtained by Phokion Kolaitis.
[17]This result is due to Arenas [2006].
[18]The class DP consists of all decision problems that can be written as the intersection of an NP problem and a coNP problem.

The fourth line deals with the same problem as the third line, but where $\mathcal{M}_{12}$ is defined by a finite set of full tgds, and the problem is then coNP-complete. In fact, this problem is coNP-complete even when $\mathcal{M}_{12}$ is also LAV, that is, when the tgds that define $\mathcal{M}_{12}$ are not only full but each has a singleton premise. The third line inherits its lower bound from the full case (the fourth line). There is a large complexity gap in the third line, since it is open as to whether this problem is even decidable.

Let $\mathcal{M}_{12}$ be a schema mapping specified by a finite set of s-t tgds. Note that we do not consider the natural problem of *computing* a global inverse of $\mathcal{M}_{12}$ specified by a finite set of s-t tgds, but instead we consider the problem of simply *deciding* whether $\mathcal{M}_{12}$ has a global inverse specified by a finite set of s-t tgds. This is good enough, since we already know how to compute such a global inverse if one exists: namely, such a global inverse is the canonical tgd global inverse. A similar comment applies to local inverses.

Although the focus in this article is on schema mappings and their inverses when both are specified by a finite set of s-t tgds, we obtain "for free" a complexity result even when we do not restrict possible inverses to be specified by a finite set of s-t tgds. Specifically, let G be the problem of deciding whether a schema mapping specified by a finite set of full s-t tgds has a global inverse specified by a finite set of s-t tgds, and let G' be the problem of deciding whether a schema mapping specified by a finite set of full s-t tgds has a global inverse (not necessarily specified by a finite set of s-t tgds). A minor modification of our proof that G is a coNP-complete problem shows that G' is a coNP-hard problem.

In the remainder of the section, we prove our complexity results.

14.1 Local Complexity

In this subsection, we shall consider the complexity of two problems:

(1) Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. What is the complexity of deciding, given an instance I of $\mathbf{S_1}$, whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I?

(2) Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ is a schema mapping where Σ_{12} is a finite set of s-t tgds. What is the complexity of deciding, given an instance I of $\mathbf{S_1}$, whether $\mathcal{M}_{12}$ has an inverse for I that is specified by a finite set of s-t tgds?

We shall consider these problems also when we restrict to *full* tgds.

We begin with a lemma. This lemma was proven in Fagin et al. [2005c], with the proof we now give, although it was not stated explicitly (it appeared in the middle of another proof).

LEMMA 14.1. *Assume that* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *and* $\mathcal{M}_{23} = (\mathbf{S_2}, \mathbf{S_3}, \Sigma_{23})$ *are schema mappings, where* $\mathbf{S_1}$, $\mathbf{S_2}$, *and* $\mathbf{S_3}$ *are pairwise disjoint, and where* Σ_{12} *and* Σ_{23} *are finite sets of s-t tgds. Then* $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$ *if and only if there is an instance* I_2 *of size polynomial in the size of* I_1, *where the degree of the polynomial depends only on* Σ_{12}, *such that* $\langle I_1, I_2\rangle \models \Sigma_{12}$ *and* $\langle I_2, I_3\rangle \models \Sigma_{23}$.

PROOF. The proof of the "if" direction follows immediately from the definition of composition (where the size of I_2 is irrelevant). We now prove the "only if" direction. Assume that $\langle I_1, I_3\rangle \models \Sigma_{12} \circ \Sigma_{23}$. Then there is J such that $\langle I_1, J\rangle \models \Sigma_{12}$ and $\langle J, I_3\rangle \models \Sigma_{12}$. It was shown in Fagin et al. [2005a] that there is a universal solution U for I with respect to Σ_{12} that is of size polynomial in the size of I, where the degree of the polynomial depends only on Σ_{12} (in fact, $chase_{12}(I)$ is such a universal solution U). So there is a homomorphism $h : U \rightarrow J$. Let $I_2 = h(U)$. Clearly, I_2 has size at most the size of U, and $I_2 \subseteq J$. Now $\langle I_1, I_2\rangle \models \Sigma_{12}$, since the homomorphic image I_2 of a solution U is a solution. Also, since $\langle J, I_3\rangle \models \Sigma_{23}$ and $I_2 \subseteq J$, it follows from Lemma 4.2 that $\langle I_2, I_3\rangle\rangle \models \Sigma_{23}$. This concludes the proof. □

THEOREM 14.2. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. The problem of deciding, given I, whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I is in NP, and can be NP-complete.*

PROOF. By Theorem 7.3, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. Doing the chase is a polynomial-time procedure (for a fixed finite set of s-t tgds). So there is a polynomial-time procedure for deciding if $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. Therefore, we need only show that the problem of deciding, given I, whether $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ is in NP. By Lemma 14.1, we know that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ if and only if there is an instance J of size polynomial in the size of I such that $\langle I, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I}\rangle \models \Sigma_{21}$. Therefore, the problem of deciding, given I, whether $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ is in NP: intuitively, we simply "guess" the intermediate instance J and verify that $\langle I, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I}\rangle \models \Sigma_{21}$.

We now show that there is a choice of $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ where the problem is NP-complete.[19]

The source schema $\mathbf{S_1}$ has binary relation symbols P and D. The target schema $\mathbf{S_2}$ has binary relation symbols P′, D′, and C. The set Σ_{12} consists of the tgds $\mathtt{P}(x, y) \rightarrow \mathtt{P}'(x, y)$, $\mathtt{D}(u, v) \rightarrow \mathtt{D}'(u, v)$, and $\mathtt{P}(x, y) \rightarrow \exists u \mathtt{C}(x, u)$.

The set Σ_{21} consists of the tgds $\mathtt{P}'(x, y) \rightarrow \widehat{\mathtt{P}}(x, y)$, $\mathtt{D}'(u, v) \rightarrow \widehat{\mathtt{D}}(u, v)$, and $\mathtt{P}'(x, y) \wedge \mathtt{C}(x, u) \wedge C(y, v) \rightarrow \widehat{\mathtt{D}}(u, v)$.

Let H be an undirected graph (thus, (x, y) is an edge of H precisely if (y, x) is an edge of H). Let V_H be the vertices of H. Let g, b, r be three distinct symbols (which intuitively represent green, blue, and red). Define the ground instance I_H by letting P contain all of the edges of H, and letting D be the inequality relation on g, b, r. Thus, the D relation consists of the six tuples $(g, b), (g, r), (b, r), (b, g), (r, g), (r, b)$. We now show that H is 3-colorable if and only if $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I_H. Since 3-colorability is an NP-complete problem, this is sufficient to prove the NP-hardness result. We assume without loss of generality that every member of V_H lies on an edge of H.

Because of the first two tgds of Σ_{12} and the first two tgds of Σ_{21}, necessarily $\widehat{I_H} \subseteq chase_{21}(chase_{12}(I_H))$ for every H. So by Theorem 7.3, it follows that $\mathcal{M}_{21}$

[19]This proof is due to Phokion Kolaitis. His proof inspired our (similar) proof of Theorem 13.3.

is an inverse of $\mathcal{M}_{12}$ for I_H if and only if $\langle I_H, \widehat{I_H}\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Therefore, we need only show that H is 3-colorable if and only if $\langle I_H, \widehat{I_H}\rangle \models \Sigma_{12} \circ \Sigma_{21}$.

Assume first that H is 3-colorable. Let $c : V_H \to \{g, b, r\}$ be a 3-coloring of H. Define J to be the $\mathbf{S_2}$ instance whose $\mathtt{P}'$ relation is the $\mathtt{P}$ relation of I_H, whose $\mathtt{D}'$ relation is the $\mathtt{D}$ relation of I_H, and where $\mathtt{C}(x, c(x))$ holds for each node x of H. It is easy to see that $\langle I_H, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I_H}\rangle \models \Sigma_{21}$. So $\langle I_H, \widehat{I_H}\rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired.

Conversely, assume that $\langle I_H, \widehat{I_H}\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Then there is J such that $\langle I_H, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I_H}\rangle \models \Sigma_{21}$. Since $\langle I_H, J\rangle \models \Sigma_{12}$, and since by assumption every member of V_H lies on an edge of H, it follows that, for each member x of V_H, there is some u such that $\mathtt{C}(x, u)$ holds in J. For each member x of V_H, define $c(x)$ to be some such u. When (x, y) is an edge of P^I (and hence of P'^J), the third tgd in Σ_{21} tells us that $(c(x), c(y))$ is in $\widehat{D}^{\widehat{I}}$, which equals D^I. Hence, $c(x)$ and $c(y)$ are each in $\{g, b, r\}$ and are distinct. Therefore, c is a 3-coloring of H, as desired. □

We now consider the case when Σ_{12} and Σ_{21} are restricted to being finite sets of *full* tgds.

THEOREM 14.3. *Assume that* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *and* $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$ *are schema mappings where* Σ_{12} *and* Σ_{21} *are each finite sets of full s-t tgds. The problem of deciding, given* I*, whether* $\mathcal{M}_{21}$ *is an inverse of* $\mathcal{M}_{12}$ *for* I *is polynomial-time solvable.*

PROOF. The theorem follows from Corollary 7.4, since there is a polynomial-time test for deciding if $\widehat{I} = chase_{21}(chase_{12}(I))$. □

Corollary 12.7 actually gives a stronger result than Theorem 14.3. Corollary 12.7 says that the problem is not only polynomial-time solvable, but even expressible in first-order logic (and therefore is solvable in logspace).

We now consider the second problem, where $\mathcal{M}_{21}$ is not given. Note that Corollary 6.5 and Theorem 8.2 give us a decision procedure for deciding if there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I. We simply check whether the active domain of $chase_{12}(I)$ contains the active domain of I. If not, we reject. If so, we produce the canonical candidate tgd local inverse $\mathcal{M}_{21}$, and verify whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. We can decide whether $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I by using the same procedure as in the proof of Theorem 14.2. However, unlike Theorem 14.2, it does not follow that this problem is in NP, since, unlike the situation in Theorem 14.2, the schema mapping $\mathcal{M}_{21}$ is not fixed, but depends on I. We now obtain an upper bound on the complexity.

THEOREM 14.4. *Assume that* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *is a schema mapping where* Σ_{12} *is a finite set of s-t tgds. The problem of deciding, given* I*, whether there is a schema mapping specified by a finite set of s-t tgds that is an inverse of* $\mathcal{M}_{12}$ *for* I *is in the complexity class* Σ_2^P *in the polynomial-time hierarchy.*

PROOF. Let $J^* = chase_{12}(I)$. Compute J^* (this is a polynomial-time procedure), and check whether the active domain of J^* contains the active domain of I. If not, then we know by Corollary 6.5 that there is no inverse of $\mathcal{M}_{12}$ for I. So assume that the active domain of J^* contains the active domain of I. Let $\mathcal{M}_{21}$ be the canonical candidate tgd local inverse. By Theorem 8.2, we know that there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I precisely if $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. Therefore, the theorem is proven if we show that the statement

$$\text{“}\mathcal{M}_{21} \text{ is an inverse of } \mathcal{M}_{12} \text{ for } I\text{”} \tag{17}$$

is in Σ_2^P. By Theorem 7.3, we know that (17) holds precisely if $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. So we need only show that the problem of deciding if $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ and the problem of deciding if $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$ are each in Σ_2^P.

Lemma 14.1 tells us that $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ if and only if there exists a polynomial-size J such that $\langle I, J\rangle \models \Sigma_{12}$ and $\langle J, \widehat{I}\rangle \models \Sigma_{21}$. Given J, the statement that $\langle I, J\rangle \models \Sigma_{12}$ is a Π_1^P statement, since $\langle I, J\rangle \models \Sigma_{12}$ precisely if there is no application of a member of Σ_{12} to selected facts in I that obtains a result outside of J. Similarly, the statement that $\langle J, \widehat{I}\rangle \models \Sigma_{21}$ is a Π_1^P statement. So the problem of deciding if $\langle I, \widehat{I}\rangle \models \Sigma_{12} \circ \Sigma_{21}$ is in Σ_2^P.

We now show that deciding if $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$ is a problem in NP, and hence in Σ_2^P. Let k be the maximum number of conjuncts in premises of members of Σ_{21}. To verify that $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$, we do the following NP procedure. For each fact F in $\widehat{I}$, we guess k applications of members of Σ_{12} to facts of I to produce k facts of $chase_{12}(I)$. We then guess an application of a member of Σ_{21} to these k facts in $chase_{12}(I)$ to produce the fact F in $chase_{21}(chase_{12}(I))$. □

If φ_1 and φ_2 are each conjunctions of atoms over the schema $\mathbf{S}$, then a *homomorphism from φ_1 to φ_2* is a homomorphism $h : I_1 \to I_2$, where I_1 is the instance of $\mathbf{S}$ whose facts are the conjuncts of φ_1, and I_2 is the instance of $\mathbf{S}$ whose facts are the conjuncts of φ_2.[20] Each value in I_1 is treated as being a member of Var, and so the restriction on homomorphisms that $h(c) = c$ for $c \in$ Const does not arise. Similarly, if I is an instance of $\mathbf{S}$, then a *homomorphism from φ_1 to I* is a homomorphism from I_1 to I. We shall make use of the following standard lemma.

LEMMA 14.5. *Assume that Σ_{12} consists of the full s-t tgd $\gamma_1 \to \gamma_2$. Then $chase_{12}(I)$ consists precisely of facts $\mathtt{S}(h(\mathbf{x}))$ such that h is a homomorphism from γ_1 to I and $\mathtt{S}(\mathbf{x})$ is a conjunct of γ_2.*

We now give a technically useful characterization in terms of homomorphisms and the chase, of when $\mathcal{M}_{12}$ has a local inverse specified by s-t tgds, when $\mathcal{M}_{12}$ is specified by full tgds.

Define a *weak endomorphism* of an instance K to be a mapping $h : K \to K$ such that, for every fact $\mathtt{R}(\mathbf{t})$ of K, we have that $\mathtt{R}(h(\mathbf{t}))$ is a fact of K. Thus, intuitively, a weak endomorphism is a homomorphism from an instance into itself that is not required to map constants into themselves.

[20]As before, these are not actual instances, but "canonical instances."

THEOREM 14.6. *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ be a schema mapping where Σ_{12} is a finite set of full s-t tgds, and let I be an instance of $\mathbf{S_1}$. Then $\mathcal{M}_{12}$ has an inverse for I by a schema mapping specified by a finite set of s-t tgds if and only if (1) $\mathcal{M}_{12}$ has the constant-propagation property for I and (2) every weak endomorphism of $chase_{12}(I)$ is a weak endomorphism of I.*

PROOF. Let $J^* = chase_{12}(I)$. Assume first that there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I. Corollary 6.5 tells us that condition (1) of the theorem must hold, and Theorem 8.2 tells us that the canonical candidate tgd local inverse $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. Since $\mathcal{M}_{21}$ is specified by a finite set Σ_{21} of full s-t tgds, it follows from Corollary 7.4 that $\widehat{I} = chase_{21}(J^*)$, and in particular $chase_{21}(J^*) \subseteq \widehat{I}$.

Let h be a weak endomorphism of J^*. Let us write the full tgd $\beta_{J^*,\widehat{I}}$ as $\gamma_1 \rightarrow \gamma_2$, where γ_1 is a conjunction of the facts of J^*, and γ_2 is a conjunction of the facts of $\widehat{I}$. Let γ_1' be the result of replacing each constant x by $h(x)$ in γ_1, and similarly for γ_2'. For example, if $\mathtt{P}(x, y)$ is a conjunct of γ_1, and $h(x) = x'$ and $h(y) = y'$, then $\mathtt{P}(x', y')$ is a conjunct of γ_1'. Since h is a weak endomorphism of J^*, we know that γ_1' is a conjunction of (some of the) facts of J^*. Since $chase_{21}(J^*) \subseteq \widehat{I}$, it follows that each conjunct of γ_2' is a fact of $\widehat{I}$. But this means that h is a weak endomorphism of $\widehat{I}$, and hence of I, as desired.

Conversely, assume that conditions (1) and (2) of the theorem hold. Since condition (1) holds, it follows that $\beta_{J^*,\widehat{I}}$ is a full tgd. Define Σ_{21} to be a set consisting only of this tgd, and, as before, let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. When we chase J^* with $\beta_{J^*,\widehat{I}}$, we clearly obtain at least $\widehat{I}$. So $\widehat{I} \subseteq chase_{21}(J^*)$. It follows easily from Lemma 14.5 that $chase_{21}(J^*)$ consists precisely of facts $\widehat{\mathtt{R}}(h(\mathbf{x}))$ such that h is a weak endomorphism of J^* and $\widehat{\mathtt{R}}(\mathbf{x})$ is a fact of $\widehat{I}$. But by condition (2) of the theorem, $\widehat{\mathtt{R}}(h(\mathbf{x}))$ is then in $\widehat{I}$. So $chase_{21}(J^*) \subseteq \widehat{I}$. Hence, $\widehat{I} = chase_{21}(J^*)$. So by Corollary 7.4, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I. □

THEOREM 14.7. *Assume that $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ is a schema mapping where Σ_{12} is a finite set of full s-t tgds. The problem of deciding, given I, whether there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I is in coNP, and can be coNP-complete.*

PROOF. The fact that the problem is in coNP follows easily from Theorem 14.6.

We now give a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ where Σ_{12} is a finite set of full tgds and where the problem is coNP-complete. We shall reduce non-3-colorability, a coNP-complete problem, to our problem. The source schema $\mathbf{S_1}$ has a binary relation symbol P and three unary relation symbols G, B, and R (which stand for green, red, and blue). The target schema $\mathbf{S_2}$ has a binary relation symbol P′ and three unary relation symbols G′, B′, and R′. The set Σ_{12} consists of the tgds $\mathtt{P}(x, y) \rightarrow \mathtt{P}'(x, y)$, $\mathtt{G}(x) \rightarrow \mathtt{G}'(x)$, $\mathtt{B}(x) \rightarrow \mathtt{B}'(x)$, $\mathtt{R}(x) \rightarrow \mathtt{R}'(x)$, $\mathtt{G}(x) \wedge \mathtt{R}(y) \rightarrow \mathtt{P}'(x, y)$, $\mathtt{G}(x) \wedge \mathtt{R}(y) \rightarrow \mathtt{P}'(y, x)$, $\mathtt{G}(x) \wedge \mathtt{B}(y) \rightarrow \mathtt{P}'(x, y)$, $\mathtt{G}(x) \wedge \mathtt{B}(y) \rightarrow \mathtt{P}'(y, x)$, $\mathtt{R}(x) \wedge \mathtt{B}(y) \rightarrow \mathtt{P}'(x, y)$, and $\mathtt{R}(x) \wedge \mathtt{B}(y) \rightarrow \mathtt{P}'(y, x)$.

Let H be an undirected graph that is connected and that has at least one edge (note that non-3-colorability is coNP-complete even when we restrict our attention to such graphs H). Let g, b, r be new symbols not appearing as vertices of H. Let V denote the nodes of H and let $C = \{g, b, r\}$. Define the ground instance I_H by letting $\texttt{P}$ contain all of the edges of H, and letting $\texttt{G}$ contain the singleton tuple (g), letting $\texttt{B}$ contain the singleton tuple (b), and letting $\texttt{R}$ contain the singleton tuple (r). We now show that there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I_H if and only if H is not 3-colorable. This is sufficient to prove the theorem.

Let $J = chase_{12}(I_H)$. Thus, the $\texttt{P}'$ relation of J consists of the tuples that are edges of H, along with the six tuples $(g, b), (g, r), (b, r), (b, g), (r, g), (r, b)$. Assume first that H is 3-colorable, under a coloring c that maps each vertex of H to $\{g, b, r\}$. Extend c to have domain $V \cup C$ by letting $c(x) = x$ for $x \in C$.

We now show that c is a weak endomorphism of J. If (x) is a tuple of the $\texttt{G}'$ relation of J, then $h(x)$ is a tuple of the $\texttt{G}'$ relation of J, since x is necessarily g, and $h(g) = g$. The same hold for the $\texttt{B}'$ relation of J and the $\texttt{R}'$ relation of J. Finally, whenever (x, y) is a tuple of the $\texttt{P}'$ relation of J, then $(c(x), c(y))$ is one of the six tuples $(g, b), (g, r), (b, r), (b, g), (r, g), (r, b)$, all of which are tuples of the $\texttt{P}'$ relation of J. So indeed, c is a weak endomorphism of J.

We now show that c is not a weak endomorphism of I_H. Let (x, y) be an edge of H (by assumption, H has at least one edge). Then $\texttt{P}(x, y)$ is a fact of I_H. However, $\texttt{P}(c(x), c(y))$ is not a fact of I_H, since $c(x)$ is either g, b, or r, and the $\texttt{P}$ relation of I_H has no tuples that contains the values g, b, or r. Therefore, c is not a weak endomorphism of I_H,

Since c is a weak endomorphism of J, but not a weak endomorphism of I_H, it follows from Theorem 14.6 that there is no schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I_H.

Assume now that H is not 3-colorable. Let c be a weak endomorphism of J. We first show that c must map V to V. Assume not. Then c maps some member of V to C. If (x, y) is an edge of H, and if c maps x to C, then necessarily c maps y to C, since there is no tuple (a, b) of the $\texttt{P}'$ relation of J where a is in V and b is in C. So since H is connected, it follows that c maps every member of V to C. For each edge $(x.y)$ of H, we have that $\texttt{P}'(x, y)$ is a fact of J, and so $\texttt{P}'(c(x), c(y))$ is a fact of J (because c is a weak endomorphism of J). Since c maps every member of V to C, it follows that the tuple $(c(x), c(y))$ is one of the six tuples $(g, b), (g, r), (b, r), (b, g), (r, g), (r, b)$, and so $c(x) \neq c(y)$. This implies that H is 3-colorable, which is a contradiction. So c maps every member of V to V. Furthermore, $c(x) = x$ for $x \in C$, since the $\texttt{G}'$ relation of J contains only (g), and similarly for $\texttt{B}'$ and $\texttt{R}'$. Therefore, c is a weak endomorphism of I_H. We have shown that every weak endomorphism of J is a weak endomorphism of I_H. Since also $\mathcal{M}_{12}$ has the constant-propagation property for I, it follows from Theorem 14.6 that there is a schema mapping specified by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I_H. □

We have now proven the complexity bounds in Table I (the coNP-hardness result in the third line is inherited from the coNP-completeness result in the fourth line).

14.2 Global Complexity

In this subsection, we shall consider the complexity of two problems:

(1) What is the complexity of deciding, given schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{12} and Σ_{21} are finite sets of s-t tgds. whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$?

(2) What is the complexity of deciding, given a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where Σ_{12} is a finite sets of s-t tgds, whether $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of s-t tgds?

We shall consider these problems also when we restrict to *full* tgds.

LEMMA 14.8. *Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where Σ_{12} is a finite set of s-t tgds. Assume that $\mathcal{M}_{12}$ has the constant-propagation property. Let $\mathcal{M}_{21}^c = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^c)$ be the canonical candidate tgd global inverse of $\mathcal{M}_{12}$. If $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}^c$, then $\widehat{I} \subseteq J$. In particular, $\widehat{I} \subseteq chase_{21}^c(chase_{12}(I))$, where $chase_{21}^c$ represents the result of chasing with Σ_{21}^c.*

PROOF. Assume that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}^c$. Then there is J' such that $\langle I, J'\rangle \models \Sigma_{12}$ and $\langle J', J\rangle \models \Sigma_{21}^c$. Let $\mathtt{R}(a_1, \ldots, a_k)$ be an arbitrary fact of I. Since $\langle I, J'\rangle \models \Sigma_{12}$, it follows that J' contains a homomorphic image J'' of the result of chasing $\mathtt{R}(a_1, \ldots, a_k)$ with Σ_{12}. Since $\langle J', J\rangle \models \Sigma_{21}^c$, it follows that J contains a homomorphic image of the result of chasing J'' with Σ_{21}^c. Therefore, by construction of Σ_{21}^c, we see that J contains the fact $\widehat{\mathtt{R}}(a_1, \ldots, a_k)$. Since $\mathtt{R}(a_1, \ldots, a_k)$ is an arbitrary fact of I, it follows that $\widehat{I} \subseteq J$, as desired.

As for the "in particular," let $J = chase_{21}^c(chase_{12}(I))$. So $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}^c$, since by Lemma 6.8, we know that $\langle I, chase_{12}(I)\rangle \models \Sigma_{12}$ and $\langle chase_{12}(I), J\rangle \models \Sigma_{21}^c$. □

THEOREM 14.9. *The problem of deciding, given a schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ with Σ_{12} a finite set of full s-t tgds, whether $\mathcal{M}_{12}$ has a global inverse that is specified by a finite set of s-t tgds, is coNP-complete. It is coNP-complete even if $\mathcal{M}_{12}$ is also LAV, that is, the members of Σ_{12} are full tgds each of which has a singleton premise.*

PROOF. We first show membership in coNP. As before, for each relational symbol $\mathtt{R}$ of $\mathbf{S_1}$, let $I_{\mathtt{R}}$ be a one-tuple instance that contains only the fact $\mathtt{R}(\mathbf{x})$, where the variables in $\mathbf{x}$ are distinct. Do a polynomial-time check to verify that for each relational symbol $\mathtt{R}$ of $\mathbf{S_1}$, we have that $\mathcal{M}_{12}$ has the constant-propagation property for $I_{\mathtt{R}}$, and if not, then reject (it is safe to reject, since by Corollary 6.5, we know that there is no global inverse of $\mathcal{M}_{12}$). Since we have that $\mathcal{M}_{12}$ has the constant-propagation property for $I_{\mathtt{R}}$ for each relational symbol $\mathtt{R}$ of $\mathbf{S_1}$, it follows easily that $\mathcal{M}_{12}$ has the constant-propagation property. Let $\mathcal{M}_{21}^c = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21}^c)$ be the canonical candidate tgd global inverse. We know from Theorem 9.3 that there is a schema mapping specified by a finite set of s-t tgds that is a global inverse of $\mathcal{M}_{12}$ if and only if $\mathcal{M}_{21}^c$ is a global inverse of $\mathcal{M}_{12}$.

We now show that the problem of determining if $\mathcal{M}_{21}^c$ is *not* a global inverse of $\mathcal{M}_{12}$ is in NP. By Theorem 13.2, we know that $\mathcal{M}_{21}^c$ is not a global inverse of $\mathcal{M}_{12}$ if and only if there is a small model I such that $\mathcal{M}_{21}^c$ is not an inverse of $\mathcal{M}_{12}$ for I. Since $\mathcal{M}_{12}$ has the constant-propagation property, it follows from Lemma 14.8 that $\widehat{I} \subseteq chase_{21}^c(chase_{12}(I))$. Let us denote $chase_{21}^c(chase_{12}(I))$ by $\bar{I}$. Hence, by Corollary 7.4, we know that $\mathcal{M}_{21}^c$ is not an inverse of $\mathcal{M}_{12}$ for I if and only if $\bar{I}$ is a proper superset of $\widehat{I}$.

We have shown that $\mathcal{M}_{21}^c$ is not a global inverse of $\mathcal{M}_{12}$ if and only if there is a small model I such that $\bar{I}$ is a proper superset of $\widehat{I}$. But this property is in NP: we simply guess the small model I and verify that $\bar{I}$ is a proper superset of $\widehat{I}$ by guessing applications of members of Σ_{12} and Σ_{21}^c that generate a member of $\bar{I}$ that is not in $\widehat{I}$.

So indeed, the problem of determining if $\mathcal{M}_{12}$ has a global inverse is in coNP. We now show that the problem is coNP-hard. We shall show that the problem of determining if $\mathcal{M}_{12}$ does *not* have a global inverse is NP-hard. We shall reduce SAT to this problem.

Let φ be a propositional formula involving n propositional symbols $A_1, \dots, A_n$ that is in conjunctive normal form $C_1 \wedge \cdots \wedge C_k$, where each C_i is a disjunction of propositional literals (either a propositional symbol A_j or the negation $\neg A_j$ of a propositional symbol). Let us write C_i as $Y_{i,1} \vee \cdots \vee Y_{i,f(i)} \vee \neg Z_{i,1} \vee \cdots \vee \neg Z_{i,g(i)}$, where each $Y_{i,j}$ and each $Z_{i,j}$ is in $\{A_1, \dots, A_n\}$. Thus, there are $f(i)$ propositional symbols that appear positively in C_i and $g(i)$ propositional symbols that appear negatively in C_i, for $1 \leq i \leq k$.

We now show how to construct in polynomial time from φ a LAV schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where the s-t tgds in Σ_{12} are full, that we shall prove has a global inverse specified by a finite set of s-t tgds if and only if φ is not satisfiable. Since SAT is NP-complete, this is sufficient to prove coNP-hardness, The source schema $\mathbf{S_1}$ consists of a $2n$-ary relational symbol $\mathtt{R}$, along with $2n$-ary relational symbols $\mathtt{U}_{i,1}, \dots, \mathtt{U}_{i,j}$ for $1 \leq i \leq k$ and $1 \leq j \leq f(i)$. The target schema $\mathbf{S_2}$ contains $2n$-ary relational symbols $\mathtt{U}'_{i,1}, \dots, \mathtt{U}'_{i,j}$ for $1 \leq i \leq k$ and $1 \leq j \leq f(i)$. Thus, for each of the relational symbols $\mathtt{U}_{i,j}$ of $\mathbf{S_1}$, there is a corresponding relational symbol $\mathtt{U}'_{i,j}$ in $\mathbf{S_2}$. Furthermore, $\mathbf{S_2}$ has $2n$-ary relational symbols $\mathtt{S}_1, \dots, \mathtt{S}_k$.

Let $\mathbf{x}$ be the $2n$-tuple $(x_1, \dots, x_{2n})$. For $1 \leq m \leq n$, let $\mathbf{x}^{=m}$ be the tuple that is the result of replacing x_{2m} in $\mathbf{x}$ by x_{2m-1}. Thus, $\mathbf{x}^{=\mathbf{m}}$ is the tuple

$$(x_1, \dots, x_{2m-1}, x_{2m-1}, x_{2m+1}, \dots x_{2n}).$$

For $1 \leq m \leq n$, let $\mathbf{x}^{\mathbf{m}}$ be the tuple that is the result of interchanging x_{2m-1} and x_{2m} in $\mathbf{x}$. Thus, $\mathbf{x}^{\mathbf{m}}$ is $(y_1, \dots, y_{2n})$, where y_{2m-1} is x_{2m} and y_{2m} is x_{2m-1}, and where y_r is x_r for $r \notin \{2m-1, 2m\}$.

Assume that $Y_{i,j}$ is the propositional symbol A_m. Let $y_{i,j}$ be the full tgd $\mathtt{U}_{i,j}(\mathbf{x}^{=\mathbf{m}}) \rightarrow \mathtt{S}_i(\mathbf{x}^{=\mathbf{m}})$. Assume that $Z_{i,j}$ is the propositional symbol A_m. Let $z_{i,j}$ be the full tgd $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x}^{\mathbf{m}})$. The set Σ_{12} consists of the following full tgds. First, it has the "copying" tgds $\mathtt{U}_{i,j}(\mathbf{x}) \rightarrow \mathtt{U}'_{i,j}(\mathbf{x})$ for each $\mathtt{U}_{i,j}$ in $\mathbf{S_1}$. It also contains the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ for $1 \leq i \leq k$. Finally, it contains the tgds $y_{i,j}$ (for $1 \leq i \leq k$ and $1 \leq j \leq f(i)$) and $z_{i,j}$ (for $1 \leq i \leq k$ and $1 \leq j \leq g(i)$).

We now prove that if φ is satisfiable, then there is no global inverse of $\mathcal{M}_{12}$. Assume that φ is satisfiable. Let us fix a truth assignment $\mathbf{Tr}$ to the propositional variables that satisfies φ, and let P be the subset of the propositional variables $A_1, \ldots, A_{2n}$ that are assigned true under the truth assignment $\mathbf{Tr}$. Let $a_1, \ldots, a_{2n}$ be distinct values. Let $\mathbf{t}$ be the $2n$-ary tuple $(t_1, \ldots, t_{2n})$, defined as follows. If $1 \leq m \leq n$, and if A_m is in P, then let t_{2m-1} and t_{2m} both be a_{2m-1}. If $1 \leq m \leq n$, and if A_m is not in P, then let t_{2m-1} be a_{2m-1} and let t_{2m} be a_{2m}. For $1 \leq m \leq n$, let $\mathbf{t^m}$ be the tuple that is the result of interchanging t_{2m-1} and t_{2m} in $\mathbf{t}$.

Define the ground instance I_1 to consist of the facts $\mathtt{U}_{i,j}(\mathbf{t})$ for each of the relational symbols $\mathtt{U}_{i,j}$ of $\mathbf{S_1}$. Furthermore, if A_m is not in P, then let I_1 also contain the fact $\mathtt{R}(\mathbf{t^m})$, for $1 \leq m \leq n$. Note that I_1 does not contain the fact $\mathtt{R}(\mathbf{t})$, since $\mathbf{t^m} \neq \mathbf{t}$ when A_m is not in P.

Let $I_\mathbf{t}$ be the one-tuple instance that contains only the fact $\mathtt{R}(\mathbf{t})$. Define the ground instance I_2 to consist of the union of the facts of I_1 and $I_\mathbf{t}$. Since I_2 contains the fact $\mathtt{R}(\mathbf{t})$ but I_1 does not, we see that $I_1 \neq I_2$. We now show that $\mathit{chase}_{12}(I_1) = \mathit{chase}_{12}(I_2)$. It then follows from Corollary 6.4 that there is no global inverse of $\mathcal{M}_{12}$. Since every member of Σ_{12} has a singleton premise, we need only show that $\mathit{chase}_{12}(I_\mathbf{t}) \subseteq \mathit{chase}_{12}(I_1)$.

The only tgds that may generate a tuple when we chase $I_\mathbf{t}$ with Σ_{12} are the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ and the tgds $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x^m})$. Consider first the result of applying the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ to $I_\mathbf{t}$. The result is $S_i(\mathbf{t})$. We must show that $S_i(\mathbf{t})$ is in $\mathit{chase}_{12}(I_1)$. Since the clause C_i is satisfied by the truth assignment $\mathbf{Tr}$, either there is some j such that $Y_{i,j} \in P$, or there is some j such that $Z_{i,j} \notin P$. Assume first that $Y_{i,j} \in P$. Say $Y_{i,j}$ is the propositional symbol A_m. So $A_m \in P$. Therefore, $t_{2m-1} = t_{2m}$. Since $Y_{i,j}$ is a disjunct of C_i, the tgd $\mathtt{U}_{i,j}(\mathbf{x^{=m}}) \rightarrow \mathtt{S}_i(\mathbf{x^{=m}})$ is in Σ_{12}. The result of applying this tgd to the fact $\mathtt{U}_{i,j}(\mathbf{t})$ of I_1 generates $S_i(\mathbf{t})$, as desired. Assume now that $Z_{i,j} \notin P$. Say $Z_{i,j}$ is the propositional symbol A_m. So $A_m \notin P$. Therefore, I_1 contains the fact $\mathtt{R}(\mathbf{t^m})$. Since $Z_{i,j}$ is a disjunct of C_i, the tgd $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x^m})$ is in Σ_{12}. Then the result of applying this tgd to the fact $\mathtt{R}(\mathbf{t^m})$ of I_1 generates $S_i(\mathbf{t})$, as desired.

Consider now the result of applying the tgd $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x^m})$ to $I_\mathbf{t}$. The result is $\mathtt{S}_i(\mathbf{t^m})$. We must show that $S_i(\mathbf{t^m})$ is in $\mathit{chase}_{12}(I_1)$. Since the clause C_i is satisfied by the truth assignment $\mathbf{Tr}$, either there is some j such that $Y_{i,j} \in P$, or there is some j such that $Z_{i,j} \notin P$. Assume first that $Y_{i,j} \in P$. Say $Y_{i,j}$ is the propositional symbol A_m. So $A_m \in P$. Therefore, $t_{2m-1} = t_{2m}$. Since $Y_{i,j}$ is a disjunct of C_i, the tgd $\mathtt{U}_{i,j}(\mathbf{x^{=m}}) \rightarrow \mathtt{S}_i(\mathbf{x^{=m}})$ is in Σ_{12}. Since $t_{2m-1} = t_{2m}$, the result of applying this tgd to the fact $\mathtt{U}_{i,j}(\mathbf{t})$ of I_1 generates $S_i(\mathbf{t})$, which is the same as the fact $S_i(\mathbf{t^m})$, as desired. Assume now that $Z_{i,j} \notin P$. Say $Z_{i,j}$ is the propositional symbol A_m. So $A_m \notin P$. Therefore, I_1 contains the fact $\mathtt{R}(\mathbf{t^m})$. Then the result of applying the tgd $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x})$ to the fact $\mathtt{R}(\mathbf{t^m})$ of I_1 generates $S_i(\mathbf{t^m})$, as desired. This concludes the proof that if φ is satisfiable, then there is no global inverse of $(\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$.

We now prove that, if φ is not satisfiable, then there is a schema mapping specified by a finite set of s-t tgds that is a global inverse of $\mathcal{M}_{12}$. Assume that φ is not satisfiable. Let Σ_{21} contain precisely the copying tgds $\mathtt{U}'_{i,j}(\mathbf{x}) \rightarrow \widehat{\mathtt{U}}_{i,j}(\mathbf{x})$ for each $U_{i,j}$ in $\mathbf{S_1}$ and the tgd $S_1(\mathbf{x}) \wedge \cdots \wedge S_k(\mathbf{x}) \rightarrow \widehat{\mathtt{R}}(\mathbf{x})$. Let $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$. We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. Assume not; we shall derive a contradiction.

Since by assumption $\mathcal{M}_{21}$ is not a global inverse of $\mathcal{M}_{12}$, it follows from Corollary 7.4 that there is a ground instance I such that $\widehat{I} \neq chase_{21}(chase_{12}(I))$. Because of the copying tgds $\mathtt{U}_{i,j}(\mathbf{x}) \rightarrow \mathtt{U}'_{i,j}(\mathbf{x})$ of Σ_{12} and the copying tgds $\mathtt{U}'_{i,j}(\mathbf{x}) \rightarrow \widehat{\mathtt{U}}_{i,j}(\mathbf{x})$ of Σ_{21}, and because of the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ (for $1 \leq i \leq k$) of Σ_{12} and the tgd $S_1(\mathbf{x}) \wedge \cdots \wedge S_k(\mathbf{x}) \rightarrow \widehat{\mathtt{R}}(\mathbf{x})$ of Σ_{21}, it follows that $\widehat{I} \subseteq chase_{21}(chase_{12}(I))$. So $\widehat{I}$ is a proper subset of $chase_{21}(chase_{12}(I))$. It is clear that every fact $\widehat{\mathtt{U}_{i,j}}(\mathbf{t})$ in $chase_{21}(chase_{12}(I))$ is in $\widehat{I}$. Hence, there is $\mathbf{t}$ such that the fact $\widehat{\mathtt{R}}(\mathbf{t})$ is in $chase_{21}(chase_{12}(I))$ but not in $\widehat{I}$. Now the only way that this fact can arise in $chase_{21}(chase_{12}(I))$ is for $S_i(\mathbf{t})$ to be a fact of $chase_{12}(I)$, for $1 \leq i \leq k$. Since $\widehat{\mathtt{R}}(\mathbf{t})$ is not a fact of $\widehat{I}$, it follows that $R(\mathbf{t})$ is not a fact of I. Therefore, the only way that $S_i(\mathbf{t})$ can be a fact of $chase_{12}(I)$ is for it to arise as a result of applying one of the tgds we denoted by $y_{i,j}$ or one of the tgds we denoted by $z_{i,j}$. Let $\mathcal{Y}_{i,j}$ be the event that $S_i(\mathbf{t})$ arises in $chase_{12}(I)$ as a result of applying the tgd $y_{i,j}$, and let $\mathcal{Z}_{i,j}$ be the event that $S_i(\mathbf{t})$ arises in $chase_{12}(I)$ as a result of applying the tgd $z_{i,j}$. From what we have said, it follows that for each i (with $1 \leq i \leq k$), either there is j such that the event $\mathcal{Y}_{i,j}$ holds, or there is j such that the event $\mathcal{Z}_{i,j}$ holds.

Assume that $\mathbf{t} = (t_1, \ldots, t_{2n})$. Define a truth assignment $\mathbf{Tr}$ to the propositional symbols $A_1, \ldots, A_n$ by letting A_m be assigned true if and only if $t_{2m-1} = t_{2m}$. Let us temporarily fix i (for $1 \leq i \leq k$). We now show that the clause C_i is true under $\mathbf{Tr}$. Recall that C_i is $Y_{i,1} \vee \cdots \vee Y_{i,f(i)} \vee \neg Z_{i,1} \vee \cdots \vee \neg Z_{i,g(i)}$, where each $Y_{i,j}$ and each $Z_{i,j}$ is in $\{A_1, \ldots, A_n\}$. We know that either there is j such that the event $\mathcal{Y}_{i,j}$ holds, or there is j such that the event $\mathcal{Z}_{i,j}$ holds. Assume that the event $\mathcal{Y}_{i,j}$ holds. Now $Y_{i,j}$ is one of $A_1, \ldots, A_n$. Assume that $Y_{i,j}$ is A_m. Then $y_{i,j}$ is the tgd $\mathtt{U}_{i,j}(\mathbf{x}^{=\mathbf{m}}) \rightarrow \mathtt{S}_i(\mathbf{x}^{=\mathbf{m}})$. Since the event $\mathcal{Y}_{i,j}$ holds, that is, $S_i(\mathbf{t})$ arises in $chase_{12}(I)$ as a result of applying the tgd $y_{i,j}$, necessarily $t_{2m-1} = t_{2m}$, and so A_m is true under $\mathbf{Tr}$. Therefore, $Y_{i,j}$ is true under $\mathbf{Tr}$, and so C_i is true under $\mathbf{Tr}$, as desired. Assume now that the event $\mathcal{Z}_{i,j}$ holds. Now $Z_{i,j}$ is one of $A_1, \ldots, A_n$. Assume that $Z_{i,j}$ is A_m. Then $z_{i,j}$ is the tgd $\mathtt{R}(\mathbf{x}) \rightarrow \mathtt{S}_i(\mathbf{x}^{\mathbf{m}})$. Now the event $\mathcal{Z}_{i,j}$ holds, that is, $S_i(\mathbf{t})$ arises in $chase_{12}(I)$ as a result of applying the tgd $z_{i,j}$. To obtain $S_i(\mathbf{t})$ by chasing with the tgd $z_{i,j}$, this tgd must be applied to the fact $\mathtt{R}(\mathbf{t}^{\mathbf{m}})$, which must therefore be in I. Since $\mathtt{R}(\mathbf{t}^{\mathbf{m}})$ is in I but $\mathtt{R}(\mathbf{t})$ is not in I, it follows that $t_{2m-1} \neq t_{2m}$, and so A_m is false under $\mathbf{Tr}$. Therefore, $Z_{i,j}$ is false under $\mathbf{Tr}$, and so C_i is true under $\mathbf{Tr}$, as desired.

We have shown that the clause C_i is true under the truth assignment $\mathbf{Tr}$. Since i is arbitrary, it follows that φ, which is the conjunction of the clauses C_i, is true under $\mathbf{Tr}$. But this is impossible, since by assumption φ is not satisfiable. □

Although the focus in this article is on schema mappings and their inverses when both are specified by a finite set of s-t tgds, we note that the proof of Theorem 14.9 gives us a hardness result even when we do not restrict possible inverses to be specified by a finite set of s-t tgds. Specifically, we have the following corollary of the proof.

COROLLARY 14.10. *The problem of deciding, given a schema mapping* $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ *with* Σ_{12} *a finite set of full s-t tgds, whether* $\mathcal{M}_{12}$ *has a global inverse (not necessarily specified by a finite set of s-t tgds) is coNP-hard. It is coNP-hard* *even if Σ_{12} is also LAV, that is, the members of Σ_{12} are full tgds each of which has a singleton premise.*

PROOF. In the proof of coNP-hardness in the proof of Theorem 14.9, we showed, given a propositional formula φ in conjunctive normal form, how to construct in polynomial time a LAV schema mapping $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$, where the s-t tgds in Σ_{12} are full, with the following properties: (1) if φ is satisfiable, then there is no global inverse of $\mathcal{M}_{12}$, and (2) if φ is not satisfiable, then there is a schema mapping specified by a finite set of s-t tgds that is a global inverse of $\mathcal{M}_{12}$. We concluded that $\mathcal{M}_{12}$ has a global inverse specified by a finite set of s-t tgds if and only if φ is not satisfiable. But we can also conclude that $\mathcal{M}_{12}$ has a global inverse (not necessarily specified by a finite set of s-t tgds) if and only if φ is not satisfiable. Since SAT is NP-complete, this latter result is sufficient to prove coNP-hardness of the problem stated in the corollary. □

We now give two technical lemmas that will be useful in establishing the complexity of the problem of deciding, given two schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each specified by a finite set of s-t tgds, whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$,

LEMMA 14.11. *Assume that $\mathcal{M}'_{12} = (\mathbf{S}'_1, \mathbf{S}'_2, \Sigma'_{12})$ and $\mathcal{M}''_{12} = (\mathbf{S}''_1, \mathbf{S}''_2, \Sigma''_{12})$ are schema mappings, where $\mathbf{S}'_1$, $\mathbf{S}'_2$, $\mathbf{S}''_1$, $\mathbf{S}''_2$ are pairwise disjoint, and where Σ'_{12} and Σ''_{12} are finite sets of s-t tgds. Let $\mathbf{S_1} = \mathbf{S}'_1 \cup \mathbf{S}''_1$ and $\mathbf{S_2} = \mathbf{S}'_2 \cup \mathbf{S}''_2$. Let $\Sigma_{12} = \Sigma'_{12} \cup \Sigma''_{12}$ and $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$. Let $I', I,'' J', J''$ be instances of $\mathbf{S}'_1, \mathbf{S}''_1, \mathbf{S}'_2, \mathbf{S}''_2$, respectively. Then $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma'_{12} \cup \Sigma''_{12}$ if and only if $\langle I', J'\rangle \models \Sigma'_{12}$ and $\langle I'', J''\rangle \models \Sigma''_{12}$.*

PROOF. Clearly $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma'_{12} \cup \Sigma''_{12}$ if and only if $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma'_{12}$ and $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma''_{12}$. Since Σ'_{12} does not contain any relational symbols in $\mathbf{S}''_1$ or $\mathbf{S}''_2$, it follows easily that $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma'_{12}$ if and only if $\langle I', J'\rangle \models \Sigma'_{12}$. Similarly, $\langle (I' \cup I''), (J' \cup J'')\rangle \models \Sigma''_{12}$ if and only if $\langle I'', J''\rangle \models \Sigma''_{12}$. The lemma then follows. □

LEMMA 14.12. *Assume that $\mathcal{M}'_{12} = (\mathbf{S}'_1, \mathbf{S}'_2, \Sigma'_{12})$, $\mathcal{M}'_{21} = (\mathbf{S}'_2, \widehat{\mathbf{S}'_1}, \Sigma'_{21})$, $\mathcal{M}''_{12} = (\mathbf{S}''_1, \mathbf{S}''_2, \Sigma''_{12})$, and $\mathcal{M}''_{21} = (\mathbf{S}''_2, \widehat{\mathbf{S}''_1}, \Sigma''_{21})$, are schema mappings, where $\mathbf{S}'_1$, $\mathbf{S}'_2$, $\widehat{\mathbf{S}'_1}$, $\mathbf{S}''_1$, $\mathbf{S}''_2$, and $\widehat{\mathbf{S}''_1}$ are pairwise disjoint, and where Σ'_{12}, Σ'_{21}, Σ''_{12}, and Σ''_{21} are finite sets of s-t tgds. Let $\mathbf{S_1} = \mathbf{S}'_1 \cup \mathbf{S}''_1$ and $\mathbf{S_2} = \mathbf{S}'_2 \cup \mathbf{S}''_2$. Let $\Sigma_{12} = \Sigma'_{12} \cup \Sigma''_{12}$ and $\Sigma_{21} = \Sigma'_{21} \cup \Sigma''_{21}$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, Then $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if $\mathcal{M}'_{21}$ is a global inverse of $\mathcal{M}'_{12}$ and $\mathcal{M}''_{21}$ is a global inverse of $\mathcal{M}''_{12}$.*

PROOF. Each instance I of $\mathbf{S_1}$ can be written uniquely in the form $I' \cup I''$, where I' is an instance of $\mathbf{S}'_1$ and I'' is an instance of $\mathbf{S}''_1$. Similarly, each instance J of $\mathbf{S_2}$ can be written uniquely in the form $J' \cup J''$, where J' is an instance of $\mathbf{S}'_2$ and J' is an instance of $\mathbf{S}''_2$. We now show that

$$\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21} \text{ if and only if } \langle I', J'\rangle \models \Sigma'_{12} \circ \Sigma'_{21} \text{ and } \langle I'', J''\rangle \models \Sigma''_{12} \circ \Sigma''_{21}. \quad (18)$$

Now $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$ if and only if there is an instance J_1 of $\mathbf{S_2}$ such that $\langle I, J_1\rangle \models \Sigma_{12}$ and $\langle J_1, J\rangle \models \Sigma_{21}$. This holds if and only if there is an instance J_1' of $\mathbf{S_2'}$ and an instance J_1'' of $\mathbf{S_2''}$ such that $\langle I, (J_1' \cup J_1'')\rangle \models \Sigma_{12}$ and $\langle (J_1' \cup J_1''), J\rangle \models \Sigma_{21}$.

By Lemma 14.11, we know that $\langle I, (J_1' \cup J_1'')\rangle \models \Sigma_{12}$ if and only if $\langle I', J_1'\rangle \models \Sigma_{12}'$ and $\langle I'', J_1''\rangle \models \Sigma_{12}''$. Similarly, $\langle (J_1' \cup J_1''), J\rangle \models \Sigma_{21}$ if and only if $\langle J_1', J'\rangle \models \Sigma_{21}'$ and $\langle J_1'', J''\rangle \models \Sigma_{21}''$.

It follows that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$ if and only if there are J_1' and J_1'' such that $\langle I', J_1'\rangle \models \Sigma_{12}'$, $\langle I'', J_1''\rangle \models \Sigma_{12}''$, $\langle J_1', J'\rangle \models \Sigma_{21}'$ and $\langle J_1'', J''\rangle \models \Sigma_{21}''$. Then (18) follows easily.

We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$ and $\mathcal{M}_{21}''$ is a global inverse of $\mathcal{M}_{12}''$. Assume first that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$; we shall show that $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$ (a similar argument shows that $\mathcal{M}_{21}''$ is a global inverse of $\mathcal{M}_{12}''$). Since $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, it follows that (6) holds. We must show that

$$\langle I', J'\rangle \models \Sigma_{12}' \circ \Sigma_{21}' \text{ if and only if } \widehat{I'} \subseteq J'. \tag{19}$$

Let I'', J_1'', and J'' be empty instances. Let $I = I'$ and $J = J'$. Then $\langle I'', J_1''\rangle \models \Sigma_{12}''$ and $\langle J_1'', J''\rangle \models \Sigma_{21}''$. Hence, $\langle I'', J''\rangle \models \Sigma_{12}'' \circ \Sigma_{21}''$. So from (18) we see that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$ if and only if $\langle I', J'\rangle \models \Sigma_{12}' \circ \Sigma_{21}'$. Furthermore, $\widehat{I} \subseteq J$ if and only if $\widehat{I'} \subseteq J'$, since $I = I'$ and $J = J'$. From (6), it therefore follows that (19) holds, as desired.

Assume now that $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$, and $\mathcal{M}_{21}''$ is a global inverse of $\mathcal{M}_{12}''$; we shall show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. Thus, we shall show that (6) holds.

Assume that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$. Let I' and I'' be the facts of I involving $\mathbf{S_1'}$ and $\mathbf{S_1''}$, respectively. Similarly, let J' and J'' be the facts of J involving $\mathbf{S_2'}$ and $\mathbf{S_2''}$, respectively. By (18) we know that $\langle I', J'\rangle \models \Sigma_{12}' \circ \Sigma_{21}'$. Since $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$, it follows that $\widehat{I'} \subseteq J'$. Similarly, $\widehat{I''} \subseteq J''$. Since $\widehat{I} = \widehat{I'} \cup \widehat{I''}$ and $J = J' \cup J''$, it follows that $\widehat{I} \subseteq J$, as desired.

Assume that $\widehat{I} \subseteq J$. As before, let I' and I'' be the facts of I involving $\mathbf{S_1'}$ and $\mathbf{S_1''}$, respectively, and let J' and J'' be the facts of J involving $\mathbf{S_2'}$ and $\mathbf{S_2''}$, respectively. Since $\widehat{I} \subseteq J$, it follows that $\widehat{I'} \subseteq J'$. Since $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$, it follows that $\langle I', J'\rangle \models \Sigma_{12}' \circ \Sigma_{21}'$. Similarly, $\langle I'', J''\rangle \models \Sigma_{12}'' \circ \Sigma_{21}$. So by (18), it follows that $\langle I, J\rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. □

THEOREM 14.13. *The problem of deciding, given schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{12} and Σ_{21} are finite sets of full s-t tgds, whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, is DP-complete.*

PROOF. We first show that the problem is in DP. As in the definition of the canonical candidate tgd $\mathcal{S}$-inverse, let R be an arbitrary relational symbol of $\mathbf{S_1}$, and let $I_{\mathtt{R}}$ be a one-tuple instance that contains only the fact $\mathtt{R}(x_1, \ldots, x_k)$, where $x_1, \ldots, x_k$ are distinct. We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if the following two properties hold:

(*) $\widehat{I_{\mathtt{R}}} \subseteq chase_{21}(chase_{12}(I_{\mathtt{R}}))$ for every relational symbol R of $\mathbf{S_1}$.
(**) $chase_{21}(chase_{12}(I)) \subseteq \widehat{I}$ for every small instance I of $\mathbf{S_1}$.

By "small" in (**), we mean that I contains at most $k_{12}k_{21}$ facts, as in the proof of Theorem 13.1.

We now show that, if $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, then (*) and (**) hold. If (*) did not hold, then by Corollary 7.4, it would follow that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for $I_{\mathtt{R}}$. If (**) did not hold, then by Corollary 7.4, it would follow that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for some (small) instance I.

We now show that if (*) and (**) hold, then $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. Assume that, (*) and (**) hold. Let I be an arbitrary instance of $\mathbf{S_1}$. By Corollary 7.4, we need only show that $\widehat{I} = \mathit{chase}_{21}(\mathit{chase}_{12}(I))$. Let $\mathtt{R}(c_1, \ldots, c_k)$ be an arbitrary fact of I, and let I' be a one-tuple instance that contains only this fact. It follows easily from (*) that $\widehat{I'} \subseteq \mathit{chase}_{21}(\mathit{chase}_{12}(I'))$. Therefore, $\widehat{I} \subseteq \mathit{chase}_{21}(\mathit{chase}_{12}(I))$. As for the reverse inclusion, assume that it were to fail, so that $\mathit{chase}_{21}(\mathit{chase}_{12}(I))$ is not a subinstance of $\widehat{I}$. By the argument given in the proof of Case 2 of Theorem 13.1, we would have a violation of (**).

We have shown that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if (*) and (**) hold. For each $\mathtt{R}$ in $\mathbf{S_1}$, we see that (*) is an NP property (we simply guess the chase steps and verify that the result of these chase steps is not in $\widehat{I}$). So (*) is an NP property. To show that (**) is a coNP property, we must show that its negation is an NP property. The negation of (**) says that there is a small instance I such that $\mathit{chase}_{21}(\mathit{chase}_{12}(I))$ is not a subset of $\widehat{I}$. This is indeed an NP property: we simply guess the small instance I and guess the chase steps. This shows that the property that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ is a DP property.

We now show that the problem is DP-hard. Let CLIQUE be the problem where the input is an undirected graph G with no self-loops, along with an integer k, and the question is whether G has a clique of size at least k (a *clique* is a set of nodes such that there is an edge between every pair of distinct nodes). Let $\overline{\text{SAT}}$ be the problem where the input is a propositional formula φ in conjunctive normal form, and the question is whether the formula is *not* satisfiable. We shall show that, given G, k, and φ as above, we can construct, in polynomial time, schema mappings $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, where Σ_{12} and Σ_{21} are finite sets of full s-t tgds, such that $\mathcal{M}_{12}$ is a global inverse of $\mathcal{M}_{12}$ if and only if (a) G has a clique of size at least k and (b) φ is not satisfiable. Since CLIQUE is an NP-complete problem and $\overline{\text{SAT}}$ is a coNP-complete problem, this is sufficient to prove DP-hardness.

The way our proof will proceed is that we will show the following:

(A) Given G and k, we can construct, in polynomial time, schema mappings $\mathcal{M}'_{12} = (\mathbf{S}'_1, \mathbf{S}'_2, \Sigma'_{12})$ and $\mathcal{M}'_{21} = (\mathbf{S}'_2, \widehat{\mathbf{S}'_1}, \Sigma'_{21})$, where Σ'_{12} and Σ'_{21} are finite sets of full s-t tgds, such that $\mathcal{M}'_{21}$ is a global inverse of $\mathcal{M}'_{12}$ if and only if G has a clique of size at least k.

(B) Given φ, we can construct, in polynomial time, schema mappings $\mathcal{M}''_{12} = (\mathbf{S}''_1, \mathbf{S}''_2, \Sigma''_{12})$ and $\mathcal{M}''_{21} = (\mathbf{S}''_2, \widehat{\mathbf{S}''_1}, \Sigma''_{21})$, where Σ''_{12} and Σ''_{21} are finite sets of full s-t tgds, such that $\mathcal{M}''_{21}$ is a global inverse of $\mathcal{M}''_{12}$ if and only if φ is not satisfiable.

If necessary, rename the relational symbols so that $\mathbf{S_1'}, \mathbf{S_2'}, \widehat{\mathbf{S_1'}}, \mathbf{S_1''}, \mathbf{S_2''}$, and $\widehat{\mathbf{S_1''}}$ are pairwise disjoint. Let $\mathbf{S_1} = \mathbf{S_1'} \cup \mathbf{S_1''}$ and $\mathbf{S_2} = \mathbf{S_2'} \cup \mathbf{S_2''}$. Let $\Sigma_{12} = \Sigma_{12}' \cup \Sigma_{12}''$ and $\Sigma_{21} = \Sigma_{21}' \cup \Sigma_{21}''$. Let $\mathcal{M}_{12} = (\mathbf{S_1}, \mathbf{S_2}, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S_2}, \widehat{\mathbf{S_1}}, \Sigma_{21})$, By Lemma 14.12, we know that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$ and $\mathcal{M}_{21}''$ is a global inverse of $\mathcal{M}_{12}''$. Therefore, if we fulfill (A) and (B), then it follows that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$ if and only if G has a clique of size at least k and φ is not satisfiable. So fulfilling (A) and (B) is sufficient to prove DP-hardness.

For part (B), we take $\mathcal{M}_{12}''$ and $\mathcal{M}_{21}''$ to be the schema mappings defined in the proof of Theorem 14.9 under the names $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$, respectively. We showed in that proof that $\mathcal{M}_{21}''$ is a global inverse of $\mathcal{M}_{12}''$ if and only if φ is not satisfiable.

So we need only prove part (A). Assume that the graph G has n distinct nodes $a_1, \ldots, a_n$. Let $\mathbf{S_1'}$ consist of the n-ary relational symbol D, and let $\mathbf{S_2'}$ consist of the the n-ary relational symbol D' and the binary relation symbol E. Let $x_1, \ldots, x_n$ be distinct variables, and let ψ be the conjunction of all formulas $E(x_i, x_j)$ such that (a_i, a_j) is an edge of the graph G. Let Σ_{12} consist of the copying tgd $D(x_1, \ldots, x_n) \to D'(x_1, \ldots, x_n)$ and the tgd $D(x_1, \ldots, x_n) \to \psi$. Let $y_1, \ldots, y_k$ be k distinct new variables, and let $\alpha(y_1, \ldots, y_k)$ be the conjunction of all formulas $E(y_i, y_j)$ where $1 \le i \le k$, $1 \le j \le k$, and $i \ne j$. Let Σ_{21}' consist of the tgd $D'(x_1, \ldots, x_n) \wedge \alpha(y_1, \ldots, y_k) \to \widehat{D}(x_1, \ldots, x_n)$. Note that each of the tgds in Σ_{12}' and Σ_{21}' are full. We now show that G has a clique of size at least k if and only if $\mathcal{M}_{21}'$ is a global inverse of $\mathcal{M}_{12}'$. By Corollary 7.4, we need only show that G has a clique of size at least k if and only if, for every ground instance I, we have $\widehat{I} = \mathit{chase}_{21}'(\mathit{chase}_{12}'(I))$. By chase_{12}', we mean the result of chasing with Σ_{12}', and similarly for chase_{21}'.

Note first that, if h is a homomorphism from $\alpha(y_1, \ldots, y_k)$ to the edge relation E of G, then h is one-to-one. This is because if $i \ne j$, then $(h(y_i), h(y_j))$ is necessarily an edge of G, since h is a homomorphism and since $E(y_i, y_j)$ is a conjunct of $\alpha(y_1, \ldots, y_k)$. Since G has no self-loops, it follows that $h(y_i) \ne h(y_j)$, as desired. Therefore if h is a homomorphism from $\alpha(y_1, \ldots, y_k)$ to the edge relation E of G, then G must have a clique of size at least k.

Assume that G does not have a clique of size at least k; we must show that there is a ground instance I such that $\widehat{I} \ne \mathit{chase}_{21}'(\mathit{chase}_{12}'(I))$. Let I be the ground instance with the single fact $\mathtt{D}(a_1, \ldots, a_n)$. Then the E relation of $\mathit{chase}_{12}'(I)$ is the edge relation of G. So $\mathit{chase}_{21}'(\mathit{chase}_{12}'(I))$ is empty: this is because, as we noted, if there is a homomorphism from $\alpha(y_1, \ldots, y_k)$ to the edge relation E of G, then G must have a clique of size at least k. Hence $\widehat{I} \ne \mathit{chase}_{21}'(\mathit{chase}_{12}'(I))$, as desired.

Assume that G has a clique of size at least k; we must show that for every ground instance I, we have $\widehat{I} = \mathit{chase}_{21}'(\mathit{chase}_{12}'(I))$. Let I be an arbitrary ground instance. Let $J^* = \mathit{chase}_{12}'(I)$ and let $I^* = \mathit{chase}_{21}'(J^*)$. We must show that $\widehat{I} = I^*$. Thus, we must show that $D^I = \widehat{D}^{I^*}$. Clearly $\widehat{D}^{I^*} \subseteq D'^{J^*} = D^I$, so $\widehat{D}^{I^*} \subseteq D^I$, Therefore, we need only show that $D^I \subseteq \widehat{D}^{I^*}$. Assume that $(b_1, \ldots, b_n) \in D^I$; we must show that $(b_1, \ldots, b_n) \in \widehat{D}^{I^*}$. Now $(b_1, \ldots, b_n) \in D'^{J^*}$.

Furthermore, $(b_i, b_j) \in E^{J^*}$ whenever (a_i, a_j) is an edge of G. Since G has a clique of size at least k, there is a set $C \subseteq \{1, \ldots, n\}$ of size k such that (a_i, a_j) is an edge of G whenever i and j are distinct members of C. So $(b_i, b_j) \in E^{J^*}$ whenever i and j are distinct members of C. Assume that $i_1, \ldots, i_k$ are the distinct members of C. By letting $b_1, \ldots, b_n$ play the roles of $x_1, \ldots, x_n$, respectively, and letting $b_{i_1}, \ldots, b_{i_k}$ play the roles of $y_1, \ldots y_k$, respectively, in the tgd $D'(x_1, \ldots, x_n) \wedge \alpha(y_1, \ldots, y_k) \rightarrow \widehat{D}(x_1, \ldots, x_n)$ of Σ'_{21}, it follows that $(b_1, \ldots, b_n) \in \widehat{D}^{I^*}$, as desired. □

We have now proven the complexity bounds in Table II (the coNP-hardness result in the third line is inherited from the coNP-completeness result in the fourth line).

15. SUBSEQUENT WORK

Fagin et al. [2007] is a followup to this article. Its main focus is a relaxed notion of the global inverse, called a *quasi-inverse*. In addition, Fagin et al. [2007] resolves two problems that were left open in the preliminary version [Fagin 2006] of this article. One of these problems was the "language of inverse," as we now discuss. In this article, we focus on inverses that are specified by a finite set of s-t tgds. For example, let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds, and let $\mathcal{M}'$ be the canonical candidate tgd global inverse of $\mathcal{M}$. Then $\mathcal{M}'$ is specified by a finite set of s-t tgds. Furthermore, $\mathcal{M}'$ is a global inverse of $\mathcal{M}$ if there is any global inverse of $\mathcal{M}$ that is specified by a finite set of s-t tgds. This leaves open the question as to how rich a language is needed to specify a global inverse of $\mathcal{M}$. The answer, as shown in Fagin et al. [2007], is *full tgds with constants and inequalities*, which we now define.

Let *Constant* be a new relation symbol. A *tgd (from* $\mathbf{T}$ *to* $\widehat{\mathbf{S}}$*) with constants and inequalities* is a first-order formula of the form $\forall \mathbf{x}(\varphi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$, where

—the formula $\varphi(\mathbf{x})$ is a conjunction of
  (1) atoms over $\mathbf{T}$, such that every variable in $\mathbf{x}$ occurs in at least one of them;
  (2) formulas *Constant*(x), where x is a variable in $\mathbf{x}$;
  (3) inequalities $x \neq x'$, where x and x' are variables in $\mathbf{x}$;
—The formula $\psi(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms over $\widehat{\mathbf{S}}$ (not all of the variables in $\mathbf{x}$ need to appear in $\psi(\mathbf{x}, \mathbf{y})$).

Naturally, an atomic formula *Constant*(x) evaluates to true if and only if x is interpreted by a value in Const. A tgd with constants and inequalities is *full* if there are no existential quantifiers. An example of a tgd with constants and inequalities is $\forall x \forall y \forall z ((\mathtt{P}(x, y, z) \wedge Constant(x) \wedge Constant(y) \wedge (x \neq y)) \rightarrow \exists w \mathtt{Q}(x, y, w))$.

THEOREM 15.1 [FAGIN ET AL. 2007]. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. If $\mathcal{M}$ has a global inverse then the following hold.*

1. *$\mathcal{M}$ has a global inverse $\mathcal{M}'$ specified by a finite set of full tgds with constants and inequalities.*
2. *There is an exponential-time algorithm for producing $\mathcal{M}'$.*

3. *Statement (1) is not necessarily true if we disallow either constants or inequalities in the premise, even if we allow existential quantifiers in the conclusion (and so allow nonfull dependencies to specify $\mathcal{M}'$).*

The second problem that was left open in the preliminary version [Fagin 2006] and was resolved in Fagin et al. [2007] was the question of whether the unique-solutions property is a necessary and sufficient condition for the existence of a global inverse of those schema mappings that are specified by a finite set of s-t tgds. (By Theorem 6.7, it is a necessary condition, and by Theorem 6.10, it is a necessary and sufficient condition for LAV schema mappings.) It was shown in Fagin et al. [2007] that the condition is not sufficient. Instead, it was shown that a strictly stronger condition, called the *subset property* (or, under the notation of Fagin et al. [2007], the $(=, =)$-*subset property*), is necessary and sufficient. This condition says that if I_1 and I_2 are two ground instances such that every solution for I_2 is a solution for I_1, then $I_1 \subseteq I_2$.

16. SUMMARY AND OPEN PROBLEMS

We have given a formal definition for one schema mapping to be an inverse of another schema mapping for a class $\mathcal{S}$ of ground instances. We have obtained a number of results about our notion of inverse, and some of these results are surprising.

In Section 15, we discussed two problems that were left open in the preliminary version [Fagin 2006] of this article and were resolved in Fagin et al. [2007]. There are still many open problems, as we would expect from a "first step" article like this.

— In Section 14, there are two complexity gaps (in line 3 of Table I and in line 3 of Table II). These complexity gaps deal with inverses that are specified by a finite set of s-t tgds. Furthermore, in Corollary 14.10, it was shown that the problem of deciding if a schema mapping specified by a finite set of s-t tgds has a global inverse (not necessarily specified by a finite set of s-t tgds) is coNP-hard. As we noted in Section 3, it is not even known whether this problem is undecidable! This is perhaps the most natural and interesting remaining open problem. Another problem, as mentioned also in Section 3, is to find a natural class of schema mappings that arise in practice where this complexity is greatly reduced.
— We have focused most of our attention on schema mappings $\mathcal{M}_{12}$ specified by a finite set of s-t tgds. What about more general schema mappings than those specified by s-t tgds? What if we allow target dependencies, such as functional dependencies?
— We have focused on right inverses, where we are given $\mathcal{M}_{12}$ and want to find a right inverse $\mathcal{M}_{21}$. It might be interesting to study the left inverse, where we are given $\mathcal{M}_{21}$ and we wish to find $\mathcal{M}_{12}$.
— In Example 5.5 we showed an example of a schema mapping specified by s-t tgds that has a unique inverse specified by s-t tgds. It might be interesting

to examine the question of when there is a unique inverse mapping specified in a given language.

—Our next open problem is somewhat imprecise, but is important in practice. Assume that we are given $\mathcal{M}_{12}$. How do we find a large class $\mathcal{S}$ and a schema mapping $\mathcal{M}_{21}$ such that $\mathcal{M}_{21}$ is an $\mathcal{S}$-inverse of $\mathcal{M}_{12}$? In fact, there might be several such large classes $\mathcal{S}$ and corresponding inverse mappings. How do we find them? This problem is imprecise, because it is not clear what we mean by a "large class" $\mathcal{S}$. We should not necessarily restrict our attention to classes $\mathcal{S}$ defined by a finitely chasable set Γ of tgds and egds.

—It might be interesting to explore more fully the unique-solutions property, which is an interesting notion in its own right.

—We might explore the property of $\widehat{I}$ and $chase_{21}(chase_{12}(I))$ being homomorphically equivalent. By Propositions 7.5 and 7.6, this notion is implied by but strictly weaker than $\mathcal{M}_{21}$ being an inverse of $\mathcal{M}_{12}$ for I.

This article is, we think, simply the first step in a fascinating journey!

ACKNOWLEDGMENTS

The author is grateful to Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan for numerous invaluable discussions and suggestions. The author is also grateful to Catriel Beeri for simplifying the proof of Theorem 6.10, and to Alan Nash for helpful comments.

REFERENCES

ARENAS, M. 2006. Private communication.

BEERI, C. AND VARDI, M. Y. 1984. A proof procedure for data dependencies. *J. Assoc. Comput. Mach. 31,* 4, 718–741.

BERNSTEIN, P. A. 2003. Applying model management to classical meta-data problems. In *Proceedings of the Conference on Innovative Data Systems Research* (CIDR). 209–220.

FAGIN, R. 1974. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of Computation, SIAM-AMS Proceedings*, Vol. 7, R. M. Karp, Ed. SIAM, Philadelphia, PA, 43–73.

FAGIN, R. 1977. Multivalued dependencies and a new normal form for relational databases. *ACM Trans. Database Syst. 2,* 3 (Sept.), 262–278.

FAGIN, R. 2006. Inverting schema mappings. In *Proceedings of the ACM Symposium on Principles of Database Systems* (PODS). 50–59.

FAGIN, R., KOLAITIS, P. G., MILLER, R. J., AND POPA, L. 2005a. Data exchange: Semantics and query answering. *Theoret. Comput. Sci.* 89–124.

FAGIN, R., KOLAITIS, P. G., AND POPA, L. 2005b. Data exchange: Getting to the core. *ACM Trans. Database Sys. 30,* 1, 174–210.

FAGIN, R., KOLAITIS, P. G., POPA, L., AND TAN, W.-C. 2005c. Composing schema mappings: Second-order dependencies to the rescue. *ACM Trans. Database Sys. 30,* 4, 994–1055.

FAGIN, R., KOLAITIS, P. G., POPA, L., AND TAN, W.-C. 2007. Quasi-inverses of schema mappings. In *Proceedings of the ACM Symposium on Principles of Database Systems* (PODS). 123–132.

HULL, R. 1986. Relative information capacity of simple relational database schemata. *SIAM J. Comput. 15*, 856–886.

LENZERINI, M. 2002. Data integration: A theoretical perspective. In *Porceedings of the ACM Symposium on Principles of Database Systems* (PODS). 233–246.

MELNIK, S. 2004. *Generic Model Management: Concepts and Algorithms*. Springer, Berlin, Germany.

MELNIK, S., BERNSTEIN, P., HALEVY, A., AND RAHM, E. 2005. Supporting executable mappings in model management. In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (SIGMOD). 167–178.

SHU, N. C., HOUSEL, B. C., TAYLOR, R. W., GHOSH, S. P., AND LUM, V. Y. 1977. EXPRESS: A data EXtraction, Processing, and REStructuring System. *ACM Trans. Database Syst. 2,* 2, 134–174.

YU, C. AND POPA, L. 2005. Semantic adaptation of schema mappings when schemas evolve. In *Proceedings of the International Conference on Very Large Data Bases* (VLDB). 1006–1017.

Received September 2006; revised April 2007; accepted July 2007