

Inverting Schema Mappings

Ronald Fagin
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120
fagin@almaden.ibm.com

ABSTRACT

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 instances, in which case an $\mathcal{S}$ -inverse is a global inverse. We focus on the important and practical case of schema mappings defined by source-to-target tuple-generating dependencies, and uncover a rich theory. When $\mathcal{S}$ is defined 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 [Heterogeneous Databases]: Data translation; H.2.4 [Systems]: Relational databases

General Terms: Algorithms, theory

Keywords: Data exchange, chase, computational complexity, data integration, dependencies, inverse, metadata model management, schema mapping, second-order logic

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 [13] that arises in many tasks where data must be 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 [2] has introduced such a framework, called *model management*. Melnik et al. [12] 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 [5, 6, 7], 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 schemas 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 *defines* the schema $\mathcal{M}$. As in [5, 6, 7], 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 atomic formulas over $\mathbf{S}$, and where $\psi(\mathbf{x}, \mathbf{y})$ is a conjunction of atomic formulas over $\mathbf{T}$.¹ They have been used to formalize data exchange [5]. They have also been used in data integration scenarios under the name of GLAV (global-and-local-as-view) assertions [10].

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 seminal research on information capacity of relational database schemas [9]. Although our schema mappings are not actually functions, they have the advantage

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

PODS’06, June 26–28, 2006, Chicago, Illinois, USA.
Copyright 2006 ACM 1-59593-318-2/06/0003 ...\$5.00.

¹There 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 .

of being simpler and more flexible. In fact, LAV mappings, which have been widely used in data integration, are special cases of schema mappings defined by s-t tgds, where the left-hand side of each tgd is a single atomic formula rather than a conjunction of atomic formulas.

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 $\langle I, J \rangle$ satisfy 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 [11] and Melnik et al. [12] give for the inverse. In those papers, 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 defined 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 $\langle I', J' \rangle \in S_{12}$.² 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 defined by a set of s-t tgds.

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 defined 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 defined 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. [7], 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 paper is the first step in exploring the very rich theory that arises.

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 equiv-

alent 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 defined by a finite set of s-t tgds, and I is a source instance. We derive a canonical local inverse, which is a schema mapping defined by a finite set of s-t tgds that is an inverse of $\mathcal{M}$ for I if there is any such inverse. If $\mathcal{S}$ is a class of source instances defined by a set of tgds and egds that always have a finite chase, then we derive a *canonical $\mathcal{S}$ -inverse*, which is a schema mapping defined 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 $\mathcal{S}$ is the class of all source instances, we refer to the canonical $\mathcal{S}$ -inverse as the *canonical global inverse*. On the face of it, the canonical local inverse seems to be 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 local inverse plays a key role in the proof of correctness of the canonical $\mathcal{S}$ -inverse.

Our canonical inverses are each defined 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 defined 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 defined 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 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 defined 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 defined by a finite set of s-t tgds, an analyst might want to investigate under what conditions $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$. (We give an example later, where $\mathcal{M}_{12}$ does a projection 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 [3] 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 defines the composition of the schema mappings. Furthermore, when $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are defined by *full* s-t tgds, this formula is first-order.

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

The proofs of theorems whose proofs are not in-line appear in the appendix.

1.1 Applications of inverse mappings

There are potentially a number of applications for inverse mappings, especially in schema evolution. For example, as-

²This is why a schema mapping may associate many target instances with each source instance, and many source instances with each target instance.

sume 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 $\mathbf{T}'_1$).

2. BACKGROUND

We now review basic concepts from data exchange.

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

Let $\mathbf{S} = \langle \mathbf{S}_1, \dots, \mathbf{S}_n \rangle$ and $\mathbf{T} = \langle \mathbf{T}_1, \dots, \mathbf{T}_m \rangle$ 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 $\mathbf{S}_i^K = \mathbf{S}_i^I$ and $\mathbf{T}_j^K = \mathbf{T}_j^J$, for $1 \leq i \leq n$ and $1 \leq j \leq m$.

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 $\sigma \in \Sigma$.

We will often drop the universal quantifiers in front of a *tg*d, and implicitly assume such quantification. However, we will write down all existential quantifiers.

Given a tuple $(t_1, \dots, t_r)$ occurring in a relation $\mathbf{R}$, we denote by $\mathbf{R}(t_1, \dots, t_r)$ the association between $(t_1, \dots, t_r)$ and $\mathbf{R}$, and call it a *fact*. We will identify an instance with its set of facts. We call each t_i in the tuple $(t_1, \dots, t_r)$ a *value*. We denote by $\mathbf{Const}$ the set of all values that appear in source instances (instances of the schema $\mathbf{S}$) and we call them *constants*. In addition, we assume an infinite set $\mathbf{Var}$ of values, which we call *nulls*, such that $\mathbf{Var} \cap \mathbf{Const} = \emptyset$.

If K is an instance with values in $\mathbf{Const} \cup \mathbf{Var}$, then $\mathbf{Var}(K)$ 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 $\mathbf{Const} \cup \mathbf{Var}$. A *homomorphism* $h : K_1 \rightarrow K_2$ is a mapping from $\mathbf{Const} \cup \mathbf{Var}(K_1)$ to $\mathbf{Const} \cup \mathbf{Var}(K_2)$ such that: (1) $h(c) = c$, for every $c \in \mathbf{Const}$; and (2) for every fact $\mathbf{R}(\mathbf{t})$ of K_1 , we have that $\mathbf{R}(h(\mathbf{t}))$ is a fact of K_2 (where, if $\mathbf{t} = (t_1, \dots, t_s)$, then $h(\mathbf{t}) = (h(t_1), \dots, h(t_s))$).

Consider a schema mapping $(\mathbf{S}, \mathbf{T}, \Sigma)$, as defined in the introduction. Recall that if I is a source instance, and J is a target instance, then J is a *solution* for I if $\langle I, J \rangle \models \Sigma$. If I is a source 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 $h : J \rightarrow J'$. When Σ is a finite set

of s-t *tg*ds, and I is a source instance, then there is always a universal solution for I [5].

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* [7], 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 is proven in [7] that when Σ_{12} and Σ_{23} are finite sets of s-t *tg*ds, then the composition formula is given by a *second-order tg*d (*SO tg*d). We give the definition of SO *tg*ds later (Definition 10.1). We now give an example (from [7]) of an SO *tg*d that defines the composition formula.

EXAMPLE 2.1. 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 $\mathbf{Emp}$ of employees. Schema $\mathbf{S}_2$ consists of a single binary relation symbol $\mathbf{Mgr}_1$, that associates each employee with a manager. Schema $\mathbf{S}_3$ consists of a similar binary relation symbol $\mathbf{Mgr}$, that is intended to provide a copy of $\mathbf{Mgr}_1$, and an additional unary relation symbol $\mathbf{SelfMgr}$, 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 *tg*d $\forall e (\mathbf{Emp}(e) \rightarrow \exists m \mathbf{Mgr}_1(e, m))$, and Σ_{23} consists of the two *tg*ds $\forall e \forall m (\mathbf{Mgr}_1(e, m) \rightarrow \mathbf{Mgr}(e, m))$ and $\forall e (\mathbf{Mgr}_1(e, e) \rightarrow \mathbf{SelfMgr}(e))$. Then the composition formula $\Sigma_{12} \circ \Sigma_{23}$ is defined by the following second-order *tg*d:

$$\begin{aligned} & \exists f (\forall e (\mathbf{Emp}(e) \rightarrow \mathbf{Mgr}(e, f(e))) \wedge \\ & \forall e (\mathbf{Emp}(e) \wedge (e = f(e)) \rightarrow \mathbf{SelfMgr}(e))). \quad \square \end{aligned} \quad (1)$$

3. 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 $\mathbf{R}$ of $\mathbf{S}_1$, let $\widehat{\mathbf{R}}$ be a new relation symbol (different from any relation symbol in $\mathbf{S}_1$ or $\mathbf{S}_2$) of the same arity as $\mathbf{R}$. Define $\widehat{\mathbf{S}}_1$ to be $\{\widehat{\mathbf{R}} : \mathbf{R} \in \mathbf{S}_1\}$. 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, $\widehat{I}^{\widehat{\mathbf{R}}} = \mathbf{R}^I$ for every $\mathbf{R}$ 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 $\mathcal{M}_{12}$, where the target schema is $\widehat{\mathbf{S}}_1$, and where Σ_{21} is a finite set of s-t *tg*ds (with source $\mathbf{S}_2$ and target $\widehat{\mathbf{S}}_1$). 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 *tg*ds. We now introduce an example that we shall use as a running example to demonstrate some of the issues that arise.

EXAMPLE 3.1. Let $\mathbf{S}_1$ consist of the ternary relation symbol $\mathbf{EDL}$ (“Employee-Department-Location”). Let $\mathbf{S}_2$ consist of the binary relation symbol $\mathbf{ED}$ (“Employee-Department”) and the binary relation symbol $\mathbf{DL}$ (“Department-Location”). Let Σ_{12} consist of the s-t *tg*d $\mathbf{EDL}(x, y, z) \rightarrow \mathbf{ED}(x, y) \wedge \mathbf{DL}(y, z)$, that corresponds to projecting $\mathbf{EDL}$ onto $\mathbf{ED}$ and $\mathbf{DL}$. Let Σ_{21} consist of the s-t *tg*d $(\mathbf{ED}(x, y) \wedge \mathbf{DL}(y, z)) \rightarrow \mathbf{EDL}(x, y, z)$, 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 dependency³

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

It is known [4] 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. 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 source instances I that satisfy Γ . $\square$

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 3.1, the class $\mathcal{S}$ would consist of those source instances that satisfy Γ .) Define Σ_{Id} (where Id stands for “identity”) to consist of the tgds $\mathbf{R}(x_1, \dots, x_k) \rightarrow \widehat{\mathbf{R}}(x_1, \dots, x_k)$, where $x_1, \dots, x_k$ are distinct variables, when $\mathbf{R}$ 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 the same source schema and the same target schema are *equivalent on I* if they have the same solutions for I .

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 an 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 I precisely if for every J ,

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

If $\mathcal{S}$ is a class of source 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 source instances. In that case, we say that $\mathcal{M}_{21}$ is a *global inverse* of $\mathcal{M}_{12}$.

EXAMPLE 3.2. Let us return to Example 3.1. We said there 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 source 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 10.7, we shall show that Γ is also a necessary condition.

If we apply the composition algorithm of [7], we find that the composition formula $\Sigma_{12} \circ \Sigma_{21}$, which we denote by σ , is

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

Let I be a source instance of $\mathbf{S}_1$ satisfying Γ . We must show that (3) holds. Assume first that $\langle I, J \rangle \models \sigma$; we must

³Note that Γ is not an s-t tgd, since the left-hand side and right-hand side use the same relation symbol EDL. Of course, Γ is a tgd in the classical sense of [1].

show that $\widehat{I} \subseteq J$. Now Σ_{Id} consists of the tgd $\text{EDL}(x, y, z) \rightarrow \widehat{\text{EDL}}(x, y, z)$. 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 $\langle I, J \rangle \models \sigma$. Thus, we must show that if $\text{EDL}(x, y, z')$ and $\text{EDL}(x', y, z)$ hold in I , then $\widehat{\text{EDL}}(x, y, z)$ holds in J . So assume that $\text{EDL}(x, y, z')$ and $\text{EDL}(x', y, z)$ hold in I . Since $I \models \Gamma$, it follows that $\text{EDL}(x, y, z)$ holds in I . Since $\widehat{I} \subseteq J$, it follows that $\widehat{\text{EDL}}(x, y, z)$ 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 10.7). $\square$

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 3.3. Let $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$, where $\mathbf{S}_1$ consists of the unary relation symbol $\mathbf{R}$, where $\mathbf{S}_2$ consists of the binary relation symbol $\mathbf{S}$, and where Σ_{12} consists of the tgd $\mathbf{R}(x) \rightarrow \mathbf{S}(x, x)$. Let Σ_{21} consist of the tgd $\mathbf{S}(x, y) \rightarrow \mathbf{R}(x)$, and let Σ'_{21} consist of the tgd $\mathbf{S}(x, y) \rightarrow \widehat{\mathbf{R}}(y)$. 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 $\mathbf{R}(x) \rightarrow \widehat{\mathbf{R}}(x)$, which defines the identity mapping. So both $\mathcal{M}_{21}$ and $\mathcal{M}'_{21}$ are global inverses of $\mathcal{M}_{12}$. $\square$

4. THE UNIQUE SOLUTIONS PROPERTY

Unlike the rest of this paper, 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 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 source 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 source 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 $\mathcal{M}_{12}$. The next theorem says that this intuition is correct.

THEOREM 4.1. *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 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 $\mathcal{M}_{12}$.*

As a corollary of Theorem 4.1, we obtain a necessary condition for $\mathcal{M}_{12}$ to have an inverse for a fixed source instance. (The proof depends on our assumption of preservation under isomorphism.)

COROLLARY 4.2. *Let $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ be schema mappings, and let I_1 and I_2 be distinct but isomorphic source 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 $\mathcal{M}_{12}$.

The next corollary, which we shall find quite useful later, applies to schema mappings defined by tgds, and makes 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 $\text{chase}_{12}(I)$ to be J .⁴ We may say loosely that J is the result of chasing I with Σ_{12} .

COROLLARY 4.3. *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 defined by s-t tgds), then every value that appears in a tuple of I appears in a tuple of $\text{chase}_{12}(I)$.*

The next proposition is an interesting application of Theorem 4.1.

PROPOSITION 4.4. *There is a schema mapping defined by a finite set of full s-t tgds that has an inverse for every source instance with a schema mapping defined by a finite set of s-t tgds, but has no global inverse.*

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 $\Sigma_{12} = \{P(x) \wedge Q(y) \rightarrow R(x, y), P(x) \rightarrow S(x), Q(x) \rightarrow S(x)\}$. Let $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$.

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

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

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 4.1 that $\mathcal{M}_{12}$ does not have a global inverse. $\square$

We now give a simple example of the use of Corollary 4.2.

EXAMPLE 4.5. 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 $\Sigma_{12} = \{R(x) \rightarrow S(x), R'(x) \rightarrow S(x)\}$. 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 source 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 source instances with the same set of solutions, it follows from Corollary 4.2 that $\mathcal{M}_{12}$ does not have an inverse for I_1 . $\square$

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 source

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 [5] 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 solutions property*, which says that whenever I_1 and I_2 are distinct source instances, then no universal solution for I_1 is a universal solution for I_2 .

Theorem 4.1 implies that the unique solutions property is a necessary condition for global invertibility. 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 left-hand side. 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.

THEOREM 4.6. *A LAV schema mapping has a global inverse if and only if it has the unique solutions property.*

The schema mapping that is a global inverse in our proof of Theorem 4.6 is rather complex (it is not defined in terms of tgds). For the rest of this paper, we shall consider only “practical” schema mappings—specifically, schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each defined by a finite set of s-t tgds.

5. CHARACTERIZING INVERTIBILITY

In this section, we give useful characterizations, in terms of the chase, of invertibility.

For the next theorem, 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 5.1. *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. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$.*

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

COROLLARY 5.2. *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 full s-t tgds. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\hat{I} = \text{chase}_{21}(\text{chase}_{12}(I))$.*

The next result⁵ gives a version of Corollary 5.2 that holds even when the tgds are not full. 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 .

COROLLARY 5.3. *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. If $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I , then $\hat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ are homomorphically equivalent.*

⁴For definiteness, we use the version of the chase as defined in [7], although it does not really matter.

⁵This result is due to Lucian Popa.

The next theorem implies the falsity of the converse of Corollary 5.3, that if $\hat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ are homomorphically equivalent, then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I . In fact, the negative result we prove is even stronger: it says that even if $\hat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ 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 .

THEOREM 5.4. *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 $\hat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ 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 $\mathbf{S}_1$.*

6. THE CANONICAL LOCAL INVERSE

Let $\mathcal{M}$ be a schema mapping defined by a finite set of s-t tgds, and let I be a source instance. In this section, we give a schema mapping that is guaranteed to be an inverse of $\mathcal{M}$ for I if there is any inverse at all that is defined by a finite set of s-t tgds.

We begin with a definition. Assume that I and J are instances (of different schemas) where every value that appears in a tuple of I also appears in a tuple of J . Define $\beta_{J,I}$ to be the full tgd where the left-hand side is the conjunction of the facts of J , and the right-hand side 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 $J^* = \text{chase}_{12}(I)$. It follows from Corollary 4.3 that every value that appears in a tuple of I (and hence in a tuple of $\hat{I}$) appears in a tuple of J^* , and so $\beta_{J^*, \hat{I}}$ is a full tgd. Define the *canonical local inverse* of $\mathcal{M}_{12}$ for I to be $(\mathbf{S}_2, \widehat{\mathbf{S}}_1, \{\beta_{J^*, \hat{I}}\})$ (we shall show that it is actually a local inverse). We call $\mathcal{M}_{21}$ the *most general s-t 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.

THEOREM 6.1. *Let $\mathcal{M}$ be a schema mapping defined by a finite set of s-t tgds, and let I be a source instance. Assume that $\mathcal{M}$ has an inverse for I that is defined by a finite set of s-t tgds. Then the canonical local inverse of $\mathcal{M}$ for I is indeed an inverse of $\mathcal{M}$ for I , and in fact the most general s-t inverse of $\mathcal{M}$ for I .*

Of course, we are much more interested in an $\mathcal{S}$ -inverse for a large class $\mathcal{S}$, rather than an inverse for a single instance I . However, the canonical local inverse is important as a tool in proving correctness of the canonical $\mathcal{S}$ -inverse (including the canonical global inverse) in the next section. In fact, even the fact that the canonical local inverse is most general is needed for the proof.

7. THE CANONICAL $\mathcal{S}$ -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 consider certain classes $\mathcal{S}$ of source instances, and show how to define a *canonical $\mathcal{S}$ -inverse* of $\mathcal{M}_{12}$, which is a schema mapping defined 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. We shall consider certain sets Γ of constraints on source instances, and let $\mathcal{S}$ be the class of source instances that satisfy Γ . When Γ is the empty set, then $\mathcal{S}$ is the class of all source instances, and so an $\mathcal{S}$ -inverse is a global inverse. In this case, we shall refer to the canonical $\mathcal{S}$ -inverse as the *canonical 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) source 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 is not hard to see that Γ is finitely chasable if and only if for every (finite) source instance I , some result of chasing I with Γ is a (finite) instance, where we allow values in I to be equated in the chase. It follows from results in [5] that when Γ is the disjoint union of a set of egds with a weakly acyclic set of tgds (as defined in [5]), then Γ is finitely chasable. We now give a simple example where the converse fails.

EXAMPLE 7.1. Let Γ' consist of the single tgd $\mathbf{R}(x, y) \rightarrow \exists z \mathbf{R}(y, z)$. It is easy to see that Γ' is not weakly acyclic, and in fact not finitely chasable. Let Γ'' consist of the single egd $\mathbf{R}(x, y) \rightarrow (x = y)$. Now let Γ be $\Gamma' \cup \Gamma''$. Then Γ is finitely chasable (since in this case, we need only chase with Γ'' alone). However, Γ is not weakly acyclic, since Γ' is not weakly acyclic. $\square$

Let Γ be a finitely chasable set of tgds and egds, and let $\mathcal{S}$ be the class of all source 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 $\mathbf{R}$ of $\mathbf{S}_1$, let $I_{\mathbf{R}}$ be a one-tuple instance that contains only the fact $\mathbf{R}(\mathbf{x})$, where the variables in $\mathbf{x}$ are distinct. Let $I_{\mathbf{R}}^{\Gamma}$ be a finite instance that is a result of chasing $I_{\mathbf{R}}$ 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 $I_{\mathbf{R}}^{\Gamma} = I_{\mathbf{R}}$. Let $J_{\mathbf{R}}^{\Gamma}$ be $\text{chase}_{12}(I_{\mathbf{R}}^{\Gamma})$, a result of chasing $I_{\mathbf{R}}^{\Gamma}$ with Σ_{12} .⁶

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 $I_{\mathbf{R}}^{\Gamma}$ (this is because $I_{\mathbf{R}}^{\Gamma}$ is a member of $\mathcal{S}$). It follows from Corollary 4.3 that every value that appears in a tuple of $I_{\mathbf{R}}^{\Gamma}$ (and hence in a tuple of $I_{\mathbf{R}}^{\Gamma}$) appears in a tuple of $J_{\mathbf{R}}^{\Gamma}$. Therefore, $\beta_{J_{\mathbf{R}}^{\Gamma}, I_{\mathbf{R}}^{\Gamma}}$ is a full tgd, where I is $I_{\mathbf{R}}^{\Gamma}$, and J is $J_{\mathbf{R}}^{\Gamma}$, with $\beta_{\cdot, \cdot}$ as defined in Section 6. Let us denote this full tgd by $\delta_{\mathbf{R}}^{\Gamma}$, and let $\Sigma_{12}^{\mathcal{S}}$ consist of all of the tgds $\delta_{\mathbf{R}}^{\Gamma}$, one for every relation symbol $\mathbf{R}$ of $\mathbf{S}_1$. Define the *canonical $\mathcal{S}$ -inverse* of $\mathcal{M}_{12}$ to be $\mathcal{M}_{12}^{\mathcal{S}} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{12}^{\mathcal{S}})$ (we shall show that it is actually an $\mathcal{S}$ -inverse). In the case where $\mathcal{S}$ is the class of all source instances, we may write $\Sigma_{12}^{\mathcal{S}}$ for Σ_{12} , and $\mathcal{M}_{12}^{\mathcal{S}}$ for $\mathcal{M}_{12}$, to honor the fact that we are then dealing with a schema mapping that is a global inverse). We call $\mathcal{M}_{21}$ the *most general s-t $\mathcal{S}$ -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 7.2. Let $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$. Assume that $\mathbf{S}_1$ consists of the binary relation symbol $\mathbf{R}$ and the unary relation symbol $\mathbf{S}$, and that $\mathbf{S}_2$ consists of the binary relation symbols $\mathbf{T}$ and $\mathbf{U}$. Let Σ_{12} consist of the s-t tgds $\mathbf{R}(x_1, x_2) \rightarrow \exists y (\mathbf{T}(x_1, y) \wedge \mathbf{U}(y, x_2))$, $\mathbf{R}(x, x) \rightarrow \mathbf{U}(x, x)$, and $\mathbf{S}(x) \rightarrow \exists y \mathbf{U}(x, y)$. Let Γ consist of the egd $\mathbf{R}(x_1, x_2) \rightarrow (x_1 = x_2)$.

⁶Even though $J_{\mathbf{R}}^{\Gamma}$ depends not just on $\mathbf{R}$ and Γ , but also on Σ_{12} , for simplicity we do not reflect the dependency on Σ_{12} in the notation $J_{\mathbf{R}}^{\Gamma}$.

Now I_R consists of the fact $R(x_1, x_2)$, and so I_R^Γ consists of the fact $R(x_1, x_1)$. Then J_R^Γ consists of the facts $T(x_1, y)$, $U(y, x_1)$, and $U(x_1, x_1)$. So δ_R^Γ is the tgd

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

Also, I_S and I_S^Γ each consist of the fact $S(x_1)$, and J_S^Γ consists of the fact $U(x_1, y)$. So δ_S^Γ is the tgd $U(x_1, y) \rightarrow \widehat{S}(x_1)$. Finally, $\mathcal{M}_{12}^S = (S_2, \widehat{S}_1, \Sigma_{12}^S)$, where Σ_{12}^S consists of the tgds δ_R^Γ and δ_S^Γ . $\square$

THEOREM 7.3. *Let $\mathcal{M}$ be a schema mapping defined 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 source instances that satisfy Γ . Assume that $\mathcal{M}$ has an $\mathcal{S}$ -inverse that is defined by a finite set of s-t tgds. Then the canonical $\mathcal{S}$ -inverse of $\mathcal{M}$ is indeed an $\mathcal{S}$ -inverse of $\mathcal{M}$, and in fact the most general s-t $\mathcal{S}$ -inverse of $\mathcal{M}$.*

8. FULL TGDS SUFFICE FOR AN INVERSE

The canonical local inverse and the canonical $\mathcal{S}$ -inverse are each defined 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 defined 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 $\mathcal{M}_{21}^f$ defined by a finite set of *full* s-t tgds and that is an $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$. While the canonical local inverse is tailored to a particular instance I , the mapping $\mathcal{M}_{21}^f$ is, as we shall see, constructed only from $\mathcal{M}_{21}$. From a technical point of view, this contrasts also with the canonical global inverse, which is constructed only from $\mathcal{M}_{12}$.

We begin with some definitions. Let γ be an s-t tgd . Assume that γ is $\forall \mathbf{x}(\varphi_S(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi_T(\mathbf{x}, \mathbf{y}))$, where $\varphi_S(\mathbf{x})$ is a conjunction of atomic formulas over $\mathbf{S}$ and $\psi_T(\mathbf{x}, \mathbf{y})$ is a conjunction of atomic formulas over $\mathbf{T}$. Let $\psi_T^f(\mathbf{x})$ be the conjunction of all atomic formulas in $\psi_T(\mathbf{x}, \mathbf{y})$ 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 $\forall \mathbf{x}(\varphi_S(\mathbf{x}) \rightarrow \psi_T^f(\mathbf{x}))$. If $\psi_T^f(\mathbf{x})$ is an empty conjunction, then γ^f is a *dummy tgd* where the right-hand side is “Truth” (and so the dummy tgd itself is “Truth”).

Let $\psi_T^n(\mathbf{x})$ be the conjunction of all atomic formulas in $\psi_T(\mathbf{x}, \mathbf{y})$ that contain some variable in $\mathbf{y}$ (the n stands for “non-full”). Define γ^n (the *non-full part* of γ) to be the tgd $\forall \mathbf{x}(\varphi_S(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi_T^n(\mathbf{x}, \mathbf{y}))$. As before, if $\psi_T^n(\mathbf{x})$ is an empty conjunction, then γ^n is a *dummy tgd* where the right-hand side 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 8.1. *Assume that $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (S_2, \widehat{S}_1, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. Let $\mathcal{M}_{21}^f = (S_2, \widehat{S}_1, \Sigma_{21}^f)$. If $\mathcal{M}_{21}$ is an $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$, then so is $\mathcal{M}_{21}^f$.*

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

COROLLARY 8.2. *Assume that $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (S_2, \widehat{S}_1, \Sigma_{21})$ are schema mappings where Σ_{12} and Σ_{21} are finite sets of s-t tgds. 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 $\mathcal{M}_{21}' = (S_2, \widehat{S}_1, \Sigma_{21}')$ is an $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$.*

9. REVERSING THE ARROWS (NOT!)

It is folk wisdom that simply “reversing the arrows” gives an inverse. What does this mean in our context?

Let us call a full tgd *reversible* if the same variables appear in the left-hand side as the right-hand side. If γ is a reversible tgd $\varphi \rightarrow \psi$, define $\text{rev}(\gamma)$ to be the full tgd $\psi \rightarrow \widehat{\varphi}$, where $\widehat{\varphi}$ is the result of replacing every relational symbol R by $\widehat{R}$. Since γ is reversible, $\text{rev}(\gamma)$ is indeed a full tgd . We think of $\text{rev}(\gamma)$ as the result of “reversing the arrow” of γ .

EXAMPLE 9.1. We now give a simple example that shows that $(S_2, \widehat{S}_1, \{\text{rev}(\gamma) : \gamma \in \Sigma_{12}\})$ is not necessarily a global inverse of $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$, even when Σ_{12} consists of a finite set of reversible tgds and $\mathcal{M}_{12}$ has a global inverse that is defined by a finite set of s-t tgds. Let S_1 consist of the unary relation symbols R_1 and R_2 . Let S_2 consist of the unary relation symbols S_1, S_2 , and S_3 . Let $\Sigma_{12} = \{R_1(x) \rightarrow S_1(x), R_2(x) \rightarrow S_2(x), R_1(x) \rightarrow S_3(x), R_2(x) \rightarrow S_3(x)\}$. Let $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$. Let $\Sigma_{21} = \{S_1(x) \rightarrow \widehat{R}_1(x), S_2(x) \rightarrow \widehat{R}_2(x)\}$. Let $\mathcal{M}_{21} = (S_2, \widehat{S}_1, \Sigma_{21})$. It is easy to see that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$.

Now let $\Sigma_{21}' = \{\text{rev}(\gamma) : \gamma \in \Sigma_{12}\}$. Thus $\Sigma_{21}' = \{S_1(x) \rightarrow \widehat{R}_1(x), S_2(x) \rightarrow \widehat{R}_2(x), S_3(x) \rightarrow \widehat{R}_1(x), S_3(x) \rightarrow \widehat{R}_2(x)\}$. Let $\mathcal{M}_{21}' = (S_2, \widehat{S}_1, \Sigma_{21}')$. 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. $\square$

Note that although $\{\text{rev}(\gamma) : \gamma \in \Sigma_{12}\}$ in Example 9.1 does not define 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 $\{\text{rev}(\gamma) : \gamma \in \Sigma_{12}\}$ that defines a global inverse.

THEOREM 9.2. *There is a schema mapping $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$ where each member of Σ_{12} is a reversible tgd with a singleton right-hand side, that has a global inverse defined by a finite set of s-t tgds, but where there is no subset X of Σ_{12} such that $(S_2, \widehat{S}_1, \{\text{rev}(\gamma) : \gamma \in X\})$ is a global inverse of $\mathcal{M}_{12}$.*

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

Given schema mappings $\mathcal{M}_{12} = (S_1, S_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (S_2, \widehat{S}_1, \Sigma_{21})$, we would like to know the class $\mathcal{S}$ of source instances I such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I . It is easy to see that this class is precisely the largest class $\mathcal{S}$ such that $\mathcal{M}_{21}$ is an $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$. Let $\mathcal{M}_{11}$ be the schema mapping $(S_1, \widehat{S}_1, \sigma)$, where σ is the composition formula $\Sigma_{12} \circ \Sigma_{21}$. So the class $\mathcal{S}$ we are seeking is the class of all 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 . We shall 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.

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 $t_1, \dots, t_k$ are terms, then $f(t_1, \dots, t_k)$ is a term. We now define a second-order tgd [7].

DEFINITION 10.1. 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:

$$\exists \mathbf{f}((\forall \mathbf{x}_1(\varphi_1 \rightarrow \psi_1)) \wedge \dots \wedge (\forall \mathbf{x}_n(\varphi_n \rightarrow \psi_n))),$$

where (1) each member of $\mathbf{f}$ is a function symbol; (2) each φ_i is a conjunction of (a) atomic formulas $S(y_1, \dots, y_k)$, where S is a k -ary relation symbol of schema $\mathbf{S}$, and $y_1, \dots, y_k$ 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 atomic formulas $T(t_1, \dots, t_l)$, where T is an l -ary relation symbol of schema $\mathbf{T}$ and $t_1, \dots, t_l$ are terms based on $\mathbf{x}_i$ and $\mathbf{f}$; and (4) each variable in $\mathbf{x}_i$ appears in some atomic formula of φ_i .

As noted in [7], every finite set of s-t tgds is logically equivalent to an SO tgd (but not conversely).

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

PROPOSITION 10.2. Let γ be an SO tgd, or a set of s-t tgds with source $\mathbf{S}$ and target $\widehat{\mathbf{S}}$, and let I be an instance of $\mathbf{S}$. Then $I \models \gamma^\#$ if and only if $\langle I, \widehat{I} \rangle \models \gamma$.

We now need some more definitions. Let γ be an SO tgd. The *equality-free reduction*⁷ γ^* of γ is obtained in multiple steps. First, we recursively replace each equality $f(t_1, \dots, t_k) = f(t'_1, \dots, t'_k)$ by $(t_1 = t'_1) \wedge \dots \wedge (t_k = t'_k)$. 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”. Intuitively, only those equalities that are “forced” remain. 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 $\exists f(\forall e(\text{Emp}(e) \rightarrow \text{Mgr}(e, f(e))))$, the result of dropping the second clause of (1). As another example, consider the following SO tgd:

$$\exists f(\forall x \forall y(\mathbf{R}(x, y) \wedge (f(x) = f(y)) \rightarrow \mathbf{S}(x, f(x)) \wedge \mathbf{T}(x, y))) \quad (5)$$

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

We now define $\text{fulltgd}(\gamma)$, which is a set of full tgds that we associate with the SO tgd γ . To obtain $\text{fulltgd}(\gamma)$, we first find the equality-free reduction γ^* of γ . We then rewrite

γ^* so that each right-hand side is a singleton. Thus, we replace $\varphi \rightarrow (\psi_1 \wedge \dots \wedge \psi_r)$, where $\psi_1, \dots, \psi_r$ are atomic formulas, by $(\varphi \rightarrow \psi_1) \wedge \dots \wedge (\varphi \rightarrow \psi_r)$. We then delete each “tgd” $\varphi \rightarrow \psi$ where the right-hand side ψ contains a function symbol. Then $\text{fulltgd}(\gamma)$ is the set of s-t tgds that remain. These are real tgds, since there are no function symbols present. By construction, $\text{fulltgd}(\gamma)$ is a set of full tgds with singleton right-hand sides.

As an example, when γ is the SO tgd (1), then $\text{fulltgd}(\gamma)$ is the empty set. As another example, when γ is the SO tgd (5), then $\text{fulltgd}(\gamma)$ contains the single tgd $\mathbf{R}(x, x) \rightarrow \mathbf{T}(x, x)$.

For each SO tgd γ where the source is $\mathbf{S}$ and the target 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 $\sigma^\#$. For each k and each k -ary relational symbol $\mathbf{R}$ of $\mathbf{S}$, take $x_1, \dots, x_k$ to be k distinct variables that do not appear in $\text{fulltgd}(\gamma)$. Let $A_{\mathbf{R}}$ be the set of all tgds of $\text{fulltgd}(\gamma)$ where the relational symbol in the right-hand side is $\widehat{\mathbf{R}}$. For each $\alpha \in A_{\mathbf{R}}$, assume that α is $\nu(\mathbf{y}) \rightarrow \widehat{\mathbf{R}}(y_1, \dots, y_k)$, where $y_1, \dots, y_k$ are not necessarily distinct (since α is full, every y_i appears in $\mathbf{y}$). Define μ_α to be the first-order formula $\exists \mathbf{y}(\nu(\mathbf{y}) \wedge (x_1 = y_1) \wedge \dots \wedge (x_k = y_k))$. Define $\psi_{\mathbf{R}}$ to be $\mathbf{R}(x_1, \dots, x_k) \rightarrow \bigvee \{\mu_\alpha : \alpha \in A_{\mathbf{R}}\}$. Since the empty disjunction represents “False”, it follows that if $A_{\mathbf{R}} = \emptyset$, then $\psi_{\mathbf{R}}$ is equivalent to $\neg \mathbf{R}(x_1, \dots, x_k)$. Now define $\gamma^\dagger$ to be the conjunction of the formulas $\psi_{\mathbf{R}}$ (over all relational symbols $\mathbf{R}$ of $\mathbf{S}$). Note that $\gamma^\dagger$ is a first-order formula.

EXAMPLE 10.3. Assume that there are two source relation symbols $\mathbf{R}$ and $\mathbf{S}$, and assume that $\text{fulltgd}(\gamma)$ consists of the following tgds, which we denote by α_1, α_2 :

$$\begin{aligned} (\alpha_1) : \quad & \mathbf{R}(y_2, y_1, y_3) \wedge \mathbf{S}(y_2, y_3, y_3) \rightarrow \widehat{\mathbf{R}}(y_1, y_1, y_2) \\ (\alpha_2) : \quad & \mathbf{S}(y_1, y_2, y_2) \rightarrow \widehat{\mathbf{R}}(y_1, y_2, y_1) \end{aligned}$$

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

$$\begin{aligned} (\mu_{\alpha_1}) : \quad & \exists y_1 \exists y_2 \exists y_3 \quad (\mathbf{R}(y_2, y_1, y_3) \wedge \mathbf{S}(y_2, y_3, y_3) \wedge \\ & (x_1 = y_1) \wedge (x_2 = y_1) \wedge (x_3 = y_2)) \\ (\mu_{\alpha_2}) : \quad & \exists y_1 \exists y_2 \exists y_3 \quad (\mathbf{S}(y_1, y_2, y_2) \wedge \\ & (x_1 = y_1) \wedge (x_2 = y_2) \wedge (x_3 = y_1)) \end{aligned}$$

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

We now have the following proposition.

PROPOSITION 10.4. Let γ be an SO tgd with source $\mathbf{S}$ and target $\widehat{\mathbf{S}}$, and let I be an instance of $\mathbf{S}$. Then $I \models \gamma^\dagger$ if and only if for every J such that $\langle I, J \rangle \models \gamma$, necessarily $\widehat{I} \subseteq J$.

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

THEOREM 10.5. 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. Let σ be $\Sigma_{12} \circ \Sigma_{21}$. Then $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $I \models \sigma^\# \wedge \sigma^\dagger$.

⁷A similar notion appears in [14] under the name “mapping reduction”.

<i>Held fixed</i>	<i>Complexity</i>
$\mathcal{M}_{12}$ 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

Figure 1: Local invertibility: input is source instance I

From our earlier Theorem 5.1, we can prove that if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each defined 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 11.) So by Fagin’s Theorem [3], the class $\mathcal{S}$ of source instances I such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I can be defined by a formula Γ in existential second-order logic. What is remarkable is that, as Theorem 10.5 says, there is such a formula Γ , namely $\sigma^\# \wedge \sigma^\dagger$, 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 10.6. *Assume that $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are schema mappings that are each defined by a finite set of full s-t tgds. There is a first-order formula φ such that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $I \models \varphi$.*

EXAMPLE 10.7. We continue with our running example (from Examples 3.1 and 3.2). We shall fulfill our promise to show that the schema mapping $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for precisely those source 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 3.2 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 $I \models \sigma^\dagger$. By Proposition 10.4, we need only show that for every J such that $\langle I, J \rangle \models \sigma$, necessarily $\hat{I} \subseteq J$. Let J be arbitrary such that $\langle I, J \rangle \models \sigma$. Since $\sigma \models \Sigma_{Id}$, it follows that $\langle I, J \rangle \models \Sigma_{Id}$. Therefore, $\hat{I} \subseteq J$, as desired. So indeed, $\sigma^\dagger$ is valid.

Therefore, by Theorem 10.5, we know that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $I \models \sigma^\dagger$. But $\sigma^\dagger$ 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 source 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 10.5. $\square$

11. COMPLEXITY RESULTS

We have investigated complexity issues, dealing with both local and global invertibility. In this paper, we do not consider complexity issues for $\mathcal{S}$ -invertibility except when $\mathcal{S}$ is a singleton (local invertibility) and when $\mathcal{S}$ is the class of all source instances (global invertibility). It might be interesting to consider complexity issues for other choices of $\mathcal{S}$. Our results are summarized in the tables of Figures 1 and 2. In both tables, we consider separately the cases where the tgds that define $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are full.

<i>Input</i>	<i>Complexity</i>
$\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	DP-hard
Full $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$	DP-complete
$\mathcal{M}_{12}$	coNP-hard
Full $\mathcal{M}_{12}$	coNP-complete

Figure 2: Global invertibility

In the table of Figure 1, the input is a source instance I . The first line of the table says that if $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are fixed schema mappings each defined 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.⁸ In the second line we consider the same problem as the first line, but $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each defined by a finite set of *full* tgds. Then the complexity drops to polynomial time (in fact, by Corollary 10.6, the problem is even definable in first-order logic, which makes it logspace computable). The third line considers whether $\mathcal{M}_{12}$ has an inverse for I defined by a finite set of s-t tgds. Since we have shown how to obtain a canonical local inverse that is an inverse for I if there is any inverse for I defined 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 local inverse grows with I , unlike the situation in the first line where Σ_{21} is given and so of fixed size. 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 fourth line we consider the same problem as the third line, but $\mathcal{M}_{12}$ is defined 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.

In the first line of the table of Figure 2, the input consists of schema mappings $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ that are each defined by a finite set of s-t tgds, and the problem is deciding whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$. In the second line we consider the same problem as the first line, but $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are each defined by a finite set of full tgds, and the problem is DP-complete.⁹ The third line considers whether $\mathcal{M}_{12}$ has a global inverse defined by a finite set of s-t tgds. In the fourth line we consider the same problem as the third line, but $\mathcal{M}_{12}$ is defined by a finite set of full tgds, and the problem is coNP-complete. In fact, this problem is coNP-complete even when the tgds that define $\mathcal{M}_{12}$ all have a singleton left-hand side. The first and third lines inherit their lower bounds from the full cases (the second and fourth lines, respectively).

There is a large complexity gap in the first and third lines, since it is open as to whether these problems are even decidable. When the tgds that define $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are full, the reason the problem is decidable (and in fact, DP-complete or coNP-complete) is a small model theorem that guarantees that if $\mathcal{M}_{21}$ is not a global inverse, then there is a small (polynomial-size) counterexample I . We close this section with a discussion of small model theorems, including a reason for the difficulty in proving a small model theorem when the tgds that define $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$ are not necessarily full.

⁸The NP-hardness result was obtained by Phokion Kolaitis.

⁹The class DP consists of all decision problems that can be written as the intersection of an NP problem and a coNP problem.

11.1 Small model theorems

We begin with a small submodel theorem for the case of schema mappings that are each defined by a finite set of full tgds.

THEOREM 11.1. *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 an instance of $\mathbf{S}_1$. 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' .*

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

THEOREM 11.2. *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 next theorem implies that the small submodel theorem fails dramatically when the tgds are not necessarily full.

THEOREM 11.3. *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} and Σ_{21} are finite sets of 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 .*

Although the small submodel theorem fails when the tgds are not necessarily full, it is open as to whether the small model theorem holds in this case. If the small model theorem were to hold, then our problems in the first and third lines of the table of Figure 2 would be decidable (and as we can show, even in Π_2^P in the polynomial-time hierarchy).

12. 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 source instances. We have obtained a number of results about our notion of inverse, and some of these results are surprising.

There are many open problems, as we would expect from a “first step” paper like this. Section 11 gives us several open problems, including closing the complexity gaps and resolving whether the small model theorem holds when the tgds are not necessarily full. We now mention some other open problems.

- We have focused most of our attention on schema mappings defined by a finite set of s-t tgds. What about more general schema mappings? 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 $\mathcal{M}_{12}$.
- 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 notion of $\widehat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ being homomorphically equivalent. By Theorem 5.4, this notion is strictly weaker than $\mathcal{M}_{21}$ being an inverse of $\mathcal{M}_{12}$ for I .

This paper 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 Alan Nash for helpful comments.

13. REFERENCES

- [1] C. Beeri and M. Y. Vardi. A Proof Procedure for Data Dependencies. *Journal of the Association for Computing Machinery (JACM)*, 31(4):718–741, 1984.
- [2] P. A. Bernstein. Applying Model Management to Classical Meta-Data Problems. In *Conference on Innovative Data Systems Research (CIDR)*, pages 209–220, 2003.
- [3] R. Fagin. Generalized First-Order Spectra and Polynomial-Time Recognizable Sets. In R. M. Karp, editor, *Complexity of Computation, SIAM-AMS Proceedings, Vol. 7*, pages 43–73, 1974.
- [4] R. Fagin. Multivalued dependencies and a new normal form for relational databases. *ACM Transactions on Database Systems (TODS)*, 2(3):262–278, Sept. 1977.
- [5] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. Data Exchange: Semantics and Query Answering. *Theoretical Computer Science*, pages 89–124, 2005.
- [6] R. Fagin, P. G. Kolaitis, and L. Popa. Data Exchange: Getting to the Core. *ACM Transactions on Database Systems (TODS)*, 30(1):174–210, 2005.
- [7] R. Fagin, P. G. Kolaitis, L. Popa, and W.-C. Tan. Composing Schema Mappings: Second-Order Dependencies to the Rescue. In *ACM Symposium on Principles of Database Systems (PODS)*, pages 83–94, 2004.
- [8] P. Hell and J. Nešetřil. The Core of a Graph. *Discrete Mathematics*, 109:117–126, 1992.
- [9] R. Hull. Relative Information Capacity of Simple Relational Database Schemata. *SIAM J. on Computing*, 15:856–886, 1986.
- [10] M. Lenzerini. Data Integration: A Theoretical Perspective. In *ACM Symposium on Principles of Database Systems (PODS)*, pages 233–246, 2002.
- [11] S. Melnik. *Generic Model Management: Concepts and Algorithms*. Springer, 2004.
- [12] S. Melnik, P. Bernstein, A. Halevy, and E. Rahm. Supporting Executable Mappings in Model Management. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 167–178, 2005.
- [13] N. C. Shu, B. C. Housel, R. W. Taylor, S. P. Ghosh, and V. Y. Lum. EXPRESS: A Data EXtraction, Processing, and REStructuring System. *ACM Transactions on Database Systems (TODS)*, 2(2):134–174, 1977.
- [14] C. Yu and L. Popa. Semantic Adaptation of Schema Mappings when Schemas Evolve. In *International Conference on Very Large Data Bases (VLDB)*, pages 1006–1017, 2005.

APPENDIX

A. APPENDIX: PROOFS

Proof of Theorem 4.1 Assume that $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{21})$. Let σ be the composition formula of $\mathcal{M}_{12}$ and $\mathcal{M}_{21}$. Assume 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,

$$\{J' : \langle I_1, J' \rangle \models \Sigma_{12}\} = \{J' : \langle I_2, J' \rangle \models \Sigma_{12}\}. \quad (6)$$

We shall derive a contradiction. We have

$$\begin{aligned} & \{J : \widehat{I}_1 \subseteq J\} \\ = & \{J : \langle I_1, J \rangle \models \sigma\} \text{ (by (3) with } I_1 \text{ for } I) \\ = & \{J : \text{There exists } J' \text{ such that} \\ & \langle I_1, J' \rangle \models \Sigma_{12} \text{ and } \langle J', J \rangle \models \Sigma_{21}\} \\ & \text{(by definition of the composition formula)} \\ = & \{J : \text{There exists } J' \text{ such that } \langle I_2, J' \rangle \models \Sigma_{12} \\ & \text{and } \langle J', J \rangle \models \Sigma_{21}\} \text{ (by (6))} \\ = & \{J : \langle I_2, J \rangle \models \sigma \text{ (by definition of the} \\ & \text{composition formula)}\} \\ = & \{J : \widehat{I}_2 \subseteq J\} \text{ (by (3) with } I_2 \text{ for } I) \end{aligned}$$

We just showed that $\{J : \widehat{I}_1 \subseteq J\} = \{J : \widehat{I}_2 \subseteq J\}$. Since $\widehat{I}_1 \in \{J : \widehat{I}_1 \subseteq J\}$, it follows that $\widehat{I}_1 \in \{J : \widehat{I}_2 \subseteq J\}$, 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, $I_1 = I_2$, which is a contradiction, as desired. $\square$

PROPOSITION A.1. *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 source instances I_1 and I_2 . Then $\text{chase}_{12}(I_1) \neq \text{chase}_{12}(I_2)$.*

PROOF. It follows from results in [5] that the solutions for a source instance I are exactly the homomorphic images of $\text{chase}_{12}(I)$. Therefore, if $\text{chase}_{12}(I_1) = \text{chase}_{12}(I_2)$, then I_1 and I_2 have the same solutions. But this is a contradiction, by Theorem 4.1. $\square$

Proof of Corollary 4.2 Assume that $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{21})$. 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, we know that $\mathcal{M}_{11}$ and the identity mapping are equivalent on I_2 (this follows in a straightforward way from our assumption of preservation under isomorphism). Hence, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I_2 . So by Theorem 4.1, the set of solutions for I_1 and I_2 are different. $\square$

Proof of Corollary 4.3 Assume that the value a that appears in a tuple of I does not appear in $\text{chase}_{12}(I)$; we shall derive a contradiction. Let a' be a new value 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 $\text{chase}_{12}(I) = \text{chase}_{12}(I')$. Hence, as in the proof of Proposition A.1, it follows that I_1 and I_2 have the same solutions. But this is a contradiction of Corollary 4.2. $\square$

Proof of Theorem 4.6 By Theorem 4.1, 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.

The *core* of a structure [6, 8] is the smallest substructure that is also a homomorphic image of the structure. When Σ_{12} is a finite set of s-t tgds, and I is a source instance, then every universal solution for I has the same core (up to isomorphism) [6], which is called the *core of the universal solutions*.

Define Σ_{21} to have the meaning that if J_1 is an $\mathbf{S}_2$ instance and J_2 is a $\widehat{\mathbf{S}}_1$ instance, then $\langle J_1, J_2 \rangle \models \Sigma_{21}$ precisely if there is an $\mathbf{S}_1$ instance I' such that $\widehat{I}' \subseteq J_2$, and J_1 is the core of the universal solutions of I' with respect to Σ_{12} . 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 J ,

$$\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21} \text{ if and only if } \widehat{I} \subseteq J. \quad (7)$$

We first show that if $\widehat{I} \subseteq J$, then $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Let J^* be the core of the universal solutions of I with respect to Σ_{12} . Then $\langle I, J^* \rangle \models \Sigma_{12}$, since the core of the universal solutions is a solution itself (it is even a universal solution [6]). By definition of Σ_{21} , it follows that $\langle J^*, J \rangle \models \Sigma_{21}$ (we let the roles of J_1, J_2, I' be played by J^*, J, I respectively). Since $\langle I, J^* \rangle \models \Sigma_{12}$ and $\langle J^*, J \rangle \models \Sigma_{21}$, it follows that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$, 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 $\langle J^*, J \rangle \models \Sigma_{21}$, there is I' such that $\widehat{I}' \subseteq J$, and J^* is the core of the universal solutions of I' with respect to Σ_{12} . Let $I'' = I \cup I'$ (that is, the set of facts of I'' consists of the union of the facts of I and I'). Since Σ_{12} is LAV, it follows that $\text{chase}_{12}(I'') = \text{chase}_{12}(I) \cup \text{chase}_{12}(I')$, that is, the chase of the union is the union of the chases. This follows immediately from the fact that the left-hand side of every member of Σ_{12} is a singleton. Now J^* is the core of $\text{chase}_{12}(I')$, since $\text{chase}_{12}(I')$ is a universal solution for I' [5], and since J^* is the core of the universal solutions of I' . Therefore, in particular, $J^* \subseteq \text{chase}_{12}(I')$. Since also $\text{chase}_{12}(I') \subseteq \text{chase}_{12}(I'')$, it follows that $J^* \subseteq \text{chase}_{12}(I'')$.

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. Since J^* is a solution for I and $\text{chase}_{12}(I)$ is a universal solution for I , there is a homomorphism h that maps $\text{chase}_{12}(I)$ into J^* . Since J^* is the core of the universal solutions of I' , there is a homomorphism h' such that the image of $\text{chase}_{12}(I')$ under h' is J^* . Define h'' on $\text{chase}_{12}(I'')$ by letting $h''(x) = h(x)$ for $x \in X$, letting $h''(x) = h'(x)$ for $x \in X'$, and letting $h''(x) = x$ otherwise. Since no tuple of $\text{chase}_{12}(I'')$ contains both a value in X and a value in X' , it follows easily that h'' is a homomorphism such that the image of $\text{chase}_{12}(I'')$ under h'' is J^* . Since J^* is also a core, it follows easily that J^* is the core of $\text{chase}_{12}(I'')$.

So J^* is the core of the universal solutions of both I' and I'' , and hence a universal solution for both I' and for I'' . Since Σ_{12} is a finite set of s-t tgds, it follows from results of [5] that I' and I'' share a universal solution if and only

if they have the same set of solutions. Therefore, I' and I'' have the same set of solutions. Since I' and I'' have the same set of solutions, it follows from the unique solutions property that $I' = I''$, that is, $I' = I \cup I'$. But this implies that $I \subseteq I'$, and so $\hat{I} \subseteq \hat{I}'$. Since also $\hat{I}' \subseteq J$, this implies that $\hat{I} \subseteq J$. This was to be shown. $\square$

The next lemma corresponds to our statement in the Introduction that the set $\langle I, J \rangle$ of pairs that satisfy a tgd is “closed down on the left and closed up on the right”.

LEMMA A.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. $\square$

The next lemma, which is standard, gives the key property of the chase.

LEMMA A.3. *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 source instance and J a target instance. Then $\langle I, J \rangle \models \Sigma_{12}$ if $\text{chase}_{12}(I) \subseteq J$. If the s-t tgds in Σ_{12} are all full, then $\langle I, J \rangle \models \Sigma_{12}$ if and only if $\text{chase}_{12}(I) \subseteq J$.*

The next lemma (Lemma A.4) gives us a different viewpoint of composition. It will be used to prove the subsequent lemma (Lemma A.5).¹⁰ Note that in Lemma A.4, 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 paper we do not allow target tgds in our schema mappings, they are allowed in [5].

LEMMA A.4. *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, $\langle I_1, I_3 \rangle \models \Sigma_{12} \circ \Sigma_{23}$. $\square$

LEMMA A.5. *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 $\text{chase}_{23}(\text{chase}_{12}(I))$ 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 A.4. 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 $I_2 = \text{chase}_{12}(I_1)$ and let $I_3 = \text{chase}_{23}(I_2)$. Then $I_3 = \text{chase}_{23}(\text{chase}_{12}(I_1))$. It is easy to see that $\langle I_1, \langle I_2, I_3 \rangle \rangle$ is a result of chasing $\langle I_1, \emptyset \rangle$ with $\Sigma_{12} \cup \Sigma_{23}$. It then follows from results in [5] 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 A.3,

¹⁰Lemmas A.4 and A.5 are due to Lucian Popa and Wang-Chiew Tan.

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 A.4, 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. $\square$

Proof of Theorem 5.1 Assume first that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I . Thus, (7) holds. It follows immediately that $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. By Lemma A.3, we know that

$$\langle I, \text{chase}_{12}(I) \rangle \models \Sigma_{12}$$

$$\langle \text{chase}_{12}(I), \text{chase}_{21}(\text{chase}_{12}(I)) \rangle \models \Sigma_{21}.$$

Therefore, $\langle I, \text{chase}_{21}(\text{chase}_{12}(I)) \rangle \models \Sigma_{12} \circ \Sigma_{21}$. So by (7), where the role of J is played by $\text{chase}_{21}(\text{chase}_{12}(I))$, it follows that $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$, as desired.

Conversely, assume that $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and that $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$; we must show that (7) holds. Assume first that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$; we must show that $\hat{I} \subseteq J$. Let $\mathcal{M}_{11} = (\mathbf{S}_1, \hat{\mathbf{S}}_1, \Sigma_{12} \circ \Sigma_{21})$. By Lemma A.5, we know that $\text{chase}_{21}(\text{chase}_{12}(I))$ is a universal solution for I under $\mathcal{M}_{11}$, and so there is a homomorphism from $\text{chase}_{21}(\text{chase}_{12}(I))$ to J . Since $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$, and since homomorphisms map values in I onto themselves, it follows that $\hat{I} \subseteq J$, as desired.

Assume now that $\hat{I} \subseteq J$; we must show that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Since $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\hat{I} \subseteq J$, it follows from Lemma A.2 that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. $\square$

Proof of Corollary 5.2 Assume first that

$$\hat{I} = \text{chase}_{21}(\text{chase}_{12}(I)).$$

By Lemma A.3, we know that

$$\langle I, \text{chase}_{12}(I) \rangle \models \Sigma_{12}$$

$$\langle \text{chase}_{12}(I), \text{chase}_{21}(\text{chase}_{12}(I)) \rangle \models \Sigma_{21}.$$

But $\text{chase}_{21}(\text{chase}_{12}(I)) = \hat{I}$, and so

$$\langle \text{chase}_{12}(I), \hat{I} \rangle \models \Sigma_{21}.$$

So $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Since also $\hat{I} = \text{chase}_{21}(\text{chase}_{12}(I))$, it follows from Theorem 5.1 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 5.1, $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. Since $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$, we know that there is J such that $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \hat{I} \rangle \models \Sigma_{21}$. Since $\langle I, J \rangle \models \Sigma_{12}$, it follows from Lemma A.3 that $\text{chase}_{12}(I) \subseteq J$. Since $\text{chase}_{12}(I) \subseteq J$ and $\langle J, \hat{I} \rangle \models \Sigma_{21}$, we know from Lemma A.2 that $\langle \text{chase}_{12}(I), \hat{I} \rangle \models \Sigma_{21}$. It follows from Lemma A.3 that $\text{chase}_{21}(\text{chase}_{12}(I)) \subseteq \hat{I}$. Since also $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$, we have $\hat{I} = \text{chase}_{21}(\text{chase}_{12}(I))$, as desired. $\square$

Proof of Corollary 5.3 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 5.1, $\langle I, \widehat{I} \rangle \models \sigma$ and $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. Since $\langle I, \widehat{I} \rangle \models \sigma$ and since $\text{chase}_{21}(\text{chase}_{12}(I))$ is a universal solution for I under $\mathcal{M}_{11}$ (by Lemma A.5), it follows that there is a homomorphism from $\text{chase}_{21}(\text{chase}_{12}(I))$ to $\widehat{I}$. Since $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$, we know that the identity function is a homomorphism from $\widehat{I}$ to $\text{chase}_{21}(\text{chase}_{12}(I))$. Hence, $\widehat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ are homomorphically equivalent, as desired. $\square$

Proof of Theorem 5.4 Let $\mathbf{S}_1$ consist of the binary relation symbol $\mathbf{R}$, let $\mathbf{S}_2$ consist of the binary relation symbol $\mathbf{S}$, let Σ_{12} consist of the s-t tgd $\mathbf{R}(a, b) \rightarrow \exists x(\mathbf{S}(a, x) \wedge \mathbf{S}(x, b))$, and let Σ_{21} consist of the s-t tgd $(\mathbf{S}(a, x) \wedge \mathbf{S}(x, b)) \rightarrow \widehat{\mathbf{R}}(a, b)$. 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 $\text{chase}_{21}(\text{chase}_{12}(I))$ 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

$$\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I)).$$

Therefore, the identity function is a homomorphism from $\widehat{I}$ to $\text{chase}_{21}(\text{chase}_{12}(I))$. We now show that there is a homomorphism from $\text{chase}_{21}(\text{chase}_{12}(I))$ to $\widehat{I}$. Each null value x in $\text{chase}_{21}(\text{chase}_{12}(I))$ is a null value in $\text{chase}_{12}(I)$, which is obtained by applying Σ_{12} to a fact $\mathbf{R}(a, b)$ to obtain the facts $\mathbf{S}(a, x)$ and $\mathbf{S}(x, b)$. 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 $\text{chase}_{21}(\text{chase}_{12}(I))$? First, there are the facts $\widehat{\mathbf{R}}(a, b)$ such that $\mathbf{R}(a, b)$ is a fact of I . For such facts, $\widehat{\mathbf{R}}(h(a), h(b))$ is simply $\widehat{\mathbf{R}}(a, b)$, which is consistent with h being a homomorphism. The other facts of $\text{chase}_{21}(\text{chase}_{12}(I))$ are facts $\widehat{\mathbf{R}}(x, y)$ such that $\mathbf{S}(x, a)$ and $\mathbf{S}(a, y)$ appear in $\text{chase}_{12}(I)$, where x and y are nulls, and where a and b are constants (values in I). We must show that $\widehat{\mathbf{R}}(h(x), h(y))$ is a fact of $\widehat{I}$.

Since $\mathbf{S}(x, a)$ appears in $\text{chase}_{12}(I)$, there is a constant c such that $\mathbf{R}(c, a)$ is a fact of I , and $\mathbf{S}(c, x)$ is a fact of $\text{chase}_{12}(I)$. Then h maps x to c , and maps y to a . Since $\mathbf{R}(c, a)$ is a fact of I , we know that $\widehat{\mathbf{R}}(c, a)$ is a fact of $\widehat{I}$, that is, $\widehat{\mathbf{R}}(h(x), h(y))$ is a fact of $\widehat{I}$, as desired. This completes the proof that $\widehat{I}$ and $\text{chase}_{21}(\text{chase}_{12}(I))$ 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 $\mathbf{R}(A, B)$ and $\mathbf{R}(B, A)$, where A and B are distinct. 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 $\mathbf{S}(A, X)$ and $\mathbf{S}(X, B)$ are facts of J . Similarly, there is some Y (either a null or a constant) such that $\mathbf{S}(B, Y)$ and $\mathbf{S}(Y, A)$ are facts of J . Since $\mathbf{S}(X, B)$ and $\mathbf{S}(B, Y)$ are facts of J , and since $\langle J, \widehat{I} \rangle \models \Sigma_{21}$, it follows that $\widehat{\mathbf{R}}(X, Y)$ is a fact of $\widehat{I}$, that is, $\mathbf{R}(X, Y)$ is a fact of I . But I contains only the facts $\mathbf{R}(A, B)$ and $\mathbf{R}(B, A)$. so either $X = A$ and $Y = B$, or $X = B$ and $Y = A$. Assume that $X = A$ and $Y = B$ (a symmetric proof works if $X = B$ and $Y = A$). Now $\mathbf{S}(A, X)$ is a fact of J , that is, $\mathbf{S}(A, A)$ is a fact of J . Since $\langle J, \widehat{I} \rangle \models \Sigma_{21}$, it follows that $\widehat{\mathbf{R}}(A, A)$ is a fact of $\widehat{I}$ (this is because we apply

the tgd $(\mathbf{S}(a, x) \wedge \mathbf{S}(x, b)) \rightarrow \widehat{\mathbf{R}}(a, b)$ to J where the roles of a, x , and b are all played by A). But this is a contradiction, since $\widehat{\mathbf{R}}(A, A)$ is not a fact of $\widehat{I}$. $\square$

LEMMA A.6. 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. $\square$

Proof of Theorem 6.1 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_{21} 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 local inverse of $\mathcal{M}_{12}$ for I . We first prove that Σ'_{21} logically implies Σ_{21} . Let $J^* = \text{chase}_{12}(I)$. By Lemma A.3, we know that $\langle I, J^* \rangle \models \Sigma_{12}$. Let J' be the result of chasing J^* with Σ'_{21} . By Lemma A.3, we know that $\langle J^*, J' \rangle \models \Sigma'_{21}$. 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 $\widehat{I} \subseteq J'$. That is, 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 A.6 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,

$$\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I)).$$

So by Theorem 5.1, it follows that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I , as desired. $\square$

Proof of Theorem 7.3 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 a finite set of s-t tgds. Let $\mathcal{M}_{12}^{\mathcal{S}} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{12}^{\mathcal{S}})$ be the canonical $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$.

Since $\mathcal{M}'_{21}$ is an $\mathcal{S}$ -inverse of $\mathcal{M}_{12}$, certainly $\mathcal{M}'_{21}$ is an inverse of $\mathcal{M}_{12}$ for $I_{\mathbf{R}}^{\Gamma}$ as defined in Section 7. Now $(\mathbf{S}_2, \widehat{\mathbf{S}}_1, \{\delta_{\mathbf{R}}^{\Gamma}\})$ is the canonical local inverse of $\mathcal{M}_{12}$ for $I_{\mathbf{R}}^{\Gamma}$, so by Theorem 6.1, we know that Σ'_{21} logically implies $\delta_{\mathbf{R}}^{\Gamma}$. Therefore, Σ'_{21} logically implies $\Sigma_{12}^{\mathcal{S}}$, 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

$$\langle I, J \rangle \models \Sigma_{12} \circ \Sigma'_{21} \text{ if and only if } \widehat{I} \subseteq J. \quad (8)$$

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

$$\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{12}^{\mathcal{S}} \text{ if and only if } \widehat{I} \subseteq J. \quad (9)$$

Assume first that I satisfies Γ , and $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{12}^{\mathcal{S}}$. Then there is J' such that $\langle I, J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{12}^{\mathcal{S}}$. Let $\mathbf{R}(\mathbf{c})$ 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 $\mathbf{R}(\mathbf{c})$ 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_R^Γ . 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 left-hand side of δ_R^Γ into J' , that maps $\mathbf{x}$ onto $\mathbf{c}$. Hence, since $\langle J', J \rangle \models \delta_R^\Gamma$ (because δ_R^Γ is in Σ_{12}^S), it follows that $\hat{R}(\mathbf{c})$ is in J . Since $R(\mathbf{c})$ is an arbitrary fact of I , this implies that $\hat{I} \subseteq J$, as desired.

Assume now that $\hat{I} \subseteq J$. By (8), 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 Σ_{12}^S , it follows from Lemma A.6 that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{12}^S$. This was to be shown. $\square$

Proof of Theorem 8.1 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 (7) holds. We must show that

$$\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}^f \text{ if and only if } \hat{I} \subseteq J. \quad (10)$$

Assume that $\hat{I} \subseteq J$. From (7), 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_{21}^f , it follows that $\langle J', J \rangle \models \Sigma_{21}^f$. Since $\langle I, J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{21}^f$, it follows that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}^f$. This proves the “if” direction of (10).

Conversely, assume that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}^f$; we must prove that $\hat{I} \subseteq J$. Since $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}^f$, there is J' such that $\langle I, J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{21}^f$. Let J'' be the result of chasing J' with Σ_{21}^f . Then $\langle J', J'' \rangle \models \Sigma_{21}^f$. Since $\langle J', J \rangle \models \Sigma_{21}^f$ and $\langle J', J'' \rangle \models \Sigma_{21}^f$, it follows that $\langle J', J \cup J'' \rangle \models \Sigma_{21}^f \cup \Sigma_{21}^f$. Since Σ_{21} is logically equivalent to $\Sigma_{21}^f \cup \Sigma_{21}^f$, 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 (7), where the role of J is played by $J \cup J''$, we know that $\hat{I} \subseteq J \cup J''$.

Every tuple introduced by doing the chase of J' with Σ_{21}^f necessarily contains nulls, by construction of Σ_{21}^f . That is, every tuple in J'' contains nulls. But $\hat{I} \subseteq J \cup J''$ and $\hat{I}$ contains no nulls. It follows that $\hat{I} \subseteq J$, which was to be shown. $\square$

Proof of Theorem 9.2 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

$$R(x, y) \rightarrow S(x, y)$$

$$R(x, y) \rightarrow T(x, y)$$

$$V(x) \rightarrow U(x)$$

$$R(x, x) \wedge R(y, y) \rightarrow S(x, y)$$

$$V(x) \rightarrow T(x, x)$$

Let $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$.

Let Σ_{21} consist of the tgds

$$S(x, y) \wedge T(x, y) \rightarrow \hat{R}(x, y)$$

$$U(x) \rightarrow \hat{V}(x)$$

Let $\mathcal{M}_{21} = (\mathbf{S}_2, \hat{\mathbf{S}}_1, \Sigma_{21})$. We now show that $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$.¹¹

¹¹This mapping $\mathcal{M}_{21}$ is not the same as the canonical global inverse in Theorem 7.3.

It is straightforward to verify that when we apply the composition algorithm of [7] to compute $\Sigma_{12} \circ \Sigma_{21}$, we obtain

$$R(x, y) \rightarrow \hat{R}(x, y)$$

$$R(x, x) \wedge V(x) \rightarrow \hat{R}(x, x)$$

$$R(x, x) \wedge R(y, y) \wedge R(x, y) \rightarrow \hat{R}(x, y)$$

$$V(x) \rightarrow \hat{V}(x).$$

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)

$$R(x, y) \rightarrow \hat{R}(x, y)$$

$$V(x) \rightarrow \hat{V}(x).$$

But this tgd defines 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 $\{rev(\gamma) : \gamma \in X\}$ is a global inverse of Σ_{12} . Define Σ_{21}^X to be $\{rev(\gamma) : \gamma \in X\}$. Let $\mathcal{M}_{21}^X = (\mathbf{S}_2, \hat{\mathbf{S}}_1, \Sigma_{21}^X)$. Assume that $\mathcal{M}_{21}^X$ is a global inverse of $\mathcal{M}_{12}$.

We first show that the set X cannot contain $R(x, y) \rightarrow S(x, y)$. Assume that it does. Then Σ_{21}^X contains $S(x, y) \rightarrow \hat{R}(x, y)$. Let I contain only the facts $R(0, 0)$ and $R(1, 1)$. Then $chase_{12}(I)$ contains the fact $S(0, 1)$. So $chase_{21}^X(chase_{12}(I))$, the result of chasing $chase_{12}(I)$ with Σ_{21}^X , contains the fact $\hat{R}(0, 1)$. Therefore, $\hat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 5.2 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 $R(x, y) \rightarrow T(x, y)$. Assume that it does. Then Σ_{21}^X contains $T(x, y) \rightarrow \hat{R}(x, y)$. Let I contain only the fact $V(0)$. Then $chase_{12}(I)$ contains the fact $T(0, 0)$. So $chase_{21}^X(chase_{12}(I))$ contains the fact $\hat{R}(0, 0)$. Therefore, $\hat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 5.2 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 following three tgds:

$$V(x) \rightarrow U(x)$$

$$R(x, x) \wedge R(y, y) \rightarrow S(x, y)$$

$$V(x) \rightarrow T(x, x).$$

In this case, Σ_{21}^X consists of some subset of the following three tgds:

$$U(x) \rightarrow \hat{V}(x).$$

$$S(x, y) \rightarrow \hat{R}(x, x) \wedge \hat{R}(y, y)$$

$$T(x, x) \rightarrow \hat{V}(x)$$

Let I contain only the fact $R(0, 1)$. Then $chase_{21}^X(chase_{12}(I))$ (and in fact, $chase_{21}^X(J)$ for an arbitrary J) does not contain the fact $\hat{R}(0, 1)$, since the only facts about $\hat{R}$ that can be generated by chasing with Σ_{21}^X are of the form $\hat{R}(x, x)$. Therefore, $\hat{I} \neq chase_{21}^X(chase_{12}(I))$. It then follows from Corollary 5.2 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. $\square$

LEMMA A.7. *Let γ be an SO tgd. Then*

1. *each tgd in $\text{fulltgd}(\gamma)$ is logically implied by γ , and*
2. *if I is a source instance, and if $\mathbf{S}(\mathbf{a})$ is a fact obtained by chasing¹² I with γ , where each member of $\mathbf{a}$ is a constant (a member of I), then $\mathbf{S}(\mathbf{a})$ is obtained by chasing I with some member of $\text{fulltgd}(\gamma)$.*

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

We now prove (2). It follows from the definition of the chase with SO tgds in [7] 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 $\mathbf{S}(\mathbf{a})$ where each member of $\mathbf{a}$ is a constant is by chasing with members of $\text{fulltgd}(\gamma)$. $\square$

Proof of Proposition 10.4 Assume first that $I \models \gamma^\dagger$ and $\langle I, J \rangle \models \gamma$; we must show that $\hat{I} \subseteq J$. Let $\mathbf{R}(a_1, \dots, a_k)$ be a fact of I . Since $I \models \gamma^\dagger$, we know that $I \models \psi_{\mathbf{R}}$. Since $I \models \psi_{\mathbf{R}}$ and $\mathbf{R}(a_1, \dots, a_k)$ is a fact of I , we know that $\psi_{\mathbf{R}}$ is not equivalent to $\neg \mathbf{R}(x_1, \dots, x_k)$. So $A_{\mathbf{R}} \neq \emptyset$. Since $I \models \psi_{\mathbf{R}}$, there is α in $A_{\mathbf{R}}$ such that I satisfies μ_α when the roles of $x_1, \dots, x_k$ are played by $a_1, \dots, a_k$ respectively. Assume that α is $\nu(\mathbf{y}) \rightarrow \hat{\mathbf{R}}(y_1, \dots, y_k)$. So I satisfies $\exists \mathbf{y}(\nu(\mathbf{y}) \wedge (x_1 = y_1) \wedge \dots \wedge (x_k = y_k))$ when the roles of $x_1, \dots, x_k$ are played by $a_1, \dots, a_k$ 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 A.7) $\langle I, J \rangle \models \text{fulltgd}(\gamma)$. Therefore, $\hat{\mathbf{R}}(a_1, \dots, a_k)$ is a fact of J . So $\hat{I} \subseteq J$, as desired.

Conversely, assume that for every J such that $\langle I, J \rangle \models \gamma$, necessarily $\hat{I} \subseteq J$; we must show that $I \models \gamma^\dagger$. Let $\langle I, J^* \rangle$ be the result of chasing $\langle I, \emptyset \rangle$ with γ . Then $\langle I, J^* \rangle \models \gamma$, and so by assumption $\hat{I} \subseteq J^*$. Let $\mathbf{R}(a_1, \dots, a_k)$ be a fact of I . So $\hat{\mathbf{R}}(a_1, \dots, a_k)$ is a fact of J^* . By part (2) of Lemma A.7, $\hat{\mathbf{R}}(a_1, \dots, a_k)$ is obtained by chasing I with some member α of $\text{fulltgd}(\gamma)$. So μ_α holds in I when the roles of $x_1, \dots, x_k$ are played by $a_1, \dots, a_k$ respectively. So $\psi_{\mathbf{R}}$ holds for I when the roles of $x_1, \dots, x_k$ are played by $a_1, \dots, a_k$ respectively. Since $\mathbf{R}(a_1, \dots, a_k)$ is an arbitrary fact of I of the form $\mathbf{R}(x_1, \dots, x_k)$, it follows that $\psi_{\mathbf{R}}$ holds for I . So $I \models \gamma^\dagger$, as desired. $\square$

Proof of Theorem 10.5 Assume first that $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I ; we must show that $I \models \sigma^\# \wedge \sigma^\dagger$. By definition of inverse, (3) holds. So $\langle I, \hat{I} \rangle \models \sigma$. By Proposition 10.2 it follows that $I \models \sigma^\#$. Furthermore, by (3), we know that whenever $\langle I, J \rangle \models \sigma$, necessarily $\hat{I} \subseteq J$. Therefore, by Proposition 10.4, it follows that $I \models \sigma^\dagger$. So $I \models \sigma^\# \wedge \sigma^\dagger$, as desired.

Conversely, assume that $I \models \sigma^\# \wedge \sigma^\dagger$; we must show that (3) holds. Since $I \models \sigma^\#$, we know from Proposition 10.2 that $\langle I, \hat{I} \rangle \models \sigma$. Since $I \models \sigma^\dagger$, it follows from Proposition 10.4 that for every J such that $\langle I, J \rangle \models \sigma$, necessarily $\hat{I} \subseteq J$. So (3) holds. $\square$

Proof of Corollary 10.6 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

¹²There is a notion of *chasing* with SO tgds [7] that is very similar to the notion of chasing with s-t tgds.

Σ_{21} are finite sets of full s-t tgds. Since Σ_{12} and Σ_{21} are finite sets of full tgds, so is the composition formula $\Sigma_{12} \circ \Sigma_{21}$ [7]. So $\sigma^\#$ is first-order. Also, $\sigma^\dagger$ is always first-order. So $\sigma^\# \wedge \sigma^\dagger$ is first-order. The result now follows from Theorem 10.5. $\square$

B. APPENDIX: COMPLEXITY RESULTS

B.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 defined by a finite set of s-t tgds?

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

The next result was proven in [7], with the proof we now give, although it was not stated as a theorem (it appeared in the middle of another proof).

LEMMA B.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 $\langle J, I_3 \rangle \models \Sigma_{23}$. It was shown in [5] 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, $\text{chase}_{12}(I)$ 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 A.2 that $\langle I_2, I_3 \rangle \models \Sigma_{23}$. This concludes the proof. $\square$

THEOREM B.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 5.1, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I if and only if $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. 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 $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. 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 B.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.¹³

The source schema $\mathbf{S}_1$ has binary relation symbols $\mathbf{P}$ and $\mathbf{D}$. The target schema $\mathbf{S}_2$ has three binary relation symbols $\mathbf{P}'$, $\mathbf{D}'$, and $\mathbf{C}$. The set Σ_{12} consists of the tgds

$$\mathbf{P}(x, y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{D}(u, v) \rightarrow \mathbf{D}'(u, v)$$

$$\mathbf{P}(x, y) \rightarrow \exists u \exists v (\mathbf{C}(x, u) \wedge \mathbf{C}(y, v))$$

The set Σ_{21} consists of the tgds

$$\mathbf{P}'(x, y) \rightarrow \widehat{\mathbf{P}}(x, y)$$

$$\mathbf{D}'(u, v) \rightarrow \widehat{\mathbf{D}}(u, v)$$

$$(\mathbf{P}'(x, y) \wedge \mathbf{C}(x, u) \wedge \mathbf{C}(y, v)) \rightarrow \widehat{\mathbf{D}}(u, v)$$

Let H be an undirected graph. 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 source instance I_H by letting $\mathbf{P}$ contain all of the edges of H , and letting $\mathbf{D}$ be the inequality relation on g, b, r . Thus, the $\mathbf{D}$ 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 $\widehat{I}_H \subseteq \text{chase}_{21}(\text{chase}_{12}(I_H))$ for every H . So by Theorem 5.1, it follows that $\mathcal{M}_{21}$ 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 $c : V_H \rightarrow \{g, b, r\}$ be a 3-coloring of H . Define J to be the $\mathbf{S}_2$ instance whose $\mathbf{P}'$ relation is the $\mathbf{P}$ relation of I_H , whose $\mathbf{D}'$ relation is the $\mathbf{D}$ relation of I_H , and where $\mathbf{C}(x, c(x))$ 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 of V_H , there is some u in $\{g, b, r\}$ such that $\mathbf{C}(x, u)$ holds in J . Define $c : V_H \rightarrow \{g, b, r\}$ by letting $c(x)$ be some such u . Since $\langle J, \widehat{I}_H \rangle \models \Sigma_{21}$, it follows easily that c is a 3-coloring of H , as desired. $\square$

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

¹³This proof is due to Phokion Kolaitis.

THEOREM B.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 5.2, since deciding if $\widehat{I} = \text{chase}_{21}(\text{chase}_{12}(I))$ is a polynomial-time test. $\square$

Corollary 10.6 actually gives a stronger result than Theorem B.3. Corollary 10.6 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 Σ_{21} is not given. Note that Corollary 4.3 and Theorem 6.1 give us a decision procedure for deciding if there is a schema mapping defined by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I . We simply check whether $\text{chase}_{12}(I)$ contains every value that appears in I . If not, we reject. If so, we produce the canonical 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 B.2. However, unlike Theorem B.2, it does not follow that this problem is in NP, since unlike the situation in Theorem B.2, the schema mapping $\mathcal{M}_{21}$ is not fixed, but depends on I . We now obtain an upper bound on the complexity.

THEOREM B.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 defined 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 $J^* = \text{chase}_{12}(I)$. Compute J^* (this is a polynomial-time procedure), and check whether J^* contains every value that appears in I . If not, then we know by Corollary 4.3 that there is no inverse of $\mathcal{M}_{12}$ for I . So assume that J^* contains every value that appears in I . Let $\mathcal{M}_{21}$ be the canonical local inverse. By Theorem 6.1, we know that there is a schema mapping defined 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 . The theorem is proven if we show that the statement

$$“\mathcal{M}_{21} \text{ is an inverse of } \mathcal{M}_{12} \text{ for } I” \quad (11)$$

is in Σ_2^P . By Theorem 5.1, we know that (11) holds precisely if $\langle I, \widehat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$ and $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. 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 $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$ are each in Σ_2^P .

Lemma B.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 $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$ is a problem in NP, and hence in Σ_2^P . Let k be the maximum

number of conjuncts in left-hand sides of members of Σ_{21} . To verify that $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$, we do the following NP procedure. For each fact F in $\hat{I}$, we guess k applications of members of Σ_{12} to facts of I to produce k facts of $\text{chase}_{12}(I)$. We then guess an application of a member of Σ_{21} to these k facts in $\text{chase}_{12}(I)$ to produce the fact F in $\text{chase}_{21}(\text{chase}_{12}(I))$. $\square$

If φ_1 and φ_2 are each conjunctions of atomic formulas over the schema $\mathbf{S}$, then a homomorphism from φ_1 to φ_2 is a homomorphism $h : I_1 \rightarrow I_2$, 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 . Each value in I_1 is treated as being a member of $\underline{\text{Var}}$, and so the restriction on homomorphisms that $h(c) = c$ for $c \in \underline{\text{Const}}$ 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 B.5. *Assume that Σ_{12} consists of the full s-t tgd $\gamma_1 \rightarrow \gamma_2$. Then $\text{chase}_{12}(I)$ consists precisely of facts $\mathbf{S}(h(\mathbf{x}))$ such that $\mathbf{S}(\mathbf{x})$ is a conjunct of γ_2 and h is a homomorphism from γ_1 to I .*

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

Define a *weak endomorphism* of an instance K to be a mapping $h : K \rightarrow K$ such that for every fact $\mathbf{R}(\mathbf{t})$ of K , we have that $\mathbf{R}(h(\mathbf{t}))$ 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.

THEOREM B.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 defined by a finite set of s-t tgds if and only if (1) every value that appears in a tuple of I appears in a tuple of $\text{chase}_{12}(I)$ and (2) every weak endomorphism of $\text{chase}_{12}(I)$ is a weak endomorphism of I .*

PROOF. Let $J^* = \text{chase}_{12}(I)$. Assume first that there is a schema mapping defined by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I . Corollary 4.3 tells us that condition (1) of the theorem must hold, and Theorem 6.1 tells us that the canonical local inverse $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I . By Corollary 5.2, we know that $\hat{I} = \text{chase}_{21}(J^*)$, and in particular $\text{chase}_{21}(J^*) \subseteq \hat{I}$.

Let h be a weak endomorphism of J^* . Let us write the full tgd $\beta_{J^*, \hat{I}}$ 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 $\hat{I}$. 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 $\mathbf{P}(x, y)$ is a conjunct of γ_1 , and $h(x) = x'$ and $h(y) = y'$, then $\mathbf{P}(x', y')$ 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 $\text{chase}_{21}(J^*) \subseteq \hat{I}$, it follows that each conjunct of γ'_2 is a fact of $\hat{I}$. But this means that h is a weak endomorphism of $\hat{I}$, and hence of I , as desired.

Conversely, assume that conditions (1) and (2) of the theorem hold. Since condition (1) holds, it follows that $\beta_{J^*, \hat{I}}$ is a full tgd. Define Σ_{21} to be a set consisting only of this tgd, and, as before, let $\mathcal{M}_{21} = (\mathbf{S}_2, \mathbf{S}_1, \Sigma_{21})$. When we chase J^*

with $\beta_{J^*, \hat{I}}$ we clearly obtain at least $\hat{I}$. So $\hat{I} \subseteq \text{chase}_{21}(J^*)$. It follows easily from Lemma B.5 that $\text{chase}_{21}(J^*)$ consists precisely of facts $\hat{\mathbf{R}}(h(\mathbf{x}))$ such that h is a weak endomorphism of J^* and $\hat{\mathbf{R}}(\mathbf{x})$ is a fact of $\hat{I}$. But by condition (2) of the theorem, $\hat{\mathbf{R}}(h(\mathbf{x}))$ is then in $\hat{I}$. So condition (2) implies that $\text{chase}_{21}(J^*) \subseteq \hat{I}$. Hence, $\hat{I} = \text{chase}_{21}(J^*)$. So by Corollary 5.2, $\mathcal{M}_{21}$ is an inverse of $\mathcal{M}_{12}$ for I . $\square$

THEOREM B.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 defined 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 B.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 $\mathbf{P}$ and three unary relation symbols $\mathbf{G}$, $\mathbf{B}$, and $\mathbf{R}$ (which stand for green, red, and blue). The target schema $\mathbf{S}_2$ has a binary relation symbol $\mathbf{P}'$ and three unary relation symbols $\mathbf{G}'$, $\mathbf{B}'$, and $\mathbf{R}'$. The set Σ_{12} consists of the tgds

$$\mathbf{P}(x, y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{G}(x) \rightarrow \mathbf{G}'(x)$$

$$\mathbf{B}(x) \rightarrow \mathbf{B}'(x)$$

$$\mathbf{R}(x) \rightarrow \mathbf{R}'(x)$$

$$\mathbf{G}(x) \wedge \mathbf{R}(y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{G}(x) \wedge \mathbf{R}(y) \rightarrow \mathbf{P}'(y, x)$$

$$\mathbf{G}(x) \wedge \mathbf{B}(y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{G}(x) \wedge \mathbf{B}(y) \rightarrow \mathbf{P}'(y, x)$$

$$\mathbf{R}(x) \wedge \mathbf{B}(y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{R}(x) \wedge \mathbf{B}(y) \rightarrow \mathbf{P}'(y, x)$$

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 source instance I_H by letting $\mathbf{P}$ contain all of the edges of H , and letting $\mathbf{G}$ contain the singleton tuple (g) , letting $\mathbf{B}$ contain the singleton tuple (b) , and letting $\mathbf{R}$ contain the singleton tuple (r) . We now show that there is a schema mapping defined 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 $J = \text{chase}_{12}(I_H)$. Thus, the $\mathbf{P}'$ 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 $\mathbf{G}'$ relation of J , then $h(x)$ is a tuple of the $\mathbf{G}'$ relation of J , since x is necessarily g , and $h(g) = g$. The same hold for the $\mathbf{B}'$ relation of J and the $\mathbf{R}'$ relation of J . Finally, whenever (x, y) is a tuple of the $\mathbf{P}'$ 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 $\mathbf{P}'$ 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 $\mathbf{P}(x, y)$ is a fact of I_H . However, $\mathbf{P}(c(x), c(y))$ is not a fact of I_H , since $c(x)$ is either g , b , or r , and the $\mathbf{P}$ 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 B.6 that there is no schema mapping defined 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 $\mathbf{P}'$ 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 $\mathbf{P}'(x, y)$ is a fact of J , and so $\mathbf{P}'(c(x), c(y))$ 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 $\mathbf{G}'$ relation of J contains only (g) , and similarly for $\mathbf{B}'$ and $\mathbf{R}'$. 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 every value in I_H appears in J , it follows from Theorem B.6 that there is a schema mapping defined by a finite set of s-t tgds that is an inverse of $\mathcal{M}_{12}$ for I_H . $\square$

B.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 $\mathcal{M}_{21} = (\mathbf{S}_2, \mathbf{S}_1, \Sigma_{21})$, 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 defined by a finite set of s-t tgds?

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

Proof of Theorem 11.1 By Corollary 5.2,

$$\hat{I} \neq \text{chase}_{21}(\text{chase}_{12}(I)).$$

Let us denote $\text{chase}_{21}(\text{chase}_{12}(I))$ by J . There are two cases.

Case 1: $\hat{I}$ is not a subinstance of J . Then there is a fact $\mathbf{R}(\mathbf{t})$ of I such that $\hat{\mathbf{R}}(\mathbf{t})$ is not in J . Let I' be the one-tuple

subinstance that contains only the fact $\mathbf{R}(\mathbf{t})$. By monotonicity of the chase, $\hat{\mathbf{R}}(\mathbf{t})$ is not in $\text{chase}_{21}(\text{chase}_{12}(I'))$. Hence, $\hat{I}' \neq \text{chase}_{21}(\text{chase}_{12}(I'))$, and so by Corollary 5.2, we know that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I' .

Case 2: $\hat{I}$ is a subinstance of J . So $\hat{I}$ is a proper subset of J . Let k_{12} be the maximum, over all members ψ of Σ_{12} , of the number of conjuncts that appear in the left-hand side of ψ . We call this quantity k_{12} the *maximum left arity* of members of Σ_{12} . Similarly, define k_{21} to be the maximum left arity of members of Σ_{21} . Since $\hat{I}$ is a proper subset of J , there is a fact $\hat{\mathbf{R}}(\mathbf{t})$ that is in J but not in $\hat{I}$. Note that the members of $\mathbf{t}$ are all constants, since all of the tgds under consideration are full. Now $\hat{\mathbf{R}}(\mathbf{t})$ was generated in one chase step when chasing $\text{chase}_{12}(I)$ with a tgd γ in Σ_{21} . There is a set W of at most k_{21} facts of $\text{chase}_{12}(I)$ such that the conjuncts in the left-hand side of γ map to members of W in order to obtain $\hat{\mathbf{R}}(\mathbf{t})$ 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 left-hand side 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 $\text{chase}_{21}(\text{chase}_{12}(I'))$ contains the fact $\hat{\mathbf{R}}(\mathbf{t})$. Since $\hat{\mathbf{R}}(\mathbf{t})$ is not in $\hat{I}$, it follows that $\hat{\mathbf{R}}(\mathbf{t})$ is not in $\hat{I}'$. Since $\hat{\mathbf{R}}(\mathbf{t})$ is not in $\hat{I}'$ but is in $\text{chase}_{21}(\text{chase}_{12}(I'))$, it follows that $\hat{I}' \neq \text{chase}_{21}(\text{chase}_{12}(I'))$. By Corollary 5.2, we know that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for I' . $\square$

Proof of Theorem 11.3 We define Σ_{12} and Σ_{21} similarly to the construction in the proof of Theorem B.2, except that here we use 2-colorability rather than 3-colorability.

The source schema $\mathbf{S}_1$ has binary relation symbols $\mathbf{P}$ and $\mathbf{D}$. The target schema $\mathbf{S}_2$ has three binary relation symbol $\mathbf{P}'$, $\mathbf{D}'$, and $\mathbf{C}$. The set Σ_{12} consists of the tgds

$$\mathbf{P}(x, y) \rightarrow \mathbf{P}'(x, y)$$

$$\mathbf{D}(u, v) \rightarrow \mathbf{D}'(u, v)$$

$$\mathbf{P}(x, y) \rightarrow \exists u \exists v (\mathbf{C}(x, u) \wedge \mathbf{C}(y, v))$$

The set Σ_{21} consists of the tgds

$$\mathbf{P}'(x, y) \rightarrow \hat{\mathbf{P}}(x, y)$$

$$\mathbf{D}'(u, v) \rightarrow \hat{\mathbf{D}}(u, v)$$

$$(\mathbf{P}'(x, y) \wedge \mathbf{P}'(y, x) \wedge \mathbf{D}'(w, z) \wedge \mathbf{D}'(z, w) \wedge \mathbf{C}(x, u) \wedge \mathbf{C}(y, v)) \rightarrow \hat{\mathbf{D}}(u, v)$$

We now define I . We let $\mathbf{P}^I$, the $\mathbf{P}$ relation of I , be an undirected cycle of odd length $2m + 1$. Specifically, define $i \oplus 1$ to be $i + 1 \bmod 2m + 1$, and let $\mathbf{P}^I$ consist of all tuples $(i, i \oplus 1)$ and $(i \oplus 1, i)$, for $0 \leq i \leq 2m$. We let $\mathbf{D}^I$ consist of the two tuples $(g, b), (b, g)$, where g and b are distinct. 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 $\langle I, \hat{I} \rangle \models \Sigma_{12} \circ \Sigma_{21}$. Hence, there is J such that $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \hat{I} \rangle \models \Sigma_{21}$. Since $\langle I, J \rangle \models \Sigma_{12}$, it follows that for each i with $0 \leq i \leq$

$2m$, there is u in $\{g, b\}$ such that $C(i, u)$ holds in J . Define $c : \{0, \dots, 2m\} \rightarrow \{g, b\}$ by letting $c(i)$ be some such u . Since $\langle J, \hat{I} \rangle \models \Sigma_{21}$, it follows easily that c is a 2-coloring of the odd cycle 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 $\hat{I}' \subseteq \text{chase}_{21}(\text{chase}_{12}(I'))$. Therefore, by Theorem 5.1, we need only show that $\langle I', \hat{I}' \rangle \models \Sigma_{12} \circ \Sigma_{21}$. There are two cases.

Case 1: $D^{I'}$ is a proper subset of D^I . Define the $\mathbf{S}_2$ instance J by letting $P^{J'} = P^{I'}$ and $D^{J'} = D^{I'}$, and defining 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 $\langle J, \hat{I}' \rangle \models \Sigma_{21}$. Clearly $\langle J, \hat{I}' \rangle$ satisfies the first two tgds of Σ_{21} . Let γ be the third tgd of Σ_{21} . Then $\langle J, \hat{I}' \rangle$ satisfies γ also, since there are no w, z such that (w, z) and (z, w) are each tuples of $D^{J'}$.

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 P^I that is not in $P^{I'}$. Define $c : \{0, \dots, 2m\} \rightarrow \{g, b\}$ 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 $P^{J'} = P^{I'}$ and $D^{J'} = D^{I'}$, and defining 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 $\langle J, \hat{I}' \rangle \models \Sigma_{21}$. Clearly $\langle J, \hat{I}' \rangle$ satisfies the first two tgds of Σ_{21} . We now show that $\langle J, \hat{I}' \rangle$ 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 left-hand side of γ fails (since $P'(x, y) \wedge P'(y, x)$ 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. $\square$

As before, if R is a relational symbol, let I_R be a one-tuple instance that contains only the fact $R(\mathbf{x})$, where the variables in $\mathbf{x}$ are distinct.

LEMMA B.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 for each relational symbol R of $\mathbf{S}_1$, every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$. Let $\mathcal{M}_{12}^{-1} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{12}^{-1})$ be the canonical global inverse of $\mathcal{M}_{12}$. If $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{12}^{-1}$, then $\hat{I} \subseteq J$. In particular, $\hat{I} \subseteq \text{chase}_{12}^{-1}(\text{chase}_{12}(I))$, where chase_{12}^{-1} represents the result of chasing with Σ_{12}^{-1} .*

PROOF. Assume that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{12}^{-1}$. Then there is J' such that $\langle I, J' \rangle \models \Sigma_{12}$ and $\langle J', J \rangle \models \Sigma_{12}^{-1}$. Let $R(a_1, \dots, a_k)$ be an arbitrary fact of I . So J' contains a homomorphic image J'' of the result of chasing $R(a_1, \dots, a_k)$ with Σ_{12} . Hence, J contains a homomorphic image of the result of chasing J'' with δ_R . Therefore, J contains the fact $\hat{R}(a_1, \dots, a_k)$. Since $R(a_1, \dots, a_k)$ is an arbitrary fact of I , it follows that $\hat{I} \subseteq J$, as desired.

As for the “in particular”, we know that

$$\langle I, \text{chase}_{12}^{-1}(\text{chase}_{12}(I)) \rangle \models \Sigma_{12} \circ \Sigma_{12}^{-1},$$

since by Lemma A.3, we know that $\langle I, \text{chase}_{12}(I) \rangle \models \Sigma_{12}$ and $\langle \text{chase}_{12}(I), \text{chase}_{12}^{-1}(\text{chase}_{12}(I)) \rangle \models \Sigma_{12}^{-1}$. $\square$

THEOREM B.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 defined by a finite set of s-t tgds, is coNP-complete. It is coNP-complete even if Σ_{12} is restricted to being a finite set of full s-t tgds all with a singleton left-hand side.

PROOF. We first show membership in coNP. As before, for each relational symbol R of $\mathbf{S}_1$, let I_R be a one-tuple instance that contains only the fact $R(\mathbf{x})$, where the variables in $\mathbf{x}$ are distinct. Do a polynomial-time check to verify that every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$ (otherwise, by Corollary 4.3, we know that there is no global inverse of $\mathcal{M}_{12}$). So assume that every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$, for each source relational symbol R . Let $\mathcal{M}_{12}^{-1} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{12}^{-1})$ be the canonical global inverse. We know from Theorem 7.3 that there is a schema mapping defined by a finite set of s-t tgds that is a global inverse of $\mathcal{M}_{12}$ if and only if $\mathcal{M}_{12}^{-1}$ is a global inverse of $\mathcal{M}_{12}$.

We now show that the problem of determining if $\mathcal{M}_{12}^{-1}$ is *not* a global inverse of $\mathcal{M}_{12}$ is in NP. By Theorem 11.2, we know that $\mathcal{M}_{12}^{-1}$ is not a global inverse of $\mathcal{M}_{12}$ if and only if there is a small model I such that $\mathcal{M}_{12}^{-1}$ is not an inverse of $\mathcal{M}_{12}$ for I . By assumption, every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$, for each source relational symbol R . So by Lemma B.8, we know that $\hat{I} \subseteq \text{chase}_{12}^{-1}(\text{chase}_{12}(I))$. Let us denote $\text{chase}_{12}^{-1}(\text{chase}_{12}(I))$ by $\bar{I}$. Hence, by Corollary 5.2, we know that $\mathcal{M}_{12}^{-1}$ is not an inverse of $\mathcal{M}_{12}$ for I if and only if $\bar{I}$ is a proper superset of $\hat{I}$.

We have shown that $\mathcal{M}_{12}^{-1}$ 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 $\hat{I}$. But this property is in NP: we simply guess the small model I and verify that $\bar{I}$ is a proper superset of $\hat{I}$ by guessing applications of members of Σ_{12} and Σ_{12}^{-1} that generate a member of $\bar{I}$ that is not in $\hat{I}$.

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 $C_1 \wedge \dots \wedge C_k$, 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 $Y_{i,1} \vee \dots \vee Y_{i,f(i)} \vee \neg Z_{i,1} \vee \dots \vee \neg Z_{i,g(i)}$, 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 define a schema mapping $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$, whose size is polynomial in the size of φ , that we shall prove has a global inverse defined 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 R , along with $2n$ -ary relational symbols $U_{i,1}, \dots, U_{i,j}$ for $1 \leq j \leq f(i)$ and $1 \leq i \leq k$. The target schema $\mathbf{S}_2$ contains $2n$ -ary relational symbols $U'_{i,1}, \dots, U'_{i,j}$ for $1 \leq j \leq f(i)$ and $1 \leq i \leq k$. Thus, for each of the relational symbols $U_{i,j}$ of $\mathbf{S}_1$, there is a corresponding relational symbol $U'_{i,j}$ in $\mathbf{S}_2$. Furthermore, $\mathbf{S}_2$ has $2n$ -ary relational symbols $S_1, \dots, S_k$.

Let $\mathbf{x}$ be the $2n$ -tuple $(x_1, \dots, x_{2m})$. 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, $\mathbf{x}^{\text{m}}$ 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 $\mathbf{x}^{\text{m}}$ be the tuple that is the result of interchanging x_{2m-1} and x_{2m} in $\mathbf{x}$. Thus, $\mathbf{x}^{\text{m}}$ 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 $U_{i,j}(\mathbf{x}^{\text{m}}) \rightarrow S_i(\mathbf{x}^{\text{m}})$. Assume that $Z_{i,j}$ is the propositional symbol A_m . Let $z_{i,j}$ be the full tgd $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$. The set Σ_{12} consists of the following full tgds. First, it has the “copying” tgds $U_{i,j}(\mathbf{x}) \rightarrow U'_{i,j}(\mathbf{x})$ for each $U_{i,j}$ 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 $1 \leq j \leq f(i)$ and $1 \leq i \leq k$) and $z_{i,j}$ (for $1 \leq j \leq g(i)$ and $1 \leq i \leq k$).

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 $A_1, \dots, A_{2n}$ that are assigned true under the truth assignment $\mathbf{Tr}$. Let $a_1, \dots, a_{2n}$ be distinct values. Let $\mathbf{t}$ be the $2n$ -ary tuple $(t_1, \dots, t_{2n})$, 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 $\mathbf{t}^{\text{m}}$ be the tuple that is the result of interchanging t_{2m-1} and t_{2m} in $\mathbf{t}$.

Define the source instance I_1 to consist of the facts $U_{i,j}(\mathbf{t})$ for each of the relational symbols $U_{i,j}$ of $\mathbf{S}_1$. Furthermore, if A_m is not in P , then let I_1 also contain the fact $R(\mathbf{t}^{\text{m}})$, for $1 \leq m \leq n$. Note that I_1 does not contain the fact $R(\mathbf{t})$, since $\mathbf{t}^{\text{m}} \neq \mathbf{t}$ when A_m is not in P .

Let I_t be the one-tuple instance that contains only the fact $R(\mathbf{t})$. Define the source instance I_2 to consist of the union of the facts of I_1 and I_t . Since I_2 contains the fact $R(\mathbf{t})$ but I_1 does not, we see that $I_1 \neq I_2$. We now show that $\text{chase}_{12}(I_1) = \text{chase}_{12}(I_2)$. It then follows from Proposition A.1 that there is not global inverse of $\mathcal{M}_{12}$. Since every member of Σ_{12} has a singleton left-hand side, we need only show that $\text{chase}_{12}(I_t) \subseteq \text{chase}_{12}(I_1)$.

The only tgds that may generate a tuple when we chase I_t with Σ_{12} are the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ and the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$. Consider first the result of applying the tgds $R(\mathbf{x}) \rightarrow S_i(\mathbf{x})$ to I_t . The result is $S_i(\mathbf{t})$. We must show that $S_i(\mathbf{t})$ is in $\text{chase}_{12}(I_1)$. 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 $U_{i,j}(\mathbf{x}^{\text{m}}) \rightarrow S_i(\mathbf{x}^{\text{m}})$ is in Σ_{12} . The result of applying this tgd to the fact $U_{i,j}(\mathbf{t})$ 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 $R(\mathbf{t}^{\text{m}})$. Since $Z_{i,j}$ is a disjunct of C_i , the tgd $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$ is in Σ_{12} . Then the result of applying this tgd to the fact $R(\mathbf{t}^{\text{m}})$ of I_1 generates $S_i(\mathbf{t})$, as desired.

Consider now the result of applying the tgd $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$ to I_t . The result is $S_i(\mathbf{t}^{\text{m}})$. We must show that $S_i(\mathbf{t}^{\text{m}})$ is in $\text{chase}_{12}(I_1)$. 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 $U_{i,j}(\mathbf{x}^{\text{m}}) \rightarrow S_i(\mathbf{x}^{\text{m}})$ is in Σ_{12} . The result

of applying this tgd to the fact $U_{i,j}(\mathbf{t})$ of I_1 generates $S_i(\mathbf{t})$, which (since $t_{2m-1} = t_{2m}$) is the same as the fact $S_i(\mathbf{t}^{\text{m}})$, 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 $R(\mathbf{t}^{\text{m}})$. Then the result of applying the tgd $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$ to the fact $R(\mathbf{t}^{\text{m}})$ of I_1 generates $S_i(\mathbf{t}^{\text{m}})$, 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 defined 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 $U'_{i,j}(\mathbf{x}) \rightarrow \hat{U}_{i,j}(\mathbf{x})$ for each $U_{i,j}$ in $\mathbf{S}_1$ and the tgd $S_1(\mathbf{x}) \wedge \dots \wedge S_k(\mathbf{x}) \rightarrow \hat{R}(\mathbf{x})$. Let $\mathcal{M}_{21} = (\mathbf{S}_2, \hat{\mathbf{S}}_1, \Sigma_{21})$. 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 5.2 that there is a source instance I such that $\hat{I} \neq \text{chase}_{21}(\text{chase}_{12}(I))$. Because of the copying tgds $U_{i,j}(\mathbf{x}) \rightarrow U'_{i,j}(\mathbf{x})$ of Σ_{12} and the copying tgds $U'_{i,j}(\mathbf{x}) \rightarrow \hat{U}_{i,j}(\mathbf{x})$ 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 $S_1(\mathbf{x}) \wedge \dots \wedge S_k(\mathbf{x}) \rightarrow \hat{R}(\mathbf{x})$ of Σ_{21} , it follows that $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. So $\hat{I}$ is a proper subset of $\text{chase}_{21}(\text{chase}_{12}(I))$. It is clear that every fact $\hat{U}_{i,j}(\mathbf{t})$ in $\text{chase}_{21}(\text{chase}_{12}(I))$ is in $\hat{I}$. Hence, there is $\mathbf{t}$ such that the fact $\hat{R}(\mathbf{t})$ is in $\text{chase}_{21}(\text{chase}_{12}(I))$ but not in $\hat{I}$. The only way that this fact can arise in $\text{chase}_{21}(\text{chase}_{12}(I))$ is for $S_i(\mathbf{t})$ to be a fact of $\text{chase}_{12}(I)$, for $1 \leq i \leq k$. Since $R(\mathbf{t})$ is not a fact of I , the only way that $S_i(\mathbf{t})$ can be a fact of $\text{chase}_{12}(I)$ 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 $\text{chase}_{12}(I)$ 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 $\text{chase}_{12}(I)$ 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 $\mathbf{t} = (t_1, \dots, t_{2n})$. Define a truth assignment $\mathbf{Tr}$ to the propositional symbols $A_1, \dots, A_n$ 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 $Y_{i,1} \vee \dots \vee Y_{i,f(i)} \vee \neg Z_{i,1} \vee \dots \vee \neg Z_{i,g(i)}$, where each $Y_{i,j}$ and each $Z_{i,j}$ is in $\{A_1, \dots, A_n\}$. 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 $A_1, \dots, A_n$. Assume that $Y_{i,j}$ is A_m . Then $y_{i,j}$ is the tgd $U_{i,j}(\mathbf{x}^{\text{m}}) \rightarrow S_i(\mathbf{x}^{\text{m}})$. Since the event $\mathcal{Y}_{i,j}$ holds, that is, $S_i(\mathbf{t})$ arises in $\text{chase}_{12}(I)$ 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 $A_1, \dots, A_n$. Assume that $Z_{i,j}$ is A_m . Then $z_{i,j}$ is the tgd $R(\mathbf{x}) \rightarrow S_i(\mathbf{x}^{\text{m}})$. Since the event $\mathcal{Z}_{i,j}$ holds, that is, $S_i(\mathbf{t})$ arises in $\text{chase}_{12}(I)$ as a result of applying the tgd $z_{i,j}$, necessarily the tgd $z_{i,j}$ is applied to a fact $R(\mathbf{x})$ for a tuple $\mathbf{x}$ other than $\mathbf{t}$. Hence, $t_{2m-1} \neq t_{2m}$, 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. $\square$

LEMMA B.10. 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. $\square$

LEMMA B.11. 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 $\Sigma'_{12}, \Sigma'_{21}, \Sigma''_{12}$, 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

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

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 B.10, 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 (12) 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 (7) holds. We must show that

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

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 (12) 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 (7), it therefore follows that (13) 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 (7) 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 (12) 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 $\widehat{I} \subseteq J$. So $\widehat{I'} \subseteq J'$. 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 (12), it follows that $\langle I, J \rangle \models \Sigma_{12} \circ \Sigma_{21}$, as desired. $\square$

THEOREM B.12. 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 proof of Theorem 7.3, let $\mathbf{R}$ be an arbitrary relational symbol of $\mathbf{S}_1$, and let $I_{\mathbf{R}}$ be a one-tuple instance that contains only the fact $\mathbf{R}(x_1, \dots, x_k)$, where $x_1, \dots, x_k$ 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:

- (*) $\widehat{I_{\mathbf{R}}} \subseteq \text{chase}_{21}(\text{chase}_{12}(I_{\mathbf{R}}))$ for every relational symbol $\mathbf{R}$ of $\mathbf{S}_1$.
- (**) $\text{chase}_{21}(\text{chase}_{12}(I)) \subseteq \widehat{I}$ 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 11.1.

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 5.2, it would follow that $\mathcal{M}_{21}$ is not an inverse of $\mathcal{M}_{12}$ for $I_{\mathbf{R}}$. If (**) did not hold, then by Corollary 5.2, 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 5.2, we need only show that $\widehat{I} = \text{chase}_{21}(\text{chase}_{12}(I))$. Let $\mathbf{R}(c_1, \dots, c_k)$ be an arbitrary fact of I , and let I' be a one-tuple instance that contains only this fact. It follows easily from (*) that $\widehat{I'} \subseteq \text{chase}_{21}(\text{chase}_{12}(I'))$. Therefore, $\widehat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$. As for the reverse inclusion, assume that it were to fail, so that $\text{chase}_{21}(\text{chase}_{12}(I))$ is not a subinstance of $\widehat{I}$. By the argument given in the proof of Theorem 11.1, 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 $\mathbf{R}$ 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. And (**) is a coNP property, since its negation is an NP property (again, we simply 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 the nodes). Let 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:

- (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 B.11, 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 $\mathcal{M}''_{12}$. From what we have said, this 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 B.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 nodes. Let $\mathbf{S}'_1$ consist of the n -ary relational symbol D , and let $\mathbf{S}'_2$ consist of the n -ary relational symbol D' and the binary relation symbol E . Label the nodes of G as $x_1, \dots, x_n$, and let ψ be the conjunction of all formulas $E(x_i, x_j)$ such that (x_i, x_j) is an edge of the graph G . Let Σ_{12} consist of the copying tgd $D(x_1, \dots, x_n) \rightarrow D'(x_1, \dots, x_n)$ and the tgd $D(x_1, \dots, x_n) \rightarrow \psi$. Let $y_1, \dots, y_k$ be k new variables, and let α be the conjunction of all formulas $E(y_i, y_j)$ where $i \neq j$. Let Σ'_{21} consist of the tgd $D'(x_1, \dots, x_n) \wedge \alpha \rightarrow \widehat{D}(x_1, \dots, x_n)$. 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 5.2, we need only show that G has a clique of size at least k if and only if for every source instance I , we have $\widehat{I} = \text{chase}'_{21}(\text{chase}'_{12}(I))$. 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 α to the edge relation E of G , then h is one-to-one. This is because if $i \neq j$, 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 α . Since G has no self-loops, it follows that $h(y_i) \neq h(y_j)$, as desired. Therefore if h is a homomorphism from α to the edge relation E of G , then G must have a clique of size at least k .

Assume that G has a clique of size at least k ; we must show that $\widehat{I} = \text{chase}'_{21}(\text{chase}'_{12}(I))$ for every source instance I . If I is the empty instance, then $\widehat{I}$ and $\text{chase}'_{21}(\text{chase}'_{12}(I))$ are both the empty instance and so $\widehat{I} = \text{chase}'_{21}(\text{chase}'_{12}(I))$. Let I be an arbitrary nonempty source instance. Then the

D' relation of $\text{chase}'_{12}(I)$ is a copy of the D relation of I , and the E relation of $\text{chase}'_{12}(I)$ is the edge relation of G . Since G has a clique of size at least k , it is easy to see that $\widehat{I} = \text{chase}'_{21}(\text{chase}'_{12}(I))$, as desired.

Assume that G does not have a clique of size at least k ; we must show that there is a source instance I such that $\widehat{I} \neq \text{chase}'_{21}(\text{chase}'_{12}(I))$. Let I be an arbitrary nonempty source instance. Then the E relation of $\text{chase}'_{12}(I)$ is the edge relation of G . So $\text{chase}'_{21}(\text{chase}'_{12}(I))$ is empty: this is because, as we noted, if there is a homomorphism from α to the edge relation E of G , then G must have a clique of size at least k . Hence $\widehat{I} \neq \text{chase}'_{21}(\text{chase}'_{12}(I))$, as desired.

□

The next two theorems give an upper bound on the complexity if the small model theorem were to hold even when we no longer restrict the tgds to be full.

THEOREM B.13. *Assume that Theorem 11.2 were to hold even when we no longer restrict the tgds to be full. Then 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 s-t tgds, whether $\mathcal{M}_{12}$ has a global inverse that is defined by a finite set of s-t tgds, would be in the complexity class Π_2^P in the polynomial-time hierarchy.*

PROOF. The proof begins like the proof of Theorem B.9. Do a polynomial-time check to verify that every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$ (otherwise, by Corollary 4.3, we know that there is no global inverse of $\mathcal{M}_{12}$). So assume that every value that appears in a tuple of I_R appears in a tuple of $\text{chase}_{12}(I_R)$, for each source relational symbol R . Let $\mathcal{M}_{12}^{-1} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{12}^{-1})$ be the canonical global inverse. We know from Theorem 7.3 that $\mathcal{M}_{12}$ has a global inverse that is defined by a finite set of s-t tgds if and only if $\mathcal{M}_{12}^{-1}$ is a global inverse of $\mathcal{M}_{12}$. We now show that the problem of determining if $\mathcal{M}_{12}$ does *not* have a global inverse that is defined by a finite set of s-t tgds is in the complexity class Σ_2^P in the polynomial-time hierarchy. So we must show that the following property is in Σ_2^P : “ $\mathcal{M}_{12}^{-1}$ is not a global inverse of $\mathcal{M}_{12}$ ”. By assumption, this holds precisely if the following statement holds:

(***) “There is a small I such that $\mathcal{M}_{12}^{-1}$ is not a global inverse of $\mathcal{M}_{12}$ for I ”.

Here “small” means polynomial in the size of Σ_{12} .

By Lemma B.8, we know that $\widehat{I} \subseteq \text{chase}_{12}^{-1}(\text{chase}_{12}(I))$. So by Theorem 5.1, we know that (***) holds precisely if $\langle I, \widehat{I} \rangle \not\models \Sigma_{12} \circ \Sigma_{12}^{-1}$. By Lemma B.1, this holds if and only if there is no small instance J such that $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \widehat{I} \rangle \models \Sigma_{12}^{-1}$. Here “small” means polynomial in the size of I , where the degree of the polynomial depends only on Σ_{12} . This gives us an Σ_2^P statement (“There exists a small I such that for every small J , we do not have $\langle I, J \rangle \models \Sigma_{12}$ and $\langle J, \widehat{I} \rangle \models \Sigma_{12}^{-1}$ ”), as desired. □

THEOREM B.14. *Assume that Theorem 11.2 were to hold even when we no longer restrict the tgds to be full. Then the problem of deciding, given schema mappings $\mathcal{M}_{12} = (\mathbf{S}_1, \mathbf{S}_2, \Sigma_{12})$ and $\mathcal{M}_{21} = (\mathbf{S}_2, \widehat{\mathbf{S}}_1, \Sigma_{21})$, with Σ_{12} and Σ_{21} finite sets of s-t tgds, whether $\mathcal{M}_{21}$ is a global inverse of $\mathcal{M}_{12}$, would be in the complexity class Π_2^P in the polynomial-time hierarchy.*

PROOF. We use the same argument as in the proof of Theorem B.13, except that we must also show that the property that “ $\hat{I} \subseteq \text{chase}_{21}(\text{chase}_{12}(I))$ ” holds for every $\mathbf{S}_1$ instance I ” is in Π_2^P . But as shown in the proof of Theorem B.12, this is equivalent to $(*)$ holding, which in turn is an NP property. $\square$