Quasi-Inverses of Schema Mappings

RONALD FAGIN, PHOKION G. KOLAITIS, and LUCIAN POPA
IBM Almaden Research Center
and
WANG-CHIEW TAN
University of California, Santa Cruz

Schema mappings are high-level specifications that describe the relationship between two database schemas. Two operators on schema mappings, namely the composition operator and the inverse operator, are regarded as especially important. Progress on the study of the inverse operator was not made until very recently, as even finding the exact semantics of this operator turned out to be a fairly delicate task. Furthermore, this notion is rather restrictive, since it is rare that a schema mapping possesses an inverse.

In this article, we introduce and study the notion of a quasi-inverse of a schema mapping. This notion is a principled relaxation of the notion of an inverse of a schema mapping; intuitively, it is obtained from the notion of an inverse by not differentiating between instances that are equivalent for data-exchange purposes. For schema mappings specified by source-to-target tuple-generating dependencies (s-t tgds), we give a necessary and sufficient combinatorial condition for the existence of a quasi-inverse, and then use this condition to obtain both positive and negative results about the existence of quasi-inverses. In particular, we show that every LAV (local-as-view) schema mapping has a quasi-inverse, but that there are schema mappings specified by full s-t tgds that have no quasi-inverse. After this, we study the language needed to express quasi-inverses of schema mappings specified by s-t tgds, and we obtain a complete characterization. We also characterize the language needed to express inverses of schema mappings, and thereby solve a problem left open in the earlier study of the inverse operator. Finally, we show that quasi-inverses can be used in many cases to recover the data that was exported by the original schema mapping when performing data exchange.

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

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Schema mapping, data exchange, data integration, metadata model management, inverse, quasi-inverse, dependencies, chase

P. G. Kolaitis is on leave from the University of California, Santa Cruz.
W.-C. Tan was supported in part by NSF CAREER Award IIS-0347065 and NSF grant IIS-0430994. Work partially done while at IBM Almaden Research Center.
Author's addresses: email: {fagin, kolaitis, lucian}@almaden.ibm.com; wctan@cse.ucsc.edu.

ACM Reference Format:
Fagin, R., Kolaitis, P. G., Popa L., and Tan, W.-C. 2008. Quasi-inverses of schema mappings. ACM Trans. Datab. Syst. 33, 2, Article 11 (June 2008), 52 pages. DOI = 10.1145/1366102.1366108 http://doi.acm.org/10.1145/1366102.1366108

1. INTRODUCTION

Schema mappings are high-level specifications that describe the relationship between two database schemas. More precisely, a schema mapping is a triple $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ consisting of a source schema $\mathbf{S}$, a target schema $\mathbf{T}$, and a set Σ of database dependencies that specify the relationship between the source schema and the target schema. Since schema mappings form the essential building blocks of such crucial data interoperability tasks as data integration and data exchange (see the surveys [Kolaitis 2005; Lenzerini 2002]), several different operators on schema mappings have been singled out as deserving study in their own right [Bernstein 2003]. The composition operator and the inverse operator have emerged as two of the most fundamental operators on schema mappings. Intuitively, the composition operator takes two schema mappings $\mathcal{M}$ and $\mathcal{M}'$ and produces a schema mapping that has the effect of applying first $\mathcal{M}$ and then $\mathcal{M}'$. The inverse operator takes a schema mapping $\mathcal{M}$ and produces a schema mapping $\mathcal{M}'$ such that, intuitively, if after applying $\mathcal{M}$ we then apply $\mathcal{M}'$, the resulting effect of $\mathcal{M}'$ is to "undo" the effect of $\mathcal{M}$.

By now, the composition operator has been investigated in depth [Fagin et al. 2005c; Madhavan and Halevy 2003; Melnik 2004; Nash et al. 2005]; however, progress on the study of the inverse operator was not made until very recently, as even finding the exact semantics of this operator turned out to be a delicate task. In Fagin [2007] the concept of an inverse of a schema mapping was rigorously defined and its basic properties were studied. The definition of an inverse was given by first defining the concept of the *identity* schema mapping Id and then stipulating that a schema mapping $\mathcal{M}'$ is an *inverse* of a schema mapping $\mathcal{M}$ if the composition of $\mathcal{M}$ with $\mathcal{M}'$ yields the identity schema mapping Id, in symbols $\mathcal{M} \circ \mathcal{M}' = \text{Id}$.

Unfortunately, the notion of an inverse of a schema mapping turned out to be rather restrictive, since it is rare that a schema mapping possesses an inverse. Indeed, as shown in Fagin [2007], if a schema mapping $\mathcal{M}$ is invertible, then $\mathcal{M}$ satisfies the *unique-solutions property*, which asserts that different source instances must have different *spaces of solutions* (that is, different sets of target instances satisfying the specifications of $\mathcal{M}$). The failure of this necessary condition for invertibility can be used as a simple, yet powerful, sufficient condition for non-invertibility. In particular, none of the following natural schema mappings possesses an inverse, because it is easy to see that none of them has the unique-solutions property:

Projection: This is the schema mapping specified by the dependency $P(x, y) \rightarrow Q(x)$.

Union: This is the schema mapping specified by the dependencies $P(x) \rightarrow S(x)$ and $Q(x) \rightarrow S(x)$.

Decomposition: This is the schema mapping specified by the dependency $P(x, y, z) \to Q(x, y) \wedge R(y, z)$.

Moreover, the invertibility of a schema mapping is not robust, as it is affected by changes to the source schema, even when the dependencies remain intact. Specifically, assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is an invertible schema mapping. If the source schema $\mathbf{S}$ is augmented with a new relation symbol R, then the new schema mapping $\mathcal{M}^* = (\mathbf{S} \cup \{R\}, \mathbf{T}, \Sigma)$ is no longer invertible.

In view of these limitations of the notion of an inverse of a schema mapping, it is natural to ask: is there a good alternative notion of an inverse that is not as restrictive as the original notion in Fagin [2007], but is still useful in data exchange? In what follows, we address this question by formulating the notion of a *quasi-inverse* of a schema mapping, by exploring its properties in depth, and by making a case for its usefulness.

Conceptual contributions. We introduce the notion of a *quasi-inverse* of a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ as a principled relaxation of the notion of an inverse mapping of $\mathcal{M}$. Intuitively, the notion of a quasi-inverse is obtained from the notion of an inverse by not differentiating between *ground* instances (null-free source instances) that are equivalent for data-exchange purposes. Formally, we first consider the equivalence relation $\sim_{\mathcal{M}}$ between ground instances such that $I_1 \sim_{\mathcal{M}} I_2$ holds if I_1 and I_2 have the same *space of solutions*, that is, for every target instance J, we have that $(I_1, J) \models \Sigma$ if and only if $(I_2, J) \models \Sigma$. We then say that a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \mathcal{M}')$ is a *quasi-inverse* of $\mathcal{M}$ if, in a precise technical sense, $\mathcal{M} \circ \mathcal{M}' = \mathrm{Id}$ holds modulo the equivalence relation $\sim_{\mathcal{M}}$.

We show that the concept of a quasi-inverse of a schema mapping is actually part of a unifying framework in which different relaxations of the notion of an inverse of a schema mapping can be obtained by using different equivalence relations that are *refinements* of the equivalence relation $\sim_{\mathcal{M}}$ (i.e., they are contained in $\sim_{\mathcal{M}}$). This framework captures, as special cases, both inverses and quasi-inverses. In fact, the notion of an inverse is the strictest one, while the notion of a quasi-inverse is the most relaxed one; all other relaxations of the notion of an inverse lie in between.

Numerous non-invertible schema mappings possess natural and useful quasi-inverses. Indeed, let us revisit the preceding examples of non-invertible schema mappings.

Projection: The schema mapping specified by $P(x, y) \to Q(x)$ has a quasi-inverse specified by $Q(x) \to \exists y P(x, y)$. Intuitively, this quasi-inverse describes the "best" you can do to recover ground instances.

Union: The schema mapping specified by the dependencies $P(x) \to S(x)$ and $Q(x) \to S(x)$ has a quasi-inverse specified by $S(x) \to P(x) \vee Q(x)$. Quasi-inverses need not be unique up to logical equivalence (the same holds true for inverses as well). Indeed, the schema mapping specified by $S(x) \to P(x)$ is also a quasi-inverse, as are the schema mapping specified by $S(x) \to Q(x)$ and the schema mapping specified by $S(x) \to P(x) \wedge Q(x)$.

Decomposition: The schema mapping specified by the dependency $P(x, y, z) \to Q(x, y) \wedge R(y, z)$ has a quasi-inverse specified by $Q(x, y) \wedge R(y, z) \to$

$P(x, y, z)$. Another quasi-inverse of this schema mapping is specified by $Q(x, y) \rightarrow \exists z P(x, y, z)$ and $R(y, z) \rightarrow \exists x P(x, y, z)$.

Finally, if $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is an invertible schema mapping and we augment $\mathbf{S}$ with a new relation symbol R, then *every* inverse of $\mathcal{M}$ is a quasi-inverse of the resulting non-invertible schema mapping $\mathcal{M}^* = (\mathbf{S} \cup \{R\}, \mathbf{T}, \Sigma)$. Moreover, if a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a quasi-inverse of a schema mapping $\mathcal{M}$, then the schema mapping $\mathcal{M}'' = (\mathbf{T}, \mathbf{S} \cup \{R\}, \Sigma')$ is a quasi-inverse of $\mathcal{M}^*$. Thus, unlike the notion of an inverse, the notion of a quasi-inverse is robust when relation symbols are added to the source schema.

Technical contributions. Our results span three different directions: an exact criterion for the existence of quasi-inverses, complete characterizations of languages needed to express quasi-inverses and inverses, and results on the use of quasi-inverses in data exchange.

Existence of quasi-inverses. For schema mappings specified by source-to-target tuple-generating dependencies (s-t tgds), we give a necessary and sufficient combinatorial condition, called the *subset property*, for the existence of a quasi-inverse. We apply this condition to obtain both positive and negative results about the existence of quasi-inverses. On the positive side, we use the subset property as a sufficient condition for quasi-invertibility to show that every LAV (local-as-view) schema mapping has a quasi-inverse; this result provides a unifying explanation for the quasi-invertibility of the *Projection*, *Union*, and *Decomposition* schema mappings. On the negative side, we use the subset property as a necessary condition for quasi-invertibility to show that there are simple schema mappings specified by full s-t tgds that have no quasi-inverse. We also show that a variation of the subset property is necessary and sufficient for the existence of an inverse.

The language of inverses and quasi-inverses. We investigate the language needed to express quasi-inverses of schema mappings specified by s-t tgds, and we obtain a complete characterization. Specifically, we show that if a schema mapping specified by a finite set of s-t tgds is quasi-invertible, then it has a quasi-inverse specified by a finite set of target-to-source disjunctive tgds with constants and inequalities (in fact, inequalities among constants suffice). Moreover, we give an exponential-time algorithm QuasiInverse (exponential in the size of the input schema mapping) for constructing such a quasi-inverse. The premise of a target-to-source disjunctive tgd with constants and inequalities is a conjunction of target atoms, formulas of the form *Constant*(x) that evaluate to true only when x is instantiated to a constant (nonnull) value, and inequalities $x_i \neq x_j$; the conclusion is a disjunction of conjunctive queries over the source. We show that our expressibility result is optimal by proving that *no* proper fragment of the language of disjunctive tgds with constants and inequalities suffices to express quasi-inverses; that is, both constants and inequalities in the premise of dependencies are needed, as are both disjunctions and existential quantifiers in the conclusion of dependencies. For schema mappings specified by a finite set of full s-t tgds, we show that if such a schema mapping is quasi-invertible, then it has a quasi-inverse specified by a finite set of target-to-source disjunctive tgds with inequalities; in other words, the predicate *Constant* is not needed in this case. We also show that every LAV schema mapping has a quasi-inverse specified by a finite set of target-to-source tgds with inequalities and constants; thus, in this case, there is no need for disjunctions in the conclusion of dependencies.

For schema mappings specified by s-t tgds, Fagin [2007] focused only on inverses specified by target-to-source tgds, and left open the problem of characterizing the language needed to express inverses of schema mappings. We settle this problem by showing that if a schema mapping specified by a finite set of s-t tgds is invertible, then it has an inverse specified by a finite set of target-to-source tgds with constants and inequalities. This turns out to be an optimal result as well.

Although we have characterized the language needed to express quasi-inverses and inverses of schema mappings specified by a finite set of s-t tgds, the complexity of deciding the existence of a quasi-inverse remains open. In fact, even decidability is open. The complexity of deciding the existence of an inverse, which was left open in Fagin [2007], was resolved in Fagin and Nash [2007], where it was shown to be coNP-complete.

Using quasi-inverses in data exchange. Since schema mappings rarely have inverses, we cannot hope to always obtain an exact copy of the original ground instance from target instances. The notion of a quasi-inverse is motivated by the idea that "similarity up to the space of solutions" is often good enough for data-exchange applications; hence, the definition of a quasi-inverse of a schema mapping $\mathcal{M}$ relaxes exact equality between ground instances to $\sim_{\mathcal{M}}$-equivalence. We show that, even though it is not possible to recover an exact copy of the original source instance, in many cases quasi-inverses allow us to recover a source instance that has "equivalent" properties from the data-exchange point of view.

More concretely, we introduce two additional notions, of *sound* and of *faithful*, which are relevant for data exchange, and then show how quasi-inverses relate to the two notions. Formally, assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping specified by a finite set of s-t tgds and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a "reverse" schema mapping specified by a finite set of target-to-source disjunctive tgds with constants and inequalities. We say that $\mathcal{M}'$ is sound with respect to $\mathcal{M}$ if the following property holds. Let I be an arbitrary ground instance and let U be the result of chasing I with Σ. Suppose we chase U back from target to source with the disjunctive dependencies in Σ' to obtain a set $\mathcal{V}$ of source instances and then we chase every member of $\mathcal{V}$ with the dependencies in Σ to obtain a set $\mathcal{U}'$ of target instances. Then $\mathcal{U}'$ contains a target instance U' that can be mapped homomorphically into U; thus, the target instance U' contains (up to a homomorphism) only facts of U, although not necessarily all of them. Furthermore, we say that $\mathcal{M}'$ is faithful with respect to $\mathcal{M}$ if the set $\mathcal{U}'$ contains a target instance U' that is homomorphically equivalent to U (thus, there are homomorphisms in both directions). In other words, there is a source instance V in $\mathcal{V}$ whose chase with Σ is homomorphically equivalent to U. This instance V is thus "data-exchange equivalent" to the original source instance I.

We then give two main results that relate quasi-inverses to the notions of soundness and faithfulness. The first main result asserts that if $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping specified by a finite set of s-t tgds and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a quasi-inverse of $\mathcal{M}$ specified by a finite set of target-to-source disjunctive tgds with constants and inequalities among constants, then $\mathcal{M}'$ is always sound. We also give an example that shows that not every such quasi-inverse is faithful, even though, by the previous result, it is sound. The second main result, however, asserts that if $\mathcal{M}'$ is obtained by applying our QuasiInverse algorithm for constructing a quasi-inverse of $\mathcal{M}$, then $\mathcal{M}'$ is faithful. In particular, if $\mathcal{M}$ is a schema mapping (not necessarily quasi-invertible) specified by a finite set of s-t tgds, there is always a schema mapping $\mathcal{M}'$ that is faithful with respect to $\mathcal{M}$.

Relation to Earlier Conference Version. This article is the full, extended version of Fagin et al. [2007]. It differs from Fagin et al. [2007] in the following ways: (a) It contains proofs of all results. Many of these proofs are technically challenging and introduce new techniques. (b) Section 3 contains several new results, namely Theorem 3.3 and all the results in Section 3.3. (c) Section 6 contains a number of new results. In particular, Theorem 6.12 extends Theorem 6.8 in Fagin et al. [2007]. In addition, Theorems 6.9 and 6.10 are new.

2. PRELIMINARIES

A *schema* $\mathbf{R}$ is a finite sequence $(R_1, \ldots, R_k)$ of relation symbols, each of a fixed arity. An *instance* I over $\mathbf{R}$ or, simply, an $\mathbf{R}$-*instance* is a sequence $(R_1^I, \ldots, R_k^I)$, where each R_i^I is a finite relation of the same arity as R_i and with entries from some set of values. We shall often use R_i to denote both the relation symbol and the relation R_i^I that interprets it. An *atom (over* $\mathbf{R}$*)* is a formula $P(v_1, \ldots, v_m)$, where P is a relation symbol in $\mathbf{R}$ and $v_1, \ldots, v_m$ are variables, not necessarily distinct. A *fact (over* $\mathbf{R}$*)* is a formula $P(w_1, \ldots, w_m)$, where P is a relation symbol in $\mathbf{R}$ and $w_1, \ldots, w_m$ are values, not necessarily distinct. It is often convenient to identify an instance with its set of facts.

Schema mappings. A *schema mapping* is a triple $\mathcal{M} = (\mathbf{R}_1, \mathbf{R}_2, \Sigma)$, where $\mathbf{R}_1$, $\mathbf{R}_2$ are schemas and Σ is a set of constraints describing the relationship between $\mathbf{R}_1$ and $\mathbf{R}_2$. We say that $\mathcal{M}$ is *specified by* Σ. Typically, constraints are formulas in some logical formalism; in this sense, a schema mapping is a syntactic object. The set of *instances of* $\mathcal{M}$ is

$$\text{Inst}(\mathcal{M}) = \{(I, J) : I \text{ is an } \mathbf{R}_1\text{-instance, } J \text{ is an } \mathbf{R}_2\text{-instance, and } (I, J) \models \Sigma\}.$$

For all practical purposes, a schema mapping $\mathcal{M} = (\mathbf{R}_1, \mathbf{R}_2, \Sigma)$ can be identified with the triple $\mathcal{M} = (\mathbf{R}_1, \mathbf{R}_2, \text{Inst}(\mathcal{M}))$; in this sense, a schema mapping is a semantic object specified by $\text{Inst}(\mathcal{M})$. Sometimes we give a schema mapping $\mathcal{M}$ as a semantic object by giving the two schemas $\mathbf{R}_1$, $\mathbf{R}_2$, and describing the (I, J) pairs that constitute $\text{Inst}(\mathcal{M})$.

Data Exchange. Let Const be a fixed infinite set of *constants* and let Var be an infinite set of *nulls* that is disjoint from Const. From now on, we assume that $\mathbf{S}$ and $\mathbf{T}$ are two fixed schemas. We call $\mathbf{S}$ the *source schema* and $\mathbf{T}$ the *target* *schema*. We assume that all **S**-instances have individual values from the set Const of constants only, while all target instances have individual values from Const ∪ Var. We refer to **S**-instances as *ground* **S**-*instances* or, simply, *ground instances* to emphasize the fact that all individual values in such instances are constants. Intuitively, schema mappings of the form $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ model the situation in which we perform data exchange from **S** to **T**: the individual values of source instances are known, while incomplete information in the specification of data exchange may give rise to null values in the target instances.

We review several notions from Fagin et al. [2005a] that will be needed in this article. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping. If I is a ground instance, then a *solution for* I under $\mathcal{M}$ is a target instance J such that $(I, J) \models \Sigma$. The set of all solutions for I under $\mathcal{M}$ is denoted by $\mathrm{Sol}(\mathcal{M}, I)$.

Let J, J' be two target instances. A function h from Const∪Var to Const∪Var is a *homomorphism* from J to J' if for every c in Const, we have that $h(c) = c$, and for every relation symbol R in **T** and every tuple $(a_1, \ldots, a_n) \in R^J$, we have that $(h(a_1), \ldots, h(a_n)) \in R^{J'}$. The instances J and J' are said to be *homomorphically equivalent* if there are homomorphisms from J to J' and from J' to J.

Given a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ and a ground instance I, a *universal solution* for I under $\mathcal{M}$ is a solution J for I under $\mathcal{M}$ such that for every solution J' for I under $\mathcal{M}$, there is a homomorphism $h : J \to J'$. Intuitively, universal solutions are the "most general" solutions among the space of all solutions for I.

A *source-to-target tuple-generating dependency* (or *s-t tgd*) is a first-order formula of the form $\forall \mathbf{x}(\varphi(\mathbf{x}) \to \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$, where $\varphi(\mathbf{x})$ is a conjunction of atoms over **S**, $\psi(\mathbf{x}, \mathbf{y})$ is a conjunction of atoms over **T**, and every variable in $\mathbf{x}$ occurs in an atom in $\varphi(\mathbf{x})$. (Not every variable in $\mathbf{x}$ needs to occur in $\psi(\mathbf{x}, \mathbf{y})$.) If there are no existential quantifiers, then the s-t tgd is *full*.

If $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping specified by a finite set Σ of s-t tgds, then *chasing*[1] I with Σ produces a target instance U such that U is a universal solution for I under $\mathcal{M}$. We often write $U = chase_\Sigma(I)$ and say that U is *the* result of the chase. (In general, there may be several such instances U but they are all homomorphically equivalent.)

Our goal in this paper is to investigate inverses and quasi-inverses of schema mappings $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds. In particular, we will identify the languages needed for expressing such inverses and quasi-inverses, and will show that these languages must be richer than the language of target-to-source tgds. The following definition introduces the richer classes of dependencies needed.

Definition 2.1. Let *Constant* be a relation symbol that is different from all relation symbols in **S** and **T**.

1. A *disjunctive tgd with constants and inequalities from* **T** *to* **S** is a first-order formula of the form $\forall \mathbf{x}(\varphi(\mathbf{x}) \to \bigvee_{i=1}^{n} \exists \mathbf{y}_i \psi_i(\mathbf{x}, \mathbf{y}_i))$, where:
   - the formula $\varphi(\mathbf{x})$ is a conjunction of
     (1) atoms over **T**, such that every variable in $\mathbf{x}$ occurs in one of them;
     (2) formulas of the form $Constant(x)$, where x is a variable in $\mathbf{x}$;

[1] A standard reference for the chase procedure is [Abiteboul et al. 1995].

(3) inequalities $x \neq x'$, where x and x' are variables in $\mathbf{x}$.

- Each formula $\psi_i(\mathbf{x}, \mathbf{y}_i)$ is a conjunction of atoms over $\mathbf{S}$.

Naturally, a formula *Constant*(x) evaluates to true if and only if x is interpreted by a value in Const.

2. A *disjunctive tgd with constants and inequalities among constants* is a disjunctive tgd with inequalities and constants where the formulas *Constant*(x) and *Constant*(x') occur as conjuncts of $\varphi(\mathbf{x})$ whenever the inequality $x \neq x'$ is a conjunct of $\varphi(\mathbf{x})$.

Clearly, disjunctive tgds with constants and inequalities extend the language of tgds with three features: (1) formulas of the form *Constant*(x) in the premise; (2) inequalities in the premise; and (3) disjunctions in the conclusion. If the conclusion consists of a single disjunct, then we talk about *tgds with constants and inequalities*. The concepts of *disjunctive tgds with inequalities*, *tgds with inequalities*, and other such special cases of Definition 2.1 are defined in an analogous way. For example,

$$P(x, y, z) \wedge Constant(x) \wedge x \neq y \rightarrow \exists w Q(x, w) \vee Q(x, y)$$

is a disjunctive tgd with constants and inequalities,

$$P(x, y, z) \wedge x \neq y \rightarrow \exists w Q(x, w) \vee Q(x, y)$$

is a disjunctive tgd with inequalities, and

$$P(x, y, z) \wedge x \neq y \rightarrow Q(x, y)$$

is a tgd with inequalities. Note that, for convenience, we have dropped the universal quantifiers in the front.

Composing and Inverting Schema Mappings. We recall the concept of the *composition* of two schema mappings, introduced in Fagin et al. [2004] and Melnik [2004], and the concept of an *inverse* of a schema mapping, introduced in Fagin [2007].

—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. The *composition* $\mathcal{M}_{12} \circ \mathcal{M}_{23}$ is a schema mapping $(\mathbf{S}_1, \mathbf{S}_3, \Sigma_{13})$ such that for every $\mathbf{S}_1$-instance I and every $\mathbf{S}_3$-instance K, we have that $(I, K) \models \Sigma_{13}$ if and only if there is an $\mathbf{S}_2$-instance J such that $(I, J) \models \Sigma_{12}$ and $(J, K) \models \Sigma_{23}$. When the schemas are understood from the context, we will often write $\Sigma_{12} \circ \Sigma_{23}$ for the composition $\mathcal{M}_{12} \circ \mathcal{M}_{23}$.

—Let $\widehat{\mathbf{S}}$ be a replica of the source schema $\mathbf{S}$, that is, for every relation symbol R of $\mathbf{S}$, the schema $\widehat{\mathbf{S}}$ contains a relation symbol $\widehat{R}$ that is not in $\mathbf{S}$ and has the same arity as R; moreover, $\widehat{\mathbf{S}}$-instances have individual values from the set Const of constants only. We also assume that $\widehat{R}$ and $\widehat{S}$ are distinct when R and S are distinct. Clearly, every ground instance I has a replica $\widehat{\mathbf{S}}$-instance $\widehat{I}$ that is also ground.

—The *identity schema mapping* is, by definition, the schema mapping Id $= (\mathbf{S}, \widehat{\mathbf{S}}, \Sigma_{\text{Id}})$, where Σ_{Id} consists of the dependencies $R(\mathbf{x}) \rightarrow \widehat{R}(\mathbf{x})$ as R ranges over the relation symbols in $\mathbf{S}$. Thus, Inst(Id) consists of all pairs (I_1, I_2) of a ground $\mathbf{S}$-instance I_1 and a ground $\widehat{\mathbf{S}}$-instance I_2 such that $\widehat{I_1} \subseteq I_2$.

—Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping. We say that a schema mapping $\mathcal{M}' = (\mathbf{T}, \widehat{\mathbf{S}}, \Sigma')$ is an *inverse* of $\mathcal{M}$ if Inst(Id) = Inst($\mathcal{M} \circ \mathcal{M}'$). This means that, for every pair (I_1, I_2) of a ground $\mathbf{S}$-instance I_1 and a ground $\widehat{\mathbf{S}}$-instance I_2, we have that $\widehat{I_1} \subseteq I_2$ if and only if there is a target instance J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$.

From now on and for notational simplicity, we will write $\mathbf{S}$ to also denote its replica $\widehat{\mathbf{S}}$; it will be clear from the context if we refer to $\mathbf{S}$ or to its replica. Moreover, we will use the same symbol to denote both a ground $\mathbf{S}$-instance I and its replica $\widehat{\mathbf{S}}$-instance $\widehat{I}$.

3. QUASI-INVERSES OF SCHEMA MAPPING: BASIC NOTIONS AND FACTS

In this section, we develop a unifying framework for defining and studying a spectrum of notions that relax the notion of an inverse of a schema mapping in a principled manner. The key idea is to group together ground instances that are *equivalent for data-exchange purposes*, and not differentiate between such equivalent instances. This idea is formalized by first introducing an equivalence relation on ground instances and then using it to define relaxations of the notion of inverse.

3.1 Data-Exchange Equivalent Ground Instances

Definition 3.1. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping and let I_1, I_2 be two ground instances.

If Sol($\mathcal{M}, I_1$) = Sol($\mathcal{M}, I_2$), then we say that *I_1 and I_2 are data-exchange equivalent with respect to $\mathcal{M}$*, and we write $I_1 \sim_{\mathcal{M}} I_2$ to denote this. When $\mathcal{M}$ is understood from the context, we may write $\sim$ in place of $\sim_{\mathcal{M}}$.

As an example, consider the *Union* schema mapping $\mathcal{M}$ specified by the dependencies $P(x) \rightarrow S(x)$ and $Q(x) \rightarrow S(x)$. It is easy to verify that if I_1 and I_2 are two ground instances, then $I_1 \sim_{\mathcal{M}} I_2$ if and only if $P^{I_1} \cup Q^{I_1} = P^{I_2} \cup Q^{I_2}$.

On the face of Definition 3.1, testing whether $I_1 \sim_{\mathcal{M}} I_2$ appears to be a difficult task, since the space of solutions of a ground instance may be an infinite set. We will show, however, that this task can be carried out in polynomial time if the schema mapping $\mathcal{M}$ is specified by a finite set of s-t tgds. The starting point is Part 2 of Proposition 2.6 in Fagin et al. [2005a], which yields the following characterization of the equivalence relation $\sim_{\mathcal{M}}$.

PROPOSITION 3.2. [Fagin et al. 2005a]. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping in which Σ is a finite set of s-t tgds. Let I_1, I_2 be two ground instances, J_1 a universal solution for I_1, and J_2 a universal solution for I_2. Then* Sol($\mathcal{M}, I_1$) $\subseteq$ Sol($\mathcal{M}, I_2$) *if and only if there is a homomorphism $h : J_2 \rightarrow J_1$. Consequently,* Sol($\mathcal{M}, I_1$) = Sol($\mathcal{M}, I_2$) *if and only if J_1 and J_2 are homomorphically equivalent.*

We are now ready to establish the first result of this article.

THEOREM 3.3. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping in which Σ is a finite set of s-t tgds. Then the equivalence relation $\sim_{\mathcal{M}}$ is decidable in polynomial time, that is to say, the following decision problem is solvable in polynomial time: given two source instances I_1 and I_2, is* Sol($\mathcal{M}, I_1$) = Sol($\mathcal{M}, I_2$)*?*

PROOF. To show that the equivalence relation $\sim_{\mathcal{M}}$ is decidable in polynomial time, it suffices to give a polynomial-time algorithm for the following problem: given ground instances I_1 and I_2, is $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$? In turn, by Proposition 3.2, this problem is equivalent to: given ground instances I_1 and I_2, is there a homomorphism from the canonical universal solution J_1 for I_1 to the canonical universal solution J_2 for I_2? As in Fagin et al. [2005a], the *canonical universal solution* of a source instance I is the target instance obtained by chasing I with the dependencies Σ of the schema mapping $\mathcal{M}$.

Let J be a target instance. The *Gaifman graph of the nulls of J* is defined as follows: (i) the nodes of this graph are the nulls of J; (ii) there is an edge between two nulls of J if and only if they appear together in some fact of J. A *block of J* is a connected component of the Gaifman graph of the nulls of J. These concepts were introduced in Fagin et al. [2005b] and used there to design a polynomial time algorithm for computing the core of the universal solutions in schema mappings specified by s-t tgds. In what follows, we use properties of blocks to design a greedy, backtrack-free algorithm for the problem: given ground instances I_1 and I_2, is there a homomorphism from the canonical universal solution J_1 for I_1 to the canonical universal solution J_2 for I_2?

Algorithm GREEDY BLOCKS ALGORITHM for the schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$.

Input: Ground instances I_1 and I_2.
Output: A homomorphism $h : J_1 \to J_2$ from the canonical universal solution J_1 for I_1 to the canonical universal solution J_2 for I_2, if such a homomorphism exists ; "No," if no such homomorphism exists.

(1) Compute the canonical universal solutions J_1 and J_2 for I_1 and I_2, respectively.
(2) Compute the blocks of J_1.
(3) For every block B of J_1, use exhaustive search to determine if there is a homomorphism from $J_1[B]$ to J_2; here, $J_1[B]$ is the subinstance of J_1 consisting of the facts $P(w_1, \ldots, w_m)$ of J_1, where each w_i is a constant of J_1 or a null of B.
   (a) If such a homomorphism h_B exists, then keep this homomorphism and proceed to the next block of J_1.
   (b) If no such homomorphism is found, then exit and return "No".
(4) Return the union of all homomorphisms h_B found in Step 3. More formally, return the following function $h : J_1 \to J_2$:

$$h(x) = \begin{cases} h_B(x) & \text{if } x \in B; \\ x & \text{if } x \text{ is a constant.} \end{cases}$$

Note that the function h is well defined, since blocks are disjoint; moreover, it extends every function h_B found by the algorithm, since homomorphisms map each constant to itself. To prove the correctness of the algorithm, assume first that the algorithm terminates by returning a function $h : J_1 \to J_2$. Let $R(w_1, \ldots, w_k)$ be a fact of J_1. Then all the nulls occurring in $\{w_1, \ldots, w_k\}$ must be in the same block, say, B of J_1. Consider the homomorphism $h_B : J_1[B] \to J_2$ found by the algorithm. It follows that $h(w_i) = h_B(w_i)$, for $1 \leq i \leq k$, and hence $R(h(w_1), \ldots, h(w_k))$ is a fact of J_2, since h_B is a homomorphism from $J_1[B]$ to J_2. For the other direction of correctness, observe that if $g : J_1 \to J_2$ is a homomorphism and B is a block of J_1, then the restriction of g on $J_1[B]$ is a homomorphism from $J_1[B]$ to J_2. Thus, if a homomorphism from J_1 to J_2 exists, then the algorithm will indeed produce such a homomorphism.

It remains to show that the running time of the algorithm is polynomial in the sizes n_1 and n_2 of I_1 and I_2, respectively, where the *size* of an instance is the number of its facts. Let a be the maximum number of universally quantified variables over all s-t tgds in Σ, and let b be the maximum number of existentially quantified variables over all s-t tgds in Σ. Since Σ is fixed, both a and b are constants. It is not hard to show that the size of J_i is $O(n_i^a)$, $i = 1, 2$. Moreover, each block in J_1 has at most b elements. It follows that J_1 has $O(n_1^a)$ many blocks. For every block B of J_1, the algorithm takes $O((n_2^a)^b) = O(n_2^{ab})$ steps to exhaustively search over all potential homomorphisms from $J_1[B]$ to J_2, and, for each such function it takes $O(n_1^a n_2^a)$ steps to test if the function is indeed a homomorphism. So, altogether the algorithm takes $O(n_1^a n_1^a n_2^a n_2^{ab}) = O(n_1^{2a} n_2^{a+ab})$ steps, which shows that it is of polynomial running time. □

Note that the preceding Theorem 3.3 is about the *data complexity* of the equivalence relation $\sim_{\mathcal{M}}$, because the schema mapping $\mathcal{M}$ is kept fixed. One can also consider the *combined complexity* of the equivalence relation $\sim_{\mathcal{M}}$, that is, the complexity of the following decision problem: given a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds, and two source instances I_1 and I_2, is $I_1 \sim_{\mathcal{M}} I_2$? The analysis of the running time of the algorithm in the proof of Theorem 3.3 shows that the combined complexity of $\sim_{\mathcal{M}}$ is in exponential time.

3.2 $(\sim_1, \sim_2)$-Inverses: A Unifying Framework

Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping. Recall that a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is an *inverse* of $\mathcal{M}$ if $\text{Inst(Id)} = \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. This means that, for every pair (I_1, I_2) of a ground $\mathbf{S}$-instance I_1 and a ground $\mathbf{S}$-instance I_2, we have that $I_1 \subseteq I_2$ if and only if there is a target instance J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$.

Suppose now that $\sim_1$ and $\sim_2$ are two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. We introduce the notion of a $(\sim_1, \sim_2)$-*inverse* of $\mathcal{M}$, which, intuitively, formalizes the idea that the equation $\text{Inst(Id)} = \text{Inst}(\mathcal{M} \circ \mathcal{M}')$ holds *modulo* the equivalence relations $\sim_1$ and $\sim_2$. In what follows, we will use the notation $\sim_{(1,2)}$ to denote the *product* equivalence relation of two equivalence relations $\sim_1$ and $\sim_2$. Thus, $(I_1, I_2) \sim_{(1,2)} (I_1', I_2')$ if and only if $I_1 \sim_1 I_1'$ and $I_2 \sim_2 I_2'$.

Definition 3.4. Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping and $\sim_1$, $\sim_2$ are two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. We say that a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a $(\sim_1, \sim_2)$-*inverse* of $\mathcal{M}$ if, for every pair (I_1, I_2) of ground instances, the following statements are equivalent:

1. There are ground instances I_1' and I_2' such that $(I_1, I_2) \sim_{(1,2)} (I_1', I_2')$ and $I_1' \subseteq I_2'$.
2. There are ground instances I_1'' and I_2'' and a target instance J such that $(I_1, I_2) \sim_{(1,2)} (I_1'', I_2'')$, $(I_1'', J) \models \Sigma$, and $(J, I_2'') \models \Sigma'$.

Thus, $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ precisely if

$$\sim_1 \circ \text{Inst(Id)} \circ \sim_2 \;=\; \sim_1 \circ \text{Inst}(\mathcal{M} \circ \mathcal{M}') \circ \sim_2,$$

where the occurrence of the symbol $\circ$ in the expression $\mathcal{M} \circ \mathcal{M}'$ denotes the composition operator on schema mappings, while all other occurrences of this symbol denote the composition of two binary relations.

Clearly, a schema mapping $\mathcal{M}'$ is an inverse of a schema mapping $\mathcal{M}$ if and only if $\mathcal{M}'$ is a $(=,=)$-inverse of $\mathcal{M}$. In this paper, our main focus will be on $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-inverses, which from now on will be referred to as *quasi-inverses*. Intuitively, the concept of a quasi-inverse relaxes the equation $\text{Inst(Id)} = \text{Inst}(\mathcal{M} \circ \mathcal{M}')$ by not distinguishing between ground instances that are data-exchange equivalent with respect to $\mathcal{M}$.

Definition 3.5. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping. A schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a *quasi-inverse* of $\mathcal{M}$ if $\mathcal{M}'$ is a $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-inverse of $\mathcal{M}$, that is, if, for every pair (I_1, I_2) of ground instances, the following statements are equivalent:

1. There are ground instances I_1' and I_2' such that $I_1 \sim_{\mathcal{M}} I_1'$, $I_2 \sim_{\mathcal{M}} I_2'$, and $I_1' \subseteq I_2'$.
2. There are ground instances I_1'' and I_2'' and a target instance J such that $I_1 \sim_{\mathcal{M}} I_1''$, $I_2 \sim_{\mathcal{M}} I_2''$, $(I_1'', J) \models \Sigma$, and $(J, I_2'') \models \Sigma'$.

So as before, $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a quasi-inverse of $\mathcal{M}$ precisely if

$$\sim_{\mathcal{M}} \circ \text{Inst(Id)} \circ \sim_{\mathcal{M}} \; = \; \sim_{\mathcal{M}} \circ \text{Inst}(\mathcal{M} \circ \mathcal{M}') \circ \sim_{\mathcal{M}} .$$

We say that $\mathcal{M}$ is *quasi-invertible* if it has a quasi-inverse, and *invertible* if it has an inverse.

Example 3.6. To illustrate these concepts, consider again the *Union* schema mapping $\mathcal{M}$ specified by the tgds $P(x) \to S(x)$ and $Q(x) \to S(x)$. As mentioned in the Introduction, this mapping is not invertible; it is not hard to verify, however, that the schema mapping $\mathcal{M}'$ specified by the dependency $S(x) \to P(x) \vee Q(x)$ is a quasi-inverse of $\mathcal{M}$. For instance, assume that I_1, I_2 are ground instances for which there are ground instances I_1'', I_2'' and a target instance J such that $I_1 \sim_{\mathcal{M}} I_1''$, $I_2 \sim_{\mathcal{M}} I_2''$, (I_1'', J) satisfies the tgds $P(x) \to S(x)$ and $Q(x) \to S(x)$, and (J, I_2'') satisfies the dependency $S(x) \to P(x) \vee Q(x)$. It follows that $P^{I_1''} \cup Q^{I_1''} \subseteq S^J \subseteq P^{I_2''} \cup Q^{I_2''}$. If we take $I_1' = I_1''$ and $I_2' = I_1'' \cup I_2''$, then clearly $I_1 \sim_{\mathcal{M}} I_1'$, $I_2 \sim_{\mathcal{M}} I_2'$, and $I_1' \subseteq I_2'$, as desired.

Example 3.7. Let $\mathcal{M}$ be the *Decomposition* schema mapping specified by the tgd $P(x, y, z) \to Q(x, y) \wedge R(y, z)$. As mentioned in the Introduction, this mapping is not invertible. However, it is not hard to show that the schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ with Σ' consisting of the tgd $Q(x, y) \wedge R(y, z) \to P(x, y, z)$ is a quasi-inverse of $\mathcal{M}$. It is also not hard to show that another quasi-inverse of $\mathcal{M}$ is the schema mapping $\mathcal{M}'' = (\mathbf{T}, \mathbf{S}, \Sigma'')$, where Σ'' consists of the tgds $Q(x, y) \to \exists z P(x, y, z)$ and $R(y, z) \to \exists x P(x, y, z)$. This also shows that a quasi-inverse of a schema mapping need not be unique up to logical equivalence. The same is true for inverses [Fagin 2007].

Next, we establish certain elementary, but quite useful, facts about $(\sim_1, \sim_2)$-inverses. We begin by introducing an auxiliary concept.

Definition 3.8. Let $\sim_1$ and $\sim_2$ be two equivalence relations on ground instances, and let D be a binary relation between ground instances. We set

$$D[\sim_1, \sim_2] \;=\; \sim_1 \circ D \circ \sim_2,$$

where the symbol $\circ$ denotes the composition of two binary relations. This means that

$$D[\sim_1, \sim_2] \;=\; \{(I_1, I_2) : \exists I_1' I_2'((I_1, I_2) \sim_{(1,2)} (I_1', I_2') \wedge (I_1', I_2') \in D)\}.$$

Note that, using this notation, $\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ if and only if

$$\mathrm{Inst}(\mathrm{Id})[\sim_1, \sim_2] \;=\; \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2].$$

Definition 3.8 can be viewed as a transformation that takes as input a triple $(D, \sim_1, \sim_2)$ and returns the binary relation $D[\sim_1, \sim_2]$ as output. The next proposition, which follows easily from the definitions and the properties of equivalence relations, states some basic properties of this transformation.

PROPOSITION 3.9. *Let D and D' be binary relations on ground instances, and let $\sim_1$, $\sim_2$, $\sim_3$, $\sim_4$ be equivalence relations on ground instances. The following statements hold.*

1. $D \subseteq D[\sim_1, \sim_2]$.
2. Monotonicity Property I: *If $D \subseteq D'$, then* $D[\sim_1, \sim_2] \subseteq D'[\sim_1, \sim_2]$.
3. Monotonicity Property II: *If $\sim_1 \subseteq \sim_3$ and $\sim_2 \subseteq \sim_4$, then* $D[\sim_1, \sim_2] \subseteq D[\sim_3, \sim_4]$.
4. Idempotence Property: $(D[\sim_1, \sim_2])[\sim_1, \sim_2] = D[\sim_1, \sim_2]$.
5. $D[\sim_1, \sim_2] \subseteq D'[\sim_1, \sim_2]$ *if and only if* $D \subseteq D'[\sim_1, \sim_2]$.
6. $D[\sim_1, \sim_2] = D'[\sim_1, \sim_2]$ *if and only if* $D \subseteq D'[\sim_1, \sim_2]$ *and* $D' \subseteq D[\sim_1, \sim_2]$.

The following fact is an immediate consequence of Part (6) of Proposition 3.9. In what follows, it will be used repeatedly in the proofs of several theorems.

COROLLARY 3.10. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be two schema mappings and let $\sim_1$, $\sim_2$ be two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. Then the following are equivalent.*

1. *$\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, i.e.,* $\mathrm{Inst}(\mathrm{Id})[\sim_1, \sim_2] = \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$.
2. $\mathrm{Inst}(\mathrm{Id}) \subseteq \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$ *and* $\mathrm{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \mathrm{Inst}(\mathrm{Id})[\sim_1, \sim_2]$.

By varying the equivalence relations $\sim_1$ and $\sim_2$, we can obtain a variety of $(\sim_1, \sim_2)$-inverses. The next proposition provides a tool for comparing them.

PROPOSITION 3.11. *Let $\mathcal{M}$ be a schema mapping and let $\sim_1$, $\sim_2$, $\sim_3$, $\sim_4$ be four equivalence relations on ground instances such that $\sim_1 \subseteq \sim_3 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_4 \subseteq \sim_{\mathcal{M}}$. Every $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ is also a $(\sim_3, \sim_4)$-inverse of $\mathcal{M}$.*

PROOF. Assume that $\mathrm{Inst}(\mathrm{Id})[\sim_1, \sim_2] = \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$. We must show that $\mathrm{Inst}(\mathrm{Id})[\sim_3, \sim_4] = \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_3, \sim_4]$. By Corollary 3.10, it suffices to show that $\mathrm{Inst}(\mathrm{Id}) \subseteq \mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_3, \sim_4]$ and $\mathrm{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \mathrm{Inst}(\mathrm{Id})[\sim_3, \sim_4]$. These hold because, using the hypothesis and Parts (1) and (3) of Proposition 3.9, we have that

$$\text{Inst}(\text{Id}) \subseteq \text{Inst}(\text{Id})[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2] \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_3, \sim_4]$$

and

$$\text{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2] = \text{Inst}(\text{Id})[\sim_1, \sim_2] \subseteq \text{Inst}(\text{Id})[\sim_3, \sim_4].$$

□

From Proposition 3.11, we see that the spectrum of $(\sim_1, \sim_2)$-inverses has both "strongest" and "weakest" elements. Indeed, if $\mathcal{M}'$ is an $(=, =)$-inverse of $\mathcal{M}$ (i.e., $\mathcal{M}'$ is an inverse of $\mathcal{M}$), then $\mathcal{M}'$ is also a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, for every two equivalence relations $\sim_1$ and $\sim_2$ contained in $\sim_{\mathcal{M}}$. At the other end of the spectrum, if $\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, then $\mathcal{M}'$ is also a $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-inverse of $\mathcal{M}$.

3.3 The Most General $(\sim_1, \sim_2)$-Inverse and the Candidate $(\sim_1, \sim_2)$-Inverse

Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping from the source schema $\mathbf{S}$ to the target schema $\mathbf{T}$. In this section, we first point out that if $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse, then it has a "most general" $(\sim_1, \sim_2)$-inverse. After this, we give a schema mapping $\mathcal{M}^* = (\mathbf{T}, \mathbf{S}, \Sigma^*)$, which we shall refer to as the *candidate* $(\sim_1, \sim_2)$*-inverse of* $\mathcal{M}$, that has the following property: if $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse, then $\mathcal{M}^*$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ and, in fact, is the "most general one". Note that if $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a schema mapping from the target schema $\mathbf{T}$ to the source schema $\mathbf{S}$, then

$$\begin{aligned}\text{Inst}(\mathcal{M}') = \{&(J, I) : J \text{ is a } \mathbf{T}\text{-instance}, I \text{ is a ground } \mathbf{S}\text{-instance},\\ &\text{and } (J, I)) \models \Sigma'\}\end{aligned}$$

Definition 3.12. Let $\mathcal{M}_1 = (\mathbf{T}, \mathbf{S}, \Sigma_1)$ and $\mathcal{M}_2 = (\mathbf{T}, \mathbf{S}, \Sigma_2)$ be two schema mappings from $\mathbf{T}$ to $\mathbf{S}$.

—We say that $\mathcal{M}_2$ *contains* $\mathcal{M}_1$, or that $\mathcal{M}_2$ *is more general than* $\mathcal{M}_1$, if
$$\text{Inst}(\mathcal{M}_1) \subseteq \text{Inst}(\mathcal{M}_2).$$
—We say that a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is the *union* of $\mathcal{M}_1$ and $\mathcal{M}_2$ if
$$\text{Inst}(\mathcal{M}') = \text{Inst}(\mathcal{M}_1) \cup \text{Inst}(\mathcal{M}_2).$$
In other words, for every $\mathbf{T}$-instance J and every ground $\mathbf{S}$-instance I, we have that $(J, I) \models \Sigma'$ if and only if $(J, I) \models \Sigma_1$ or $(J, I) \models \Sigma_2$.

—In general, if $\{\mathcal{M}_k : k \in K\}$ is a set of schema mappings from $\mathbf{T}$ to $\mathbf{S}$, then the *union* of the schema mappings in this set is a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ such that
$$\text{Inst}(\mathcal{M}') = \bigcup_{k \in K} \text{Inst}(\mathcal{M}_k).$$
We write $\bigcup_{k \in K} \mathcal{M}_k$ to denote the union of all schema mappings in $\{\mathcal{M}_k : k \in K\}$.

Note that "$\mathcal{M}_2$ contains $\mathcal{M}_1$" means that whenever J is a target instance and I is a ground instance such that $(J, I) \models \Sigma_1$, then $(J, I) \models \Sigma_2$. This is a notion of *weak implication*, a slightly weaker notion than logical implication, which would make the same requirement even when I is a source instance that is not ground. For $(\sim_1, \sim_2)$-inverses, weak implication is a natural notion, since we care only about ground source instances. Note also that, when referring to *the* union of two schema mappings or to *the* union of a set of schema mappings, we are viewing schema mappings as semantic objects described by giving two schemas and a set of pairs of instances from these two schemas.

PROPOSITION 3.13. *Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping and $\sim_1$, $\sim_2$ are two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. If $\{\mathcal{M}_k : k \in K\}$ is a nonempty set of schema mappings from $\mathbf{T}$ to $\mathbf{S}$ such that each $\mathcal{M}_k$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, then the union $\bigcup_{k \in K} \mathcal{M}_k$ is also a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$.*

PROOF. Use Corollary 3.10 and the easily verifiable fact that $\text{Inst}(\mathcal{M} \circ (\bigcup_{k \in K} \mathcal{M}_k)) = \bigcup_{k \in K} \text{Inst}(\mathcal{M} \circ \mathcal{M}_k)$. □

COROLLARY 3.14. *Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping and $\sim_1$, $\sim_2$ are two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. If $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse, then the union $\mathcal{M}^*$ of all $(\sim_1, \sim_2)$-inverses of $\mathcal{M}$ is also a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$. Moreover, $\mathcal{M}^*$ contains every $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, and hence $\mathcal{M}^*$ is the most general $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$.*

Next, given a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, we describe a schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \mathcal{M}')$ such that if $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse, then $\mathcal{M}'$ is the most general $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$. Let I be a ground instance. Let $\tau_{(\sim_1, \sim_2)}(\Sigma, I)$ define the set of all pairs (J, I_2) such that J is a target instance, I_2 is a ground instance, and such that if $(I, J) \models \Sigma$, then there are ground instances I' and I_2' such that $(I, I_2) \sim_{(1,2)} (I', I_2')$ and $I' \subseteq I_2'$. Later, we refer to $\tau_{(\sim_1, \sim_2)}(\Sigma, I)$ as *the constraint associated with I*. Let $cand_{(\sim_1, \sim_2)}(\Sigma)$ be the set of all constraints $\tau_{(\sim_1, \sim_2)}(\Sigma, I)$ (one for every ground instance I). Thus, $(J, I_2) \models cand_{(\sim_1, \sim_2)}(\Sigma)$ precisely if I_2 is a ground instance and for every ground instance I such that $(I, J) \models \Sigma$, there are ground instances I' and I_2' such that $(I, I_2) \sim_{(1,2)} (I', I_2')$ and $I' \subseteq I_2'$. Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, cand_{(\sim_1, \sim_2)}(\Sigma))$; in what follows, we refer to $\mathcal{M}'$ as the *candidate* $(\sim_1, \sim_2)$*-inverse* of $\mathcal{M}$.

THEOREM 3.15. *Assume that $\mathcal{M}$ is a schema mapping that has a $(\sim_1, \sim_2)$-inverse. Then the candidate $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$; moreover, it contains every $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, and hence is the most general $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$.*

PROOF. Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be the candidate $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ we have defined. Let $\mathcal{M}'' = (\mathbf{T}, \mathbf{S}, \Sigma'')$ be a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$. Since $\mathcal{M}''$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$, we know that $\text{Inst(Id)}[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}'')[\sim_1, \sim_2]$. We must show that $\text{Inst(Id)}[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$ holds, and also that $\mathcal{M}'$ contains $\mathcal{M}''$.

We first show that $\mathcal{M}'$ contains $\mathcal{M}''$. Assume not; we shall derive a contradiction. Since $\mathcal{M}'$ does not contain $\mathcal{M}''$, there is some pair (J, I_2) such that I_2 is a ground instance and $(J, I_2) \models \Sigma''$ but $(J, I_2) \not\models \Sigma'$. Since $(J, I_2) \not\models \Sigma'$, there is some ground instance I such that $(I, J) \models \Sigma$, but there is no choice of I', I'_2 such that $(I, I_2) \sim_{(1,2)} (I', I'_2)$ and $I' \subseteq I'_2$. Since $(I, J) \models \Sigma$ and $(J, I_2) \models \Sigma''$, it follows that (I, I_2) is in Inst($\mathcal{M} \circ \mathcal{M}''$)$[\sim_1, \sim_2]$. However, since there is no choice of I', I'_2 such that $(I, I_2) \sim_{(1,2)} (I', I'_2)$ and $I' \subseteq I'_2$, it follows that (I, I_2) is not in Inst(Id)$[\sim_1, \sim_2]$. This contradicts the hypothesis that $\mathcal{M}''$ is a $(\sim_1, \sim_2)$-inverse-inverse of $\mathcal{M}$. So $\mathcal{M}'$ contains $\mathcal{M}''$.

Since $\mathcal{M}'$ contains $\mathcal{M}''$, we have that Inst($\mathcal{M} \circ \mathcal{M}''$)$[\sim_1, \sim_2] \subseteq$ Inst($\mathcal{M} \circ \mathcal{M}'$)$[\sim_1, \sim_2]$. Since Inst(Id)$[\sim_1, \sim_2] =$ Inst($\mathcal{M} \circ \mathcal{M}''$)$[\sim_1, \sim_2]$, we also have that Inst(Id)$[\sim_1, \sim_2] \subseteq$ Inst($\mathcal{M} \circ \mathcal{M}'$)$[\sim_1, \sim_2]$. So, to prove Inst(Id)$[\sim_1, \sim_2] =$ Inst($\Sigma \circ \Sigma'$)$[\sim_1, \sim_2]$, we need only show that Inst($\mathcal{M} \circ \mathcal{M}'$)$[\sim_1, \sim_2] \subseteq$ Inst(Id)$[\sim_1, \sim_2]$. By Part (5) of Proposition 3.9, it is sufficient to show Inst($\mathcal{M} \circ \mathcal{M}'$) $\subseteq$ Inst(Id) $[\sim_1, \sim_2]$.

If $(I_1, I_2) \in$ Inst($\mathcal{M} \circ \mathcal{M}'$), there is J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. If we let I_1 play the role of I in the definition of Σ', we see that there are I'_1, I'_2 such that $(I_1, I_2) \sim_{(1,2)} (I'_1, I'_2)$ and $I'_1 \subseteq I'_2$. So $(I_1, I_2) \in$ Inst(Id)$[\sim_1, \sim_2]$, as desired. □

In Section 5.1, we apply Theorem 3.15 to inverses (i.e., $(=, =)$-inverses), to prove that if a schema mapping specified by a finite set of s-t tgds has an inverse, then it has an inverse specified by a finite set of full s-t tgds with constants and inequalities.

3.4 The $(\sim_1, \sim_2)$-Subset Property and Its Applications

In this section, we introduce a combinatorial property, which we call the $(\sim_1, \sim_2)$-*subset property*, and show that it gives an exact criterion for the existence of $(\sim_1, \sim_2)$-inverses; in particular, we obtain necessary and sufficient conditions for the existence of quasi-inverses and for the existence of inverses.

Definition 3.16. Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping and $\sim_1$, $\sim_2$ are two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$.

- We say that $\mathcal{M}$ has the $(\sim_1, \sim_2)$-*subset property* if for every pair (I_1, I_2) of ground instances such that Sol($\mathcal{M}, I_2$) $\subseteq$ Sol($\mathcal{M}, I_1$), there is a pair (I'_1, I'_2) of ground instances such that $(I_1, I_2) \sim_{(1,2)} (I'_1, I'_2)$ and $I'_1 \subseteq I'_2$.
- In particular, a schema mapping $\mathcal{M}$ has the $(=, =)$-*subset property* if for every pair (I_1, I_2) of ground instances such that Sol($\mathcal{M}, I_2$) $\subseteq$ Sol($\mathcal{M}, I_1$), we have that $I_1 \subseteq I_2$.

Before stating any technical results, let us give some insight to the $(=, =)$-subset property and to the $(\sim_1, \sim_2)$-subset property. Using Proposition 3.2, it is easy to see that if Σ is a set of s-t tgds and I_1, I_2 are two ground instances such that $I_1 \subseteq I_2$, then Sol($\mathcal{M}, I_2$) $\subseteq$ Sol($\mathcal{M}, I_1$). Thus, the $(=, =)$-subset property is the converse of this fact. More generally, the $(\sim_1, \sim_2)$-subset property is the converse of the following fact, which follows also from Proposition 3.2: if I_1, I_2 are two ground instances such that there are two ground instances I'_1, I'_2 with $(I_1, I_2) \sim_{(1,2)} (I'_1, I'_2)$ and $I'_1 \subseteq I'_2$, then Sol($\mathcal{M}, I_2$) $\subseteq$ Sol($\mathcal{M}, I_1$).

We now give two examples: one of a schema mapping $\mathcal{M}$ that has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property, and one of a (full) schema mapping $\mathcal{M}$ that does not have the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property.

Example 3.17. Let $\mathcal{M}$ be the *Decomposition* schema mapping of Example 3.7. We now show that $\mathcal{M}$ has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property. To see this, let I_1 and I_2 be two ground instances such that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. Let J be the solution for I_2 obtained by taking $Q^J = \pi_{12}(P^{I_2})$ and $R^J = \pi_{23}(P^{I_2})$.[2] Since $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$, we have that J is also in $\text{Sol}(\mathcal{M}, I_1)$, so $\pi_{12}(P^{I_1}) \subseteq \pi_{12}(P^{I_2})$ and $\pi_{23}(P^{I_1}) \subseteq \pi_{23}(P^{I_2})$. Let $I_2' = I_1 \cup I_2$. From the two inclusions we have just established, it follows that $I_2' \sim_\mathcal{M} I_2$; moreover, we have that $I_1 \subseteq I_2'$. This shows that $\mathcal{M}$ has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property (actually, this shows that $\mathcal{M}$ has the stronger $(=, \sim_\mathcal{M})$-subset property). □

Example 3.18. Let $\mathcal{M}$ be the schema mapping specified by the full s-t tgd $E(x, z) \wedge E(z, y) \rightarrow F(x, y) \wedge M(z)$. We claim that $\mathcal{M}$ does not have the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property. Towards the claim, consider the following two ground instances I_1 and I_2:

$$I_1 = \{E(1,4), E(4,3), E(1,2), E(2,5), E(4,2)\} \quad \text{and} \quad I_2 = \{E(1,2), E(2,3), E(1,3)\}.$$

For $i = 1, 2$, let J_i be the result of the chase of I_i with Σ. It is easy to see that

$$J_1 = \{F(1,3), F(1,5), F(4,5), F(1,2), M(2), M(4)\} \quad \text{and} \quad J_2 = \{F(1,3), M(2)\}.$$

Since $J_2 \subseteq J_1$ and since J_i is a universal solution for I_i, $i = 1, 2$, we have that $\text{Sol}(\mathcal{M}, I_1) \subseteq \text{Sol}(\mathcal{M}, I_2)$. We will show that there are do not exist ground instances I_1' and I_2' such that $I_1 \sim_\mathcal{M} I_1'$, $I_2 \sim_\mathcal{M} I_2'$, and $I_2' \subseteq I_1'$.

Towards a contradiction, assume that such ground instances I_1' and I_2' do exist. For $i = 1, 2$, let J_i' be the result of the chase of I_i' with Σ. Since $I_i \sim_\mathcal{M} I_i'$, we must have that $J_i = J_i'$, for $i = 1, 2$. In particular, $F(1, 3) \in J_2'$, which implies that there is an element n in the active domain of I_2' such that $E(1, n) \in I_2'$ and $E(n, 3) \in I_2'$. In turn, this implies that $M(n) \in J_2' = J_2$, which forces $n = 2$. Consequently, we have that $I_2 \subseteq I_2'$, which implies that $I_2 \subseteq I_1'$ (since $I_2' \subseteq I_1'$). Let us now focus attention on I_1'. Since $J_1 = J_1'$, we have that $F(4, 5) \in J_1'$, which implies that there is an element m in the active domain of I_1' such that $E(4, m) \in I_1'$ and $E(m, 5) \in I_1'$. In turn, this implies that $M(m) \in J_1' = J_1$, which means that either $m = 2$ or $m = 4$. If $m = 2$, we have that $E(4, 2) \in I_1'$; however, $E(2, 3) \in I_1'$, since $E(2, 3) \in I_2 \subseteq I_1'$. This implies that $F(4, 3) \in J_1'$, which contradicts the assumption that $J_1' = J_1$. If $m = 4$, then we have that $E(4, 4) \in I_1'$, which implies that $F(4, 4) \in J_1'$; this contradicts the assumption that $J_1 = J_1'$. □

The next theorem asserts that the $(\sim_1, \sim_2)$-subset property is a necessary and sufficient condition for the existence of a $(\sim_1, \sim_2)$-inverse of a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ in which Σ is a finite set of s-t tgds.

[2] $\pi_{12}(P^{I_2})$ is the result of projecting P^{I_2} onto its first two columns, and similarly for $\pi_{23}(P^{I_2})$.

THEOREM 3.19. *Assume that* $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ *is a schema mapping in which* Σ *is a finite set of s-t tgds and let* $\sim_1, \sim_2$ *be two equivalence relations on ground instances such that* $\sim_1 \subseteq \sim_{\mathcal{M}}$ *and* $\sim_2 \subseteq \sim_{\mathcal{M}}$. *Then the following statements are equivalent:*

1. $\mathcal{M}$ *has a* $(\sim_1, \sim_2)$*-inverse.*
2. $\mathcal{M}$ *has the* $(\sim_1, \sim_2)$*-subset property.*

PROOF. We first show that (1) $\Rightarrow$ (2). Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be a $(\sim_1, \sim_2)$-inverse of Σ. Assume that (I_1, I_2) is a pair of ground instances such that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. Since $(I_2, I_2) \in \text{Inst(Id)}$, it follows that $(I_2, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$. Therefore, there is a pair (I_3, I_4) of ground instances and a target instance J such that $(I_2, I_2) \sim_{(1,2)} (I_3, I_4)$, and such that $(I_3, J) \models \Sigma$ and $(J, I_4) \models \Sigma'$. Since $I_2 \sim_1 I_3$ and $\sim_1 \subseteq \sim_{\mathcal{M}}$, we obtain that $I_2 \sim_{\mathcal{M}} I_3$. Therefore, since $(I_3, J) \models \Sigma$, it follows that $(I_2, J) \models \Sigma$. Since $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$, we have that $(I_1, J) \models \Sigma$. Consequently, $(I_1, I_4) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, which implies that $(I_1, I_4) \in \text{Inst(Id)}[\sim_1, \sim_2]$. In turn, this implies that there is a pair (I_1', I_4') of ground instances such that $(I_1, I_4) \sim_{(1,2)} (I_1', I_4')$ and $I_1' \subseteq I_4'$. Since $I_2 \sim_2 I_4$ and $I_4 \sim_2 I_4'$, we have that $I_2 \sim_2 I_4'$. Therefore, $(I_1, I_2) \sim_{(1,2)} (I_1', I_4')$ where $I_1' \subseteq I_4'$. This establishes the implication (1) $\Rightarrow$ (2).

Towards the reverse implication, assume that $\mathcal{M}$ has the $(\sim_1, \sim_2)$-subset property. Put

$$D = \{(J, I) : J \text{ is a universal solution for } I \text{ with respect to } \mathcal{M}\}.$$

Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be the schema mapping such that $(J, I) \models \Sigma'$ if and only if $(J, I) \in D$. We claim that $\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of Σ. By Corollary 3.10, it suffices to prove that $\text{Inst(Id)} \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$ and $\text{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \text{Inst(Id)}[\sim_1, \sim_2]$.

Assume that $(I_1, I_2) \in \text{Inst(Id)}$, that is, $I_1 \subseteq I_2$. We must show that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$. Let $J = \mathit{chase}_{\Sigma}(I_2)$. So J is a universal solution for I_2 with respect to $\mathcal{M}$. We thus have $(I_2, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. Since $I_1 \subseteq I_2$ and $(I_2, J) \models \Sigma$, and since Σ is a set of s-t tgds, it follows that $(I_1, J) \models \Sigma$. Since $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$, we have that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$, as desired. Observe that we did not use the $(\sim_1, \sim_2)$-subset property in proving this inclusion.

For the other inclusion, assume that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$; we must show that $(I_1, I_2) \in \text{Inst(Id)}[\sim_1, \sim_2]$. Since $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, it follows that there is a target instance J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. Since $(J, I_2) \models \Sigma'$, we know that J is a universal solution for I_2 with respect to $\mathcal{M}$. We claim that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. To see this, assume that J' is a target instance such that $J' \in \text{Sol}(\mathcal{M}, I_2)$; we must show that $J' \in \text{Sol}(\mathcal{M}, I_1)$. Since J is a universal solution for I_2, it follows that there is a homomorphism $h : J \to J'$. Since Σ is a set of s-t tgds and $(I_1, J) \models \Sigma$, we have that $(I_1, J') \models \Sigma$ (to find witnesses in J' for the existential quantifiers in the tgds in Σ, we apply h to the witnesses in J). So $J' \in \text{Sol}(\mathcal{M}, I_1)$, as desired. This completes the proof of the claim that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$.

Since $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. and since $\mathcal{M}$ has the $(\sim_1, \sim_2)$-subset property, it follows that there are I_1', I_2' such that $(I_1, I_2) \sim_{(1,2)} (I_1', I_2')$ and $I_1' \subseteq I_2'$. So $(I_1, I_2) \in \text{Inst}(\text{Id})[\sim_1, \sim_2]$, as desired. □

COROLLARY 3.20. *Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping in which Σ is a finite set of s-t tgds. Then the following statements are equivalent:*

1. *$\mathcal{M}$ has a quasi-inverse.*
2. *$\mathcal{M}$ has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property. That is, if (I_1, I_2) is a pair of ground instances such that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(M, I_1)$, then there is a pair (I_1', I_2') of ground instances such that $I_1 \sim_\mathcal{M} I_1'$, $I_2 \sim_\mathcal{M} I_2'$, and $I_1' \subseteq I_2'$.*

The $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property can be used to show that several natural schema mappings that are not invertible have quasi-inverses. In particular, this holds true for the *Projection*, *Union*, and *Decomposition* schema mappings discussed in the Introduction. We showed explicitly in Example 3.17 that the *Decomposition* schema mapping has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property.

The *Projection*, *Union*, and *Decomposition* schema mappings are LAV *(local-as-view)* schema mappings, that is, the premise of each dependency is a single atom. The next result shows that *every* LAV schema mapping has a quasi-inverse. The proof generalizes the argument in Example 3.17.

PROPOSITION 3.21. *If $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a LAV schema mapping, then $\mathcal{M}$ has the $(\sim_\mathcal{M}, \sim_\mathcal{M})$-subset property. Consequently, every LAV schema mapping has a quasi-inverse.*

PROOF. Assume that I_1 and I_2 are two ground instances such that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. Let J_1 be a universal solution for I_1, and let J_2 be a universal solution for I_2. There is a homomorphism from J_1 to J_2, since J_2 is a solution for I_1. Let $I_2' = I_1 \cup I_2$. Clearly, $I_1 \subseteq I_2'$. It remains to show that $I_2 \sim_\mathcal{M} I_2'$, that is, $\text{Sol}(\mathcal{M}, I_2) = \text{Sol}(\mathcal{M}, I_2')$. Since $I_2 \subseteq I_2'$, we have that $\text{Sol}(\mathcal{M}, I_2') \subseteq \text{Sol}(\mathcal{M}, I_2)$. For the other inclusion, let J be a solution for I_2. We have to show that J is also a solution for I_2'. Consider an arbitrary s-t tgd in Σ; since Σ is LAV, the tgd must be of the form $\forall \mathbf{x}(R(\mathbf{x}) \rightarrow \exists \mathbf{y}\varphi(\mathbf{x}, \mathbf{y}))$, where $R(\mathbf{x})$ is an atom.

Let $\mathbf{a}$ be a tuple of constants such that $I_2' \models R(\mathbf{a})$. We must show that $J \models \exists \mathbf{y}\varphi(\mathbf{a}, \mathbf{y})$. Since $I_2' \models R(\mathbf{a})$, there are two possibilities: either $R(\mathbf{a})$ is a fact of I_1, or $R(\mathbf{a})$ is a fact of I_2. If $R(\mathbf{a})$ is a fact of I_1, then $J_1 \models \exists \mathbf{y}\varphi(\mathbf{a}, \mathbf{y})$, since J_1 is a solution for I_1. Since there is a homomorphism from J_1 to J_2 and a homomorphism from J_2 to J (because J_2 is a universal solution for I_2) we have that there is a homomorphism from J_1 to J. Hence, $J \models \exists \mathbf{y}\varphi(\mathbf{a}, \mathbf{y})$, as desired. If $R(\mathbf{a})$ is a fact of I_2, then again $J \models \exists \mathbf{y}\varphi(\mathbf{a}, \mathbf{y})$, since J is a solution for I_2. So in both cases, $J \models \exists \mathbf{y}\varphi(\mathbf{a}, \mathbf{y})$, as desired. □

An inspection of the proof of Proposition 3.21 reveals that actually a stronger fact holds: if $\mathcal{M}$ is a LAV schema mapping, then $\mathcal{M}$ has the $(=, \sim_\mathcal{M})$-subset property; consequently, every LAV mapping has a $(=, \sim_\mathcal{M})$-inverse.

Our next result asserts that, in contrast to LAV schema mappings, there are schema mappings specified by *full* s-t tgds that have no quasi-inverses.

PROPOSITION 3.22. *There is a schema mapping $\mathcal{M}$ that is specified by a single full s-t tgd and has no quasi-inverse.*

PROOF. Let $\mathcal{M}$ be the schema mapping in Example 3.18. It is specified by a single full s-t tgd, and we showed in Example 3.18 that $\mathcal{M}$ does not have the $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-subset property. So by Corollary 3.20, it follows that $\mathcal{M}$ has no quasi-inverse. □

The full s-t tgd $E(x, z) \wedge E(z, y) \rightarrow F(x, y) \wedge M(z)$ in Example 3.18 is logically equivalent to the set consisting of the following two full s-t tgds:

$$E(x, z) \wedge E(z, y) \rightarrow F(x, y) \quad \text{and} \quad E(x, z) \wedge E(z, y) \rightarrow M(z).$$

It is not hard to show that each of these two full s-t tgds specifies schema mappings that have a quasi-inverse. Thus, the schema mapping specified by the union of the set of constraints of two quasi-invertible schema mappings need not be quasi-invertible.

Note that the $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-subset property is used "positively" in the proof of Proposition 3.21 and "negatively" in the proof of Proposition 3.22. Indeed, the $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-subset property is used as a sufficient condition for the existence of quasi-inverses in Proposition 3.21 and as a necessary condition in Proposition 3.22.

3.5 The $(=, =)$-Subset Property and Inverses of Schema Mappings

Theorem 3.19 yields a necessary and sufficient condition for the existence of an inverse:

COROLLARY 3.23. *Assume that $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping where Σ is a finite set of s-t tgds. Then the following statements are equivalent:*

1. *$\mathcal{M}$ has an inverse.*
2. *$\mathcal{M}$ has the $(=, =)$-subset property. That is, if I_1 and I_2 are two ground instances such that* $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$*, then $I_1 \subseteq I_2$.*

As mentioned in the Introduction, the *unique-solutions property* was identified in Fagin [2007] as a necessary condition for a schema mapping $\mathcal{M}$ to have an inverse. By definition, this property says that if I_1 and I_2 are ground instances such that $I_1 \neq I_2$, then we have that $\text{Sol}(\mathcal{M}, I_1) \neq \text{Sol}(\mathcal{M}, I_2)$. It is easy to see that the $(=, =)$-subset property implies the unique-solutions property. Indeed, if $\text{Sol}(\mathcal{M}, I_1) = \text{Sol}(\mathcal{M}, I_2)$, then by applying the $(=, =)$-subset property twice, we have that $I_1 \subseteq I_2$ and $I_2 \subseteq I_1$, and so $I_1 = I_2$.

The unique solutions property implies that if a schema mapping $\mathcal{M}$ is invertible, then the equivalence relation $\sim_{\mathcal{M}}$ coincides with the equality relation $=$ on ground instances. In turn, this observation shows that the following proposition is true.

PROPOSITION 3.24. *Every quasi-inverse of an invertible schema mapping $\mathcal{M}$ is an inverse of $\mathcal{M}$.*

Thus, for invertible schema mappings, there is no distinction between inverses and quasi-inverses (and $(\sim_1, \sim_2)$-inverses as well).

The failure of the unique solutions property is typically used as a sufficient condition for proving non-invertibility. In particular, the *Projection*, *Union*, and *Decomposition* schema mappings are not invertible because none of them possesses the unique solutions property. In Fagin [2007], it was shown that for LAV schema mappings, the unique solutions property is a necessary and sufficient condition for invertibility. The question of whether the unique solutions property is not just necessary but also sufficient for invertibility of general (not necessarily LAV) schema mappings specified by a finite set of s-t tgds was left open. The next theorem resolves this problem, by showing that the unique-solutins problem is not sufficient for invertibility.

THEOREM 3.25. *There is a schema mapping specified by a finite set of s-t tgds that has the unique solutions property but does not have an inverse.*

PROOF. Let $\mathbf{S}$ consist of unary relation symbols A and B. Let $\mathbf{T}$ consist of a unary relation symbol C and a binary relation symbol R. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ where Σ consists of the tgds:

$$A(x) \rightarrow R(x, x), \quad B(x) \rightarrow \exists y R(x, y), \quad A(x) \wedge B(x) \rightarrow C(x).$$

We shall show that $\mathcal{M}$ has the unique solutions property but does not have an inverse. We begin by showing that $\mathcal{M}$ does not satisfy the $(=, =)$-subset property, and hence does not have an inverse. Let I_1 have the single fact $B(0)$, and let I_2 have the single fact $A(0)$. Then $\text{Sol}(\mathcal{M}, I_1)$ consists of those target instances that contain some fact of the form $R(0, y)$, and $\text{Sol}(\mathcal{M}, I_2)$ consists of those target instances that contain the fact $R(0, 0)$. Therefore, $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. Since $I_1 \not\subseteq I_2$, it follows that $\mathcal{M}$ does not satisfy the $(=, =)$-subset property, and hence does not have an inverse.

We now sketch a proof that $\mathcal{M}$ has the unique solutions property. Let I_1 and I_2 be ground instances that have the same set of solutions; we must show that $I_1 = I_2$.

First, it is straightforward to show that I_1 and I_2 have the same active domain. Now every member a of this active domain has one of three possible types in I_1: a is in A^{I_1} and in B^{I_1}; a is in A^{I_1} but not in B^{I_1}; or a is in B^{I_1} but not in A^{I_1}. Similarly, we define the corresponding types of a in I_2. It is straightforward to show that whatever type a has in I_1, it has the corresponding type in I_2. This implies that $I_1 = I_2$, as desired. □

4. THE LANGUAGE OF QUASI-INVERSES

In this section, we identify the language needed to express quasi-inverses of schema mappings specified by s-t tgds.

4.1 The General Case

One of our main results is the following characterization of the language for quasi-inverses of schema mappings specified by tgds.

THEOREM 4.1. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. If $\mathcal{M}$ has a quasi-inverse then the following hold.*

1. *$\mathcal{M}$ has a quasi-inverse $\mathcal{M}'$ specified by a finite set of disjunctive tgds with constants and inequalities.*
2. *There is an exponential-time algorithm for producing $\mathcal{M}'$.*
3. *Statement (1) is not necessarily true if we disallow either constants or inequalities in the premise, or disallow disjunctions or existential quantifiers in the conclusion.*

In fact, the quasi-inverse our algorithm produces has inequalities only among constants.

We illustrate the intuition behind the construction of $\mathcal{M}'$, with two examples. We begin with the *Union* example, where Σ consists of the s-t tgds $P(x) \to S(x)$ and $Q(x) \to S(x)$. There are two possible "generators" of $S(x)$, namely $P(x)$ and $Q(x)$. These possibilities are reflected by the disjunctive tgd $S(x) \to P(x) \vee Q(x)$ (we shall put a variation of this disjunctive tgd into Σ'). As another example, let Σ consist of the s-t tgds $S(x, y) \to P(x, y)$ and $T(x, y) \to P(x, x)$. There is only one possible generator of $P(x, y)$ if x and y are different, namely $S(x, y)$, and this is reflected by the tgd with inequalities $P(x, y) \wedge (x \neq y) \to S(x, y)$. But there are two possible generators of $P(x, x)$, namely $S(x, x)$ and $T(x, y)$, and this is reflected by the disjunctive tgd $P(x, x) \to S(x, x) \vee \exists y T(x, y)$. The algorithm for producing quasi-inverses systematically considers all such generators.

We now introduce the machinery behind the algorithm to produce $\mathcal{M}'$, including a formal definition of "generator." If α is a conjunction of atoms (or an instantiation of atoms), define I_α to be an instance whose facts are the conjuncts of α. Note that I_α is not an instance in the usual sense, since the active domain consists of variables. Thus, I_α is a type of canonical instance. Let $\mathbf{x}$ be a vector of distinct variables. A *complete description* $\delta(\mathbf{x})$ is a conjunction of equalities and inequalities among members of $\mathbf{x}$ such that for each x_i, x_j in $\mathbf{x}$, exactly one of the formulas $x_i = x_j$ or $x_i \neq x_j$ is a conjunct of $\delta(\mathbf{x})$, and such that $\delta(\mathbf{x})$ is satisfiable (in other words, consistent).

Let Σ be a finite set of s-t tgds. We now define a set Σ^* that includes Σ and that is logically equivalent to Σ. For each member σ of Σ, and for each complete description δ of the variables that each appear in both the premise and the conclusion of σ, select a unique representative of each equivalence class determined by δ, and let σ_δ be obtained from σ by replacing every variable in σ by the representative of its equivalence class. Let Σ^* consist of Σ and all such formulas σ_δ (for all choices of σ in Σ and all complete descriptions δ of the variables that each appear in both the premise and the conclusion of σ). For example, if σ is $R(x_1, x_2, x_3, x_4) \to \exists y(Q(x_1, y) \wedge S(y, x_2, x_3))$, and if δ is $(x_1 = x_3) \wedge (x_1 \neq x_2) \wedge (x_2 \neq x_3)$, then $\{x_1, x_3\}$ forms one equivalence class and $\{x_2\}$ is the other equivalence class, and σ_δ is $R(x_1, x_2, x_1, x_4) \to \exists y(Q(x_1, y) \wedge S(y, x_2, x_1))$.

Definition 4.2. Let $\beta(\mathbf{x}, \mathbf{z})$ be a conjunction of source atoms, and let $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ be a conjunction of target atoms, where the members of $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are all distinct, and the members of $\mathbf{x}$ are exactly the variables that appear in both $\beta(\mathbf{x}, \mathbf{z})$ and $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. Let Σ be a finite set of s-t tgds. We say that $\beta(\mathbf{x}, \mathbf{z})$ is a *generator of* $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ *(with respect to* Σ*)* if the s-t tgd $\beta(\mathbf{x}, \mathbf{z}) \rightarrow \exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ is a logical consequence of Σ.

When Σ is understood, we shall often drop the words "with respect to Σ". It follows easily from the standard theory of the chase that $\beta(\mathbf{x}, \mathbf{z})$ is a generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ with respect to Σ if and only if the chase of $I_{\beta(\mathbf{x},\mathbf{z})}$ with Σ gives at least $I_{\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y}')}$ for a substitution where some $\mathbf{y}'$ substitutes for $\mathbf{y}$.

Definition 4.3. The source formula $\beta(\mathbf{x}, \mathbf{z})$ is a *minimal generator of* $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ if $\beta(\mathbf{x}, \mathbf{z})$ is a generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ and there is no $\beta'(\mathbf{x}, \mathbf{z})$ that is a conjunction of a strict subset of the conjuncts of $\beta(\mathbf{x}, \mathbf{z})$ such that $\beta'(\mathbf{x}, \mathbf{z})$ is a generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$.

We shall make use of the following simple lemma.

LEMMA 4.4. *Let Σ be a finite set of s-t tgds, each with at most s_1 conjuncts in its premise. Let $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ be a conjunction of s_2 target atoms. Then every minimal generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ with respect to Σ has at most $s_1 s_2$ conjuncts.*

PROOF. At most s_2 chase steps are required to generate $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ from $\beta(\mathbf{x}, \mathbf{z})$ (one chase step for every conjunct of $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$). Each chase step involves at most s_1 conjuncts of $\beta(\mathbf{x}, \mathbf{z})$. So if $\beta(\mathbf{x}, \mathbf{z})$ is minimal, it has at most $s_1 s_2$ conjuncts. □

From Lemma 4.4, we see that there is a simple exhaustive-search algorithm for finding minimal generators:

Algorithm MinGen($\mathcal{M}, \exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$)

Input: A schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds, and a formula $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, where $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ is a conjunction of target atoms, and where the variables in $\mathbf{x}, \mathbf{y}$ are all distinct, and all appear in $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$.
Output: A finite set of the minimal generators of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ with respect to Σ.

1. (*Initialization.*) Initialize the set G of minimal generators of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ to $\emptyset$.
2. (*Exhaustive search.*) Let s_1 and s_2 be as in Lemma 4.4. Systematically check every conjunction $\beta(\mathbf{x}, \mathbf{z})$ (up to renaming of variables in $\mathbf{z}$) of at most $s_1 s_2$ atoms where the variables in $\mathbf{z}$ are distinct and distinct from members of $\mathbf{x}, \mathbf{y}$, to see if the chase of $I_{\beta(\mathbf{x},\mathbf{z})}$ with Σ gives at least $I_{\psi_{\mathbf{T}}(\mathbf{x},\mathbf{y}')}$ for a substitution where some $\mathbf{y}'$ substitutes for $\mathbf{y}$. If so, add $\beta(\mathbf{x}, \mathbf{z})$ to G.
3. (*Minimize.*) For each member $\beta(\mathbf{x}, \mathbf{z})$ of G, check to see if there is some other $\beta'(\mathbf{x}, \mathbf{z}')$ in G whose conjuncts are a subset of the conjuncts of $\beta(\mathbf{x}, \mathbf{z})$ (up to renaming of variables in $\mathbf{z}, \mathbf{z}'$). If so, remove $\beta(\mathbf{x}, \mathbf{z})$ from G. Continue the process until there is no more change in G. **Return** G.

The next algorithm produces a finite set of disjunctive tgds with constants and inequalities that we shall prove specifies a quasi-inverse if one exists.

Algorithm QuasiInverse($\mathcal{M}$)

Input: A schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds.
Output: A schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ' is a finite set of disjunctive tgds with constants and inequalities, that is a quasi-inverse of $\mathcal{M}$ if $\mathcal{M}$ has a quasi-inverse.

1. (*Create* Σ^*.) Create Σ^* from Σ as defined just before Definition 4.2.
2. (*Create the formulas* σ'.) For each member σ of Σ^*, create an implication σ' as follows. Assume that σ is $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u}) \rightarrow \exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, where the variables in $\mathbf{x}$ are distinct, and consist exactly of the variables that appear in both $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u})$ and $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. The premise of σ' is the conjunction of $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, along with each of the formulas $Constant(x)$ for members x of $\mathbf{x}$, along with the formulas $x_i \neq x_j$ for each pair x_i, x_j of distinct variables in $\mathbf{x}$. For each formula $\beta(\mathbf{x}, \mathbf{z})$ in the output of MinGen($\mathcal{M}, \exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$), let $\exists \mathbf{z} \beta(\mathbf{x}, \mathbf{z})$ be a disjunct in the conclusion of σ'.
3. (*Construct* Σ'.) Let Σ' consist of each of these formulas σ'. **Return** $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$.

Note that the disjunction in the conclusion that is created in Step (2) of the algorithm is nonempty, since $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u})$, the premise of σ, is a generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, and so some subset of the conjunctions of $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u})$ forms a minimal generator.

It is interesting to note that the schema mapping $\mathcal{M}'$ that is the output of the algorithm QuasiInverse($\mathcal{M}$) is the same no matter what the choice is for $(\sim_1, \sim_2)$ (as long as $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$). However, it may happen that $\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$ and yet $\mathcal{M}'$ fails to be a $(\sim'_1, \sim'_2)$-inverse of $\mathcal{M}$. This is because $\mathcal{M}'$ is guaranteed to be a $(\sim'_1, \sim'_2)$-inverse of $\mathcal{M}$ only if $\mathcal{M}$ has a $(\sim'_1, \sim'_2)$-inverse.

Example 4.5. Let Σ consist of the tgds

$$P(x_1, x_2, x_3) \rightarrow \exists y (S(x_1, x_2, y) \wedge Q(y, y)) \qquad T(x_3, x_4) \rightarrow S(x_4, x_4, x_3)$$
$$U(x_1) \rightarrow \exists y (S(x_1, x_1, y) \wedge Q(y, y) \wedge Q(x_1, y)) \qquad R(x_1, x_2, x_4) \rightarrow Q(x_1, x_2).$$

Let σ_1 be the first tgd in Σ, and let σ_2 be the result of replacing each occurrence of x_2 in σ_1 by x_1 (so σ_2 is $P(x_1, x_1, x_3) \rightarrow \exists y (S(x_1, x_1, y) \wedge Q(y, y))$). Then σ_1 and σ_2 are both in Σ^*. To show Step (2) of the algorithm QuasiInverse, in this example we shall produce σ'_1 from σ_1, and we shall produce σ'_2 from σ_2. Thus, the algorithm puts σ'_1 and σ'_2 into Σ'.

The only generator of the conclusion $\exists y (S(x_1, x_2, y) \wedge Q(y, y))$ of σ_1 is $P(x_1, x_2, x_3)$, so σ'_1 is $S(x_1, x_2, y) \wedge Q(y, y) \wedge Constant(x_1) \wedge Constant(x_2) \wedge (x_1 \neq x_2) \rightarrow \exists x_3 P(x_1, x_2, x_3)$. There are four minimal generators of $\exists y (S(x_1, x_1, y) \wedge Q(y, y))$, the conclusion of σ_2. The first is $P(x_1, x_1, x_3)$, the premise of σ_2. The second is $U(x_1)$, since its chase yields $S(x_1, x_1, y)$, $Q(y, y)$, $Q(x_1, y)$, which includes the conjuncts in the conclusion of σ_2. The third is $T(x_1, x_1) \wedge R(x_1, x_1, x_4)$, since chasing the two facts in this conjunct yields $S(x_1, x_1, x_1)$, $Q(x_1, x_1)$, where the role of y in the conclusion of σ_2 is played by the variable x_1. The fourth is $T(x_3, x_1) \wedge R(x_3, x_3, x_4)$, since the chase of the two facts in this conjunct yields

$S(x_1, x_1, x_3)$, $Q(x_3, x_3)$, where the role of y in the conclusion of σ_2 is played by the variable x_3. Then σ_2' is

$$S(x_1, x_1, y) \wedge Q(y, y) \wedge Constant(x_1) \rightarrow \big(\exists x_3 P(x_1, x_1, x_3) \vee U(x_1) \vee \exists x_4(T(x_1, x_1) \wedge R(x_1, x_1, x_4)) \vee \exists x_3 \exists x_4(T(x_3, x_1) \wedge R(x_3, x_3, x_4))\big)$$

Note that the fourth disjunct in the conclusion of σ_2' is implied by the third disjunct (by letting the role of x_3 be played by x_1). So the third disjunct could be removed, since we need only keep the more general disjunct.

We shall soon prove that that the QuasiInverse algorithm defines a quasi-inverse of $\mathcal{M}$ if one exists. We first need a lemma, which follows from the standard theory of the chase.

LEMMA 4.6. *Assume that Σ is a finite set of s-t tgds. Let $\beta(\mathbf{x}, \mathbf{z})$ be a conjunction of source atoms, and $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ a conjunction of target atoms. Assume that $\mathbf{x}$ and $\mathbf{z}$ are vectors of distinct variables, with no variables in common, and $\bar{\mathbf{x}}, \bar{\mathbf{z}}$ are an assignment of distinct constants to $\mathbf{x}, \mathbf{z}$. Assume that $\bar{\mathbf{y}}$ is an assignment where the members of $\bar{\mathbf{y}}$ may be constants in $\bar{\mathbf{x}}$, constants in $\bar{\mathbf{z}}$, or nulls. Assume that a chase of $I_{\beta(\bar{\mathbf{x}},\bar{\mathbf{z}})}$ with Σ yields at least $I_{\psi_{\mathbf{T}}(\bar{\mathbf{x}},\bar{\mathbf{y}})}$. Then $\beta(\mathbf{x}, \mathbf{z})$ is a generator of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$.*

We now show that the QuasiInverse algorithm defines a quasi-inverse of $\mathcal{M}$ if one exists. This implies Part (1) of Theorem 4.1.

THEOREM 4.7. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping, where Σ is a finite set of tgds. Let $\sim_1$, $\sim_2$ be two equivalence relations on ground instances such that $\sim_1 \subseteq \sim_{\mathcal{M}}$ and $\sim_2 \subseteq \sim_{\mathcal{M}}$. The following are equivalent.*

1. *$\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse.*
2. *The output $\mathcal{M}'$ of QuasiInverse($\mathcal{M}$) is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$. The schema mapping $\mathcal{M}'$ is specified by a finite set of disjunctive tgds with constants and inequalities among constants.*

PROOF. It is obvious that (2) implies (1). We now show that (1) implies (2). Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where $\mathcal{M}'$ is the output of QuasiInverse($\mathcal{M}$). By construction, the schema mapping $\mathcal{M}'$ is specified by a finite set of disjunctive tgds with constants and inequalities among constants. Assume that $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse. We now show that $\mathcal{M}'$ is a $(\sim_1, \sim_2)$-inverse of $\mathcal{M}$.

By Corollary 3.10, we must show that $\text{Inst(Id)} \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$ and $\text{Inst}(\mathcal{M} \circ \mathcal{M}') \subseteq \text{Inst(Id)}[\sim_1, \sim_2]$. We first show that $\text{Inst(Id)} \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$. By Part (1) of Proposition 3.9, it is sufficient to show that $\text{Inst(Id)} \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. Towards this, let I_1, I_2 be ground instances such that $I_1 \subseteq I_2$. Let J_1 be the result of doing a chase of I_1 with Σ. Clearly, $(I_1, J_1) \models \Sigma$. We claim that $(J_1, I_2) \models \Sigma'$, which implies that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, as desired. Let σ and σ' be as in Step (2) of QuasiInverse($\mathcal{M}$). To show that $(J_1, I_2) \models \Sigma'$, we need only show that $(J_1, I_2) \models \sigma'$. Assume that $\bar{\mathbf{x}}, \bar{\mathbf{y}}$ are an assignment of values to $\mathbf{x}, \mathbf{y}$ such that the premise of σ' holds in J_1. We need only show that there is a minimal generator $\beta(\mathbf{x}, \mathbf{z})$ of $\exists \mathbf{y} \psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$ where also $\exists \mathbf{z} \beta(\bar{\mathbf{x}}, \mathbf{z})$

holds in I_1, and hence in I_2. Since the entries of $\bar{\mathbf{x}}$ are constants (because of the formulas *Constant*(x) in the premise of σ'), and since J_1 is the result of doing a chase of I_1 with Σ, there is a vector $\mathbf{z}$ of distinct variables, all distinct from the variables in $\mathbf{x}$, and an assignment $\bar{\mathbf{z}}$ of distinct constants in I_1, all distinct from the constants in $\bar{\mathbf{x}}$, and a conjunction $\beta''(\bar{\mathbf{x}}, \bar{\mathbf{z}})$ of facts of I_1, such that a chase of $I_{\beta''(\bar{\mathbf{x}},\bar{\mathbf{z}})}$ with Σ yields at least $I_{\psi_{\mathbf{T}}(\bar{\mathbf{x}},\bar{\mathbf{y}})}$. The members of $\bar{\mathbf{x}}$ are distinct, because of the inequalities in the premise of σ'. So by Lemma 4.6, we see that $\beta''(\mathbf{x}, \mathbf{z})$ is a generator of $\exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. Note that $\exists\mathbf{z}\beta''(\bar{\mathbf{x}}, \mathbf{z})$ holds in I_1. Let $\beta(\mathbf{x}, \mathbf{z})$ be obtained from $\beta''(\mathbf{x}, \mathbf{z})$ by removing as many conjuncts as possible so that $\beta(\mathbf{x}, \mathbf{z})$ is still a generator of $\exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. So $\beta(\mathbf{x}, \mathbf{z})$ is a minimal generator of $\exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. Since $\exists\mathbf{z}\beta''(\bar{\mathbf{x}}, \mathbf{z})$ holds in I_1, so does $\exists\mathbf{z}\beta(\bar{\mathbf{x}}, \mathbf{z})$. So $\beta(\mathbf{x}, \mathbf{z})$ is a minimal generator of $\exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, and $\exists\mathbf{z}\beta(\bar{\mathbf{x}}, \mathbf{z})$ holds in I_1. This was to be shown.

For the other direction, assume that I_1, I_2 are ground instances such that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$; we must show that $(I_1, I_2) \in \text{Inst}(\text{Id})[\sim_1, \sim_2]$. Since $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, there is a target instance J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. We now show that $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$. Assume $J^* \in \text{Sol}(\mathcal{M}, I_2)$; we must show that $J^* \in \text{Sol}(\mathcal{M}, I_1)$. Let σ be a member of Σ; we must show that $(I_1, J^*) \models \sigma$. Let σ be $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u}) \rightarrow \exists\mathbf{y}\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$, where the variables in $\mathbf{x}$ are distinct, and are exactly the variables that appear in both $\varphi_{\mathbf{S}}(\mathbf{x}, \mathbf{u})$ and $\psi_{\mathbf{T}}(\mathbf{x}, \mathbf{y})$. Assume that $\varphi_{\mathbf{S}}(\bar{\mathbf{x}}, \bar{\mathbf{u}})$ holds in I_1; we must show that $\exists\mathbf{y}\psi_{\mathbf{T}}(\bar{\mathbf{x}}, \mathbf{y})$ holds in J^*. Let δ be a complete description such that $\delta(\bar{\mathbf{x}})$ holds. Let σ_1 be σ_δ, as defined earlier in our construction of Σ^*. Assume that σ_1 is $\varphi_1(\mathbf{x}', \mathbf{u}) \rightarrow \exists\mathbf{y}\psi_1(\mathbf{x}', \mathbf{y})$, where $\mathbf{x}'$ is obtained from $\mathbf{x}$ (as in the construction of Σ^*) by replacing each member of $\mathbf{x}$ by a representative of its equivalence class with respect to δ. Let $\bar{\mathbf{x}}'$ be the assignment of values to $\mathbf{x}'$ obtained from $\bar{\mathbf{x}}$. In particular, the members of $\bar{\mathbf{x}}'$ are distinct constants. Also, $\varphi_1(\bar{\mathbf{x}}', \bar{\mathbf{u}})$ is the same as $\varphi_{\mathbf{S}}(\bar{\mathbf{x}}, \bar{\mathbf{u}})$, and $\psi_1(\bar{\mathbf{x}}', \mathbf{y})$ is the same as $\psi_{\mathbf{T}}(\bar{\mathbf{x}}, \mathbf{y})$. So we know that $\varphi_1(\bar{\mathbf{x}}', \bar{\mathbf{u}})$ holds in I_1, and we need to show that $\exists\mathbf{y}\psi_1(\bar{\mathbf{x}}', \mathbf{y})$ holds in J^*. Since $\varphi_1(\bar{\mathbf{x}}', \bar{\mathbf{u}})$ holds in I_1, and since $(I_1, J) \models \sigma_1$, it follows that $\exists\mathbf{y}\psi_1(\bar{\mathbf{x}}', \mathbf{y})$ holds in J. So there is $\bar{\mathbf{y}}$ such that $\psi_1(\bar{\mathbf{x}}', \bar{\mathbf{y}})$ holds in J. Let σ_1' be obtained from σ_1 as in Step (2) of QuasiInverse$(\mathcal{M})$. Let $\psi_1'(\mathbf{x}', \mathbf{y})$ be the premise of σ_1'. Then $\psi_1'(\bar{\mathbf{x}}', \bar{\mathbf{y}})$ holds in J, since $\psi_1(\bar{\mathbf{x}}', \bar{\mathbf{y}})$ holds in J, and the members of $\bar{\mathbf{x}}'$ are distinct constants. Since $(J, I_2) \models \Sigma'$, we know that $(J, I_2) \models \sigma_1'$. Since $\psi_1'(\bar{\mathbf{x}}', \bar{\mathbf{y}})$ holds in J, and since $(J, I_2) \models \sigma_1'$, there is a minimal generator $\beta(\mathbf{x}', \mathbf{z})$ of $\exists\mathbf{y}\psi_1(\mathbf{x}', \mathbf{y})$ such that $\exists\mathbf{z}\beta(\bar{\mathbf{x}}', \mathbf{z})$ holds in I_2. Now (I_2, J^*) satisfies Σ, and $\exists\mathbf{z}\beta(\bar{\mathbf{x}}', \mathbf{z})$ holds in I_2. It follows from the definition of a generator that $\exists\mathbf{y}\psi_1(\bar{\mathbf{x}}', \mathbf{y})$ holds in J^*, which was to be shown. So $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$, as desired.

Since by assumption, $\mathcal{M}$ has a $(\sim_1, \sim_2)$-inverse, it follows that $\mathcal{M}$ has the $(\sim_1, \sim_2)$-subset property. Therefore, since $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$, it follows that $(I_1, I_2) \in \text{Inst}(\text{Id})[\sim_1, \sim_2]$, which was to be shown. $\square$

4.2 The Full Case

In this section, we consider the case of schema mappings specified by a finite set of *full* s-t tgds. The next proposition implies that the constant predicate then plays no role.

PROPOSITION 4.8. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ is a finite set of full s-t tgds, and Σ' is a finite set of disjunctive tgds with constants and inequalities. Let I_1 and I_2 be ground instances.*

1. $(I_1, I_2) \models \Sigma \circ \Sigma'$ *if and only if* $(chase_\Sigma(I_1), I_2) \models \Sigma'$.
2. *Let Σ'' be obtained from Σ' by removing every conjunct of the form Constant(v) from members of Σ'. Then*
   a. $(I_1, I_2) \models \Sigma \circ \Sigma'$ *if and only if* $(I_1, I_2) \models \Sigma \circ \Sigma''$.
   b. $\mathcal{M}'$ *is a* $(\sim_1, \sim_2)$*-inverse of* $\mathcal{M}$ *if and only if* $\mathcal{M}''$ *is a* $(\sim_1, \sim_2)$*-inverse of* $\mathcal{M}$.

PROOF. Let $J^* = chase_\Sigma(I_1)$. Then J^* is ground (since Σ is full), and $(I_1, J^*) \models \Sigma$. We now prove (1). Assume first that $(I_1, I_2) \models \Sigma \circ \Sigma'$. Then there is J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. Now $J^* \subseteq J$, since $J^* = chase_\Sigma(I_1)$. It is easy to see that since (a) $(J, I_2) \models \Sigma'$, (b) $J^* \subseteq J$, and (c) Σ' is a finite set of disjunctive tgds with constants and inequalities, it follows that $(J^*, I_2) \models \Sigma'$. This was to be shown. Conversely, assume that $(J^*, I_2) \models \Sigma'$. Since $(I_1, J^*) \models \Sigma$ and $(J^*, I_2) \models \Sigma'$, it follows that $(I_1, I_2) \models \Sigma \circ \Sigma'$. This was to be shown.

We now prove (2a). Assume first that $(I_1, I_2) \models \Sigma \circ \Sigma'$. By Part (1), we see that $(J^*, I_2) \models \Sigma'$. Since J^* consists only of constants, it follows that $(J^*, I_2) \models \Sigma''$ if and only if $(J^*, I_2) \models \Sigma'$. So $(J^*, I_2) \models \Sigma''$. Since $(I_1, J^*) \models \Sigma$ and $(J^*, I_2) \models \Sigma''$, it follows that $(I_1, I_2) \models \Sigma \circ \Sigma''$, as desired. The proof of the converse is exactly the same, but where the roles of Σ' and Σ'' are reversed.

We now prove (2b). It follows from Part (2a) that $\text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}'')[\sim_1, \sim_2]$. Therefore, $\text{Inst}(\text{Id})[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim_1, \sim_2]$ if and only if $\text{Inst}(\text{Id})[\sim_1, \sim_2] = \text{Inst}(\mathcal{M} \circ \mathcal{M}'')[\sim_1, \sim_2]$. This implies (2b). □

We obtain the following modified version of Theorem 4.1 in the full case:

THEOREM 4.9. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of full s-t tgds. If $\mathcal{M}$ has a quasi-inverse then the following hold.*

1. *$\mathcal{M}$ has a quasi-inverse $\mathcal{M}'$ specified by a finite set of disjunctive tgds with inequalities.*
2. *There is an exponential-time algorithm for producing $\mathcal{M}'$.*
3. *Statement (1) is not necessarily true if we disallow either disjunctions, inequalities, or existential quantifiers (disallowing existential quantifiers means requiring full dependencies).*

Part (1) follows from Part (1) of Theorem 4.1, along with Part (2b) of Proposition 4.8. We shall prove Part (3) in Section 4.4.

In the full case, we can still, of course, make use of the algorithms MinGen to obtain the minimal generators. However, we now give an algorithm that is especially tailored to the full case. We can assume without loss of generality that the conclusion of every member of Σ is a single atom. Let Σ^* be as before. We are interested in obtaining minimal generators only of formulas that are the conclusion of members of Σ^*, and hence only of atoms. The algorithm

MinGenFull that applies in this case to produce minimal generators does not need to use exhaustive search.

Algorithm MinGenFull($\mathcal{M}$, $R(\mathbf{x})$)

Input: A schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of full tgds with singleton conclusions, and an atom $R(\mathbf{x})$.
Output: A finite set G of the minimal generators of $R(\mathbf{x})$ with respect to Σ.

1. (*Create* Σ^*.) Create Σ^* from Σ as defined just before Definition 4.2.
2. (*Initialization.*) Initialize the set G of minimal generators of $R(\mathbf{x})$ to $\emptyset$.
3. (*Look for matches.*) For each member σ of Σ^*, check to see if there is a formula σ' that can be obtained from σ by renaming variables, so that the conclusion becomes $R(\mathbf{x})$. When this succeeds, and when the premise of σ' is $\beta(\mathbf{x}, \mathbf{z})$, where $\mathbf{z}$ consists of the distinct variables in the premise of σ' that do not appear in $R(\mathbf{x})$, then add $\beta(\mathbf{x}, \mathbf{z})$ to the set G.
4. (*Minimize.*) For each member $\beta(\mathbf{x}, \mathbf{z})$ of G, check to see if there is some other $\beta'(\mathbf{x}, \mathbf{z}')$ in G whose conjuncts are a subset of the conjuncts of $\beta(\mathbf{x}, \mathbf{z})$ (up to renaming of variables in $\mathbf{z}, \mathbf{z}'$). If so, remove $\beta(\mathbf{x}, \mathbf{z})$ from G. Continue the process until there is no more change in G. **Return** G.

It is straightforward to verify that MinGenFull produces all minimal generators (up to renaming of variables). In the full case, we can modify the algorithm QuasiInverse not only by making use of MinGenFull instead of MinGen, but also (because of Part (2b) of Proposition 4.8) by no longer including the formulas $Constant(x)$.

4.3 The LAV Case

Proposition 3.21 tells us that every LAV schema mapping has a quasi-inverse. The next theorem asserts that disjunctions are not needed in the language of quasi-inverses of LAV schema mappings.

THEOREM 4.10. *Every LAV schema mapping has a quasi-inverse specified by a finite set of tgds with constants and inequalities. Thus, disjunctions are not needed.*

We prove this theorem by giving an algorithm that we prove produces a quasi-inverse.

Define a *prime atom* to be one that contains precisely the variables $x_1, x_2, \ldots, x_k$ for some k, and where the initial appearance of x_i precedes the initial appearance of x_j if $i < j$. For example, $P(x_1, x_2, x_1, x_3, x_2)$ is a prime atom, but $Q(x_2, x_1)$ and $R(x_2, x_3)$ are not. Note that for every atom, there is a unique renaming of variables to obtain a prime atom. Define a *prime instance* to be an instance whose only fact is a single prime atom. As with our definition of I_α, a prime instance is not an instance in the usual sense, but is a type of canonical instance.

The algorithm for producing a quasi-inverse is as follows:

Algorithm LAV($\mathcal{M}$)

Input: A LAV schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$.
Output: A schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ' is a finite set of tgds with constants and inequalities, and $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$.

1. (*Generate all prime source atoms in lexicographic order.*) For example, if R is a ternary source relation symbol, the atomic formulas involving R, in lexicographic order, are $R(x_1, x_1, x_1)$, $R(x_1, x_1, x_2)$, $R(x_1, x_2, x_1)$, $R(x_1, x_2, x_2)$, $R(x_1, x_2, x_3)$.
2. (*Construct a tgd for each prime atom* α.) For each prime atom α generated in Step (1), let I_α be the prime instance containing only α. Let ψ_α be the conjunction of the facts of $chase_\Sigma(I_\alpha)$. Let $\mathbf{y}$ consist of the distinct variables in α that appear in ψ_α, and let $\mathbf{z}$ consist of the remaining distinct variables in α. Let ψ'_α be the conjunction of ψ_α along with the formulas $Constant(y)$ for each variable y in $\mathbf{y}$, along with inequalities $y_i \neq y_j$ for each pair y_i, y_j of distinct variables in $\mathbf{y}$. Let τ_α be $\psi'_\alpha \rightarrow \exists \mathbf{z}\alpha$, which is a tgd with constants and inequalities.
3. (*Construct* Σ'.) Let Σ' consist of each of these formulas τ_α, one for each prime atom α. **Return** $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$.

Note that as before, the quasi-inverse $\mathcal{M}'$ that our algorithm produces has inequalities only among constants.

THEOREM 4.11. *If $\mathcal{M}$ is a LAV schema mapping, then the output of the algorithm LAV($\mathcal{M}$) is a quasi-inverse of $\mathcal{M}$. Hence, every LAV schema mapping has a quasi-inverse specified by a finite set of tgds with constants and inequalities.*

PROOF. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a LAV schema mapping, and let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be the output of LAV($\mathcal{M}$). Let α be a prime atom, which we shall write as $\alpha(\mathbf{y}, \mathbf{z})$, where $\mathbf{y}, \mathbf{z}$ are as in the algorithm. Let us say that α *affects the chase* if $chase_\Sigma(I_\alpha)$ is not the empty set. The terminology is justified by the LAV assumption, which implies that if α does not affect the chase under our definition, and if $\bar{\mathbf{y}}, \bar{\mathbf{z}}$ is an assignment of values where every variable is assigned a distinct value, then $chase_\Sigma(I) = chase_\Sigma(I \setminus \alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}}))$ for every instance I (here $\setminus$ is the set difference).

We now show that $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$. Assume first that I_1 and I_2 are ground instances such that $(I_1, I_2) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. So there is J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. Let I'_1 consist of those facts $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ in I_1 such that $\alpha(\mathbf{y}, \mathbf{z})$ affects the chase and $\bar{\mathbf{y}}, \bar{\mathbf{z}}$ is an assignment of values where every variable is assigned a distinct value. It is straightforward to see that $chase_\Sigma(I_1) = chase_\Sigma(I'_1)$, and so $I_1 \sim I'_1$. In particular, $(I'_1, J) \models \Sigma$. Let $I'_2 = I_2 \cup I'_1$. Clearly $I'_1 \subseteq I'_2$. We shall show that $I_2 \sim I'_2$. This proves the second inclusion in Part (2) of Corollary 3.10.

Let $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ be a fact in I'_1, where $\bar{\mathbf{y}}, \bar{\mathbf{z}}$ is an assignment such that every variable is assigned a distinct value. Since $(I'_1, J) \models \Sigma$, it follows that J contains a homomorphic image of the result of chasing $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ with Σ (note that the members of $\bar{\mathbf{y}}$, being constants, are each mapped onto themselves by the homomorphism). Therefore, since Σ' contains τ_α, and since $(J, I_2) \models \Sigma'$, there is some assignment $\mathbf{a}$ to $\mathbf{z}$ such that $\alpha(\bar{\mathbf{y}}, \mathbf{a})$ is a fact in I_2. Now the chase of $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ with Σ is contained in the chase of $\alpha(\bar{\mathbf{y}}, \mathbf{a})$ with Σ: this is because (1) by definition of $\mathbf{z}$, it follows that no member of $\bar{\mathbf{z}}$ appears in the chase of $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ with Σ, and (2) $\bar{z_i} = \bar{z_j}$ precisely when $i = j$, whereas there may be additional equalities among the entries of in $\mathbf{a}$. So adding $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ to I_2 does not change its chase with Σ. Since $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ is an arbitrary fact in I_1', this tells us that adding I_1' to I_2 does not affect its chase with Σ (again, we are using the LAV assumption). That is, $chase_\Sigma(I_2) = chase_\Sigma(I_2 \cup I_1')$. But $I_2 \cup I_1' = I_2'$. So $chase_\Sigma(I_2) = chase_\Sigma(I_2')$. Therefore, $I_2 \sim I_2'$, as desired.

Assume now that $I_1 \subseteq I_2$. Let $J = chase_\Sigma(I_2)$, and let $I = chase_{\Sigma'}(J)$. Clearly $(I_2, J) \models \Sigma$ and $(J, I) \models \Sigma'$. Since $(I_2, J) \models \Sigma$ and $I_1 \subseteq I_2$, we see that $(I_1, J) \models \Sigma$. Let I_2' be a ground instance obtained from I by replacing each distinct null in I by a new constant. Since $(J, I) \models \Sigma'$, it follows easily that $(J, I_2') \models \Sigma'$. Let $I_2'' = I_2 \cup I_2'$. Since $(J, I_2') \models \Sigma'$ and $I_2' \subseteq I_2''$, it follows that $(J, I_2'') \models \Sigma'$. Since $(I_1, J) \models \Sigma$ and $(J, I_2'') \models \Sigma'$, we see that $(I_1, I_2'') \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. We must show $I_2 \sim I_2''$ to prove the first inclusion in Part (2) of Corollary 3.10. This will complete the proof that both inclusions in Part (2) of Corollary 3.10 hold, and so (by Corollary 3.10) $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$.

Let γ be a fact in I_2'. So γ is obtained from a fact γ' in I by replacing each distinct null in γ' by a new constant. Since γ' is a fact in I, it is obtained by chasing J with some member τ_α of Σ'. Let ψ_α and ψ_α' be as in the algorithm. Let $\mathbf{y}$ and $\mathbf{z}$ be as in the algorithm, and let $\mathbf{w}$ consist of the distinct variables in ψ_α that are not in $\mathbf{y}$. Let us write ψ_α and ψ_α' as $\psi_\alpha(\mathbf{y}, \mathbf{w})$ and $\psi_\alpha'(\mathbf{y}, \mathbf{w})$ respectively. For notational convenience, let us take $\bar{\mathbf{w}}$ to consist of distinct nulls and let the chase of I_α with Σ be $\psi_\alpha(\mathbf{y}, \bar{\mathbf{w}})$. There are assignments $\bar{\mathbf{y}}, \bar{\mathbf{z}}, \bar{\mathbf{z}}'$ such that γ' is $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}}')$ and γ is $\alpha(\bar{\mathbf{y}}, \bar{\mathbf{z}})$, and the members of $\bar{\mathbf{y}}$ are distinct constants (because of the constant predicates and inequalities in $\psi_\alpha'(\mathbf{y}, \mathbf{w})$), the members of $\bar{\mathbf{z}}'$ are distinct nulls (distinct since γ' is obtained from J by chasing J with τ_α), and the members of $\bar{\mathbf{z}}$ are distinct constants (because γ is obtained from γ' by replacing each distinct null in γ' by a new constant). Now γ' is obtained from J by chasing J with τ_α. Therefore, there is $\bar{\mathbf{w}}'$ such that $\psi_\alpha(\bar{\mathbf{y}}, \bar{\mathbf{w}}')$ holds in J. Let h be the function where $h(\bar{y}_i) = \bar{y}_i$ for each y_i in $\mathbf{y}$, and $h(\bar{w_j}) = \bar{w_j}'$ for each of the nulls $\bar{w_j}$ in $\bar{\mathbf{w}}$. Then h is a homomorphism from $I_{\psi_\alpha(\bar{\mathbf{y}}, \bar{\mathbf{w}})}$ to J (since $\psi_\alpha(\bar{\mathbf{y}}, \bar{\mathbf{w}}')$ holds in J). By construction of τ_α, we know that $chase_\Sigma(I_{\alpha(\mathbf{y},\mathbf{z})})$ is $I_{\psi_\alpha(\mathbf{y}, \bar{\mathbf{w}})}$. Since the members of $\mathbf{y}, \mathbf{z}$ are all distinct, as are the members of $\bar{\mathbf{y}}, \bar{\mathbf{z}}$, it follows that $chase_\Sigma(I_{\alpha(\bar{\mathbf{y}},\bar{\mathbf{z}})})$ is $I_{\psi_\alpha(\bar{\mathbf{y}}, \bar{\mathbf{w}})}$. Thus, the result of chasing (with Σ) the fact γ of I_2' is $I_{\psi_\alpha(\bar{\mathbf{y}}, \bar{\mathbf{w}})}$, which has a homomorphic image (under h) in J. Hence, $chase_\Sigma(I_2 \cup \{\gamma\})$ is homomorphically equivalent to $J = chase_\Sigma(I_2)$. Since γ is an arbitrary fact in I_2', and since $\mathcal{M}$ is LAV, it follows that $chase_\Sigma(I_2 \cup I_2')$ is homomorphically equivalent to $chase_\Sigma(I_2)$. That is, $chase_\Sigma(I_2'')$ is homomorphically equivalent to $chase_\Sigma(I_2)$. So $I_2 \sim I_2''$, as desired. □

4.4 Necessity of the Language

In this section, we prove Part (3) of Theorem 4.1, which says that constants, inequalities, disjunctions, and existential quantifiers are needed in general to express a quasi-inverse. We also prove Part (3) of Theorem 4.9, which says that in the full case, inequalities, disjunctions, and existential quantifiers are needed in general to express a quasi-inverse. We also prove that the result of

Theorem 4.10 is optimal, in that constants, inequalities, and existential quantifiers are needed in general to specify a quasi-inverse of a LAV mapping.

4.4.1 *Necessity of Constants.* We begin with a theorem about the necessity of constants for an inverse, and as a corollary show the necessity of constants for a quasi-inverse.

THEOREM 4.12. *There is a LAV schema mapping that has an inverse, but does not have an inverse specified by a set of disjunctive tgds with inequalities.*

To prove this, we shall prove a stronger result. Let σ be a formula that involves only source and target relation symbols. Let us say that σ is *indifferent to constants* if whenever $(J, I) \models \sigma$, and (J', I') is isomorphic to (J, I) (but where the isomorphism may map constants into either constants or nulls, and may map nulls into either constants or nulls), then $(J', I') \models \sigma$. It is easy to see that disjunctive tgds with inequalities (but not with constants) are indifferent to constants. We shall prove the following theorem, which immediately implies Theorem 4.12.

THEOREM 4.13. *There is a LAV schema mapping that has an inverse, but does not have an inverse specified by a set of formulas that are indifferent to constants.*

PROOF. Let $\mathbf{S}$ consist of a binary relation symbol P, and let $\mathbf{T}$ consist of a binary relation symbol Q. Let Σ consist of the tgd $P(x, y) \rightarrow \exists z(Q(x, z) \wedge Q(z, y))$, and let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$. We now show that $\mathcal{M}$ has an inverse (specified by a full tgd with constants), but does not have an inverse specified by a set of formulas that are indifferent to constants.

Let Σ' be $Q(x, z) \wedge Q(z, y) \wedge \mathit{Constant}(x) \wedge \mathit{Constant}(y) \rightarrow P(x, y)$, and let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$. We now show that $\mathcal{M}'$ is an inverse of $\mathcal{M}$. Assume first that I_1 and I_2 are ground instances and $(I_1, I_2) \models \Sigma \circ \Sigma'$. So there is J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$. It is easy to see by definition of Σ and Σ' that $I_1 \subseteq I_2$. Conversely, assume that I_1 and I_2 are ground instances with $I_1 \subseteq I_2$. Let $J = \mathit{chase}_\Sigma(I_1)$. It is straightforward to verify that if x and y represent constants and $Q(x, z)$ and $Q(z, y)$ hold in $J = \mathit{chase}_\Sigma(I_1)$, then $P(x, y)$ holds in I_1, and hence in I_2. It follows easily that $(J, I_2) \models \Sigma'$. Since also $(I_1, J) \models \Sigma$, it follows that $(I_1, I_2) \models \Sigma \circ \Sigma'$. So $\mathcal{M}'$ is an inverse of $\mathcal{M}$, as desired.

Assume now that $\mathcal{M}$ has an inverse $\mathcal{M}'' = (\mathbf{T}, \mathbf{S}, \Sigma'')$, where Σ'' is a set of formulas that are indifferent to constants; we shall derive a contradiction. Let α be the formula $Q(x, z) \wedge Q(z, y) \rightarrow P(x, y)$. We now prove that Σ'' logically implies α. Assume not; we shall derive a contradiction. Since Σ'' does not logically imply α, there are J, I such that $(J, I) \models \Sigma''$ but $(J, I) \not\models \alpha$. Since $(J, I) \not\models \alpha$, there are $\bar{x}, \bar{y}, \bar{z}$ such that $Q(\bar{x}, \bar{z})$ and $Q(\bar{z}, \bar{y})$ hold in J but $P(\bar{x}, \bar{y})$ does not hold in I. Since the formulas in Σ'' are indifferent to constants, we can assume (by renaming the nulls to new constants, if needed) that the active domains of J and I (including $\bar{x}$, $\bar{y}$, and $\bar{z}$) consist only of constants. Let I' consist of the single fact $P(\bar{x}, \bar{y})$. Then $(I', J) \models \Sigma$, since J contains the facts $Q(\bar{x}, \bar{z})$ and $Q(\bar{z}, \bar{y})$. Since also $(J, I) \models \Sigma''$, it follows that $(I', I) \models \Sigma \circ \Sigma''$. But it is false that $I' \subseteq I$, since $P(\bar{x}, \bar{y})$ does not hold in I. This contradicts the definition of inverse. So indeed, Σ'' logically implies α.

Let I be the ground instance consisting of the two facts $P(0, 1)$ and $P(1, 0)$. Since $\mathcal{M}''$ is an inverse of $\mathcal{M}$, it follows that $(I, I) \models \Sigma \circ \Sigma''$. So there is J such that $(I, J) \models \Sigma$ and $(J, I) \models \Sigma''$. Since Σ'' logically implies α, it follows that $(J, I) \models \alpha$. Since $(I, J) \models \Sigma$, there is z_1 such that $Q(0, z_1)$ and $Q(z_1, 1)$ hold in J, and there is z_2 such that $Q(1, z_2)$ and $Q(z_2, 0)$ hold in J. Since $Q(z_1, 1)$ and $Q(1, z_2)$ hold in J and since $(J, I) \models \alpha$, it follows that $P(z_1, z_2)$ holds in I. So either $z_1 = 0$ and $z_2 = 1$, or $z_1 = 1$ and $z_2 = 0$. By symmetry, we can assume without loss of generality that $z_1 = 0$ and $z_2 = 1$. So $Q(0, 0)$, $Q(0, 1)$, $Q(1, 1)$ and $Q(1, 0)$ hold in J. Because $Q(0, 1)$ and $Q(1, 0)$ hold in J, and $(J, I) \models \alpha$, it follows that $P(0, 0)$ holds in I, which is our desired contradiction. □

COROLLARY 4.14. *There is a LAV schema mapping that has a quasi-inverse, but does not have a quasi-inverse specified by a set of disjunctive tgds with inequalities.*

PROOF. By Theorem 4.12, there is a schema mapping $\mathcal{M}$ specified by a finite set of s-t tgds that has an inverse, but does not have an inverse specified by a set of disjunctive tgds with inequalities. Since $\mathcal{M}$ has an inverse, it has a quasi-inverse (every inverse is a quasi-inverse). Moreover, if $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$, then $\mathcal{M}'$ is an inverse of $\mathcal{M}$, by Proposition 3.24. Thus, if $\mathcal{M}$ were to have quasi-inverse specified by a set of disjunctive tgds with inequalities, then it would have an inverse specified by a set of disjunctive tgds with inequalities, which is a contradiction. □

4.4.2 *Necessity of Inequalities.* We begin this section with results about the necessity of inequalities for an inverse, and the necessity of inequalities for a quasi-inverse.

THEOREM 4.15. *There is a LAV schema mapping specified by a finite set of full s-t tgds that has an inverse, but does not have an inverse specified by a finite set of disjunctive tgds with constants.*

PROOF. Let $\mathbf{S}$ consist of the binary relation symbol P and the unary relation symbol T. Let $\mathbf{T}$ consist of the binary relation symbol P' and the unary relation symbols Q and T'. Let $\Sigma = \{P(x, y) \to P'(x, y), P(x, x) \to Q(x), T(x) \to T'(x), T(x) \to P'(x, x)\}$, and let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$.

We now show that $\mathcal{M}$ has the $(=, =)$-subset property. Assume that I_1 and I_2 are ground instances such that $\mathrm{Sol}(\mathcal{M}, I_2) \subseteq \mathrm{Sol}(\mathcal{M}, I_1)$. Let $J_1 = \mathit{chase}_\Sigma(I_1)$ and $J_2 = \mathit{chase}_\Sigma(I_2)$. Since $\mathrm{Sol}(\mathcal{M}, I_2) \subseteq \mathrm{Sol}(\mathcal{M}, I_1)$, it follows that $\mathit{chase}_\Sigma(I_1) \subseteq \mathit{chase}_\Sigma(I_2)$, that is, $J_1 \subseteq J_2$. We will show that $P^{I_1} \subseteq P^{I_2}$ and $T^{I_1} \subseteq T^{I_2}$, which means that $I_1 \subseteq I_2$, as desired. If $(a, b) \in P^{I_1}$ and $a \neq b$, then $(a, b) \in P'^{J_1} \subseteq P'^{J_2}$. Since $a \neq b$, the only way for $P'^{J_2}(a, b)$ to hold is to have $(a, b) \in P^{I_2}$. If $a = b$ and $P^{I_1}(a, a)$ holds, then $Q^{J_1}(a)$ holds, and so $Q^{J_2}(a)$ also holds, hence $P^{I_2}(a, a)$ holds. This shows that $P^{I_1} \subseteq P^{I_2}$. We also have that $T^{I_1} = T'^{J_1} = T'^{J_2} = T^{I_2}$. This completes the proof that $I_1 \subseteq I_2$. So $\mathcal{M}$ has the $(=, =)$-subset property, which implies that $\mathcal{M}$ has an inverse.

Assume now that $\mathcal{M}$ has an inverse $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ' is a set of disjunctive tgds with constants but without inequalities; we shall derive a contradiction. Let Σ'' be obtained from Σ' by removing every conjunct of the form *Constant*(v). Let $\mathcal{M}'' = (\mathbf{T}, \mathbf{S}, \Sigma'')$, By Part (2b) of Proposition 4.8, it follows that $\mathcal{M}''$ is an inverse of $\mathcal{M}$. Now Σ'' is a set of disjunctive tgds, without inequalities or constants.

Let χ be a member of Σ'' such that the relation symbol Q does not occur in the premise of χ. We claim that the conclusion of χ must contain a disjunct whose conjuncts are of the form $\exists y T(y)$ or of the form $T(x)$. To see this, let I be the source instance with $P^I = \emptyset$ and $T^I = \{a\}$, for some constant a. Since $(I, I) \in \text{Inst(Id)}$ and $\mathcal{M}''$ is an inverse of $\mathcal{M}$, we know that $(I, I) \models \Sigma \circ \Sigma''$, so there is a target instance J such that $(I, J) \models \Sigma$ and $(J, I) \models \Sigma''$. Clearly, J contains the facts $T'(a)$ and $P'(a, a)$. Since $(J, I) \models \chi$ and Q does not occur in the premise of χ, it follows that the premise of χ becomes true when all universally quantified variables in χ take value a. Consequently, there is a disjunct $\exists \mathbf{y}\varphi(\mathbf{x}, \mathbf{y})$ in the conclusion of χ, such that $\exists \mathbf{y}\varphi(\bar{\mathbf{x}}, \mathbf{y})$ holds in I, where $\bar{\mathbf{x}}$ assigns a to every variable in $\mathbf{x}$. Since $P^I = \emptyset$, it must be the case that P does not occur in $\varphi(\mathbf{x}, \mathbf{y})$, which means that each conjunct of $\varphi(\mathbf{x}, \mathbf{y})$ is of the form $\exists y T(y)$ or of the form $T(x)$.

Let $I_1 = \{P(a_1, a_2), T(a_1), T(a_2)\}$ and $I_2 = \{T(a_1), T(a_2)\}$. Clearly, $I_1 \not\subseteq I_2$ and $(I_1, chase_\Sigma(I_1)) \models \Sigma$. We claim that $(chase_\Sigma(I_1), I_2) \models \Sigma''$, which would imply that $(I_1, I_2) \in \text{Inst}(\Sigma \circ \Sigma'') = \text{Inst(Id)}$, thus arriving at a contradiction. Observe that $chase_\Sigma(I_1)$ consists of the facts $\{P'(a_1, a_2), P'(a_1, a_1), P'(a_2, a_2)\}$; in particular, it contains no fact involving the relation symbol Q. Let us consider an arbitrary dependency χ in Σ''. If Q occurs in the premise of χ, then $(chase_\Sigma(I_1), I_2) \models \chi$, since the premise of χ never becomes true. So, assume that Q does not occur in the premise of χ. If there is an assignment of values to the universally quantified variables in χ so that $(chase_\Sigma(I_1), I_2)$ satisfies the premise of χ, then each variable must be assigned value a_1 or a_2. Since, by the claim in the preceding paragraph, the conclusion of χ contains a disjunct consisting of conjunctions of the form $\exists y T(y)$ or of the form $T(x)$, we have that the conclusion of χ becomes true in I_2, since I_2 contains the facts $T(a_1)$ and $T(a_2)$. This completes the proof of the claim that $(chase_\Sigma(I_1), I_2) \models \Sigma''$, which was to be shown. □

COROLLARY 4.16. *There is a LAV schema mapping specified by a finite set of full s-t tgds that has a quasi-inverse, but does not have a quasi-inverse specified by a finite set of disjunctive tgds with constants.*

PROOF. This follows from Theorem 4.15 just as Corollary 4.14 follows from Theorem 4.12. □

4.4.3 *Necessity of Disjunctions.*

THEOREM 4.17. *There is a schema mapping specified by a finite set of full tgds that has a quasi-inverse, but no quasi-inverse that is specified by a set of tgds with constants and inequalities.*

PROOF. Let **S** consist of four unary relation symbols P_1, P_2, P_3, P_4. and let **T** consist of six unary relation symbols S_1, S_2, R_{13}, R_{14}, R_{23}, R_{24}.

Let Σ consist of the following eight full tgds:

$$P_1(x) \to S_1(x) \qquad P_2(x) \to S_1(x) \qquad P_3(x) \to S_2(x) \qquad P_4(x) \to S_2(x)$$
$$P_i(x) \wedge P_j(x) \to R_{ij}(x), \text{ for } i \in \{1,2\} \text{ and } j \in \{3,4\}$$

We now show that $\mathcal{M}$ satisfies the $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-subset property, and hence has a quasi-inverse. In fact, we shall show that $\mathcal{M}$ satisfies the $(\sim_{\mathcal{M}}, =)$-subset property. Thus, we shall show that if (I_1, I_2) is a pair of ground instances, where $\text{Sol}(\mathcal{M}, I_1) \subseteq \text{Sol}(\mathcal{M}, I_2)$, then there is a ground instance I_2' such that $I_2 \sim I_2'$ and $I_2' \subseteq I_1$.

Let us refer to members of $P_1^{I_1}$ or $P_2^{I_1}$ as members of the *top half* of I_1. Similarly, let us refer to members of $P_3^{I_1}$ or $P_4^{I_1}$ as members of the *bottom half* of I_1. Likewise, define the top half and bottom half of I_2. Each constant in the active domain of I_1 must be in the top half of I_1, the bottom half of I_1, or both, and similarly for I_2.

We now define a ground instance I_2' with the same active domain as I_2. We shall show that $I_2 \sim I_2'$, and $I_2' \subseteq I_1$, as desired. We shall define I_2' by saying, for each constant c in the active domain of I_2, for which relations P_i we have $P_i^{I_2'}(c)$. There are three cases.

Case 1: c is a constant in both the top half and the bottom half of I_2. For definiteness, assume that $c \in P_i^{I_2}$, where i is either 1 or 2, and $c \in P_j^{I_2}$, where j is either 3 or 4. We now show that $c \in P_i^{I_1}$ and $c \in P_j^{I_1}$. Assume not. Let $J = chase_\Sigma(I_1)$. Then R_{ij}^J does not contain c. But every solution J' for I_2 has $c \in R_{ij}^{J'}$. So J is a solution for I_1 but not for I_2. This is a contradiction, since $\text{Sol}(\mathcal{M}, I_1) \subseteq \text{Sol}(\mathcal{M}, I_2)$. Put c in the P_i relation of I_2' precisely if c is in the P_i relation of I_2, for $i \in \{1, 2, 3, 4\}$, Therefore, by what we just showed, if c is in the P_i relation of I_2', then c is in the P_i relation of I_1, for $i \in \{1, 2, 3, 4\}$.

Case 2: c is a constant in the top half of I_2 but not in the bottom half of I_2. We show that c is in the top half of I_1. Assume not. Let $J = chase_\Sigma(I_1)$. Then S_1^J does not contain c. But every solution J' for I_2 has $c \in S_1^{J'}$. So J is a solution for I_1 but not for I_2. This is a contradiction, since $\text{Sol}(\mathcal{M}, I_1) \subseteq \text{Sol}(\mathcal{M}, I_2)$. So c is in the top half of I_1. Therefore, c is in either $P_1^{I_1}$ or $P_2^{I_1}$. If c is in $P_1^{I_1}$, then put c in the P_1 relation of I_2'; otherwise, put c in the P_2 relation of I_2'.

Case 3: c is a constant in the bottom half of I_2 but not in the top half of I_2. By a similar argument to what we gave in Case 2, we see that c is in the bottom half of I_1. Since c is in the bottom half of I_1, we know that c is in either $P_3^{I_1}$ or $P_4^{I_1}$. If c is in $P_3^{I_1}$, then put c in the P_3 relation of I_2'; otherwise, put c in the P_4 relation of I_2'.

This completes the definition of I_2'. By considering, for each constant c in the active domain of I_2, which of the three possible cases c satisfies, it is straightforward to verify that $chase_\Sigma(I_2) = chase_\Sigma(I_2')$, and so $I_2 \sim I_2'$. By the comments we made during our construction of I_2', it follows easily that $I_2' \subseteq I_1$. This completes the proof that $\mathcal{M}$ satisfies the $(\sim_{\mathcal{M}}, =)$-subset property, and so satisfies the $(\sim_{\mathcal{M}}, \sim_{\mathcal{M}})$-subset property. Therefore,$\mathcal{M}$ has a quasi-inverse.

Assume now that $\mathcal{M}$ has a quasi-inverse $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ where Σ' is a set of tgds with constants and inequalities. We shall derive a contradiction. Let J_1 be the target instance $\{S_1(0)\}$. We now show that $chase_{\Sigma'}(J_1)$ contains either the fact $P_1(0)$ or the fact $P_2(0)$. Assume not. Let $I_1 = \{P_1(0)\}$. Then $J_1 = chase_{\Sigma}(I_1)$. Let I_2 be the result of replacing each null in $chase_{\Sigma'}(J_1)$ by a constant other than 0. Then $(J_1, I_2) \models \Sigma'$, and I_2 is a ground instance that contains neither the fact $P_1(0)$ nor the fact $P_2(0)$. Since $(I_1, J_1) \models \Sigma$ and $(J_1, I_2) \models \Sigma'$, it follows that $(I_1, I_2) \models \Sigma \circ \Sigma'$. So by the definition of quasi-inverse, there are I_1', I_2' where $I_1 \sim I_1'$ and $I_2 \sim I_2'$, and $I_1' \subseteq I_2'$. Since $I_1 \sim I_1'$, and since I_1 contains the fact $P_1(0)$, it follows that I_1' contains either the fact $P_1(0)$ or the fact $P_2(0)$. Since $I_2 \sim I_2'$, and since I_2 contains neither the fact $P_1(0)$ nor the fact $P_2(0)$, it follows that I_2' contains neither the fact $P_1(0)$ nor the fact $P_2(0)$. But this is impossible, since $I_1' \subseteq I_2'$. So there is a member σ_1 of Σ' such that the chase of J_1 with σ_1 contains either the fact $P_1(0)$ or the fact $P_2(0)$. Without loss of generality, assume that the chase of J_1 with σ_1 contains the fact $P_1(0)$.

Let J_2 be the target instance $\{S_2(0)\}$. By the same argument, but where we replace P_1 by P_3, and replace P_2 by P_4, we see that there is a member σ_2 of Σ' such that the chase of J_2 with σ_2 contains either the fact $P_3(0)$ or the fact $P_4(0)$. Without loss of generality, assume that the chase of J_2 with σ_2 contains the fact $P_3(0)$.

Let $I_1 = \{P_2(0), P_4(0)\}$. Since $(I_1, I_1) \in \text{Inst(Id)}$, it follows from the definition of a quasi-inverse that $(I_1, I_1) \in \text{Inst}(\Sigma \circ \Sigma')[\sim_{\mathcal{M}}, \sim_{\mathcal{M}}]$. So there are I_1', I_1'' with $I_1 \sim I_1'$ and $I_1 \sim I_1''$ such that $(I_1', I_1'') \models \Sigma \circ \Sigma'$. Hence, there is J such that $(I_1', J) \models \Sigma$ and $(J, I_1'') \models \Sigma'$. Since $I_1 \sim I_1'$, it follows that $chase_{\Sigma}(I_1) = chase_{\Sigma}(I_1')$. Now $chase_{\Sigma}(I_1)$ contains the facts $S_1(0), S_2(0)$. So $chase_{\Sigma}(I_1')$ contains the facts $S_1(0), S_2(0)$, and hence J contains the facts $S_1(0), S_2(0)$. We showed that there is σ_1 in Σ' such that the chase of $\{S_1(0)\}$ with σ_1 contains the fact $P_1(0)$. So the result of chasing J with σ_1 contains the fact $P_1(0)$. Since every fact about constants in the result of chasing J with σ_1 is forced to hold for all solutions for J with respect to σ_1, it follows that I_1'' contains the fact $P_1(0)$. We also showed that there is σ_2 in Σ' such that the chase of $\{S_2(0)\}$ with σ_2 contains the fact $P_3(0)$, and so by a similar argument it follows that I_1'' contains the fact $P_3(0)$. Since I_1'' contains the facts $P_1(0), P_3(0)$, it follows that every solution of I_1'' with respect to Σ contains $R_{13}(0)$. Let $J^* = chase_{\Sigma}(I_1) = \{S_1(0), S_2(0), R_{24}(0)\}$. Then J^* is a solution of I_1 with respect to Σ, and hence of I_1'' with respect to Σ, since $I_1 \sim I_1''$. But J^* does not contain $R_{13}(0)$. This is our desired contradiction. □

4.4.4 *Necessity of Existential Quantifiers.* In Theorem 5.1, the inverse is specified by *full* tgds with constants and inequalities. By contrast, in Theorem 4.1, the disjunctive tgds with constants and inequalities that specify the quasi-inverse are not required to be full. In this section, we show that this is inevitable. We begin with an example where it seems at first that there is no *full* disjunctive tgd that specifies a quasi-inverse, but there actually is.

Example 4.18. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ where $\Sigma = \{P(x, y) \to R(x)\}$. It is straightforward to verify that the (nonfull) s-t tgd $R(x) \to \exists y P(x, y)$

specifies a quasi-inverse of $\mathcal{M}$. Since R is unary and P is binary, it is natural to believe that existential quantification is necessary. However, it is straightforward to verify that $R(x) \rightarrow P(x, x)$ also specifies a quasi-inverse of $\mathcal{M}$, and is full.

We now show that, in fact, we cannot always avoid the existential quantifiers.

THEOREM 4.19. *There is a LAV schema mapping specified by a finite set of full tgds that has a quasi-inverse, but does not have a quasi-inverse specified by a set of full disjunctive tgds with constants and inequalities.*

PROOF. Let $\mathbf{S}$ consist of a single binary relation symbol P, and let $\mathbf{T}$ consist of two unary relation symbols R and S. Let $\Sigma = \{P(x, y) \rightarrow R(x), P(x, x) \rightarrow S(x)\}$, and let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$. Since $\mathcal{M}$ is a LAV mapping, it has a quasi-inverse.

Assume now that $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is a quasi-inverse of $\mathcal{M}$, where Σ' is a set of full disjunctive tgds with constants and inequalities; we shall derive a contradiction. There are two cases.

Case 1: Every member of Σ' has either S or an inequality in the premise. Let I_1 consist of the fact $P(0, 1)$, and let $J = \mathit{chase}_\Sigma(I_1) = \{R(0)\}$. Clearly $(I_1, J) \models \Sigma$. Let I_2 be the empty instance. Let σ be an arbitrary member of Σ'. Since by assumption σ has either S or an inequality in the premise, there is no assignment of values to the variables in the premise of σ that is true in J. So $(J, I_2) \models \sigma$. Since σ is an arbitrary member of Σ', it follows that $(J, I_2) \models \Sigma'$. Therefore, $(I_1, I_2) \models \Sigma \circ \Sigma'$. Since $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$, it follows that there are I_1', I_2' with $I_1 \sim I_1'$ and $I_2 \sim I_2'$, and with $I_1' \subseteq I_2'$. Since $I_1 \sim I_1'$, we see that I_1' is nonempty. Since $I_2 \sim I_2'$, we see that I_2' is empty. So it is not possible to have $I_1' \subseteq I_2'$. This is our desired contradiction.

Case 2: Some member of Σ' does not have S in the premise and does not have an inequality in the premise. Again, let I_1 consist of the fact $P(0, 1)$. Since $\mathcal{M}'$ is a quasi-inverse of $\mathcal{M}$, it follows that there are I_1', I_1'' with $I_1 \sim I_1'$ and $I_1 \sim I_1''$, where $(I_1', I_1'') \models \Sigma \circ \Sigma'$. Since $I_1 \sim I_1'$, it is easy to see that I_1' consists of a nonempty set of facts all of the form $P(0, x)$, where $x \neq 0$. Identically, I_1'' consists of a nonempty set of facts all of the form $P(0, x)$, where $x \neq 0$.

Since $(I_1', I_1'') \models \Sigma \circ \Sigma'$, there is J such that $(I_1', J) \models \Sigma$ and $(J, I_1'') \models \Sigma'$. Since $(I_1', J) \models \Sigma$, clearly J contains $\mathit{chase}_\Sigma(I_1) = \{R(0)\}$. By assumption, there is a member σ of Σ' that does not have S in the premise and does not have an inequality in the premise. By assumption, σ is a full disjunctive tgd with constants and inequalities. Take a truth assignment that assigns 0 to every variable in the premise of σ (and hence to every variable in the conclusion of σ, since σ is full). The premise is then a conjunction of formulas $R(0)$ and $\mathit{Constant}(0)$, all of which hold in J. The conclusion is simply a disjunction of conjunctions of formulas $P(0, 0)$, since P is the only relation symbol in the conclusion, and every variable is assigned the value 0. So the conclusion is equivalent to $P(0, 0)$. Since $(J, I_1'') \models \Sigma'$, it follows that $P(0, 0)$ holds in I_1''. This is our desired contradiction. □

4.4.5 *Putting Together the Necessity Results.* Part (3) of Theorem 4.1 follows from Corollary 4.14 (necessity of constants), Corollary 4.16 (necessity of inequalities), Theorem 4.17 (necessity of disjunctions), and Theorem 4.19 (necessity of existential quantifiers). Part (3) of Theorem 4.9 follows from the fact that for each of the examples used to prove Part (3) of Theorem 4.1 (except the necessity of constants), the s-t tgds in Σ are full. Note also that Corollary 4.14 (necessity of constants), Corollary 4.16 (necessity of inequalities), and Theorem 4.19 (necessity of existential quantifiers) all use LAV mappings. Hence, the result of Theorem 4.10 is optimal, in that constants, inequalities, and existential quantifiers are needed in general to specify a quasi-inverse of a LAV mapping.

5. THE LANGUAGE OF INVERSES

The focus in Fagin [2007] is on inverses that are specified by a finite set of tgds. For example, given a schema mapping $\mathcal{M}$ specified by a finite set of s-t tgds, Fagin [2007] gives an algorithm for constructing a schema mapping specified by finite set of tgds that is an inverse of $\mathcal{M}$ if and only if there is an inverse of $\mathcal{M}$ that is specified by a finite set of tgds. If there is an inverse $\mathcal{M}'$ but there is no inverse specified by a finite set of tgds, then the algorithm in Fagin [2007] will not find $\mathcal{M}'$. The "language of inverses" is left as an open problem in Fagin [2007]. This is the question as to what language is needed to specify the inverse of $\mathcal{M}$, when $\mathcal{M}$ is specified by a finite set of s-t tgds. The next theorem resolves this open problem.

THEOREM 5.1. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. If $\mathcal{M}$ has an inverse then the following hold.*

1. *$\mathcal{M}$ has an inverse $\mathcal{M}'$ specified by a finite set of full tgds with constants and inequalities.*
2. *There is an exponential-time algorithm for producing $\mathcal{M}'$.*
3. *Statement (1) is not necessarily true if we disallow either constants or inequalities in the premise, even if we allow existential quantifiers in the conclusion (and so allow nonfull dependencies to specify $\mathcal{M}'$).*

In fact, the inverse our algorithm produces has inequalities only among constants.

We now discuss the machinery used to prove Theorem 5.1.

Definition 5.2. A schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds, satisfies the *constant-propagation property* if for every ground instance I, every member of the active domain of I is in the active domain of $chase_\Sigma(I)$.

It is straightforward to see that $\mathcal{M}$ satisfies the constant-propagation property precisely if, for each relation symbol R in $\mathbf{S}$, the chase of $R(x_1, \ldots, x_m)$ with Σ includes each of the m distinct variables $x_1, \ldots, x_m$, where m is the arity of R.

We shall use the following proposition from [Fagin 2007].

PROPOSITION 5.3. *[Fagin 2007] Every invertible schema mapping that is specified by a finite set of s-t tgds satisfies the constant-propagation property.*

We now give an algorithm that produces an inverse if one exists.

Algorithm Inverse($\mathcal{M}$)

Input: A schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where Σ is a finite set of s-t tgds.
Output: A schema mapping $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ' is a finite set of full tgds with constants and inequalities, and $\mathcal{M}'$ is an inverse of $\mathcal{M}$ if $\mathcal{M}$ has an inverse. There is no output if $\mathcal{M}$ does not satisfy the constant-propagation property.

1. (*Verify that $\mathcal{M}$ satisfies the constant-propagation property.*) Check to see if, for each relation symbol R in $\mathbf{S}$, the chase of $R(x_1, \ldots, x_m)$ with Σ includes each of the m distinct variables $x_1, \ldots, x_m$, where m is the arity of R. If not, halt without output. If so, continue to the next step.
2. (*Generate all prime source atoms in lexicographic order.*) For example, if R is a ternary source relation symbol, the atoms for R, in lexicographic order, are $R(x_1, x_1, x_1)$, $R(x_1, x_1, x_2)$, $R(x_1, x_2, x_1)$, $R(x_1, x_2, x_2)$, $R(x_1, x_2, x_3)$.
3. (*Construct a full tgd $\omega(\Sigma, I)$ with constants and inequalities for each prime instance I.*) For each prime source atom α generated in Step (2), let I_α be the prime instance containing only α. Let ψ_α be the conjunction of the facts of $chase_\Sigma(I_\alpha)$. Form a full tgd $\omega(\Sigma, I_\alpha)$ with constants and inequalities whose premise is the conjunction of ψ_α with the formulas $Constant(x)$ for each variable x that appears in α, along with inequalities $x_i \neq x_j$ for each pair x_i, x_j of distinct variables that appear in α, and whose conclusion is α.
4. (*Construct Σ'.*) Let Σ' consist of each of these formulas $\omega(\Sigma, I)$, one for each prime instance I. **Return** $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$.

Assume that $\mathcal{M}$ satisfies the constant-propagation property. Then the algorithm gives an output. Furthermore, $\omega(\Sigma, I_\alpha)$, as formed in Step (3), is then a well-defined full tgd with constants and inequalities, since every variable in the conclusion of $\omega(\Sigma, I_\alpha)$ necessarily appears in the premise.

Note that in the case when $\mathcal{M}$ is a schema mapping (LAV or otherwise) that satisfies the constant-propagation property, the output of Inverse($\mathcal{M}$) and LAV($\mathcal{M}$) is the same. The tgds in Σ' that are constructed in Step (2) of LAV($\mathcal{M}$) are automatically full when $\mathcal{M}$ satisfies the constant-propagation property.

Example 5.4. Let $\mathbf{S}$ consist of a binary relation symbol R. Let $\mathbf{T}$ consist of a binary relation symbol Q, ternary relation symbol S, and unary relation symbol U. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ where Σ consists of the tgds:

$$R(x_1, x_2) \wedge R(x_2, x_1) \rightarrow \exists y\, Q(x_1, y)$$
$$R(x_1, x_2) \rightarrow \exists y\, S(x_1, x_2, y)$$
$$R(x_1, x_1) \rightarrow U(x_1)$$

Then $\mathcal{M}$ satisfies the constant-propagation property, since the chase of $R(x_1, x_2)$ is $S(x_1, x_2, y)$, which contains both of the variables x_1 and x_2 of $R(x_1, x_2)$. The two prime source atoms are $R(x_1, x_1)$ and $R(x_1, x_2)$. The two prime instances are $I_{R(x_1,x_1)} = \{R(x_1, x_1)\}$ and $I_{R(x_1,x_2)} = \{R(x_1, x_2)\}$. Then $\omega(\Sigma, I_{R(x_1,x_1)})$ is

$$Q(x_1, y_1) \wedge S(x_1, x_1, y_2) \wedge U(x_1) \wedge Constant(x_1) \rightarrow R(x_1, x_1). \qquad (1)$$

Also, $\omega(\Sigma, I_{R(x_1,x_2)})$ is

$$S(x_1, x_2, y) \wedge Constant(x_1) \wedge Constant(x_2) \wedge (x_1 \neq x_2) \rightarrow R(x_1, x_2). \qquad (2)$$

The output of Inverse($\mathcal{M}$) is $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where Σ' consists of (1) and (2).

Later (Theorem 5.7), we show that if $\mathcal{M}$ is invertible, then the output $\mathcal{M}'$ of the algorithm is an inverse of $\mathcal{M}$.

Proposition 3.24 tells us that every quasi-inverse of an invertible schema mapping $\mathcal{M}$ is an inverse of $\mathcal{M}$. The reader might therefore wonder why we need both the algorithms QuasiInverse and Inverse, since QuasiInverse will necessarily produce an inverse if the input is an invertible schema mapping. The answer is that in this case, QuasiInverse will produce an inverse specified by disjunctive tgds with constants and equalities where disjunctions may actually appear, even though there is an inverse specified by full (and non-disjunctive) tgds with constants and equalities that Inverse will find.

Example 5.5. We now give an example of an invertible schema mapping specified by full s-t tgds where the output of the QuasiInverse algorithm is an inverse where disjunctions actually appear. Let **S** consist of two unary relation symbols P and Q, and let **T** consist of three unary relation symbols P', Q', and R. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$, where $\Sigma = \{P(x) \rightarrow P'(x), Q(x) \rightarrow Q'(x), P(x) \rightarrow R(x), Q(x) \rightarrow R(x)\}$. Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$, where $\Sigma' = \{P'(x) \rightarrow P(x), Q'(x) \rightarrow Q(x)\}$. It is easy to see that $\mathcal{M}'$ is an inverse of $\mathcal{M}$.

The algorithm Inverse produces a slightly weaker inverse, namely one that is specified by $\{P'(x) \wedge R(x) \wedge Constant(x) \rightarrow P(x), Q'(x) \wedge R(x) \wedge Constant(x) \rightarrow Q(x)\}$. Of course, we see directly from Part (2b) of Proposition 4.8 that the constant predicates are redundant.

The algorithm QuasiInverse produces a schema mapping that is specified by the set $\{P'(x) \wedge Constant(x) \rightarrow P(x), Q'(x) \wedge Constant(x) \rightarrow Q(x), R(x) \wedge Constant(x) \rightarrow P(x) \vee Q(x)\}$ of disjunctive tgds with constants. In particular, the last member actually has a disjunction appear.

5.1 Candidate Inverse

In Section 3.3, we introduced a candidate $(\sim_1, \sim_2)$-inverse, and showed that if a schema mapping specified by a finite set of s-t tgds has a $(\sim_1, \sim_2)$-inverse, then this candidate $(\sim_1, \sim_2)$-inverse is a $(\sim_1, \sim_2)$-inverse. In the case where $(\sim_1, \sim_2)$ is $(=, =)$, we refer to it as a *candidate inverse*. Let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ and $\mathcal{M}'' = (\mathbf{T}, \mathbf{S}, \Sigma'')$ be target to source schema mappings. We say that Σ'' *weakly implies* Σ' if whenever J is a target instance and I is a ground instance such that $(J, I) \models \Sigma''$, then $(J, I) \models \Sigma'$. (This is the notion of weak implication mentioned in Section 3.3.) Let us say that Σ' and Σ'' are *weakly equivalent* if Σ' weakly implies Σ'', and Σ'' weakly implies Σ'. Thus, Σ' and Σ'' are weakly equivalent precisely if whenever J is a target instance and I is a ground instance, then $(J, I) \models \Sigma'$ if and only if $(J, I) \models \Sigma''$. We may then say that the schema mappings $\mathcal{M}'$ and $\mathcal{M}''$ are weakly equivalent.

In this section, we show that if $\mathcal{M}$ is an invertible schema mapping that is specified by a finite set of s-t tgds, then our candidate inverse from Section 3.3

is weakly equivalent to the schema mapping that is produced by Inverse($\mathcal{M}$). We make use of this to show if $\mathcal{M}$ is invertible, then the output of Inverse($\mathcal{M}$) is an inverse of $\mathcal{M}$.

Recall that $cand_{(=,=)}(\Sigma)$, which specifies the candidate inverse, consists of all of the formulas $\tau_{(=,=)}(\Sigma, I)$, the "constraint associated with I", for every ground instance I. In turn, $\tau_{(=,=)}(\Sigma, I)$ defines the set of all pairs (J, I_2) such that I_2 is a ground instance and such that if $(I, J) \models \Sigma$ then $I \subseteq I_2$. The next theorem says that $cand_{(=,=)}(\Sigma)$, which specifies the candidate inverse, is weakly equivalent to our finite set $\Omega(\Sigma)$ of full tgds with constants and inequalities in the schema mapping that is produced by Inverse($\mathcal{M}$). Thus, the theorem says that if J is a target instance, and I_2 is a ground instance, then

$$(J, I_2) \models cand_{(=,=)}(\Sigma) \text{ if and only if } (J, I_2) \models \Omega(\Sigma). \tag{3}$$

THEOREM 5.6. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping, where Σ is a finite set of s-t tgds. Assume that $\mathcal{M}$ satisfies the constant-propagation property. Then $cand_{(=,=)}(\Sigma)$ is weakly equivalent to $\Omega(\Sigma)$.*

PROOF. Assume that J is a target instance and I_2 is a ground instance. We shall hold J and I_2 fixed throughout the proof. For each ground instance I, define $T(I)$ to be the set of all of the constraints $\tau_{(=,=)}(\Sigma, I_1)$ where I_1 is a ground instance that is isomorphic to I (since Σ is fixed, we do not bother with including it in our notation $T(I)$). We first show that if I is a ground instance, then $T(I)$ is weakly equivalent to $\omega(\Sigma, I)$, as defined in Step (3) of the algorithm Inverse($\mathcal{M}$). Let us write $\omega(\Sigma, I)$ as $\psi(\mathbf{x}, \mathbf{y}) \rightarrow \varphi(\mathbf{x})$, where the members of $\mathbf{x}$ are exactly the active domain of I. (Thus, for convenience, we are treating the constants in the active domain of I as if they were variables.) Then every member of $\mathbf{x}$ actually appears in both $\psi(\mathbf{x}, \mathbf{y})$ and in $\varphi(\mathbf{x})$.

Assume first that $(J, I_2) \models T(I)$; we must show that $(J, I_2) \models \omega(\Sigma, I)$. Let $\bar{\mathbf{x}}, \bar{\mathbf{y}}$ be an assignment of values to $\mathbf{x}, \mathbf{y}$ such that $\psi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ holds in J; we must show that $\varphi(\bar{\mathbf{x}})$ holds in I_2. Since by construction $Constant(x_i)$ appears in $\psi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ for each x_i in $\mathbf{x}$, and the inequalities $x_i \neq x_j$ appear in $\psi(\mathbf{x}, \mathbf{y})$ whenever x_i and x_j are distinct members of $\mathbf{x}$, we know that $\bar{x_i}$ is a constant for each x_i in $\mathbf{x}$, and $\bar{x_i}$ and $\bar{x_j}$ are distinct when x_i and x_j are distinct variables in $\mathbf{x}$. By construction, $\varphi(\mathbf{x})$ is a conjunction of the facts of I. Let I_1 be the instance whose facts are given by $\varphi(\bar{\mathbf{x}})$. Since $\bar{x_i}$ and $\bar{x_j}$ are distinct constants when x_i and x_j are distinct variables in $\mathbf{x}$, it follows that I_1 is a ground instance and that I is isomorphic to I_1 under the isomorphism that maps x_i into $\bar{x_i}$ for each i.

Since $\psi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ holds in J, it follows that J contains the result of chasing I_1 by Σ and then replacing the nulls in $\mathbf{y}$ by the values in $\bar{\mathbf{y}}$. So $(I_1, J) \models \Sigma$. Now $(J, I_2) \models T(I)$, and so $(J, I_2) \models \tau_{(=,=)}(\Sigma, I_1)$. Since also $(I_1, J) \models \Sigma$, it follows by definition of $\tau_{(=,=)}(\Sigma, I_1)$ that $I_1 \subseteq I_2$. Hence, $\varphi(\bar{\mathbf{x}})$ holds in I_2. This was to be shown.

Conversely, assume that $(J, I_2) \models \omega(\Sigma, I)$; we must show that $(J, I_2) \models T(I)$. Assume that I_1 is a ground instance isomorphic to I; we need only show that $(J, I_2) \models \tau_{(=,=)}(\Sigma, I_1)$. To prove this, assume that $(I_1, J) \models \Sigma$; we must show that $I_1 \subseteq I_2$. Now $\omega(\Sigma, I)$ is $\psi(\mathbf{x}, \mathbf{y}) \rightarrow \varphi(\mathbf{x})$, where the members of $\mathbf{x}$ are exactly the active domain of I. Define $\bar{\mathbf{x}}$ by having the mapping that maps x_i into $\bar{x_i}$ for each i be an isomorphism from I to I_1. Since $(I_1, J) \models \Sigma$, there is an assignment $\bar{\mathbf{y}}$ of values to $\mathbf{y}$ such that $\psi(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ holds in J. Since also $(J, I_2) \models \omega(\Sigma, I)$, that is, $(J, I_2) \models (\psi(\mathbf{x}, \mathbf{y}) \rightarrow \varphi(\mathbf{x}))$, it follows that I_2 satisfies $\varphi(\bar{\mathbf{x}})$. So $I_1 \subseteq I_2$, as desired.

Let Γ be the set of all formulas $\omega(\Sigma, I)$, one for every ground instance I. We have shown that $T(I)$ is weakly equivalent to $\omega(\Sigma, I)$, for every ground instance I. Therefore, $cand_{(=,=)}(\Sigma)$ (which is the set of all of the constraints $\tau_{(=,=)}(\Sigma, I)$, one for every ground instance I) is weakly equivalent to Γ. We wish to show that $cand_{(=,=)}(\Sigma)$ is weakly equivalent to $\Omega(\Sigma)$, which is a subset of Γ. So the proof is completed if we show that $\Omega(\Sigma)$ weakly implies Γ.

Let $\omega(\Sigma, I)$ be an arbitrary member of Γ. Let $I_1, \ldots, I_k$ be one-tuple instances where $I = I_1 \cup \cdots \cup I_k$. Let Z be $\{\omega(\Sigma, I_1), \ldots, \omega(\Sigma, I_k)\}$. We need only show that Z weakly implies $\omega(\Sigma, I)$ (this makes use of the fact that for each such one-tuple instance I_j, there is a prime instance I'_j such that $\omega(\Sigma, I_j)$ is logically equivalent to $\omega(\Sigma, I'_j)$).

It is easy to see that (up to a renaming of nulls), every conjunct that appears in the premise of some member of Z is a conjunct that appears in the premise of $\omega(\Sigma, I)$ (since the union of the results of chasing each individual tuple of I is a subset of the result of chasing I). Also, the set consisting of the (singleton) conclusions of members of Z is precisely the set of conjuncts that appear in the conclusion of $\omega(\Sigma, I)$. It follows immediately that Z weakly implies $\omega(\Sigma, I)$, as desired. This concludes the proof. □

We now show that Inverse($\mathcal{M}$) produces an inverse of $\mathcal{M}$ if $\mathcal{M}$ is invertible.

THEOREM 5.7. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping, where Σ is a finite set of s-t tgds. The following are equivalent.*

1. *$\mathcal{M}$ has an inverse.*
2. *The schema mapping that is the output of Inverse($\mathcal{M}$), which is specified by a finite set of full tgds with constants and inequalities among constants, is an inverse of $\mathcal{M}$.*

PROOF. The fact that (2) implies (1) is immediate. Assume now that (1) holds; we shall show that (2) holds. Write $cand_{(=,=)}(\Sigma)$ as Σ^c, and let $\mathcal{M}^c = (\mathbf{T}, \mathbf{S}, \Sigma^c)$. Write $\Omega(\Sigma)$, the set of full tgds with constants and inequalities that specifies the output of Inverse($\mathcal{M}$), as Σ^Ω. Let $\mathcal{M}^\Omega = (\mathbf{T}, \mathbf{S}, \Sigma^\Omega)$. Since (1) holds, it follows from Theorem 3.15 that $cand_{(=,=)}(\Sigma)$ specifies an inverse of $\mathcal{M}$. So $\text{Inst(Id)}[\sim_1, \sim_2] = \text{Inst}(\Sigma \circ \Sigma^c)[\sim_1, \sim_2]$. We need only show that $\text{Inst(Id)}[\sim_1, \sim_2] = \text{Inst}(\Sigma \circ \Sigma^\Omega)[\sim_1, \sim_2]$. Let I_1 and I_2 be ground instances. It is sufficient to show that $(I_1, I_2) \models \Sigma \circ \Sigma^c$ if and only if $(I_1, I_2) \models \Sigma \circ \Sigma^\Omega$. Now $(I_1, I_2) \models \Sigma \circ \Sigma^c$ if and only if there is J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma^c$. But Theorem 5.6 tells us that $(J, I_2) \models \Sigma^c$ if and only if $(J, I_2) \models \Sigma^\Omega$. □

Part (1) of Theorem 5.1 follows from Theorem 5.7. Part (3) of Theorem 5.1 follows from Theorem 4.12 (the necessity of constants) and Theorem 4.15 (the necessity of inequalities).

5.2 The Full Case

When Σ consists of *full* tgds, we obtain the following modified version of Theorem 5.1:

THEOREM 5.8. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of full s-t tgds. If $\mathcal{M}$ has an inverse then the following hold.*

1. *$\mathcal{M}$ has an inverse $\mathcal{M}'$ specified by a finite set of full tgds with inequalities.*
2. *There is an exponential-time algorithm for producing $\mathcal{M}'$.*
3. *Statement (1) is not necessarily true if we disallow inequalities, even if we allow existential quantifiers on the conclusion (and so allow nonfull dependencies to specify $\mathcal{M}'$).*

Part (1) follows from Part (1) of Theorem 5.1, along with Part (2b) of Proposition 4.8. In the full case, we can modify the algorithm Inverse by no longer including the formulas *Constant*(x). Part (3) follows from Theorem 4.15 (the necessity of inequalities).

6. QUASI-INVERSES IN DATA EXCHANGE

Next, we shall describe two desirable properties that an "inverse" should possess for data exchange. (Here, we use the term "inverse" loosely, to mean any schema mapping $\mathcal{M}'$ that goes in the reverse direction of $\mathcal{M}$.) In particular, we define the two notions of a *sound* "inverse" and of a *faithful* "inverse," which are relevant for data exchange, and then show how quasi-inverses relate to these notions.

First, it is desirable for an "inverse" to be *sound*. Specifically, assume $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ is a schema mapping where Σ is a finite set of s-t tgds, and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ is an "inverse" schema mapping. For now, assume that Σ' is given by a finite set of tgds. Suppose that we perform data exchange with $\mathcal{M}$, by chasing a ground instance I with Σ, to obtain a target instance U, denoted by $U = chase_\Sigma(I)$. We can then perform a reverse data exchange from U with $\mathcal{M}'$ and obtain V (i.e., compute $V = chase_{\Sigma'}(U)$). Then $\mathcal{M}'$ is *sound* with respect to $\mathcal{M}$ if the following holds for every choice of ground instance I: When we redo the original exchange with Σ but this time starting from V, we obtain a *subset* of the facts that are in U (modulo homomorphic images of nulls). Intuitively, the result of the reverse data exchange with $\mathcal{M}'$, followed by a data exchange with $\mathcal{M}$ (i.e., $chase_\Sigma(V)$), does not introduce any new information that cannot be found in U). If, in addition to $\mathcal{M}'$ being sound, U can be embedded homomorphically into $chase_\Sigma(V)$, then no information that is in U has been lost. We then say that $\mathcal{M}'$ is *faithful* with respect to $\mathcal{M}$.

Example 6.1. Let us revisit the earlier *Decomposition* example with a schema mapping $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ where Σ consists of the s-t tgd $P(x, y, z) \rightarrow Q(x, y) \wedge R(y, z)$. Let us recall, from Example 3.17, that $\mathcal{M}$ has quasi-inverses $\mathcal{M}'$ and $\mathcal{M}''$ specified by the following sets Σ' and Σ'' of tgds:

$$\Sigma' = \{Q(x, y) \wedge R(y, z) \rightarrow P(x, y, z)\}$$
$$\Sigma'' = \{Q(x, y) \rightarrow \exists z P(x, y, z), \quad R(y, z) \rightarrow \exists x P(x, y, z)\}$$

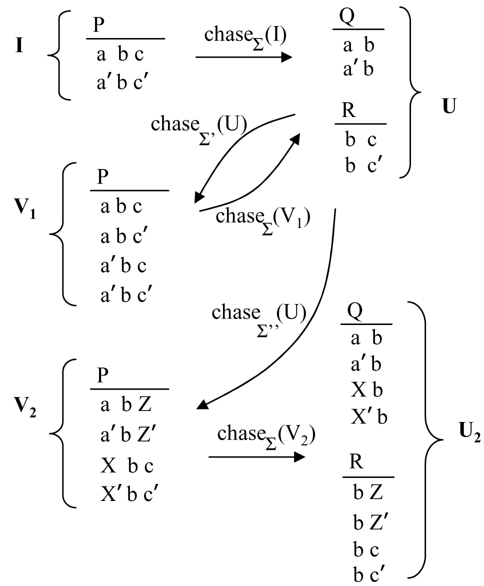

Fig. 1. $\mathcal{M}'$ and $\mathcal{M}''$ are faithful with respect to $\mathcal{M}$.

Let I be the ground instance shown in Figure 1. The result of chasing I with Σ (i.e., the result of the data exchange with $\mathcal{M}$) is the instance U shown in the figure. If we now chase U with Σ' (i.e., perform the reverse data exchange with $\mathcal{M}'$), we obtain the source instance V_1. Furthermore, if we now redo the original data exchange with $\mathcal{M}$ starting from V_1, the result is identical to U. In fact, it can be shown that, for every ground instance I, the result of redoing the original data exchange on V_1 is identical to U. Hence, $\mathcal{M}'$ is faithful with respect to $\mathcal{M}$.

Consider now $\mathcal{M}''$. Again, let U be the result of the first data exchange on I with $\mathcal{M}$. Let V_2 be obtained, as in the figure, by reverse data exchange with $\mathcal{M}''$ from U. If we now redo the original data exchange with $\mathcal{M}$ starting from V_2, the result is the instance U_2. Now U_2 is different from the target instance U because U_2 contains extra tuples with nulls. The instances U and U_2, however, are homomorphically equivalent. It can be shown that this is true for every ground instance I, and therefore $\mathcal{M}''$ is faithful with respect to $\mathcal{M}$.

It turns out that it is not an accident that $\mathcal{M}$ has faithful quasi-inverses. In this section, we show that if $\mathcal{M}$ is a schema mapping that is specified by a finite set of s-t tgds and has a quasi-inverse, then $\mathcal{M}$ is guaranteed to have a faithful quasi-inverse (and the algorithm QuasiInverse produces one). □

Note that nulls may arise when we chase I with a schema mapping $\mathcal{M}$, and also when we chase the result U with an "inverse" $\mathcal{M}'$. In particular, the result V of the reverse data exchange with $\mathcal{M}'$ may not necessarily be a ground instance, but rather a source instance with constants and nulls. For the purposes of "reverse" data exchange, we now expand the domain of values for source instances to be $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$. Thus, in addition to the ground instances, we shall also consider in this section nonground instances such as V_2 in Example 6.1. For the bidirectional data exchange scenario where $U = chase_{\Sigma}(I)$ and $V = chase_{\Sigma'}(U)$, we note that (I, U) satisfies the specification given by $\mathcal{M}$ and (U, V) satisfies the specification given by $\mathcal{M}'$. We point out, however, that (I, V) is outside of $\mathrm{Inst}(\mathcal{M} \circ \mathcal{M}')$ because V is not ground. Nevertheless, if $\mathcal{M}'$ is faithful, these nulls do not matter: when we redo the data exchange with $\mathcal{M}$ from V, we obtain a target instance that is homomorphically equivalent to the original result U.

In order to define soundness and faithfulness in the general case, when $\mathcal{M}'$ is expressed by a set of disjunctive tgds with constants and inequalities, we need to consider an extension of the chase for this more general language. The standard notion of the chase can be easily extended to handle the *Constant* predicate and the inequalities in the premise of the tgds in Σ'. However, when the conclusion of a tgd in Σ' contains disjunction, we need to use the *disjunctive chase*. Chasing with disjunctive dependencies has been considered before in various contexts [Deutsch and Tannen 2001; Fagin et al. 2005a]; we use a similar notion here, which we make precise via the following three definitions. When defining the disjunctive chase, we do not need to assume a separation into a source and a target schema. However, the subsequent definitions and results about soundness and faithfulness will apply the disjunctive chase in the context where such separation exists.

Definition 6.2. Let $\varphi(\mathbf{x})$ be a conjunction of formulas that may include: (1) atoms, such that every variable in $\mathbf{x}$ occurs in one of them, (2) formulas of the form $Constant(x)$, where x is a variable in $\mathbf{x}$, and, (3) inequalities of the form $x_1 \neq x_2$ where x_1 and x_2 are variables in $\mathbf{x}$. Let K be an instance over $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$. A *homomorphism h from $\varphi(\mathbf{x})$ to K* is a mapping from the variables $\mathbf{x}$ to values in $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$ such that: (1) for every atom $T(x_1, \ldots, x_k)$ in φ we have that $T(h(x_1), \ldots, h(x_k))$ is a fact in K, (2) for every formula $Constant(x)$ in φ, we have that $h(x)$ is in $\underline{\mathsf{Const}}$, and, (3) for every inequality $x_i \neq x_j$ in φ, we have that $h(x_i) \neq h(x_j)$.

Definition 6.3 (*Disjunctive Chase Step*). Let σ be a disjunctive tgd with constants and inequalities of the form:

$$\forall \mathbf{x}[\varphi(\mathbf{x}) \rightarrow (\exists \mathbf{y_1} \psi_1(\mathbf{x_1}, \mathbf{y_1}) \vee \ldots \vee \exists \mathbf{y_p} \psi_p(\mathbf{x_p}, \mathbf{y_p}))].$$

Let σ_i be the tgd with constants and inequalities that is obtained from σ by taking the ith disjunct $\forall \mathbf{x}[\varphi(\mathbf{x}) \rightarrow \exists \mathbf{y_i} \psi_i(\mathbf{x_i}, \mathbf{y_i})]$. Let K be an instance over $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$. Assume that h is a homomorphism from $\varphi(\mathbf{x})$ to K such that for each $i \in \{1, \ldots, p\}$, there is no extension of h to a homomorphism from $\varphi(\mathbf{x}) \wedge \psi_i(\mathbf{x_i}, \mathbf{y_i})$ to K. We say that *σ can be applied to K with homomorphism h*. Note that this also means that σ_i can be be applied to K with homomorphism h (this is the nondisjunctive definition of a chase step).

Let $K_1, \ldots, K_p$ be the instances that result by applying each of $\sigma_1, \ldots, \sigma_p$ to K with homomorphism h. We say that *the result of applying σ to K* is the set $\{K_1, \ldots, K_p\}$, and write $K \xrightarrow{\sigma,h} \{K_1, \ldots, K_p\}$.

Definition 6.4 (*Disjunctive Chase*). Let Σ be a finite set of disjunctive tgds with constants and inequalities. The *disjunctive chase of an instance K with* Σ is a tree (finite or infinite) that has K as a root and for each node K', if K' has children $K_1, \ldots, K_p$, then it must be the case that $K' \xrightarrow{\sigma,h} \{K_1, \ldots, K_p\}$ for some σ in Σ and some homomorphism h. Moreover, each leaf L in the tree has the requirement that there is no σ in Σ and no homomorphism h such that σ can be applied to L with h. When the chase tree is finite, we say that *the result of the disjunctive chase of K with* Σ is the set of leaves in the chase tree.

In the case when the disjunctive tgds are from a schema $\mathbf{T}$ to a schema $\mathbf{S}$, we can chase instances of the form (J, I) where J is a $\mathbf{T}$-instance and I is an $\mathbf{S}$-instance. Note that any such chase tree will be finite (since there is no recursion). Our case of interest is applying the disjunctive chase with Σ' to an instance of the form $(U, \emptyset)$ where $U = chase_\Sigma(I)$, for some ground instance I. The result of such chase is a set $\{(U, V_1), \ldots, (U, V_m)\}$ of instances where $V_1, \ldots, V_m$ are $\mathbf{S}$-instances. If $\mathcal{V}$ denotes the set $\{V_1, \ldots, V_m\}$, we shall also say that $\mathcal{V}$ is the result of chasing U with Σ' and write $\mathcal{V} = chase_{\Sigma'}(U)$. Furthermore, let us denote by $\mathcal{U}' = chase_\Sigma(\mathcal{V})$ the set of all instances U' that are obtained by chasing, in the standard way, each member V of $\mathcal{V}$ with Σ.

Definition 6.5. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping where Σ is a finite set of s-t tgds, and let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be a schema mapping where Σ' is a finite set of disjunctive tgds with constants and inequalities.

(1) We say that $\mathcal{M}'$ is *sound* with respect to $\mathcal{M}$ if:

for every ground instance I over $\mathbf{S}$, if $U = chase_\Sigma(I)$, $\mathcal{V} = chase_{\Sigma'}(U)$ and $\mathcal{U}' = chase_\Sigma(\mathcal{V})$, then there is a homomorphism from some member of $\mathcal{U}'$ into U.

(2) We say that $\mathcal{M}'$ is *faithful* with respect to $\mathcal{M}$ if:

for every ground instance I over $\mathbf{S}$, if $U = chase_\Sigma(I)$, $\mathcal{V} = chase_{\Sigma'}(U)$ and $\mathcal{U}' = chase_\Sigma(\mathcal{V})$, then there is some member of $\mathcal{U}'$ that is homomorphically equivalent to U.

Observe that by definition, if $\mathcal{M}'$ is faithful with respect to $\mathcal{M}$, then $\mathcal{M}'$ is sound with respect to $\mathcal{M}$. Also observe that when the dependencies in Σ' have no disjunction, the set $\mathcal{V}$ of source instances becomes a singleton set. Thus, if $\mathcal{M}'$ is faithful, chasing with Σ' recovers a single source instance whose chase (with Σ) is homomorphically equivalent to U. In fact, even when the dependencies in Σ' have disjunction, if $\mathcal{M}'$ is faithful, we can still recover a single source instance: we search among the instances in $\mathcal{V}$ to find the source instance whose chase (with Σ) is homomorphically equivalent to U.

We shall make use of the following lemma, which is an extension of Lemma 3.4 in Fagin et al. [2005a] to the case of disjunctive tgds with constants and inequalities among constants.

LEMMA 6.6. *Let σ be a disjunctive tgd with constants and inequalities among constants, and let K be an instance over* $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$. *Assume that σ is applicable* *to K with homomorphism h and let $K \xrightarrow{\sigma,h} \{K_1, \ldots, K_p\}$ be the corresponding disjunctive chase step.*

Let K^* be an instance over $\underline{\mathsf{Const}} \cup \underline{\mathsf{Var}}$ such that K^* satisfies σ. If there is a homomorphism g from K to K^*, then there is an instance K_m in $\{K_1, \ldots, K_p\}$ and an extension $\bar{g}$ of g such that $\bar{g}$ is a homomorphism from K_m to K^*.

PROOF. Let σ be $\forall \mathbf{x}[\varphi(\mathbf{x}) \rightarrow (\exists \mathbf{y_1} \psi_1(\mathbf{x_1}, \mathbf{y_1}) \vee \ldots \vee \exists \mathbf{y_p} \psi_p(\mathbf{x_p}, \mathbf{y_p}))]$. Since h is a homomorphism from $\varphi(\mathbf{x})$ to K, it must be the case that each atom $R(x_1, \ldots, x_k)$ of φ is embedded into a fact $R(h(x_1), \ldots, h(x_k))$ of K. Moreover, for every conjunct of the form $Constant(x)$ in φ, it must be the case that $h(x)$ is a constant in K. Furthermore, for every inequality $x_i \neq x_j$ in φ it must be the case that $h(x_i)$ and $h(x_j)$ are two distinct constants c_i and c_j in K. (We know that $Constant(x_i)$ and $Constant(x_j)$ must appear in φ, since the inequalities are among constants.)

We now show that $g \circ h$ is a homomorphism from $\varphi(\mathbf{x})$ into K^*. Indeed, each atom $R(x_1, \ldots, x_k)$ of φ is embedded into a fact $R(g(h(x_1)), \ldots, g(h(x_k)))$ of K^*. We used here the fact that g is a homomorphism from K to K^* and $R(h(x_1), \ldots, h(x_k))$ is a fact of K. Moreover, for every conjunct of the form $Constant(x)$ in φ, we have that $h(x)$ is a constant c in K, which implies $g \circ h(x) = g(h(x)) = g(c) = c$ (since g maps constants to themselves). Furthermore, for every inequality $x_i \neq x_j$ in φ, we know that $h(x_i)$ and $h(x_j)$ are two distinct constants c_i and c_j in K, which implies that $g \circ h(x_i) = g(h(x_i)) = g(c_i) = c_i$ and $g \circ h(x_j) = g(h(x_j)) = g(c_j) = c_j$ and therefore $g \circ h(x_i) \neq g \circ h(x_j)$. Again we used that g maps constants to themselves.[3]

Since $g \circ h$ is a homomorphism from $\varphi(\mathbf{x})$ to K^*, and K^* satisfies σ, there must be some $m \in \{1, \ldots, p\}$ and some values $b_1, \ldots, b_l$ playing the role of the variables $\mathbf{y}$ such that $K^* \models \psi_m(g \circ h(\mathbf{x}), b_1, \ldots, b_l)$. Let K_m be the instance that is obtained from K, in the chase step with σ and h, by using the mth disjunct of σ. Thus, K_m is obtained by adding the facts $\psi_m(h(\mathbf{x}), n_1, \ldots, n_l)$ to K, where $n_1, \ldots, n_l$ are fresh nulls chosen to instantiate the existentially quantified variables $\mathbf{y}$. Define $\bar{g}$ so that $\bar{g}(v) = g(v)$ for every v in $d$om(g) and $\bar{g}(n_i) = b_i$, for every $1 \leq i \leq l$. Clearly, $\bar{g}(c) = c$ for every constant c. Moreover, $\bar{g}$ maps the new facts in $\psi_m(h(\mathbf{x}), n_1, \ldots, n_l)$ to facts in $\psi_m(g \circ h(\mathbf{x}), b_1, \ldots, b_l)$, which in turn are facts of K^*. Thus, $\bar{g}$ is a homomorphism from K_m to K^*. □

The following proposition states an important property of the disjunctive chase in the context of bidirectional data exchange, when the disjunctive tgds (with constants and inequalities among constants) are part of the "reverse" mapping.

PROPOSITION 6.7. [UNIVERSALITY OF "CHASE OF THE CHASE"] *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping where Σ is a finite set of s-t tgds and let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be a schema mapping where Σ' is a finite set of disjunctive tgds with constants and inequalities among constants. Moreover, let I be a ground instance over $\mathbf{S}$. If $U = chase_{\Sigma}(I)$ and $\mathcal{V} = chase_{\Sigma'}(U)$, then for every ground instance K such*

[3]Note that g in general can collapse nonconstant elements. Thus, it is essential in the above argument that the inequalities involve variables that appear inside the *Constant* predicate.

that $(I, K) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$*, there is* $V \in \mathcal{V}$ *such that there is a homomorphism from* V *to* K*.*

Proof. Let K be a ground instance such that $(I, K) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. This implies that there is J over $\mathbf{T}$ such that $(I, J) \models \Sigma$ and $(J, K) \models \Sigma'$. Let $\mathcal{K} = \mathit{chase}_{\Sigma'}(J)$. The proof proceeds in two parts. First, we show that there is an instance K' in $\mathcal{K}$ such that there is a homomorphism from K' to K. Then we show that there is an instance $V \in \mathcal{V}$ such that there is a homomorphism from V to K'. It follows that there is a homomorphism from V to K.

To prove the first part, let g be the identity mapping on the set of values of J unioned with the set Const of all the constants. Clearly, g is a homomorphism from $(J, \emptyset)$ to (J, K). We know that (J, K) satisfies all the dependencies in Σ'. Thus, by repeatedly applying Lemma 6.6 for the successive steps in the chase with Σ' from $(J, \emptyset)$, we find an instance $K' \in \mathcal{K}$ and an extension $\bar{g}$ of g such that $\bar{g}$ is a homomorphism from (J, K') to (J, K). In particular, $\bar{g}$ is a homomorphism from K' to K.

For the second part, note that U is a universal solution [Fagin et al. 2005a] for I and Σ. Since J is a solution for I and Σ, it follows that there is a homomorphism h from U to J. Hence, h is a homomorphism from $(U, \emptyset)$ to (J, K') where K' is the instance constructed in the first part of the proof. By the properties of the disjunctive chase, we have that (J, K') satisfies Σ', since K' is one of the instances that result after chasing J with Σ'. Then we can apply the same Lemma 6.6 to find an instance $V \in \mathcal{V}$ and an extension $\bar{h}$ of h such that $\bar{h}$ is a homomorphism from (U, V) to (J, K'). In particular, $\bar{h}$ is a homomorphism from V to K'. □

The next theorem shows that every quasi-inverse specified by disjunctive tgds with constants and inequalities among constants is sound. We have shown earlier that this language is sufficient to express quasi-inverses of schema mappings that are specified by s-t tgds.

Theorem 6.8. *Let* $\mathcal{M}$ *be a schema mapping specified by a finite set of s-t tgds. If* $\mathcal{M}'$ *is a quasi-inverse of* $\mathcal{M}$ *that is specified by a finite set of disjunctive tgds with constants and inequalities among constants, then* $\mathcal{M}'$ *is sound with respect to* $\mathcal{M}$*.*

Proof. Let $\mathcal{M}$ be $(\mathbf{S}, \mathbf{T}, \Sigma)$ and let $\mathcal{M}'$ be $(\mathbf{T}, \mathbf{S}, \Sigma')$. Assume that I is a ground instance over S. Let $U = \mathit{chase}_{\Sigma}(I)$, $\mathcal{V} = \mathit{chase}_{\Sigma'}(U)$, and $\mathcal{U}' = \mathit{chase}_{\Sigma}(\mathcal{V})$. We need to show that there is a homomorphism from a member of $\mathcal{U}'$ into U.

Since $(I, I) \in \text{Inst}(\text{Id})[\sim, \sim]$ and $\text{Inst}(\text{Id})[\sim, \sim] \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim, \sim]$, it follows that $(I, I) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')[\sim, \sim]$. This means that there exist I_1, J_0 and K such that I_1 and K are ground, $I \sim I_1$, $I \sim K$, $(I_1, J_0) \models \Sigma$ and $(J_0, K) \models \Sigma'$. Because $(I_1, J_0) \models \Sigma$ and $I \sim I_1$, we have that $(I, J_0) \models \Sigma$. This, together with $(J_0, K) \models \Sigma'$, implies that $(I, K) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. By Proposition 6.7, this implies that there exists $V \in \mathcal{V}$ and a homomorphism $h : V \to K$.

Now, let U' be the result of chasing V with Σ. In more precise terms, this means that (V, U') is the result of chasing $(V, \emptyset)$ with Σ. Since $V \subseteq \mathcal{V}$, it follows that U' is a member of $\mathcal{U}' = \mathit{chase}_{\Sigma}(\mathcal{V})$. Let J denote the result of chasing K with Σ. Since h is a homomorphism from $(V, \emptyset)$ to (K, J), and $(K, J) \models \Sigma$, it follows from Lemma 3.4 in Fagin et al. [2005a] that there is a homomorphism $\bar{h}$, which is an extension of h, from (V, U') into (K, J). Therefore, there is a homomorphism $\bar{h} : U' \to J$. Furthermore, by definition, $I \sim K$ means that $\text{Sol}(\mathcal{M}, I) = \text{Sol}(\mathcal{M}, K)$. By Proposition 2.6 of Fagin et al. [2005a], this last equality is equivalent to the fact that every universal solution for K with respect to $\mathcal{M}$ is homomorphically equivalent to every universal solution for I with respect to $\mathcal{M}$. Since $J = \mathit{chase}_\Sigma(K)$ and $U = \mathit{chase}_\Sigma(I)$ are universal solutions for K and I, respectively, we obtain that J and U are homomorphically equivalent. In particular, there is a homomorphism $g : J \to U$. Hence, $g \circ \bar{h}$ is a homomorphism from U' to U. □

In the next two theorems, we give examples that show: (1) the restriction of allowing inequalities *only among constants* in the quasi-inverse is necessary for Theorem 6.8 to hold, and (2) we cannot strengthen Theorem 6.8 to say that a quasi-inverse specified by a finite set of disjunctive tgds with constants and inequalities among constants is *faithful* (in addition to being sound). Hence, Theorem 6.8 is optimal in two senses. Furthermore, the two examples we give are strong in that they involve not just quasi-inverses but inverses.

THEOREM 6.9. *There exist a schema mapping $\mathcal{M}$ specified by a finite set of s-t tgds and an inverse $\mathcal{M}'$ of $\mathcal{M}$ that is specified by a finite set of tgds with inequalities (not among constants) such that $\mathcal{M}'$ is not sound with respect to $\mathcal{M}$.*

PROOF. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ where Σ and Σ' are defined as follows:

$$\Sigma: \; S_1(x) \to \exists y (T(x, y) \wedge T(y, x)) \qquad S_2(x) \to S_2'(x)$$

$$\Sigma': \; T(x, y) \wedge T(y, x) \to S_1(x) \qquad S_2'(x) \to S_2(x) \qquad T(x, y) \wedge T(y, x) \wedge y \neq x \to S_2(x)$$

The third tgd in Σ' includes an inequality that is not among constants. We first show that $\mathcal{M}'$ is an inverse of $\mathcal{M}$. To do this, we show that $(I, I') \models \Sigma \circ \Sigma'$ if and only if $I \subseteq I'$.

First, it is easy to see that if $(I, I') \models \Sigma \circ \Sigma'$, then $I \subseteq I'$, because the tgds in Σ together with the first two tgds in Σ' assert that a copy of I must exist in I'. We now show that $(I, I) \models \Sigma \circ \Sigma'$, which in turn implies that if $I \subseteq I'$, then $(I, I') \models \Sigma \circ \Sigma'$. Let I be an arbitrary ground instance with m facts in S_1 and n facts in S_2. That is, I consists of $S_1(a_i)$ where $1 \leq i \leq m$, and $S_2(b_j)$ where $1 \leq j \leq n$. Let J consist of $T(a_i, a_i)$ where $1 \leq i \leq m$ and $S_2'(b_j)$ where $1 \leq j \leq n$. It is easy to see that $(I, J) \models \Sigma$ and $(J, I) \models \Sigma'$ (in particular, the third tgd in Σ' is vacuously satisfied). Hence, $(I, I) \models \Sigma \circ \Sigma'$.

We now show that $\mathcal{M}'$ is not sound. Let I consist of the fact $S_1(0)$. Let $U = \mathit{chase}_\Sigma(I)$, $V = \mathit{chase}_{\Sigma'}(U)$, and $U' = \mathit{chase}_\Sigma(V)$. The instance U consists of two facts $T(0, X)$ and $T(X, 0)$ where X is a null introduced by the chase. Note that S_2' is empty in U. Furthermore, the instance V consists of two facts $S_1(0)$ and $S_2(0)$. (Note that $S_2(0)$ is the result of applying the third tgd in Σ', since $X \neq 0$.) Finally, the instance U' consists of three facts: $T(0, Y)$, $T(Y, 0)$, and $S_2'(0)$, where Y is a null introduced by the chase. Since S_2' is empty in U but not in U', there is no homomorphism from U' to U. □

THEOREM 6.10. *There exist a schema mapping $\mathcal{M}$ specified by a finite set of s-t tgds and an inverse $\mathcal{M}'$ of $\mathcal{M}$ specified by a finite set of tgds with constants such that $\mathcal{M}'$ is sound but not faithful with respect to $\mathcal{M}$.*

PROOF. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ and $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ where Σ and Σ' are defined below:

$$\Sigma:\ P(x) \rightarrow \exists y\, Q(x, y) \qquad \Sigma':\ Q(x, y) \wedge Constant(y) \rightarrow P(x)$$
$$Q(x, y) \rightarrow P(y)$$

Observe that Σ' consists of a finite set of tgds with constants. We first show that $\mathcal{M}'$ is a inverse of $\mathcal{M}$. To do this, we need to show that $(I, I') \models \Sigma \circ \Sigma'$ if and only if $I \subseteq I'$.

First, let I be a ground instance that consists of n facts $P(x_1), \ldots, P(x_n)$ and let J be $\{Q(x_i, x_i) \mid 1 \leq i \leq n\}$. It is easy to see that $(I, J) \models \Sigma$ and $(J, I) \models \Sigma'$. Hence, $(I, I) \models \Sigma \circ \Sigma'$ which implies that if $I \subseteq I'$, then $(I, I') \models \Sigma \circ \Sigma'$. Next, assume that I and I' are ground instances such that $(I, I') \models \Sigma \circ \Sigma'$; we shall show that $I \subseteq I'$. Since $(I, I') \models \Sigma \circ \Sigma'$, there is J such that $(I, J) \models \Sigma$ and $(J, I') \models \Sigma'$. Suppose I consists of n facts $P(x_1), \ldots, P(x_n)$. Since $(I, J) \models \Sigma$, we know that J contains $\{Q(x_i, y_i) \mid 1 \leq i \leq n\}$, for some choices of $y_1, \ldots, y_n$. There are two cases:

- *Case* 1. Some y_i is not a constant. Then I' contains $P(y_i)$, and so is not ground. Hence, this case is not possible.
- *Case* 2. Every y_i is a constant. Then I' contains $P(x_i)$, $1 \leq i \leq n$, and so $I \subseteq I'$, as desired.

We now show $\mathcal{M}'$ is sound but not faithful with respect to $\mathcal{M}$. Since $\mathcal{M}'$ is specified by a finite set of tgds with constants, $\mathcal{M}'$ is sound by Theorem 6.8. We now show $\mathcal{M}'$ is not faithful. Let I be a ground instance consisting of a single fact $P(0)$. Then the chase of I with Σ is $U = \{Q(0, Y)\}$ for a null Y, the chase of U with Σ' is $V = \{P(Y)\}$, and the chase of V with Σ is $U' = \{Q(Y, Y')\}$ for another null Y'. Clearly, there is no homomorphism from $\{Q(0, Y)\}$ to $\{Q(Y, Y')\}$. Hence, $\mathcal{M}'$ is not faithful. □

In the following lemma, we prove that if $\mathcal{M}'$ is the result of applying the algorithm QuasiInverse on $\mathcal{M}$, and I, U and $\mathcal{U}'$ are as in Definition 6.5, then there exists a homomorphism from U into every member U' of $\mathcal{U}'$.

LEMMA 6.11. *Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ be a schema mapping where Σ is a finite set of s-t tgds and let $\mathcal{M}' = (\mathbf{T}, \mathbf{S}, \Sigma')$ be the schema mapping returned by applying the algorithm* QuasiInverse *on $\mathcal{M}$. Assume that I is a ground instance over* $\mathbf{S}$. *Let $U = chase_\Sigma(I)$, $\mathcal{V} = chase_{\Sigma'}(U)$ and $\mathcal{U}' = chase_\Sigma(\mathcal{V})$. Then there is a homomorphism from U into every member U' of $\mathcal{U}$.*

PROOF. Let U' be a member of $\mathcal{U}'$, that is, $U' = chase_\Sigma(V)$ for some V in $\mathcal{V}$. We will show that there is a homomorphism from U to U', by adapting an argument in the proof of Theorem 4.7. In the second direction of that proof, we have shown the following fact: if I_1 and I_2 are ground instances such that there is a target instance J such that $(I_1, J) \models \Sigma$ and $(J, I_2) \models \Sigma'$ (where Σ' is the set constructed from Σ in the QuasiInverse algorithm), then $\text{Sol}(\mathcal{M}, I_2) \subseteq \text{Sol}(\mathcal{M}, I_1)$.

The main observation is that the proof of the preceding fact continues to hold even when I_2 is non-ground (the proof did not use the assumption that I_2 is ground). We apply this to our situation as follows. We have that $(I, U) \models \Sigma$ and $(U, V) \models \Sigma'$, where I is ground. Then, by the previous argument where I plays the role of I_1 and V plays the role of I_2, we have that $\text{Sol}(\mathcal{M}, V) \subseteq \text{Sol}(\mathcal{M}, I)$. Since U' is a solution for $\mathcal{M}$ and V, we thus obtain that U' is a solution for $\mathcal{M}$ and I. Since U is a universal solution for $\mathcal{M}$ and I, it follows that there is a homomorphism from U to U'. □

The final theorem says that the schema mapping obtained by applying the QuasiInverse algorithm is always faithful. This implies that, in particular, when quasi-inverses exist, the schema mapping returned by the QuasiInverse algorithm is a faithful quasi-inverse.

THEOREM 6.12. *Let $\mathcal{M}$ be a schema mapping specified by a finite set of s-t tgds. Let $\mathcal{M}'$ be the schema mapping returned by applying the algorithm* Quasi-Inverse *on $\mathcal{M}$. Then $\mathcal{M}'$ is faithful with respect to $\mathcal{M}$. In particular, if $\mathcal{M}$ has a quasi-inverse, then $\mathcal{M}'$ is a faithful quasi-inverse of $\mathcal{M}$.*

PROOF. Let $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma)$ and $\mathcal{M}' = (\mathbf{S}, \mathbf{T}, \Sigma')$, where Σ is a finite set of s-t tgds and Σ' is the finite set of disjunctive tgds with constants and inequalities among constants that is constructed by the QuasiInverse algorithm. Let $U = chase_{\Sigma}(I)$, $\mathcal{V} = chase_{\Sigma'}(U)$ and $\mathcal{U}' = chase_{\Sigma}(\mathcal{V})$. We shall show that there is a a member U' of $\mathcal{U}'$ that is homomorphically equivalent to U, by showing that there is a member U' of $\mathcal{U}'$ such that there is a homomorphism from U' to U. The existence of a reverse homomorphism from U to U', which follows from Lemma 6.11, completes the proof of the theorem.

We now show that there is a member U' of $\mathcal{U}'$ such that there is a homomorphism from U' to U. As shown in the proof of Theorem 4.7, the schema mapping $\mathcal{M}'$ produced by the QuasiInverse algorithm satisfies $\text{Inst(Id)} \subseteq \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. Since $(I, I) \in \text{Inst(Id)}$, it follows that $(I, I) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$. By the "universality of the chase of the chase" (Proposition 6.7), the following holds for every ground instance I over $\mathbf{S}$: For every ground instance K such that $(I, K) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, there is $V \in \mathcal{V}$ such that there is a homomorphism from V to K. Since $(I, I) \in \text{Inst}(\mathcal{M} \circ \mathcal{M}')$, there is $V \in \mathcal{V}$ such that there is a homomorphism h from V to I. So h is a homomorphism from $(V, \emptyset)$ to (I, U) and $(I, U) \models \Sigma$. Let $U' = chase_{\Sigma}(V)$. Clearly, $U' \in \mathcal{U}'$. By repeatedly applying Lemma 3.4 in [Fagin et al. 2005a] for the successive steps in the chase with Σ from $(V, \emptyset)$, we find an extension $\bar{h}$ of h such that $\bar{h}$ is a homomorphism from (V, U') to (I, U). In particular, $\bar{h}$ is a homomorphism from U' to U. □

It is interesting to note that the QuasiInverse algorithm, which was designed to produce quasi-inverses if they exist, turns out to have the desirable property that it always produces faithful mappings. The mappings produced by the algorithm are faithful even when the original schema mapping is not quasi-invertible. This also means, in particular, that faithful mappings (of schema mappings given by finite sets of s-t tgds) always exist, even when quasi-inverses may not exist.

7. CONCLUDING REMARKS

The notion of an inverse of a schema mapping is rather restrictive, since it is rare that a schema mapping has an inverse. We therefore introduced and studied a more relaxed notion of a quasi-inverse of a schema mapping. Both inverses and quasi-inverses are special cases of a unifying framework that we developed. We gave exact criteria for the existence of quasi-inverses and inverses, complete characterizations of the languages needed to express quasi-inverses and inverses, and results regarding the use of quasi-inverses in data exchange.

Some of the important remaining problems are decision and complexity issues. We have shown that for LAV schema mappings, a quasi-inverse always exists. However, the complexity of the decision problem for the existence of a quasi-inverse of a schema mapping specified by a finite set of s-t tgds (even in the full case) remains open. We do not know whether the problem is even decidable. Another open problem concerns the optimality of the algorithms QuasiInverse and Inverse. Given a schema mapping specified by a finite set of s-t tgds, these algorithms produce a schema mapping that is exponential in the size of the input schema mapping. We do not know whether the size of a quasi-inverse is necessarily exponential. If it turns out that there is always a polynomial-size quasi-inverse, this raises the question of finding a polynomial-time algorithm that can produce it. Similarly, the same question arises for inverses. Finally, the notion of faithfulness deserves further exploration. This is especially true since for each schema mapping $\mathcal{M}$ specified by a finite set of s-t tgds, there is always a schema mapping $\mathcal{M}'$ that is faithful with respect to $\mathcal{M}$.

REFERENCES

Abiteboul, S., Hull, R., and Vianu, V. 1995. *Foundations of Databases*. Addison-Wesley.

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

Deutsch, A. and Tannen, V. 2001. Optimization properties for classes of conjunctive regular path queries. In *Proceedings of the International Workshop on Database Programming Languages (DBPL)*. 21–39.

Fagin, R. 2007. Inverting schema mappings. *ACM Trans. Data. Syst. 32,* 4.

Fagin, R., Kolaitis, P. G., Miller, R. J., and Popa, L. 2005a. Data exchange: Semantics and query answering. *Theor. Comput. Science 336,* 1, 89–124.

Fagin, R., Kolaitis, P. G., and Popa, L. 2005b. Data exchange: Getting to the core. *ACM Trans. Data. Syst. 30,* 1, 174–210.

Fagin, R., Kolaitis, P. G., Popa, L., and Tan, W.-C. 2004. Composing schema mappings: Second-order dependencies to the rescue. In *Proceedings of the ACM Symposium on Principles of Database Systems (PODS)*. 83–94.

Fagin, R., Kolaitis, P. G., Popa, L., and Tan, W.-C. 2005c. Composing schema mappings: Second-order dependencies to the rescue. *ACM Trans. Data. Syst. 30,* 4, 994–1055.

Fagin, R., Kolaitis, P. G., Popa, L., and Tan, W.-C. 2007. Quasi-inverses of schema mappings. In *Proceedings of the ACM Symposium on Principles of Database Systems (PODS)*. 123–132.

Fagin, R. and Nash, A. 2007. The structure of inverses in schema mappings. IBM Research Report RJ10425.

Kolaitis, P. G. 2005. Schema mappings, data exchange, and metadata management. In *Proceedings of the ACM Symposium on Principles of Database Systems (PODS)*. 61–75.

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

MADHAVAN, J. AND HALEVY, A. Y. 2003. Composing mappings among data sources. In *Proceedings of the International Conference on Very Large Data Bases (VLDB)*. 572–583.

MELNIK, S. 2004. Generic model management: concepts and algorithms. Lecture Notes in Computer Science, vol. 2967. Springer.

NASH, A., BERNSTEIN, P. A., AND MELNIK, S. 2005. Composition of mappings given by embedded dependencies. In *Proceedings of the ACM Symposium on Principles of Database Systems (PODS)*. 172–183.

Received October 2007; accepted February 2008