

On the Data Complexity of Consistent Query Answering

Balder ten Cate · Gaëlle Fontaine ·
Phokion G. Kolaitis

Published online: 4 November 2014
© Springer Science+Business Media New York 2014

Abstract The framework of database repairs is a principled approach to managing inconsistency in databases. In particular, the consistent answers of a query on an inconsistent database provide sound semantics and the guarantee that the values obtained are those returned by the query on every repair of the given inconsistent database. In this paper, we carry out a systematic investigation of the data complexity of the consistent answers of conjunctive queries for set-based repairs and with respect to classes of constraints that, in recent years, have been extensively studied in the context of data exchange and data integration. Our results, which range from polynomial-time computability to undecidability, complement or improve on earlier work, and provide a fairly comprehensive picture of the data complexity of consistent query answering. We also address the problem of finding a “representative” or “useful” repair of an inconsistent database. To this effect, we introduce the notion of a universal repair, as well as relaxations of it, and then apply it to the investigation of the data complexity of consistent query answering.

Keywords Inconsistent databases · Constraints · Repairs · Consistent answers

B. ten Cate
LogicBlox, Inc., Atlanta, GA, USA

G. Fontaine (✉)
University of Chile, Santiago, Chile
e-mail: gaelle@dcc.uchile.cl

P. G. Kolaitis · B. ten Cate
University of California, Santa Cruz, CA, USA

P. G. Kolaitis
IBM Research - Almaden, San Jose, CA, USA
e-mail: kolaitis@cs.ucsc.edu

1 Introduction

An inconsistent database is a database that fails to satisfy one or more integrity constraints that the data are hand are supposed to obey. Inconsistency in databases arises in a variety of applications, including data integration and data warehousing, where the task is to bring together data distributed over different sources that may obey mutually incompatible constraints. In practice, inconsistency is handled mainly via data cleaning, which means that the inconsistent database is transformed, through deletions or additions, to a consistent one that is then used for query answering or for warehousing purposes. This process, however, forces arbitrary choices to be made since, in general, there is a multitude of ways in which an inconsistent databases can be transformed to a consistent one. Arenas, Bertossi and Chomicki [4] introduced a principled approach to the management of inconsistency by formulating the notions of a repair of an inconsistent database and of consistent query answering. Intuitively, a *repair* of an inconsistent database I is a consistent database J that differs from I in a “minimal” way. Furthermore, the *consistent answers* of a query q on an inconsistent database I are defined to be the intersection $\bigcap \{q(J) : J \text{ is a repair of } I\}$. Thus, the inconsistencies in the database are kept, but are handled at query time by considering all repairs and returning the tuples that are guaranteed to be in the result of the query on every repair.

Two algorithmic problems concerning repairs of inconsistent databases naturally arise. The first is, of course, the problem of computing the consistent answers of a query over an inconsistent database. The second is the *repair checking* problem, which can be thought of as the model-checking problem for repairs: given two database instances I and J , is J a repair of I ? Since the publication of [4] in 1999, these two problems have been extensively explored for different types of repairs (set-based repairs, cardinality-based repairs, attribute-based repairs) and for different types of constraints. As regards types of constraints, the earlier work on repair checking and consistent query answering focused on functional dependencies, inclusion dependencies, and denial constraints (see the overviews [7, 13]). More recently, broader classes of constraints, such as tuple-generating dependencies (tgds) and equality-generating dependencies (egds), have also been considered in the study of repair checking and consistent query answering. As is well known, these classes of constraints were originally investigated in the context of classical dependency theory, but in the past decade have found numerous uses in the context of data integration and data exchange. For some of these broader classes of constraints, the repair-checking problem has been studied in [2] and [24], and the consistent query answering problem in [32] and [3].

In this paper, we systematically explore the *data complexity* of consistent query answering for sets of tgds and egds. This means that for every fixed set Σ of tgds and egds and for every fixed conjunctive query q , we consider the complexity of the following algorithmic problem: given a database instance I , compute the consistent answers of q on I w.r.t. Σ . Concerning the types of repairs, we consider set-based repairs, that is, subset-repairs, superset-repairs, and \oplus -repairs (symmetric difference repairs).

Our main results about the data complexity of consistent query answering, together with previously known results, are summarized in Table 1. In this table, by an entry such as “C/C-comp.” we mean that for every fixed set of dependencies in the class considered and for every fixed conjunctive query, the problem is in C and that there is a set of dependencies in the class and a conjunctive query for which the problem is C-complete.

One finding of our investigation is that the class of LAV (local-as-view) tgds exhibits tractable algorithmic behavior as regards the data complexity of consistent query answering with respect to all three types of set-based repairs considered here; this extends earlier results about inclusion dependencies and subset-repairs [14]. Another finding is that there is a set (in fact, a singleton set) of GAV (global-as-view) tgds for which the data complexity of the consistent query answering with respect to both subset-repairs and \oplus -repairs is conP-complete; this strengthens a conP-completeness result for GAV tgds and functional dependencies with respect to \oplus -constraints in [32]. We also show that there are sets of weakly acyclic sets of tgds for which the data complexity of the consistent query answering problem is Π_2^P -complete with respect to both subset-repairs and \oplus -repairs; earlier, Π_2^P -completeness results had been obtained for the data complexity of consistent query answering for sets of functional dependencies and universal constraints [32] with respect to \oplus -repairs. Finally, we show that the assumption of weak acyclicity is of the essence for the decidability of the consistent query answering problem. Specifically, we show that there is a fixed set of tgds and a fixed conjunctive query for which the consistent query answering problem is undecidable with respect to \oplus -repairs; furthermore, a similar undecidability result holds for superset-repairs. Previously, it was known that the consistent query answering problem was undecidable in combined complexity for conjunctive queries and for sets of inclusion dependencies and functional dependencies with respect to superset-repairs [31]. It was also known that the consistent query answering problem was undecidable in combined complexity for unions of conjunctive queries and for sets of inclusion dependencies and functional dependencies with respect to \oplus -repairs [11]. Finally, it was known that there is a fixed set of universal constraints and a fixed universal query for which consistent query answering is undecidable with respect to \oplus -repairs [3].

In addition to the data complexity of consistent query answering, we also addressed the following question: For which types of dependencies is it the case that, given a database instance, there is an efficient way to compute a “representative” and “useful” repair?

We formalized and answered this question by introducing the notion of a *universal repair* and the notion of an n -universal repair, where n is a positive integer. These notions are analogous to and, in fact, are motivated from the notion of a universal solution in data exchange [17]. Furthermore, the notion of universal repair is closely related to the notion of *nucleus* in [35], since a universal repair is nucleus that is also a repair. Informally, a universal repair is a repair such that the consistent answers of an arbitrary conjunctive query can be computed by essentially evaluating the query on the universal repair. Similarly, an n -universal repair has the same property but only for conjunctive queries with at most n atoms. We study the existence of universal

Table 1 Data complexity of the repair checking problem and the consistent query answering problem for conjunctive queries

Dependencies	★-repair checking, where ★ ∈ {⊕, subset, superset}	Superset-CQA	Subset-CQA	⊕-CQA
LAV tgds	PTIME [†]	PTIME [†]	PTIME [†]	PTIME [†]
GAV tgds [‡]	PTIME [32]	PTIME [17] (implicit)	coNP[32]/coNP-comp. [†]	coNP[32]/coNP-comp. [†]
Weakly acyclic tgds [‡]	coNP/coNP-comp. [2, 24]	PTIME [17] (implicit)	Π_2^P / Π_2^P -comp. [†]	Π_2^P / Π_2^P -comp. [†]
Arbitrary tgds [‡]	coNP/coNP-comp. [2, 24]	undecidable [28] (implicit)	Π_2^P / Π_2^P -comp. [†]	undecidable [†]

[‡] Indicates that all results hold also in the presence of egds

[†] Indicates new result obtained in this paper

repairs and of n -universal repairs, as well as structural properties of such repairs and the complexity of computing them. We show that, if Σ is a set of LAV tgds, then every database instance has a unique universal \oplus -repair that is also a universal subset-repair and can be computed in polynomial time. Furthermore, while not every database instance has a universal superset-repair, we show that for every $n \geq 1$, an n -universal superset-repair exists and can be computed in polynomial time. If Σ is a weakly acyclic set of tgds and egds, things are the other way around: every database instance has a universal superset-repair that can be computed in polynomial time, but not every instance has a universal, or even a 1-universal, \oplus -repair; furthermore, similar limitations hold true for subset-repairs.

An extended abstract of the present paper has appeared earlier in [33]. The present paper augments it by providing detailed proofs and new material. In particular, Proposition 3.8 and Theorem 4.1 are new results obtained since the publication of [33].

2 Basic Notions

A *schema* \mathbf{R} is a finite sequence (R_1, \dots, R_k) of relation symbols, each of a fixed arity. An *instance* I over \mathbf{R} is a sequence (R_1^I, \dots, R_k^I) , where each R_i^I is a relation of the same arity as R_i . For notational simplicity, we shall write R_i to denote both the relation symbol and the relation R_i^I that interprets it. A *fact* of an instance I (over \mathbf{R}) is an expression $R_i(v_1, \dots, v_m)$, where R_i is one of the relations of I and v_1, \dots, v_m are values such that $(v_1, \dots, v_m) \in R_i$. Every instance can be identified with the set of its facts. The *size* of an instance is the number of facts it has. The *active domain* of an instance I , denoted by $\text{adom}(I)$, is the set of all values occurring in the facts of I . We assume that all instances I considered are finite, which means that every relation R_i of I is finite, for $1 \leq i \leq k$.

Definition 2.1 (Dependencies) A *tuple-generating dependency* (tgd) is a first-order sentence of the form

$$\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})),$$

where ϕ, ψ are conjunctions of atomic formulas, $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_m)$ are tuples of variables, and every universally quantified variable x_i occurs in ϕ .

An *equality-generating dependency* (egd) is a first-order sentence of the form

$$\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow x_k = x_\ell),$$

where ϕ is a conjunction of atomic formulas, $\mathbf{x} = (x_1, \dots, x_n)$ is a tuple of variables, $1 \leq k, \ell \leq n$, and each universally quantified variable x_i occurs in ϕ .

In what follows, by the term *dependency*, we will mean a tgd or an egd. Also, by a *set of dependencies*, we will mean a finite set of dependencies.

Definition 2.2 A *local-as-view* or, simply, *LAV* tgd is a tgd $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$ in which ϕ is a single atomic formula.

A *global-as-view* or, simply, *GAV* tgd is a tgd $\forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \psi(\mathbf{x}'))$ in which ψ is a single atomic formula such that the variables in \mathbf{x}' are among the variables of \mathbf{x} .

For example, the *copy* tgd $\forall x \forall y(E(x, y) \rightarrow F(x, y))$ is both a LAV tgd and a GAV tgd. In contrast,

$$\forall x \forall y(E(x, y) \rightarrow \exists z(F(x, z) \wedge F(z, y)))$$

is a LAV tgd that is not a GAV tgd, while

$$\forall x \forall y \forall z(E(x, z) \wedge E(z, y) \rightarrow F(x, y))$$

is a GAV tgd that is not a LAV tgd. Note that every inclusion dependency is a LAV tgd, but not vice-versa. Furthermore, every tgd with no existential quantifiers in its right-hand side (such tgds are called *full*) is logically equivalent to a set of GAV tgds. For example, the full tgd $\forall x \forall y \forall z(E(x, z) \wedge E(z, y) \rightarrow F(x, y) \wedge P(x))$ is logically equivalent to the set consisting of the GAV tgds $\forall x \forall y \forall z(E(x, z) \wedge E(z, y) \rightarrow F(x, y))$ and $\forall x \forall y \forall z(E(x, z) \wedge E(z, y) \rightarrow P(x))$.

From now on and for the sake of readability, we will often drop the universal quantifiers when writing dependencies.

As mentioned in the Introduction, tgds play an important role in data exchange and data integration, where they are used to specify the relationship between a source (local) schema and a target (global) schema or to express constraints in a target (global) schema. Moreover, it is known that *weakly acyclic* sets of tgds have good algorithmic properties as regards data exchange that, in general, are not possessed by arbitrary sets of tgds.

Definition 2.3 (Weak Acyclicity [15, 17]) Let Σ be a set of tgds and egds over a schema \mathbf{S} .

- The *dependency graph* of Σ is the following directed graph.
 - The nodes are the pairs (R, i) , where $R \in \mathbf{S}$ is a relation of some arity k , and $1 \leq i \leq k$. We call such pairs *positions*.
 - There is an edge from position (R, i) to position (S, j) if Σ contains a tgd $\forall \mathbf{x}(\phi \rightarrow \exists \mathbf{y} \psi)$ such that some variable from \mathbf{x} occurs in position (R, i) in ϕ and in position (S, j) in ψ .
 - There is a *special edge* from position (R, i) to position (S, j) if Σ contains a tgd $\forall \mathbf{x}(\phi \rightarrow \exists \mathbf{y} \psi)$ such that (i) some variable from \mathbf{x} occurs in position (R, i) in ϕ and also occurs in ψ , and (ii) some variable from \mathbf{y} occurs in position (S, j) in ψ .
- We say that Σ is *weakly acyclic* if the dependency graph contains no cycle going through a special edge.

Every set of GAV tgds is weakly acyclic, since the dependency graph has no special edges. It is also easy to see that every acyclic set of inclusion dependencies is weakly acyclic. However, the set $\Sigma = \{D(x, y) \rightarrow M(y), M(y) \rightarrow \exists x D(x, y)\}$ is a cyclic, yet weakly acyclic set of inclusion dependencies. Finally, the singleton set

consisting of the inclusion dependency $R(x, y) \rightarrow \exists zR(y, z)$ is not weakly acyclic because the dependency graph contains a self-loop on position $R.2$.

Intuitively, a *repair* of an instance I with respect to a set of dependencies Σ is an instance J that satisfies Σ and “differs minimally” from I . Here, we will consider three different set-based types of repairs. If I and J are instances, we will write $I \oplus J$ to denote the *symmetric difference* $(I \setminus J) \cup (J \setminus I)$ of I and J , that is, $I \oplus J$ is the set of all facts that either belong to I and not to J , or belong to J and not to I .

Definition 2.4 (Repairs) Let I, J be instances, and Σ a set of dependencies.

- We say that J is an \oplus -repair of I w.r.t. Σ if $J \models \Sigma$ and there is no instance J' such that $J' \models \Sigma$ and $I \oplus J' \subsetneq I \oplus J$.
- We say that J is a *subset-repair* of I w.r.t. Σ if J is both a subset of I and a \oplus -repair of I w.r.t. Σ . Equivalently, J is a subset-repair of I w.r.t. Σ if $J \subseteq I$, $J \models \Sigma$ and there is no instance $J' \subseteq I$ such that $J' \models \Sigma$ and $J \subsetneq J'$.
- We say that J is a *superset-repair* of I w.r.t. Σ if J is both a superset of I and a \oplus -repair of I w.r.t. Σ . Equivalently, J is a superset-repair of I w.r.t. Σ if $I \subseteq J$, $J \models \Sigma$, and there is no instance J' such that $I \subseteq J'$, $J' \models \Sigma$ and $J' \subsetneq J$.

In what follows and when the set Σ of dependencies is understood from the context, we will often drop the reference to Σ and, instead, write simply that J is an \oplus -repair of I , and similarly for subset and superset repairs.

For example, consider the key constraint $R(x, y) \wedge R(x, z) \rightarrow y = z$ and the instance $I = \{R(a, b), R(a, c), R(b, c)\}$. Then I has exactly two subset-repairs J_1 and J_2 that are also its only \oplus -repairs, namely, $J_1 = \{R(a, b), R(b, c)\}$ and $J_2 = \{R(a, c), R(b, c)\}$.

Next, let $\Sigma = \{P(x) \rightarrow Q(x)\}$ and $I = \{P(a), P(b)\}$. Then I has exactly four \oplus -repairs, namely, $J_1 = \emptyset$, $J_2 = \{P(a), Q(a)\}$, $J_3 = \{P(b), Q(b)\}$, and $J_4 = \{P(a), P(b), Q(a), Q(b)\}$. Note that J_1 is the only subset-repair of I , while J_4 is the only superset-repair of I .

For a different example, let $\Sigma = \{R(x, y) \rightarrow \exists zR(y, z)\}$ and $I = \{R(1, 2)\}$. It is easy to see that, for every $n \geq 1$ and every $k < n$, the “loop” instance

$$J_{n,k} = \{R(i, i + 1) \mid 1 \leq i < n\} \cup \{R(n, k)\}$$

is a superset-repair (hence, also a \oplus -repair) of I w.r.t. Σ . Thus, I has infinitely many, and arbitrarily large, superset-repairs w.r.t. Σ .

Next, we give the definitions of the main algorithmic problems in the study of repairs.

Definition 2.5 (Repair Checking) Let Σ be a set of dependencies, and let $\star \in \{\oplus, \text{subset}, \text{superset}\}$. The \star -repair checking problem w.r.t. Σ asks: given instances I and I' , check whether I' is a \star -repair of I w.r.t. Σ .

Definition 2.6 (Consistent Answers) Let Σ be a set of dependencies, q a conjunctive query, I an instance, and $\star \in \{\oplus, \text{subset}, \text{superset}\}$.

- The \star -consistent answers of q on I w.r.t. Σ , denoted by $\star\text{-Con}(q, I, \Sigma)$, is the intersection

$$\bigcap \{q(J) \mid J \text{ is a } \star\text{-repair of } I \text{ w.r.t. } \Sigma\}.$$

- The \star -consistent query answering problem of q w.r.t. Σ asks: given an instance I , compute the \star -consistent answer of q on I with respect to Σ .

Several remarks are in order now. First, we wish to emphasize that only finite instances and only finite repairs are considered in this paper. It is known that there exist sets of dependencies Σ , conjunctive queries q , and instances I such that the consistent answer to q on I w.r.t. Σ would yield a different result if infinite repairs would be considered as well (see [31]). Second, if Σ consists of tgds and egds, there is always a subset-repair and, hence, also an \oplus -repair. If Σ consists of tgds only, there is also always a superset-repair.

Proposition 2.7 *If Σ is a set dependencies, then every instance has a subset-repair and, hence, also an \oplus -repair w.r.t. Σ . Furthermore, Σ is a set of tgds, then every instance has a superset-repair w.r.t. Σ .*

Proof The empty instance satisfies Σ , hence, for every instance I that fails to satisfy Σ , there is a maximal (w.r.t. \subseteq) instance J such that $J \subseteq I$ and $J \models \Sigma$. Now let Σ consist of tgds only, and let I be an instance. Take J to be the instance that has the same active domain of I and in which each relation symbol denotes the complete relation (of appropriate arity) over the given active domain. Then $I \subseteq J$ and $J \models \Sigma$. Hence, there is a minimal (w.r.t. \subseteq) instance $J' \subseteq J$ such that $I \subseteq J'$ and $J' \models \Sigma$. By construction, J is a superset-repair of I w.r.t. Σ . \square

Finally, it is easy to see that, for all $\star \in \{\oplus, \text{subset}, \text{superset}\}$, it is the case that the \star -repair checking problem is in coNP . Moreover, the subset-consistent query answering is always in Π_2^P . For example, to check that J is not an \oplus -repair of I , one has to guess an instance J' of size at most $|I| + |J|$ and verify that $J' \models \Sigma$ and $I \oplus J' \subset I \oplus J$.

Note that the notion of *superset-consistent answers* is closely related to the notion of *certain answers* in incomplete information and in data exchange. To make the connection with data exchange precise, recall from [17] that a *schema mapping* (also called *data exchange setting*) is a 4-tuple $\mathcal{M} = (\mathbf{S}, \mathbf{T}, \Sigma_{st}, \Sigma_t)$, where \mathbf{S} and \mathbf{T} are disjoint schemas called the *source schema* and the *target schema*, Σ_{st} is a collection of source-to-target tgds, and Σ_t is a set of target dependencies, that is, dependencies over the target schema. Recall also that an instance J over the target schema is said to be a *solution* for an instance I over the source schema if the pair (I, J) satisfies all constraints in $\Sigma_{st} \cup \Sigma_t$. Finally, recall that, if \mathcal{M} is a schema mapping, q is a query over the target schema of \mathcal{M} , and I is an instance over the source schema of \mathcal{M} , then the certain answers of q on I w.r.t. \mathcal{M} , denoted by $\text{certain}(q, I, \mathcal{M})$, are defined as follows:

$$\text{certain}(q, I, \mathcal{M}) = \bigcap \{q(J) \mid J \text{ is solution for } I \text{ w.r.t. } \mathcal{M}\}.$$

The connection between superset-consistent answers and certain answers is given by the next proposition.

Proposition 2.8 *Let Σ is a set of dependencies and let \mathcal{M}_Σ be the schema mapping defined as follows:*

- (i) *The source-to-target tgds of \mathcal{M}_Σ are copy tgds of the form $R'(\mathbf{x}) \rightarrow R(\mathbf{x})$, where R is a relation symbol of the schema of Σ and R' is a new relation symbol whose arity is that of R .*
- (ii) *The target tgds and egds of \mathcal{M}_Σ are the tgds and egds in Σ .*

Then for every conjunctive query q and every instance I ,

$$\text{superset-Con}(q, I, \Sigma) = \text{certain}(q, I', \mathcal{M}_\Sigma),$$

where I' is the copy of the instance I obtained by changing the name of every relation R of I to R' .

The preceding proposition is proved by first observing that every superset repair of I w.r.t. Σ is a solution of I' w.r.t. \mathcal{M}_Σ and then combining the monotonicity of conjunctive queries with the fact that every solution J for I' w.r.t. \mathcal{M}_Σ contains a superset-repair of I . We will make use of this connection in some of our complexity results for the superset-consistent query answering problem.

Finally, we recall the definition of *data complexity* from [34] adapted to our context.

Definition 2.9 Let \mathcal{L} be a class of dependencies, let $\star \in \{\oplus, \text{subset}, \text{superset}\}$, and let \mathcal{C} be a complexity class.

- We say that the *data complexity of \star -consistent query answering w.r.t. \mathcal{L} is in \mathcal{C}* if for every finite subset Σ of \mathcal{L} and for every conjunctive query q , we have that the \star -consistent query answering problem of q w.r.t. Σ is in \mathcal{C} .
- We say that the *data complexity of \star -consistent query answering w.r.t. \mathcal{L} is \mathcal{C} -complete* if it is in \mathcal{C} and there are a finite subset Σ' of \mathcal{L} and a conjunctive query q' , such that the \star -consistent query answering problem of q' w.r.t. Σ' is \mathcal{C} -complete.

3 Universal Repairs

By analogy to the notion of a *universal solution* in data exchange [17], we will now introduce the notion of a *universal repair*. Since it will turn out that universal repairs often do not exist, we also introduce a weaker notion, namely, that of a n -universal repair, $n \geq 1$.

Definition 3.1 Let Σ be a set of dependencies, I an instance, and $\star \in \{\oplus, \text{subset}, \text{superset}\}$.

- A universal \star -repair of I w.r.t. Σ is a \star -repair J of I w.r.t. Σ such that if q is a conjunctive query, then

$$\star\text{-Con}(q, I, \Sigma) = q(J)_{\downarrow},$$

where $q(J)_{\downarrow}$ is the set of tuples in $q(J)$ containing only values from the active domain $\text{adom}(I)$ of I .

- For $n \geq 1$, an n -universal \star -repair of I w.r.t. Σ is a \star -repair J of I w.r.t. Σ such that if q is conjunctive query with at most n atoms, then

$$\star\text{-Con}(q, I, \Sigma) = q(J)_{\downarrow}.$$

Clearly, a universal repair is also a n -universal repair, for every $n \geq 1$. As we shall soon see, however, the converse is not always true. The next proposition, which follows immediately from the definitions and the fact that the data complexity of conjunctive queries is in PTIME, justifies the introduction of the notions of universal repairs and n -universal repairs.

Proposition 3.2 *Let Σ be a set of dependencies and let $\star \in \{\oplus, \text{subset}, \text{superset}\}$.*

- Assume that n is a positive integer. If there is a polynomial-time algorithm that, given an instance I , it returns an n -universal \star -repair of I w.r.t. Σ , then the data complexity of the \star -consistent query answering problem for conjunctive queries with at most n atoms is in PTIME.
- In particular, if there is a polynomial-time algorithm that, given an instance I , it returns a universal \star -repair for I w.r.t. Σ , then the data complexity of the \star -consistent query answering for conjunctive queries is in PTIME.

It is not hard to see that the notion of a universal repair can be characterized in terms of homomorphisms.

Proposition 3.3 *Let Σ be a set of dependencies, let $\star \in \{\oplus, \text{subset}, \text{superset}\}$, and let I and J be two instances. Then the following statements are equivalent.*

- (1) J is a universal \star -repair of I w.r.t. Σ .
- (2) J is a \star -repair of I w.r.t. Σ and for every \star -repair J' of I w.r.t. Σ , there is a homomorphism $h : J \rightarrow J'$ such that $h(a) = a$, for every a in $\text{adom}(I) \cap \text{adom}(J)$.

Proof The direction (1) \Rightarrow (2) is proved by applying the definition of universal \star -repair to the variant of canonical conjunctive query of J in which the elements in the intersection $\text{adom}(I) \cap \text{adom}(J)$ are viewed as free variables and the elements in $\text{adom}(J) \setminus \text{adom}(I)$ are existentially quantified. The direction (2) \Rightarrow (1) follows from the preservation of conjunctive queries under homomorphisms. \square

Note that a similar characterization of n -universal repairs in terms of homomorphisms does not hold; intuitively, the reason is that, in general, an n -universal repair

may contain more than n facts, so the proof of direction (1) \Rightarrow (2) does not go through.

The notion of a universal repair is closely related to that of a *nucleus*, which was introduced and studied in the context of *update* repairs by J. Wijsen in [35]. The notion of a nucleus can be adapted to set-based repairs as follows.

Let Σ be a set of dependencies, $\star \in \{\oplus, \text{subset}, \text{superset}\}$, and I an instance. A \star -*nucleus* for I w.r.t. Σ is an instance J such that if q is a conjunctive query, then $\star\text{-Con}(q, I, \Sigma) = q(J)_\downarrow$. Thus, a \star -universal repair for I w.r.t. Σ is \star -nucleus for I w.r.t. Σ that also happens to be a \star -repair of I w.r.t. Σ . It should be pointed out that, if an instance I has finitely many \star -repairs w.r.t. Σ , then a \star -nucleus for I w.r.t. Σ always exists. To see this, let J be the direct product of all \star -repairs of I w.r.t. Σ . From the definition of the direct product and the fact that the direct product has homomorphisms to each of its factors, it is easy to verify that for every Boolean conjunctive query q , we have that $\star\text{-Con}(q, I, \Sigma) = q(J)$. This does not exactly extend to non-Boolean conjunctive queries, but comes quite close. For example, if q is a unary conjunctive query and a is an element in the active domain $\text{adom}(I)$, then $a \in \star\text{-Con}(q, I, \Sigma)$ if and only if the tuple $(a, \dots, a) \in q(J)$. Therefore, one can obtain a \star -nucleus J^* for I from J by taking an isomorphic copy of J in which every tuple in J of the form (a, \dots, a) with $a \in \text{adom}(I)$ is identified with a , while every tuple that is not of this form is identified with a unique element not in $\text{adom}(I)$. This construction is also described in [35, Definition 4.4]. See also the proof of Proposition 3.8 below, where a similar construction is used.

Note that all \star -nuclei of I w.r.t. Σ are homomorphically equivalent via homomorphisms that are the identity on elements of the active domain $\text{adom}(I)$ of I . However, it is possible that every \star -nucleus of I w.r.t. Σ is of size exponentially bigger than the size of I .

We now give several examples illustrating universal repairs and n -universal repairs.

Example 3.4 Let $\Sigma = \{R(x, y) \rightarrow \exists z R(y, z)\}$.

First, consider the instance $I = \{R(1, 2), R(2, 2), R(1, 3), R(3, 4)\}$. Then the instance $J = \{R(1, 2), R(2, 2)\}$ is a universal subset-repair of I w.r.t. Σ , because it is the only subset-repair of I .

Next, consider the instance $I = \{R(1, 2)\}$. Then I has no universal superset-repair; indeed, this is so because every superset-repair of I must contain a cycle of some length n_0 (recall that all repairs must be finite), and therefore cannot be universal since, as seen earlier, not every superset-repair contains a cycle of the same length n_0 . However, for every $n \geq 1$, there are n -universal superset-repairs of I . Specifically, one such n -universal repair is the “loop” instance

$$J_{4n,2n} = \{R(i, i + 1) | 1 \leq i < 4n\} \cup \{R(4n, 2n)\}.$$

Example 3.5 Let $\Sigma = \{P(x) \wedge Q(x) \rightarrow R(x)\}$ and $I = \{P(a), Q(a)\}$. The instance $J = \{P(a), Q(a), R(a)\}$ is a universal superset-repair of I w.r.t. Σ , because it is the only superset-repair of I . However, I has no universal subset-repair, since its

subset-repairs are $\{P(a)\}$ and $\{Q(a)\}$. In fact, I does not even have a 1-universal subset-repair, since the queries $P(x)$ and $Q(x)$ return different values on the two subset-repairs of I .

Example 3.6 Let $\Sigma = \{P(x) \wedge Q(y) \rightarrow x = y\}$ and $I = \{P(a), Q(b)\}$. First, I has no superset-repair w.r.t. Σ , hence it has no 1-universal superset-repair. Furthermore, it has neither a 1-universal subset-repair, nor a 1-universal \oplus -repair w.r.t. Σ .

The next proposition describes some basic properties of the different types of universal repairs, and the relationship between them. When we consider an isomorphism between repairs J, J' of an instance I , we assume that the isomorphism h in question is constant on the active domain of I , that is, $h(a) = a$ for all $a \in \text{adom}(I)$.

Proposition 3.7 *Let Σ be a set of dependencies.*

1. *Every instance has at most one universal superset-repair, up to isomorphism.*
2. *An instance has a universal subset-repair if and only if it has exactly one subset-repair. Consequently, every instance has at most one universal subset-repair.*
3. *Every universal \oplus -repair of an instance I is a universal subset-repair of I . Consequently, every instance has at most one universal \oplus -repair. In contrast, a universal subset-repair need not be a universal \oplus -repair.*
4. *For every instance I and every instance J , the following statements are equivalent:*
 - (1) *J is a universal \oplus -repair of I .*
 - (2) *J is a universal subset-repair of I and $J \subseteq J'$, for every \oplus -repair J' of I .*
 - (3) *J is a subset repair of I and $J \subseteq J'$, for every \oplus -repair J' of I .*

Proof 1. Let J_1 and J_2 be universal superset-repairs of I . In particular, the active domain of I is contained in both the active domain $\text{adom}(J_1)$ of J_1 and the active domain $\text{adom}(J_2)$ of J_2 . Let q_1 and q_2 be the canonical conjunctive queries of J_1 and J_2 , where, as in the proof of Proposition 3.3, the values in the active domain $\text{adom}(I)$ of I are treated as free variables of the queries, and the values in $\text{adom}(J_i) \setminus \text{adom}(I)$, $i = 1, 2$, are treated as existentially quantified variables of the queries. By universality, $q_1(J_1)_{\downarrow} = q_1(J_2)_{\downarrow}$ and $q_2(J_1)_{\downarrow} = q_2(J_2)_{\downarrow}$. This implies that there are homomorphisms $h_1 : J_1 \rightarrow J_2$ and $h_2 : J_2 \rightarrow J_1$ such that $h_1(a) = a$ and $h_2(a) = a$ for all values a in the active domain of I . Consequently, J_1 and J_2 are homomorphically equivalent, where the values from the active domain of I are treated as constants. Furthermore, since J_1 and J_2 are superset-repairs of I , we claim that each of them is a *core* (with the values from the active domain of I treated as constants). This means that there is no proper *retraction* of J_i to itself, i.e., we have to show that, for $i = 1, 2$, there is no homomorphism $h : J_i \rightarrow J_i$ such that h is the identity on its range and $h(J_i)$ (as a set of facts) is properly contained in J_i . Suppose that $h : J_i \rightarrow J_i$ is a homomorphism that is the identity on its range. Then we have that $I \subseteq h(J_i) \subseteq J_i$. Moreover, since Σ consists of tgds and egds, J_i satisfies Σ , and h is the identity on its range, we have

that $h(J_i) \models \Sigma$ (this means that tgds and egds are preserved under retractions - see also cf. [23]). Since J_i is a superset repair of I , it follows that $h(J_i) = J_i$. This completes the proof that J_1 and J_2 are cores. But it is a well known fact that any two homomorphically equivalent core instances are isomorphic (cf. [18, 25]).

2. If an instance I has exactly one subset-repair J , then, trivially, J is a universal subset-repair. If, on the other hand, there are (at least) two different subset-repairs J_1 and J_2 , then, by the definition of subset repairs, neither is a subset of the other, and hence, neither can be universal. This is so because each $J_i, i = 1, 2$, contains a fact, say $R(\mathbf{a})$, that is not included in all subset-repairs; therefore, J_i does not correctly compute the consistent answers for the query $R(\mathbf{x})$.

3. Let J be a universal \oplus -repair of I w.r.t. Σ . Let J' be a subset-repair of I (recall that every instance has a subset-repair). Since J is a universal \oplus -repair, every fact $R(\mathbf{a})$ of I that belongs to J also belongs also to J' (for, if not, then evaluating the query $q(\mathbf{x}) = R(\mathbf{x})$ in J would not yield the consistent answers). Consequently, $J' \oplus I \subseteq J \oplus I$, and therefore, by the definition of \oplus -repairs, $J = J'$. Hence J is a subset-repair.

Not every universal subset-repair is a universal \oplus -repair. To see this, consider the set $\Sigma = \{P(x) \rightarrow \exists y Q(y), P(x) \wedge Q(x) \rightarrow R(x)\}$ and the instances $I = \{P(a), Q(a)\}$ and $J = \{Q(a)\}$. Then J is a universal subset-repair of I , but not a universal \oplus -repair of I , as may be seen by considering the \oplus -repair $J' = \{P(a), Q(b)\}$ and the query $q(x) = Q(x)$. In fact, this shows that J is not even a 1-universal \oplus -repair of I .

4. For the implication (1) \Rightarrow (2), if J is a universal \oplus -repair of I , then, as shown in Part 3, J is a universal subset-repair of I . Moreover, by \oplus -universality, J is contained in every \oplus -repair of I . The implication (2) \Rightarrow (3) is obvious. Finally, for the implication (3) \Rightarrow (1), assume that J is a subset repair of I that is contained in every \oplus -repair of I , and let q be a conjunctive query. Since every subset repair is also a \oplus -repair, we have that $\oplus\text{-Con}(q, I, \Sigma) \subseteq q(J)_\downarrow = q(J)$ (the equality holds because J is a subinstance of I). Moreover, if J' is an arbitrary \oplus -repair of I , then we have that $q(J) \subseteq q(J')$, since $J \subseteq J'$ and conjunctive queries are monotone. Thus, $q(J) \subseteq \oplus\text{-Con}(q, I, \Sigma)$ and so $\oplus\text{-Con}(q, I, \Sigma) = q(J)_\downarrow = q(J)$. \square

Note that Part 2 and Part 3 of Proposition 3.7 hold with “ n -universal” in place of “universal”, where n is an arbitrary positive integer. The reason is that universality was applied only to queries with a single atom. In particular, the counterexample in the proof of Part 3 shows that a universal subset-repair need not even be a 1-universal \oplus -repair. Note also that the notion of a *universal superset-repair* is intimately linked to the notion of a *core universal solution* in data exchange [18], as seen in the proof of Proposition 3.7. We will further exploit this connection when we consider superset-repairs for weakly acyclic sets of tgds and egds.

Proposition 3.8 *Let Σ be a set of dependencies. Every instance that has a superset-repair w.r.t. Σ has an n -universal superset-repair w.r.t. Σ , for each positive integer*

n .¹ Moreover, there is a set of tgds Σ for which there is no algorithm that, given an instance I , computes a 1-universal superset repair of I .

Proof Let Σ be a set of dependencies, let I be an instance, and let n be a natural number. Let \mathcal{Q}_n the set of all pairs $\langle q, \mathbf{a} \rangle$ where q is a conjunctive query with at most n atoms, and $\mathbf{a} \in \text{adom}(I)^k$ (where $k \geq 0$ is the arity of q) be such that $\mathbf{a} \notin \text{superset-Con}(q, I, \Sigma)$. For each $\langle q, \mathbf{a} \rangle \in \mathcal{Q}_n$, let $J_{\langle q, \mathbf{a} \rangle}$ be a superset-repair of I w.r.t. Σ such that $\mathbf{a} \notin q(J_{\langle q, \mathbf{a} \rangle})$. Let J_\times be the direct product $\prod\{J_{\langle q, \mathbf{a} \rangle} \mid \langle q, \mathbf{a} \rangle \in \mathcal{Q}_n\}$ (cf. [16]), and let J be an isomorphic copy of J_\times in which each “diagonal element”, i.e., each element of the form (a, a, \dots, a) is replaced by a .

Claim 1 $I \subseteq J$ and $J \models \Sigma$.

First, we prove that I is a subset of J . Every fact $R(a_1, \dots, a_m)$ of I must belong to all superset-repairs of I , from which it follows, by the definition of direct products, that $R((a_1, \dots, a_1), \dots, (a_m, \dots, a_m))$ belongs to J_\times , and therefore $R(a_1, \dots, a_m)$ belongs to J .

Next we show that Σ is true in J . This follows from the known fact that tgds and egds are preserved under taking direct products [16]: whenever I_1, \dots, I_n are instances satisfying a set of tgds and egds Σ , then the direct product $\prod_i I_i$ satisfies Σ . This finishes the proof of the Claim.

It follows from Claim 1 that there is some superset-repair J_0 of I such that $J_0 \subseteq J$. We prove that J_0 is a n -universal superset-repair of I w.r.t. Σ . That is, for all queries q with at most n atoms,

$$\text{superset-Con}(q, I, \Sigma) = q(J_0)_\downarrow.$$

One inclusion follows from the fact that J_0 is a superset-repair of I w.r.t. Σ . For the other inclusion, we reason by contraposition: suppose that $\mathbf{a} \notin \text{superset-Con}(q, I, \Sigma)$. Then, by construction, $\mathbf{a} \notin q(J_{\langle q, \mathbf{a} \rangle})$. Let $h : J_\times \rightarrow J_{\langle q, \mathbf{a} \rangle}$ be the natural map that sends each tuple to its value at the coordinate corresponding to $\langle q, \mathbf{a} \rangle$. Let $h' : J \rightarrow J_{\langle q, \mathbf{a} \rangle}$ be the composition of h with the isomorphism between J_\times and J . In other words,

$$h'(x) = \begin{cases} a & \text{if } x = a \text{ for some } a \in \text{adom}(I) \\ h(x) & \text{otherwise} \end{cases}$$

Since $h : J_\times \rightarrow J_{\langle q, \mathbf{a} \rangle}$ is a homomorphism, we have that $h' : J \rightarrow J_{\langle q, \mathbf{a} \rangle}$ is a homomorphism as well. Since conjunctive queries are preserved by homomorphisms, it follows that $\mathbf{a} \notin q(J)$. Since $J_0 \subseteq J$ and conjunctive queries are monotone, this implies that $\mathbf{a} \notin q(J_0)$.

¹Proposition 7.1 in [33] asserts the existence of a set of tgds Σ and an instance I , such that I has no 1-universal superset-repair w.r.t. Σ . Unfortunately, there was a mistake in the proof of this proposition. Proposition 3.8 clarifies the situation (note that, by Proposition 2.7, for sets of tgds Σ , every instance has a superset-repair).

We now proceed to the second part of the proposition. It was shown in [10] that there is a set of tgds Σ and a Boolean conjunctive query q consisting of a single atom, such that the following problem is undecidable:

Given an instance I , decide whether, for every instance $J \models \Sigma$ with $I \subseteq J$, we have that q is true in J .

Using the notation used in Proposition 2.8, this is precisely the problem of computing the *certain answers* of q with respect to the schema mapping \mathcal{M}_Σ . It follows that there is no effective algorithm for computing a 1-universal superset-repair with respect to Σ , since, by Proposition 2.8, such an algorithm would immediately yield an algorithm for computing the certain answers of q with respect to Σ . \square

4 LAV tgds

The main result of this section is that if Σ is a set of LAV tgds and q is a conjunctive query, then there is a polynomial-time algorithm for computing the \star -consistent answers of q on a given instance I w.r.t. Σ , where $\star \in \{\oplus, \text{subset}, \text{superset}\}$. We also show that there is a polynomial-time algorithm for the \star -repair checking problem w.r.t. Σ , where $\star \in \{\oplus, \text{subset}, \text{superset}\}$. In the special case in which Σ is a set of inclusion dependencies, the tractability of the subset-repair checking problem and the subset-consistent query answering problem for conjunctive queries was established in [14]. The tractability of the subset and the \oplus -repair checking problem for a weakly acyclic set of LAV tgds was shown in [2] and was subsequently extended to the broader class of semi-LAV sets of tgds [24] (which is still a subclass of the class of weakly acyclic sets of tgds). Finally, the \oplus -consistent query problem for sets of inclusion dependencies and universal constraints was studied in [9], where disjunctive logic programming was used to obtain the \oplus -consistent answers, but no complexity results were established.

In obtaining our tractability results, we will use extensively the notions of universal repair and of n -universal repair, $n \geq 1$, for sets of LAV tgds.

We say that a set of dependencies Σ is *closed under union* if for all instances I_1, I_2 such that $I_1 \models \Sigma$ and $I_2 \models \Sigma$, we have that $I_1 \cup I_2 \models \Sigma$. It is well known that every set of LAV tgds is closed under union (e.g., see [2, 12]). The next result shows that, in a certain sense, this property is characteristic of LAV tgds. For completeness, we also include a proof that sets of LAV tgds are closed under union.

Theorem 4.1 *Every set of LAV tgds is closed under union. Furthermore, if Σ is a set of tgds that is closed under union, then Σ is logically equivalent to a set of LAV tgds.*

Proof Suppose that Σ is a set of LAV tgds and that I_1 and I_2 are instances such that both satisfy Σ . Let $\forall \mathbf{x}(R(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$ be a tgd in Σ . Then, for every tuple of values \mathbf{a} such that $R(\mathbf{a})$ holds in $I_1 \cup I_2$, we have that $R(\mathbf{a})$ already holds in I_1 or I_2 , and hence $\exists \mathbf{y} \psi(\mathbf{a}, \mathbf{y})$ holds in I_1 or in I_2 , which implies that it also holds in $I_1 \cup I_2$. This shows that $I_1 \cup I_2 \models \Sigma$.

We now prove that the converse holds. Let Σ be a set of tgds that is closed under union, and let \mathbf{R} be the schema of Σ . We define Σ' as the (finite) set of all LAV tgds over \mathbf{R} implied by Σ and containing at most n atoms in the right-hand side, where n is the maximum number of atoms in the right-hand side of a tgd in Σ . By definition, Σ logically implies every LAV tgd in Σ' . Conversely, consider any tgd in Σ of the form

$$\sigma = \forall \mathbf{x}(\phi(\mathbf{x}) \rightarrow \exists \mathbf{y}\psi(\mathbf{x}, \mathbf{y}))$$

We will show that $\Sigma' \models \sigma$. Let I be the canonical instance of $\phi(\mathbf{x})$, that is, the instance whose active domain consists of the variables in \mathbf{x} and whose facts are the atoms in ϕ . Let g be the natural variable assignment for \mathbf{x} , i.e., the map that sends every variable in \mathbf{x} to itself.

For each fact f of I , let J_f be an n -universal superset-repair of $\{f\}$ with respect to Σ (using Proposition 3.8), using all fresh elements except for the elements belonging to the fact f itself. Finally, let $J = \bigcup\{J_f \mid f \text{ is a fact of } I\}$. It follows from the closure under union that $J \models \Sigma$. Furthermore, by construction, we have that $I \subseteq J$. In particular, we have that $J \models \phi(\mathbf{x})$ under the variable assignment g , and therefore, $\psi(\mathbf{x}, \mathbf{y})$ must be satisfied under some variable assignment g' extending g .

In what follows, the basic idea will be that we split the conjunctive query $\exists \mathbf{y}\psi(\mathbf{x}, \mathbf{y})$ into different parts, such that each part is mapped (by g') into a different “component” J_f of J . If it happens to be the case that g' maps each existential variable to a value outside of $\text{adom}(I)$, then this will turn out to be easy to do. In general, it may be the case that g' maps some of the variables \mathbf{y} to elements of $\text{adom}(I)$. However, if this is the case, then we can find a substitution instance of ψ that is satisfied, and where those existential variables have been replaced some of the universally quantified variables. Let $\exists \mathbf{y}'\psi'(\mathbf{x}, \mathbf{y}')$ be a maximal substitution instance (i.e., replacing as many existentially quantified variables by universal quantified variables as possible) for which it is the case that $J \models \exists \mathbf{y}'\psi'(\mathbf{x}, \mathbf{y}')$ under the variable assignment g , and let h be a variable assignment for \mathbf{x} and \mathbf{y}' extending g such that $\psi'(\mathbf{x}, \mathbf{y}')$ is satisfied in J under h . Note that by construction, $h(x) = x$ for all $x \in \mathbf{x}$ and $h(y) \notin \text{adom}(I)$ for all $y \in \mathbf{y}'$.

We now label each atom A of ψ by a fact f of I such that $h(A)$ belongs to J_f (note that, by construction of J , we can always find such a fact. If there are more than one, we choose one arbitrarily). Observe that, if two atoms of ψ receive a different label, they cannot share an existentially quantified variable (this is precisely because all existential variables are mapped by h to values outside $\text{adom}(I)$, which must therefore belong to a unique “component” J_q of J). For each fact f of I , let q_f be the conjunctive query that is obtained as the subquery of $\exists \mathbf{y}'\psi(\mathbf{x}, \mathbf{y}')$ consisting of the atoms labeled f . Then by construction, q_f is a conjunctive query that is satisfied in J_f under the natural variable assignment g . Since q_f has at most n atoms and J_f is an n -universal superset-repair of $\{f\}$, it follows that q_f is true in every superset-repair of $\{f\}$ with respect to Σ , under the natural variable assignment g . But this means that Σ implies the LAV tgd $\forall \mathbf{x}(f(\mathbf{x}) \rightarrow q(\mathbf{x}))$, and hence the latter belongs to Σ' . This argument applies to each fact f of I . It follows that, whenever $\phi(\mathbf{x})$ is satisfied in an instance satisfying Σ' , then the LAV tgds in Σ' ensure that $\exists \mathbf{y}'\psi(\mathbf{x}, \mathbf{y}')$, and hence also $\exists \mathbf{y}\psi(\mathbf{x}, \mathbf{y})$ is satisfied. In other words, Σ' implies σ . \square

Closure under union implies that every instance has a universal \oplus -repair. In fact, the following characterization holds.

Proposition 4.2 *Let Σ be a set of dependencies. The following statements are equivalent:*

- (1) Σ is closed under union.
- (2) Every instance has a universal \oplus -repair w.r.t. Σ .
- (3) Every instance has a universal subset-repair w.r.t. Σ .

Proof For the direction (1) \Rightarrow (2), let Σ be a set of tgds that is closed under union, and let I be an instance. Since the empty instance satisfies Σ , we have that I has at least one subset-repair w.r.t. Σ . In fact, I has a unique subset-repair.² For, if J_1 and J_2 are subset-repairs of I , then $J_1 \cup J_2$ is also a subinstance of I satisfying Σ , and hence, by the definition of repairs, $J_1 = J_1 \cup J_2 = J_2$. Clearly, the unique subset-repair J of I is a universal subset-repair of I . Furthermore, we claim that J is a universal \oplus -repair of I . Let q be a conjunctive query. We have to show that $\oplus\text{-Con}(q, I, \Sigma) = q(J)_\downarrow$. From the definition of $\oplus\text{-Con}(q, I, \Sigma)$ and since J is a \oplus -repair of I , it follows that $\oplus\text{-Con}(q, I, \Sigma) \subseteq q(J)_\downarrow$. For the other inclusion, let K be a \oplus -repair of I . We have to show that $q(J)_\downarrow \subseteq q(K)$. Since J and K are \oplus -repairs, we have $J \models \Sigma$ and $K \models \Sigma$. By closure under union, we have that also $J \cup K \models \Sigma$. As $J \subseteq I$, we have $I \oplus (J \cup K) \subseteq I \oplus K$. Since K is a \oplus -repair of I , this can only happen if $K \cup J = K$, which means that $J \subseteq K$. It follows that $q(J)_\downarrow \subseteq q(K)$. Hence, J is a universal \oplus -repair of I .

The direction (2) \Rightarrow (3) follows immediately from Proposition 3.7.

[3 \Rightarrow 1] Let I_1, I_2 be instances satisfying Σ . Towards a contradiction, suppose that $I_1 \cup I_2$ does not satisfy Σ . Let J be the universal subset-repair of $I_1 \cup I_2$. Then J must omit some fact of $I_1 \cup I_2$. Without loss of generality, we may assume that J omits a fact of I_1 . Since I_1 satisfies Σ , there must be a subset-repair of $I_1 \cup I_2$ that contains all facts in I_1 . But this subset-repair must then be incomparable to J , which means that J is not the only subset-repair of I , and hence, by Proposition 3.7, J is not a universal subset-repair of I , a contradiction. \square

Corollary 4.3 *If Σ is a set of LAV tgds, then every instance I has a unique subset repair, which is also the unique universal subset-repair and the unique universal \oplus -repair of I w.r.t. Σ .*

Corollary 4.3 implies, in particular, that if Σ is a set of LAV tgds, then, for all conjunctive queries q and instances I , $\oplus\text{-Con}(q, I, \Sigma) = \text{subset-Con}(q, I, \Sigma)$.

Theorem 4.4 *Let Σ be a set of LAV tgds. There is a polynomial-time algorithm that, given an instance I , computes the unique universal subset-repair of I w.r.t. Σ (which is also the unique universal \oplus -repair of I w.r.t. Σ).*

²For inclusion dependencies, this fact was pointed out in [14].

Proof Let I be an instance. By Corollary 4.3, I has a unique subset repair I_0 , which is also both the unique universal subset-repair and the unique universal \oplus -repair of I . We will show how to compute I_0 in polynomial time. We start with a definition. Let τ be a LAV tgd of the form $\forall \mathbf{x} (R(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}))$. We say that a fact t falsifies the LAV tgd τ in an instance J if t is in J , it is of the form $R(\mathbf{a})$, and $\exists \mathbf{y} \psi(\mathbf{a}, \mathbf{y})$ is false in J .

The algorithm works as follows. It starts with the instance I . As long as the current instance contains a fact $R(\mathbf{a})$ that falsifies some tgd in Σ , it removes it from the current instance.

Clearly, the algorithm runs in time polynomial in the size of I . It is easy to show by induction that, at any time of the run of the algorithm, the current instance is a superset of the unique subset-repair I_0 of I . Hence, if J is the instance obtained at the end of the run of the algorithm, we have that $I_0 \subseteq J$. Moreover, since there is no tuple falsifying a tgd in Σ in J , we have that J satisfies Σ . Since $J \models \Sigma$, $I_0 \subseteq J \subseteq I$, and I_0 is a subset-repair of I , we conclude that I_0 is equal to J . \square

The situation for superset-repairs is very different. Example 3.4 shows that an instance may have infinitely many superset-repairs of arbitrarily large size, and it may have no universal superset-repairs. Nevertheless, we will show that if Σ is a set of LAV tgds then, for every instance I and every $n \geq 1$, there is an n -universal superset-repair for I w.r.t. Σ that can be computed in polynomial time in the size of I . For this, we need the following lemma.

Lemma 4.5 *Let Σ be a set of LAV tgds. There is a polynomial time algorithm that, given instances I and J with $I \subseteq J$ and $J \models \Sigma$, computes a superset-repair K of I w.r.t. Σ such that $K \subseteq J$.*

Proof Figure 1 depicts an algorithm that, given I and J with $I \subseteq J$ and $J \models \Sigma$, either verifies that J is a repair of I or computes a set K such that $I \subseteq K \subsetneq J$ and $K \models \Sigma$. By applying repeatedly the algorithm until it outputs a repair, we obtain the result.

By Theorem 4.4, computing the subset-repair of an instance is in polynomial time. It follows that the algorithm in Fig. 1 runs in time polynomial in the sizes of I and J .

We prove the correctness of the algorithm. Suppose first that the algorithm outputs an instance K . We have to show that $I \subseteq K \subsetneq J$ and $K \models \Sigma$. By construction, there exists a fact $R(\mathbf{a})$ in $J \setminus I$ such that K is the subset-repair of $J \setminus \{R(\mathbf{a})\}$. This implies that $K \models \Sigma$ and $K \subsetneq J$. Moreover, K contains I , as the algorithm returned K .

Fig. 1 Algorithm for the superset-repair checking problem with respect to a set of LAV tgds

Input: Two instances I and J such that $I \subseteq J$ and $J \models \Sigma$
Output: Either “ J is a repair” or an instance K such that $I \subseteq K \subsetneq J$ and $K \models \Sigma$
For each fact $R(\mathbf{a})$ in $J \setminus I$,
 compute the subset-repair L of $J \setminus \{R(\mathbf{a})\}$
 If $I \subseteq L$, **Then Return** L
 End If
End For
Return “ J is a repair”

Next, assume that J is not a superset-repair. We have to prove that the algorithm does not output “ J is a repair”. Since J is not a superset-repair, there exists an instance J_0 such that $I \subseteq J_0 \subsetneq J$ and $J_0 \models \Sigma$. Take $R(\mathbf{a})$ in $J \setminus J_0$. Let L be the subset-repair of $J \setminus \{R(\mathbf{a})\}$.

Consider the instance $L_0 := L \cup J_0$. By Theorem 4.1, $L_0 \models \Sigma$. Moreover, $L_0 \subseteq J \setminus \{R(\mathbf{a})\}$ and $J \oplus L_0 \subseteq J \oplus K$. Since K is a subset-repair of J , this can only happen if $L_0 = K$. That is, $J_0 \subseteq L$. Together with $I \subseteq J_0$, we obtain $I \subseteq L$. Recall that L is the subset-repair of $J \setminus \{R(\mathbf{a})\}$. This means that, when the algorithm enters the loop and picks $R(\mathbf{a})$, it stops running and returns the instance L . In particular, the algorithm does not output “ J is a repair”. \square

Theorem 4.6 *Let Σ be a set of LAV tgds, and let $n \geq 1$. There is a polynomial time algorithm that, given an instance I , computes an n -universal superset-repair of I w.r.t. Σ .*

Proof Let \mathbf{R} be the schema consisting of the relations in Σ . We may safely restrict attention to instances over \mathbf{R} (because facts involving relations outside of \mathbf{R} do not require repairing).

By Propositions 3.8 and 2.7, every instance I has an n -universal superset-repair w.r.t. Σ . Using the fact that Σ consists of LAV tgds, we can furthermore show that an n -universal superset-repair of a given instance can be effectively constructed: let q be any conjunctive query over \mathbf{R} with at most n atoms, and let \mathbf{a} be a tuple from $\text{adom}(I)$ whose length is the arity of q . Recall from Proposition 2.8 that $\mathbf{a} \in \text{superset-Con}(q, I, \Sigma)$ holds if and only if \mathbf{a} belongs to $q(J)$ for all instances J extending I and satisfying Σ . It was shown in [5, 10] that the latter is decidable in double exponential time, and, moreover, when $\mathbf{a} \notin \text{superset-Con}(q, I, \Sigma)$, a counterexample (i.e., an instance $J \models \Sigma$ such that $I \subseteq J$ and $\mathbf{a} \notin q(J)$) can be effectively constructed (this was shown in [5, 10] for guarded tgds, which form a proper generalization of LAV tgds). These facts immediately turn the construction described in the proof of Proposition 3.8 into an effective procedure.

In the remainder of the proof, we show that, for fixed Σ and n , we can compute n -universal superset-solutions in polynomial time.

We first consider instances consisting of a single fact (over \mathbf{R}). Let $I = \{R(\mathbf{a})\}$, where $R(\mathbf{a})$ is some fact over \mathbf{R} . By Proposition 3.8, I has an n -universal superset-repair, which we will denote by $J_{R(\mathbf{a})}$. For each fact $R(\mathbf{a})$, we choose some suitable such $J_{R(\mathbf{a})}$. There are only finitely many distinct facts over \mathbf{R} , up to isomorphism (i.e., up to a one-to-one renaming of values), and therefore the map that sends each fact $R(\mathbf{a})$ to the chosen instance $J_{R(\mathbf{a})}$ can be viewed as a finite object (by assuming that, whenever two facts f, f' are isomorphic, then the chosen instances J_f and $J_{f'}$ are isomorphic as well). We can think of it as a lookup table that will be used by the algorithm below to efficiently construct n -universal repairs for instances consisting of a single fact. Note that the lookup table does not depend on the input instance (it depends only on Σ , which is fixed) and hence constructing the lookup table is not part of the computation performed by the algorithm.

We now proceed with the description of the algorithm. Given as input an arbitrary instance I , the algorithm constructs, in polynomial time, an instance J , namely, the

union of all instances $J_{R(\mathbf{a})}$ where $R(\mathbf{a})$ is a fact belonging to I . We assume here that each instance $J_{R(\mathbf{a})}$ contains, besides \mathbf{a} , only fresh values not occurring in any of the other instances.

Clearly, we have that $I \subseteq J$. Since LAV tgds are closed under union, we have that $J \models \Sigma$. In what follows, we will show that, for every query q of size at most n , and for every tuple \mathbf{a} from the active domain of I , if $\mathbf{a} \in q(J)$ then $\mathbf{a} \in \text{superset-Con}(q, I, \Sigma)$ in I . By Lemma 4.5, we can construct, in polynomial time, a subinstance J' of J is a superset-repair of I . We will show that J' is in fact an n -universal superset-repair of I .

Let $\mathbf{a} \in q(J)$ where q is a conjunctive query of size at most n and \mathbf{a} is a tuple from the active domain of I . For the sake of a contradiction, assume that $\mathbf{a} \notin \text{superset-Con}(q, I, \Sigma)$. We may assume that q is a minimal size query for which this happens.

Let h be a witnessing homomorphism from q to J . We know that h is injective, for, otherwise, we could replace q by a smaller query. We also know that the image of h is entirely contained in a single part J_F of J (where F is a fact of I). This is because, otherwise, q could be divided into two subqueries, one whose image is contained in some J_F and the other containing the rest of the query q . One of these two parts would then be a smaller counterexample.

But, since the image of h is entirely contained in J_F , and we know that J_F is an n -universal superset-repair of $\{F\}$, we have that $\mathbf{a} \in \text{superset-Con}(q, \{F\}, \Sigma) \subseteq \text{superset-Con}(q, I, \Sigma)$. \square

The proof of Theorem 4.6 also shows that, given a set of LAV tgds Σ , a natural number $n \geq 1$, and an instance I , we can effectively compute an n -universal superset-repair of I w.r.t. Σ .

The main result of this section is now an immediate consequence of Proposition 3.2, Theorems 4.4, and 4.6.

Theorem 4.7 *For every set Σ of LAV tgds, for every conjunctive query q , and for $\star \in \{\oplus, \text{superset}, \text{subset}\}$, the \star -consistent query answering problem w.r.t. Σ is solvable in polynomial time.*

The Repair-Checking Problem We now consider the repair checking problem for sets of LAV tgds. First, we show how the \oplus -repair checking problem can be reduced to the subset-repair checking problem and the superset-repair checking problem.

Lemma 4.8 *Let Σ be a set of LAV tgds, and let I and J be two instances. Then the following statements are equivalent:*

1. J is a \oplus -repair of I w.r.t. Σ .
2. J is a superset-repair of $I \cap J$ w.r.t. Σ and J is a subset-repair of $I \cup J$ w.r.t. Σ .

Proof Let Σ be a set of LAV tgds, and let I and J be two instances. For the direction (1) \Rightarrow (2), assume that J is a \oplus -repair of I . First, we prove that J is a superset-repair of $I \cap J$. Let K_0 be an instance such that $K_0 \models \Sigma$ and $I \cap J \subseteq K_0 \subseteq J$. We have

to show that $K_0 = J$. As $I \cap J \subseteq K_0 \subseteq J$, we also have $I \oplus K_0 \subseteq I \oplus J$. Since $K_0 \models \Sigma$ and J is a \oplus -repair of I , this can only happen if $K_0 = J$.

Next we prove that J is a subset-repair of $I \cup J$. Let K_1 be an instance such that $K_1 \models \Sigma$ and $J \subseteq K_1 \subseteq I \cup J$. We have to show that $J = K_1$. Since $J \subseteq K_1 \subseteq I \cup J$, we have $I \oplus K_1 \subseteq I \oplus J$. Putting this together with the facts that $K_1 \models \Sigma$ and J is a \oplus -repair of I , we obtain $K_1 = J$.

For the direction (2) \Rightarrow (1), assume that J is a superset-repair of $I \cap J$ w.r.t. Σ and that J is a subset-repair of $I \cap J$ w.r.t. Σ . Let K be an instance such that $K \models \Sigma$ and $I \oplus K \subseteq I \oplus J$. We have to show that $K = J$. Since $I \oplus K \subseteq I \oplus J$, we have

$$J \cap I \subseteq K \cap I \text{ and } K \setminus I \subseteq J \setminus I.$$

As $K \setminus I \subseteq J \setminus I$, we have $K \cup J \subseteq I \cup J$. Moreover, since $K \models \Sigma$ and $J \models \Sigma$, it follows from Theorem 4.1 that $K \cup J \models \Sigma$. Recall that J is a subset-repair of $I \cup J$. Since $K \cup J \models \Sigma$ and $J \subseteq K \cup J \subseteq I \cup J$, this implies that $J = K \cup J$. Hence, $K \subseteq J$.

Recall that $J \cap I \subseteq K \cap I$. Putting everything together, we have $I \cap J \subseteq K \subseteq J$. Since $K \models \Sigma$ and J is a repair of $I \cap J$, this implies $K = J$. □

Theorem 4.9 *Let Σ be a set of LAV tgds, and let $\star \in \{\oplus, \text{superset}, \text{subset}\}$. The \star -repair checking problem w.r.t. Σ is solvable in polynomial time.*

Proof By Lemma 4.8, it suffices to give polynomial-time algorithms for the superset-repair checking problem and the subset-repair checking problem. The fact that the subset-repair checking problem is solvable in polynomial time follows directly from Theorem 4.4 (since it suffices to check for equality with the universal subset-repair). Finally, the fact that the superset-repair checking problem is solvable in polynomial time is a direct consequence of Lemma 4.5 □

Adding egds The tractability results concerning LAV tgds are optimal, in the sense that if we consider sets of LAV tgds and egds, then most of them do not remain true. We discuss first the repair-checking problem for sets of LAV tgds and egds. In [14], it was shown that there is a set consisting of a cyclic inclusion dependency and a functional dependency for which the subset-repair checking problem is **conP**-complete. In [2], the intractability of subset-repair checking (and \oplus -repair checking) was shown to hold for the union of an acyclic set of inclusion dependencies with a set of egds. However, the superset-repair checking problem can still be solved in polynomial time. This follows from the next lemma combined with Theorem 4.9.

Lemma 4.10 *Let Σ be a set of tgds and egds, and let I and J be two instances such that $I \subseteq J$ and $J \models \Sigma$. Then J is a superset-repair of I w.r.t. Σ if and only if J is a superset-repair of I w.r.t. the set Σ_1 consisting of all tgds in Σ .*

Proof Assume that $\Sigma = \Sigma_1 \cup \Sigma_2$, where Σ_1 is a set of tgds and Σ_2 is a set of egds. Let I and J be two instances such that $I \subseteq J$ and $J \models \Sigma_1 \cup \Sigma_2$. It is clear that if J is a superset-repair of I w.r.t. Σ_1 , then J is a superset-repair of I w.r.t. $\Sigma_1 \cup \Sigma_2$.

For the other direction, assume that J is a superset-repair of I w.r.t. $\Sigma_1 \cup \Sigma_2$. Towards a contradiction, assume that J is not a superset-repair of I w.r.t. Σ_1 . Hence, there is an instance J' such that $I \subseteq J' \subsetneq J$ and $J' \models \Sigma_1$. Since Σ_2 is a set of egds, $J \models \Sigma_2$ and $J' \subseteq J$, we also have that $J' \models \Sigma_2$. So, J' is a superset of I such that $J' \subseteq J$ and $J' \models \Sigma_1 \cup \Sigma_2$. This contradicts the fact that J is a superset-repair of I w.r.t. $\Sigma = \Sigma_1 \cup \Sigma_2$. \square

Next, we comment on the consistent query answering problem for sets of LAV tgds and egds. Theorem 4.7 in [14] shows that there is a set Σ of inclusion dependencies and functional dependencies, and a conjunctive query q such that computing the subset-consistent answers of q w.r.t. Σ is a Π_2^p -complete problem. Furthermore, as regards the combined complexity of consistent query answering, it follows from Theorem 5 in [31] that the following problem is undecidable: given a set Σ of LAV tgds and egds, a conjunctive query q and an instance I , compute the superset-consistent answers of q on I w.r.t. Σ . We leave it as an open problem whether or not this undecidability result still holds for some fixed conjunctive query q and some fixed set Σ of LAV tgds and egds. It is also unknown whether or not the data complexity of \oplus -consistent answers for conjunctive queries is decidable. As mentioned in the Introduction, the \oplus -consistent answering problem for unions of conjunctive queries and for sets of inclusion dependencies and functional dependencies was shown to be undecidable in combined complexity [11].

Finally, if we consider LAV tgds together with egds, there might not always exist a universal (or even 1-universal) subset-repair, \oplus -repair or superset-repair, as seen in Example 3.6.

5 GAV tgds

In this section, we investigate the existence and efficient computability of universal repairs for GAV tgds, we review known complexity results for repair checking and consistent query answering for GAV tgds, and we provide new matching lower bounds. Furthermore, toward the end of the section, we will briefly consider the addition of egds, and will also discuss semi-LAV sets of tgds; as mentioned earlier, semi-LAV sets of tgds were introduced in [24].

Recall that in the case of LAV tgds, every instance has a unique subset-repair, which is also the unique universal \oplus -repair. This is far from true for GAV tgds. First of all, Example 3.5 shows that there is a set Σ of GAV tgds and an instance I such that I has no 1-universal subset repair w.r.t. Σ ; therefore, Proposition 3.7 implies that I also has no 1-universal \oplus -repair w.r.t. Σ . In addition, even if an instance happens to have a universal subset-repair, this may not be a universal \oplus -repair, as revealed by the next example.

Example 5.1 Consider the set

$$\Sigma = \{P(x) \wedge Q(y) \rightarrow M(y), M(x) \rightarrow P(x)\}$$

and the instance $I = \{M(a), Q(b)\}$. The instance $J = \{Q(b)\}$ is then the only subset-repair of I , hence it is also a universal subset-repair of I . However, J is not a universal \oplus -repair of I , as can be seen by considering the \oplus -repair $J' = \{M(a), P(a)\}$ and the query $Q(x)$. This also shows that, unlike the case of LAV tgds, the subset-consistent answers of a conjunctive query w.r.t. a class of GAV tgds do not always coincide with the \oplus -consistent answers of the same query w.r.t. the same class of GAV tgds.

Unlike the case for subset-repairs and \oplus -repairs, we will show that, for sets of GAV tgds, every instance has a universal superset-repair. In fact, every instance has a *unique* superset-repair; moreover, the latter property is characteristic for GAV tgds (this will be analogous to the characterization of LAV tgds given in the previous section). To state and prove this characterization, we need a definition and a lemma. We say that a set Σ of dependencies is *closed under intersection* if for all instances I_1, I_2 such that $I_1 \models \Sigma$ and $I_2 \models \Sigma$, we have that $I_1 \cap I_2 \models \Sigma$. Closure under intersection will turn out to be equivalent to the existence of a unique superset-repair, and this characterizes the GAV tgds among all tgds.

Lemma 5.2 *Let Σ be a set of tgds and let I be an instance that has a unique superset-repair J w.r.t. Σ . If I' is an instance and J' is a superset-repair of I' w.r.t. Σ , then every homomorphism $h : I \rightarrow I'$ can be extended to a homomorphism $h' : J \rightarrow J'$.*

Proof Here, $\text{adom}(K)$ will denote the active domain of an instance K , while $\text{rng}(h)$ will denote the range of the function h . Let J'_h be the following instance:

- $\text{adom}(J'_h) = \text{adom}(I) \cup (\text{adom}(J') \setminus \text{rng}(h))$.
- For every fact $R(b_1, \dots, b_n)$ of J' , we populate J'_h with all possible facts that can be obtained by replacing each value b_i by a value $a_i \in \text{dom}(I)$ such that $h(a_i) = b_i$, if $b_i \in \text{rng}(h)$, or leaving b_i untouched, otherwise.

Since $h(I) \subseteq I' \subseteq J'$, we have that $I \subseteq J'_h$. Let h' be a function defined on $\text{adom}(J'_h)$ in the following way: if $a \in \text{dom}(I)$, then $h'(a) = h(a)$; if $a \notin \text{dom}(I)$ (but $a \in \text{adom}(J') \setminus \text{rng}(h)$), then $h'(a) = a$. From the construction of J'_h and the definition of h' , it follows that h' is a homomorphism from J'_h to J' . To complete the proof, it suffices to show that $J'_h \models \Sigma$. Indeed, once we have shown this, it will follow that $J \subseteq J'_h$, because $I \subseteq J'_h$ and J is the unique superset-repair of I ; consequently, the restriction of h' on $\text{adom}(J)$ is a homomorphism from J to J' .

To show that $J'_h \models \Sigma$, consider an arbitrary tgd from Σ , say,

$$\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y}).$$

Suppose that J'_h satisfies $\phi(\mathbf{a})$ for some tuple of values \mathbf{a} . Then, by the construction of J'_h , we have that J' satisfies $\phi(h'(\mathbf{a}))$, and hence, it also satisfies $\psi(h'(\mathbf{a}), \mathbf{b})$ for some values $\mathbf{b} = b_1, \dots, b_m$. Let $\mathbf{a}' = a'_1, \dots, a'_m$ be a tuple of values such that $h'(a'_i) = b_i$, for $1 \leq i \leq m$. Then, by construction, J'_h satisfies $\psi(\mathbf{a}, \mathbf{a}')$. □

Theorem 5.3 *Let Σ be a set of tgds. Then the following statements are equivalent:*

1. *Every instance has a unique (universal) superset-repair w.r.t. Σ .*
2. *Σ is closed under intersection.*
3. *Σ is logically equivalent to a set of GAV tgds.*

Proof The proof goes round-robin.

[1 \Rightarrow 3] Assume that every instance I has a unique superset-repair J . It follows that J can only contain values that come from the active domain of I (otherwise, by taking an isomorphic copy in which the values outside of the active domain of I are replaced by fresh values, we would refute the uniqueness of the repair). Consider a tgd $\tau \in \Sigma$ of the form $\phi(\mathbf{x}) \rightarrow \exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$, where ψ is a conjunction of atomic formulas. Let I_ϕ be the canonical instance of $\phi(\mathbf{x})$, and let J_ϕ be the unique superset-repair of I_ϕ . Since J_ϕ is a repair of I_ϕ , it satisfies the tgd τ . It follows that $\exists \mathbf{y} \psi(\mathbf{x}, \mathbf{y})$ is satisfied in J_ϕ under the natural assignment. Since J_ϕ contains only values that are from the active domain of I_ϕ , this means that J_ϕ satisfies $\psi(\mathbf{x}, \mathbf{x}')$ for some tuple of values \mathbf{x}' from the active domain of I_ϕ . Let $\widehat{\tau}$ be the dependency

$$\phi(\mathbf{x}) \rightarrow \psi(\mathbf{x}, \mathbf{x}'),$$

which can be equivalently written as a finite set of GAV tgds, and let $\widehat{\Sigma} = \{\widehat{\tau} \mid \tau \in \Sigma\}$. We claim that Σ and $\widehat{\Sigma}$ are logically equivalent. The fact that Σ' logically implies Σ follows directly from the construction of Σ' . For the other direction, suppose that $I \models \Sigma$ and $I \models \phi(\mathbf{a})$ for some tuple of values \mathbf{a} . Let I_ϕ be again the canonical instance of $\phi(\mathbf{x})$. Then the function $h : \mathbf{x} \mapsto \mathbf{a}$ is a homomorphism from I_ϕ to I . Since J_ϕ is the unique superset-repair of I_ϕ w.r.t. Σ , Lemma 5.2 implies that h extends to a homomorphism from J_ϕ to every superset-repair of I . Since I satisfies Σ , it is its own repair, which means that I satisfies $\psi(\mathbf{a}, h(\mathbf{a}))$. Therefore, $I \models \widehat{\tau}$.

[3 \Rightarrow 2] Let Σ be a set of GAV tgds. Let $I_1 \models \Sigma$ and $I_2 \models \Sigma$, and suppose that, for some GAV tgd $t \in \Sigma$ of the form

$$\phi(x_1, \dots, x_n) \rightarrow R(x_{i_1}, \dots, x_{i_k})$$

it holds that $I_1 \cap I_2 \models \phi(a_1, \dots, a_n)$ for some values a_1, \dots, a_n . Then both I_1 and I_2 satisfy $\phi(a_1, \dots, a_n)$; since I_1 and I_2 satisfy Σ , we have that I_1 and I_2 satisfy $R(a_{i_1}, \dots, a_{i_k})$; thus, the fact $R(a_{i_1}, \dots, a_{i_k})$ belongs to the intersection of I_1 and I_2 .

[2 \Rightarrow 1] Let Σ be a set of tgds that is closed under intersection. Let I be an instance and let J be the instance consisting of all possible facts over the active domain of I . Then $I \subseteq J$ and $J \models \Sigma$, hence a superset-repair of I must exist. Suppose that I has two distinct superset-repairs J_1 and J_2 . By closure under intersection, $J_1 \cap J_2 \models \Sigma$, hence the properties of superset-repairs imply that $J_1 = J_1 \cap J_2 = J_2$, a contradiction. \square

The classical chase procedure from dependency theory (see [1]) provides a method for efficiently computing the unique superset-repair of an instance w.r.t. a set of GAV tgds.

Theorem 5.4 *Let Σ be a set of GAV tgds. There is a polynomial time algorithm that, given an instance I , computes the unique (universal) superset-repair of I w.r.t. Σ .*

We now move to consistent query answering. Let Σ be a set of GAV tgds and let q be a conjunctive query. The preceding Theorem 5.4 implies immediately that the superset-consistent query answering problem for q w.r.t. Σ is solvable in polynomial time. Staworko and Chomicki [32] showed that both the subset-consistent query answering problem and the \oplus -consistent query answering problem for q w.r.t. Σ are in **coNP**. This will also follow from Theorem 6.2 in the next section, which implies that the size of every \oplus -repair of an instance I w.r.t. a set of GAV tgds is bounded by a polynomial in the size of I . In [32], it is also shown that there is a conjunctive query q and a set Σ consisting of GAV tgds and egds such that the \oplus -consistent query answering problem for q w.r.t. Σ is **coNP**-complete. Our next result improves on this lower bound by showing that it can hold even for a set consisting of a single GAV tgd.

Theorem 5.5 *There is a set Σ consisting of a single GAV tgd and there is a conjunctive query q such that both the subset-consistent query answering problem and the \oplus -consistent query answering problem for q w.r.t. Σ are **coNP**-complete.*

Proof We produce a GAV tgd τ , a conjunctive query q , and a polynomial-time reduction from the complement of **POSITIVE 1-IN-3-SAT** to the problem of finding the \oplus -consistent answers of q w.r.t. $\{\tau\}$. Recall that **POSITIVE 1-IN-3-SAT** is the following NP-complete problem [22]: given a Boolean formula ϕ in conjunctive normal form and such that each clause is a disjunction of the form $(x_1 \vee x_2 \vee x_3)$ of three positive literals, is there a truth assignment that makes true exactly one variable in every clause?

Let τ be the following GAV tgd:

$$\forall x, u, u' (P(x, u) \wedge P(x, u') \rightarrow E(u, u'))$$

and let q be the following conjunctive query:

$$\begin{aligned} \exists x_1, x_2, x_3, u_1, u_2, u_3 & (R(x_1, x_2, x_3) \wedge P(x_1, u_1) \\ & \wedge P(x_2, u_2) \wedge P(x_3, u_3) \wedge S(u_1, u_2, u_3)). \end{aligned}$$

The intuition behind the relation symbols is as follows: P encodes truth values for the variables in a given Boolean formula, while E is used to simulate equality. The tgd τ expresses that each variable is assigned at most one truth value.

The relation R encodes every clause $(x_1 \vee x_2 \vee x_3)$ occurring in a formula in conjunctive normal form such that each clause is a disjunction of three positive literals. The relation S will consist of the triples in $\{0, 1\}^3 \setminus \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. The conjunctive query q expresses that there are a clause $(x_1 \vee x_2 \vee x_3)$ and truth values assigned to x_1, x_2 and x_3 such that it is not the case that exactly one variable is true in the clause $(x_1 \vee x_2 \vee x_3)$.

Let ϕ be a formula of the form

$$(x_{11} \vee x_{12} \vee x_{13}) \wedge \dots \wedge (x_{n1} \vee x_{n2} \vee x_{n3}),$$

where for all $1 \leq i \leq n$ and for all $1 \leq j \leq 3$, x_{ij} is a variable. We denote by Var the set $\{x_{ij} \mid 1 \leq i \leq n, 1 \leq j \leq 3\}$. We construct an instance I such that

$$\oplus\text{-Con}(q, I, \Sigma) = \text{true} \quad \text{iff} \quad \phi \notin \text{POSITIVE 1-IN-3-SAT}. \tag{1}$$

The instance I is defined as follows:

$$\begin{aligned} R^I &= \{(x_{i1}, x_{i2}, x_{i3}) \mid 1 \leq i \leq n\}, \\ P^I &= \{(x_{ij}, 0), (x_{ij}, 1) \mid 1 \leq i \leq n, 1 \leq j \leq 3\}, \\ E^I &= \{(1, 1), (0, 0)\}, \\ S^I &= \{0, 1\}^3 \setminus \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}. \end{aligned}$$

In order to prove the implication from left to right of (1), suppose that $\oplus\text{-Con}(q, I, \Sigma) = \text{true}$. Let $V : Prop \rightarrow \{0, 1\}$ be an assignment. We have to show that it is not the case that exactly one variable is true in every clause. That is, there exists a conjunct

$$x_{i1} \vee x_{i2} \vee x_{i3}$$

of ϕ such that $(V(x_{i1}), V(x_{i2}), V(x_{i3}))$ does not belong to $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

We define an instance J_0 in the following way:

$$\begin{aligned} R^{J_0} &= \{(x_{i1}, x_{i2}, x_{i3}) \mid 1 \leq i \leq n\}, \\ P^{J_0} &= \{(x_{ij}, V(x_{ij})) \mid 1 \leq i \leq n, 1 \leq j \leq 3\}, \\ E^{J_0} &= \{(1, 1), (0, 0)\}, \\ S^{J_0} &= \{0, 1\}^3 \setminus \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}. \end{aligned}$$

We can check that J_0 is a \oplus -repair of I with respect to Σ . Since the \oplus -consistent answer of q is true, q is true in J_0 . Hence, there are x_1, x_2, x_3, u_1, u_2 and u_3 such that

$$R(x_1, x_2, x_3) \wedge P(x_1, u_1) \wedge P(x_2, u_2) \wedge P(x_3, u_3) \wedge S(u_1, u_2, u_3)$$

holds in J_0 . By definition of R^{J_0} , $R^{J_0}(x_1, x_2, x_3)$ implies that $x_1 \vee x_2 \vee x_3$ is a conjunct of ϕ . It also follows from the definition of P^{J_0} that for all $1 \leq i \leq 3$, $P^{J_0}(x_i, u_i)$ implies $u_i = V(x_i)$. Next, by definition of S^{J_0} , $S^{J_0}(u_1, u_2, u_3)$ means that $(u_1, u_2, u_3) \notin \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

Putting everything together, we have a conjunct $x_1 \vee x_2 \vee x_3$ of ϕ such that $(V(x_{i1}), V(x_{i2}), V(x_{i3}))$ does not belong to $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. This finishes the proof of the implication from left to right of (1).

Next we show the implication from right to left of (1). Suppose that ϕ does not belong to POSITIVE 1-IN-3-SAT. Let J be a \oplus -repair of I .

We start by proving that $R^J = R^I$ and $S^J = S^I$. Consider the instance J_0 obtained by replacing R^J by R^I and S^J by S^I . Since t is true in J and neither R nor S occurs in Σ , t is true in K . By definition of J_0 , $I \oplus J_0 \subseteq I \oplus J$. Since J_0 is a \oplus -repair, this implies that J_0 is equal to J . Therefore, $R^J = R^I$ and $S^J = S^I$.

Since $R^J = R^I$ and $S^J = S^I$, it follows from the definitions of R^I, S^I and q that (\dagger) q is true in J iff there are a conjunct $x_{i1} \vee x_{i2} \vee x_{i3}$ of ϕ and u_1, u_2, u_3 such that $P^J(x_{i1}, u_1), P^J(x_{i2}, u_2), P^J(x_{i3}, u_3)$ and $(u_1, u_2, u_3) \notin \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$.

Claim 1 One of the followings holds:

- (a) For all $x \in Prop$ and for all $u \in \{0, 1\}$, $P(x, u)$ belongs to J .
- (b) For all $x \in Prop$, there is a unique $u \in \{0, 1\}$ such that $P^J(x, u)$ holds.

□

Proof of Claim We start by proving that

$$\forall x \in Prop, \forall u, u'(P^J(x, u) \text{ or } E^J(u, u')) \Rightarrow u, u' \in \{0, 1\} \tag{2}$$

Consider the instance J_1 obtained from J by deleting the tuples of the form $P^J(x, u)$, $E^J(u, u')$ and $E^J(u', u)$, where $x \in Prop$ and $u \notin \{0, 1\}$. The $\text{tgd } t$ is satisfied in J_1 . It also follows from the definitions of I and J_1 that $I \oplus J_1 \subseteq I \oplus J$. Since J is a \oplus -repair, this means that $J_1 = J$. That is, for all u, u' such that $P^J(x, u)$ or $E^J(u, u')$ or $E^J(u', u)$ holds, u belongs to $\{0, 1\}$.

Next we show that

$$\{(0, 0), (1, 1)\} \subseteq E^J. \tag{3}$$

Consider the instance J_2 obtained from J by adding the tuples $E(0, 0)$ and $E(0, 1)$ to J . We can check that J_2 satisfies t . Moreover, by definitions of I and J_2 , we have $I \oplus J_2 \subseteq I \oplus J$. Since J is a \oplus -repair, $J_2 = J$. That is, $E(0, 0)$ and $E(0, 1)$ belong to J .

Putting (2) and (3) together, we obtain that

$$\{(0, 0), (1, 1)\} \subseteq E^J \subseteq \{(0, 0), (1, 1), (0, 1), (1, 0)\}.$$

We make the following case distinction:

- Suppose that $E^J = \{(0, 0), (1, 1)\}$. We show that (b) holds. That is, for all $x \in Prop$, there is a unique $u \in \{0, 1\}$ such that $P^J(x, u)$ holds. Since $E^J = \{(0, 0), (1, 1)\}$ and the constraint

$$\forall x, u, u'(P(x, u) \wedge P(x, u') \rightarrow E(u, u'))$$

is true in J , we have that for all $x \in Prop$, there exists at most one element $u \in \{0, 1\}$ such that $P^J(x, u)$.

Hence, in order to prove (b), it remains to show that for all $x \in Prop$, there exists $u \in \{0, 1\}$ such that $P^J(x, u)$ holds. Let x be an element of $Prop$. Suppose for contradiction that there is no $u \in \{0, 1\}$ such that $P^J(x, u)$. By (2), this implies that there is no u such that $P^J(x, u)$. Let J_3 be the instance obtained by adding the tuple $P^J(x, 0)$ to J . We prove that t is true in J_3 . Suppose that $P^{J_3}(y, v)$ and $P^{J_3}(y, v')$ hold. We have to show $E^{J_3}(v, v')$. If $y \neq x$, this follows from the definition of J_3 and the fact that t is true in J .

Next suppose that $y = x$. That is, $P^{J_3}(x, v)$ and $P^{J_3}(x, v')$ hold. Recall that there is no u such that $P^J(x, u)$. Hence, by definition of J_3 , $P^{J_3}(x, v)$ and $P^{J_3}(x, v')$ imply $v = v' = 0$. Since $E^J = \{(0, 0), (1, 1)\}$ and $E^J = E^{J_3}$, we have $E^{J_3}(v, v')$. This finishes the proof that t is true in J_3 .

Next, it follows from the definition of J_3 that $I \oplus J_3 \subseteq I \oplus J$. Since J is a repair, $J_3 = J$. This contradicts our assumption that there is no $u \in \{0, 1\}$ such that $P^J(x, u)$. This finishes the proof that (b) holds.

- Next suppose that $(0, 1)$ or $(1, 0)$ belongs to E^J . Without loss of generality, we may assume that $(0, 1)$ belongs to E^J .

There must be an element x_0 such that $P^J(x_0, 0)$ and $P^J(x_0, 1)$ hold. Indeed, if there is no such element, then we can let J_4 be the instance obtained by removing the tuple $E(0, 1)$ from J . Then, since there is no x_0 such that $P^J(x_0, 0)$ and $P^J(x_0, 1)$, t is satisfied in J_4 . We also have that $I \oplus J_4 \subseteq J \oplus J$. As J is a \oplus -repair, this implies that $J = J_4$. That is, $E(0, 1)$ does not belong to J , which is a contradiction. Hence, there is x_0 such that $P^J(x_0, 0)$ and $P^J(x_0, 1)$.

Now, since $P(x_0, 1) \wedge P(x_0, 0)$ holds in J and the tgds

$$\forall x, u, u'(P(x, u) \wedge P(x, u') \rightarrow E(u, u'))$$

is satisfied in J , we have that $E(1, 0)$ belongs to J . Therefore, $E^J = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$.

We show that for all $x \in \text{Prop}$ and for all $u \in \{0, 1\}$, $P(x, u)$ belongs to J . Take $x \in \text{Prop}$ and $u \in \{0, 1\}$. Define J_5 as the instance obtained by adding the tuple $P(x, u)$ to J . Since t is true in J and $E^J = \{(0, 0), (1, 0), (0, 1), (1, 1)\}$, t remains true in J_5 . It is clear that $I \oplus J_5 \subseteq I \oplus J$. As J is a \oplus -repair, this implies that $J = J_5$. In other words, $P(x, u)$ belongs to J . This finishes the proof of the claim. □

Suppose first that (a) holds. That is, for all $x \in \text{Prop}$ and for all $u \in \{0, 1\}$, $P(x, u)$ belongs to J . We show that q is true in J . Let $x_{i1} \vee x_{i2} \vee x_{i3}$ be a conjunct of ϕ . Since (a) holds, $P(x_{i1}, 0)$, $P(x_{i2}, 0)$ and $P(x_{i3}, 0)$ belong to J . By (\dagger) , this implies that q is true in J .

Next suppose that (b) holds. That is, for all $x \in \text{Prop}$, there is a unique $u \in \{0, 1\}$ such that $P^J(x, u)$ holds. Hence, we may choose a valuation $V : \text{Prop} \rightarrow \{0, 1\}$ such that for all $x \in \text{Prop}$ and for all $u \in \{0, 1\}$,

$$V(x) = u \quad \text{iff} \quad P^J(x, u).$$

Since ϕ does not belong to POSITIVE 1-IN-3-SAT, it is not the case that exactly one variable is true in every clause. Therefore, there exists a conjunct $x_{i1} \vee x_{i2} \vee x_{i3}$ of ϕ such that

$$(V(x_{i1}), V(x_{i2}), V(x_{i3})) \notin \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}.$$

Moreover, by definition of V , $P(x_{i1}, V(x_{i1}))$, $P(x_{i2}, V(x_{i2}))$ and $P(x_{i3}, V(x_{i3}))$ belong to J . It follows from (\dagger) that q is true in J .

Extensions: egds and Semi-LAV All complexity upper bounds described in this section hold for the more general case of sets of GAV tgds and egds. Note that if Σ is a set of GAV tgds and egds, then an instance I may not have any superset-repair w.r.t. Σ . Still, if a superset-repair exists, then it is unique, and it can be computed in polynomial time using the chase procedure. The existence of a superset-repair can be tested in polynomial time as well (c.f. Theorem 6.1 below).

In [24], the class of *semi-LAV* sets of dependencies was introduced; it contains properly the class of all sets of GAV tgds, as well as the class of all weakly acyclic sets of LAV tgds. It was shown in [24] that the \oplus -repair checking problem for semi-LAV

sets of tgds is still solvable in polynomial time (however, this is no longer true if egds are allowed [2]). For \star -consistent query answering, where $\star \in \{\text{subset, superset, } \oplus\}$, the complexity for semi-LAV sets of tgds is the same as the complexity for sets of GAV tgds. The lower bounds naturally transfer, and the upper bounds follow from the fact that repair checking is in polynomial time, together with Theorems 6.1 and 6.2 in the next section, since every semi-LAV set of tgds is, by definition, weakly acyclic.

6 Weakly Acyclic Sets of tgds (and egds)

In this section, we study the consistent query answering problem for conjunctive queries and weakly acyclic sets of tgds and egds. We begin by considering universal superset-repairs. Recall that, according to Proposition 3.7, every instance has at most one universal superset-repair, up to isomorphism.

Theorem 6.1 *Let Σ be a weakly acyclic set of tgds and egds. If an instance has a superset-repair w.r.t. Σ , then it has a universal superset-repair. Moreover, there is a polynomial time algorithm that, given an instance I , tests whether it has a superset-repair w.r.t. Σ and if so, computes the (unique up to isomorphism) universal superset-repair of I .*

Proof We rely on known results from data exchange. We say that an instance J is a solution for an instance I w.r.t. Σ if $I \subseteq J$ and J satisfies Σ . Thus, a superset-repair of I is a solution of I such that no strict subset is a solution for I . In [23], it was shown that if Σ is a fixed weakly acyclic set of tgds and egds, then there is a polynomial-time algorithm that, given an instance I , tests whether it has a solution w.r.t. Σ , and if so, computes a solution J satisfying the following additional properties:

- (i) For each solution J' of I , there is a homomorphism $h : J \rightarrow J'$ that is the identity on values from the active domain of I .
- (ii) For each solution J'' satisfying the above condition (i), we have that $J \subseteq J''$.

This solution J is known as the “core universal solution” of I [18], and is unique up to isomorphism.

Let J be the core universal solution of I w.r.t. Σ . Since $I \subseteq J$ and J satisfies Σ , it remains to show that (a) J is a superset-repair of I , i.e., there is no $J' \subsetneq J$ that contains I and satisfies Σ , and (b) J is a *universal* superset-repair of I . The first item follows immediately from the above properties (i) and (ii): any such J' would satisfy condition (i), thereby contradicting the fact that J satisfies condition (ii) above. The second item follows from the fact that J satisfies condition (i) and from the preservation of conjunctive queries under homomorphisms. \square

Recall from Example 3.5 that, in general, an instance might not have a 1-universal subset-repair (hence, a 1-universal \oplus -repair) w.r.t. a weakly acyclic set of tgds.

Next, we move to the consistent query answering problem. One of the key observations for obtaining an upper bound for the complexity of the \oplus -consistent answers

is that the size of every \oplus -repair of an instance I is polynomial in the size of I . The proof relies on the *solution-aware chase*, which was introduced in the context of peer data exchange [32].

Theorem 6.2 *Let Σ be a weakly acyclic set of tgds and egds. There is a polynomial $p(x)$ such that for all instances I and for all \oplus -repairs J of I with respect to Σ , the size of J is bounded by $p(|I|)$, where $|I|$ is the size of I .*

Proof Let Σ be a weakly acyclic set of tgds and egds. The proof relies on the algorithm described in the proof of Lemma 3.4 in [20]. Specifically, the proof of that lemma yields a polynomial $p(x)$ and an algorithm that works as follows. Given instances K and L such that $K \subseteq L$ and $L \models \Sigma$, the algorithm computes an instance $f(K, L)$ such that $K \subseteq f(K, L) \subseteq L$, $f(K, L) \models \Sigma$, and the size of $f(K, L)$ is bounded by $p(|K|)$.

Fix an \oplus -repair I_0 of I . We show that the size of I_0 is bounded by $p(|I|)$. We can use the algorithm described in Lemma 3.4 in [20] to produce an instance $J := f(I \cap I_0, I_0)$ satisfying the following. The size of J is bounded by $p(|I \cap I_0|)$, $I \cap I_0 \subseteq J \subseteq I_0$ and $J \models \Sigma$. In particular, the size of J is bounded by $p(|I|)$.

In order to show that the size of I_0 is bounded by $p(|I|)$, it is sufficient to show that $I_0 = J$. Since I_0 is a \oplus -repair of I , this is equivalent to proving that $J \models \Sigma$ and $I \oplus J \subseteq I \oplus I_0$.

We already know that $J \models \Sigma$. It remains to show that $I \oplus J \subseteq I \oplus I_0$. Let $R(\mathbf{a})$ be a fact in $I \oplus J$. We prove that $R(\mathbf{a})$ belongs to $I \oplus I_0$. Suppose first that $R(\mathbf{a})$ belongs to $J \setminus I$. We know that $J \subseteq I_0$. Hence, if $R(\mathbf{a})$ belongs to $J \setminus I$, then $R(\mathbf{a})$ belongs to $I_0 \setminus I$. This implies that $R(\mathbf{a})$ belongs to $I \oplus I_0$. Next, assume that $R(\mathbf{a})$ belongs to $I \setminus J$. We have to show that $R(\mathbf{a}) \in I \oplus I_0$. Since $R(\mathbf{a})$ belongs to I , this means that we have to prove that $R(\mathbf{a}) \notin I_0$. Suppose for contradiction that $R(\mathbf{a})$ belongs to I_0 . Then $R(\mathbf{a})$ belongs to $I \cap I_0$. Recall that $I \cap I_0 \subseteq J$. Putting this together with $R(\mathbf{a}) \in I \cap I_0$, we obtain that $R(\mathbf{a})$ belongs to J , which is a contradiction. \square

Our main result concerning consistent query answering w.r.t. a weakly acyclic set of tgds is a Π_2^P lower bound both for subset-repairs and for \oplus -repairs. A Π_2^P lower bound had been obtained earlier for the problem of finding the \oplus -consistent answers w.r.t. a set of functional dependencies and universal constraints [32].

Theorem 6.3 *Let Σ be a weakly acyclic set of tgds and egds and let q be a conjunctive query.*

1. *The superset-consistent query answering problem for q w.r.t. Σ is in PTIME.*
2. *The \oplus -consistent (subset-consistent) query answering problem for q w.r.t. Σ is in Π_2^P .*
3. *There is a set Σ of weakly acyclic tgds and a conjunctive query q such that both the \oplus -consistent and the subset-consistent query answering problems for q w.r.t. Σ are Π_2^P -complete.*

Proof Part 1 is an immediate consequence of Theorem 6.1. It also follows from Proposition 2.8 and results in [17] concerning the tractability of the certain answers of conjunctive queries in data exchange w.r.t. weakly acyclic sets of tgds and egds. Part 2 follows from Theorem 6.2 and the fact that both the \oplus -repair and the subset-repair checking problems with respect to any set of tgds and egds is in **CONP**. For Part 3, we give a weakly acyclic set Σ of tgds and a Boolean conjunctive query q with the following properties: for every quantified Boolean formula ϕ of the form $\forall p_1 \dots \forall p_n \exists q_1 \dots \exists q_k \psi$, where ψ is a conjunction of clauses containing 3 literals, we can construct in polynomial time an instance I_ϕ so that the following statements are equivalent:

1. ϕ is true ;
2. $\text{subset-Con}(q, I_\phi, \Sigma) = \top$;
3. $\oplus\text{-Con}(q, I_\phi, \Sigma) = \top$.

Evaluating such formulas is known to be a Π_2^P -complete problem [30].

We introduce the following tgds:

$$\begin{aligned} \phi_1 &= Q(q, v) \rightarrow \exists s A(s), \\ \phi_2 &= A(s) \wedge Q'(q) \rightarrow \exists v Q(q, v), \\ \phi_3 &= Q(q, v) \wedge Q(q, v') \rightarrow E(v, v'), \\ \phi_4 &= P(p, v) \wedge P(p, v') \rightarrow E(v, v'), \\ \phi_{a_1 a_2 a_3}(X_1, X_2, X_3) &= R_{a_1 a_2 a_3}(x_1, x_2, x_3) \wedge A(s) \wedge \\ &X'_1(x_1) \wedge X'_2(x_2) \wedge X'_3(x_3) \rightarrow \\ &\exists v_1, v_2, v_3 (X_1(x_1, v_1) \wedge X_2(x_2, v_2) \\ &\wedge X_3(x_3, v_3) \wedge T_{a_1 a_2 a_3}(v_1, v_2, v_3)), \end{aligned}$$

where $a_i \in \{0, 1\}$ and $X_i \in \{P, Q\}$, for $i = 1, 2, 3$. We now let

$$\begin{aligned} \Sigma &= \{\phi_1, \phi_2, \phi_3, \phi_4\} \cup \\ &\{\phi_{a_1 a_2 a_3}(X_1, X_2, X_3) \mid a_i \in \{0, 1\}, X_i \in \{P, Q\}\}. \end{aligned}$$

We also let q be the query $\exists s A(s)$. Note that Σ is a weakly acyclic set. Indeed, the only special edges are from positions in Q' , P' and R ; however, there is no incoming edge to a position in Q' , P' or R .

The intuition behind the relation symbols is as follows. The relation P' encodes the universally quantified variables p_1, \dots, p_n of ϕ , while Q' encodes the existentially quantified variables q_1, \dots, q_k of ϕ . Furthermore, P encodes truth values for p_1, \dots, p_n , while Q encodes truth values for q_1, \dots, q_k . We use the symbol E to simulate equality; it will consist of the tuples $(0, 0)$ and $(1, 1)$. The tgds ϕ_3 and ϕ_4 express that each variable gets assigned at most one truth value.

In what follows, we will use x, x_1, x_2, x_3, \dots to denote literals. If x is a positive literal, we define the *inverse sign* of x , denoted by $is(x)$, to be 0. If x is a negative literal, we define $is(x)$ to be 1. Thus, $is(x)$ denotes the opposite of the sign of the literal x . For $a_1, a_2, a_3 \in \{0, 1\}$, we will use a relation $R_{a_1 a_2 a_3}$ to encode clauses of the form $x_1 \vee x_2 \vee x_3$ such that $is(x_i) = a_i$. The relation $T_{a_1 a_2 a_3}$ consists of those truth assignments that make true the clauses of the form $x_1 \vee x_2 \vee x_3$, where $is(x_i) = a_i$.

The symbol A acts as guard. It is activated (that is, becomes non-empty) as soon as one variable in $\{q_1, \dots, q_k\}$ gets assigned a truth value (see the tgd ϕ_1). Once A has been activated, it makes sure that each variable in $\{q_1, \dots, q_k\}$ is assigned a truth value (as expressed by the tgd ϕ_2). Hence, either no variable in $\{q_1, \dots, q_k\}$ gets assigned a truth value or all variables in $\{q_1, \dots, q_k\}$ get assigned a truth value. Moreover, when A is non-empty, the tgds of the form $\phi_{a_1 a_2 a_3}(X_1, X_2, X_3)$ express that the truth assignment is such that each clause of the form $x_1 \vee x_2 \vee x_3$ occurring in ϕ is true under the truth assignment.

Consider a quantified Boolean formula ϕ of the form

$$\forall p_1, \dots, p_n \exists q_1, \dots, q_k \psi,$$

where ψ is a conjunction of clauses of the form

$$x_{i1} \vee x_{i2} \vee x_{i3},$$

with $1 \leq i \leq l$ and $x_{ij} \in \{p_a, \neg p_a, q_b, \neg q_b \mid 1 \leq a \leq n, 1 \leq b \leq k\}$.

We construct now an instance I_ϕ so that the following statements are equivalent:

1. ϕ is true ;
2. $\text{subset-Con}(q, I_\phi, \Sigma) = \top$;
3. $\oplus\text{-Con}(q, I_\phi, \Sigma) = \top$.

We define $I_\phi = I$ by

$$\begin{aligned} A^I &= \{1\}, \\ (Q')^I &= \{q_i \mid 1 \leq i \leq k\}, \\ (P')^I &= \{p_i \mid 1 \leq i \leq n\}, \\ P^I &= \{(p_i, v) \mid 1 \leq i \leq n, v \in \{0, 1\}\}, \\ Q^I &= \{(q_i, v) \mid 1 \leq i \leq k, v \in \{0, 1\}\}, \\ E^I &= \{(0, 0), (1, 1)\}, \\ T_{a_1 a_2, a_3}^I &= \{(v_1, v_2, v_3) \mid \forall 1 \leq i \leq 3, v_i \in \{0, 1\}, \\ &\quad (v_1, v_2, v_3) \neq (a_1, a_2, a_3)\}, \\ R_{a_1 a_2 a_3}^I &= \{(x_1, x_2, x_3) \mid x_1 \vee x_2 \vee x_3 \text{ is a clause of } \psi \text{ and} \\ &\quad \forall 1 \leq i \leq 3, is(x_i) = a_i\}. \end{aligned}$$

In order to prove that the instance I has the required properties, it suffices to show that

$$\text{subset-Con}(q, I, \Sigma) = \top \Rightarrow \phi \text{ is true,} \tag{4}$$

$$\phi \text{ is true} \Rightarrow \oplus\text{-Con}(q, I, \Sigma) = \top. \tag{5}$$

We start by showing (4).

Claim 2 If the subset-consistent answer of q is true, then ϕ is true. □

Proof of Claim Suppose that the consistent answer of q is true and let $V_P : \{p_i \mid 1 \leq i \leq n\} \rightarrow \{0, 1\}$ be a valuation. We have to find a valuation $V_Q : \{q_i \mid 1 \leq i \leq k\} \rightarrow$

$\{0, 1\}$ such that ψ is true under $V_P \cup V_Q$. We define an instance J in the following way:

$$\begin{aligned} A^J &= \emptyset, \\ (Q')^J &= (Q')^I, \\ (P')^J &= (P')^I, \\ P^J &= \{(p_i, V_P(p_i)) \mid 1 \leq i \leq n\}, \\ Q^J &= \emptyset, \\ E^J &= E^I, \\ T_{a_1 a_2 a_3}^J &= T_{a_1 a_2 a_3}^I, \\ R_{a_1 a_2 a_3}^J &= R_{a_1 a_2 a_3}^I. \end{aligned}$$

The formulas of Σ are satisfied in J . Hence, there is a subset-repair J_0 of I such that $I \oplus J_0 \subseteq I \oplus J$. This implies that $(Q')^{J_0} = (Q')^I$, $(P')^{J_0} = (P')^I$, $E^{J_0} = E^I$, $T^{J_0} = T^I$, $R^{J_0} = R^I$, $Q^{J_0} \subseteq Q^I$ and $P^J \subseteq P^{J_0}$.

As the subset-consistent answer of q is true, we have $A^{J_0} \neq \emptyset$. Take $i \in \{1, \dots, k\}$. Since $A^{J_0} \neq \emptyset$, $(Q')^{J_0}(q_i)$ holds and ϕ_{2i} is true in J_0 , there exists v_i such that $Q^{J_0}(q_i, v_i)$ holds. As $Q^{J_0} \subseteq Q^I$, v_i belongs to $\{0, 1\}$.

We define a valuation $V_Q : \{q_i \mid 1 \leq i \leq k\} \rightarrow \{0, 1\}$ such that for all $i \in \{1, \dots, k\}$, we have

$$V_Q(q_i) = v_i.$$

We show that ψ is true under $V := V_P \cup V_Q$. Consider a clause $x_{i1} \vee x_{i2} \vee x_{i3}$ of ψ . For all $1 \leq j \leq 3$, define a_j by $is(x_{ij})$. By definition of R^I , we have $R_{a_1 a_2 a_3}^I(x_{i1}, x_{i2}, x_{i3})$. Since $R_{a_1 a_2 a_3}^I = R_{a_1 a_2 a_3}^{J_0}$, we also have $R_{a_1 a_2 a_3}^{J_0}(x_{i1}, x_{i2}, x_{i3})$.

If $j \in \{1, 2, 3\}$ and $x_{ij} = p_l$ for some $1 \leq l \leq n$, we define X_j as P . Otherwise, we define X_j as Q . By definition of $(P')^J$ and $(Q')^J$, we have $(X'_j)^J(x_{ij})$. Since $(X'_j)^J \subseteq (X'_j)^{J_0}$, this implies $(X'_j)^{J_0}(x_{ij})$.

Putting everything together we obtain that

$$R_{a_1 a_2 a_3}^{J_0}(x_{i1}, x_{i2}, x_{i3}) \wedge (X'_1)^{J_0}(x_{i1}) \wedge (X'_2)^{J_0}(x_{i2}) \wedge (X'_3)^{J_0}(x_{i3}). \tag{6}$$

Since Σ is true in J_0 , $\phi_{a_1 a_2 a_3}(X_1, X_2, X_3)$ holds in J_0 . Putting this together with (6) and the fact that $A^{J_0} \neq \emptyset$, there exists v_1, v_2 and v_3 such that

$$X_1^{J_0}(x_{i1}, v_1) \wedge X_2^{J_0}(x_{i2}, v_2) \wedge X_3^{J_0}(x_{i3}, v_3) \wedge T_{a_1 a_2 a_3}^{J_0}(v_1, v_2, v_3).$$

We show that for all $1 \leq j \leq 3$,

$$v_j = V(x_{ij}). \tag{7}$$

Take $j \in \{1, 2, 3\}$. Suppose first that $X_j = Q$. Since ϕ_{3i} is true in J_0 and $E^{J_0} = \{(0, 0), (1, 1)\}$, there is a unique v_j such that $Q^{J_0}(x_{ij}, v_j)$ and $v_j \in \{0, 1\}$. By definition of V , we have $Q^{J_0}(x_{ij}, V(x_{ij}))$. Hence, $v_j = V(x_{ij})$. The case where $X_j = P$ is similar and this finishes the proof of (7).

Since $T_{a_1 a_2 a_3}^{J_0}(v_1, v_2, v_3)$ and $T_{a_1 a_2 a_3}^{J_0} = T_{a_1 a_2 a_3}^I$, we have $(v_1, v_2, v_3) \neq (a_1, a_2, a_3)$. It follows from (7) that

$$(V(x_{i1}), V(x_{i2}), V(x_{i3})) \neq (a_1, a_2, a_3).$$

By definition a_1, a_2 and a_3 , this means

$$(V(x_{i1}), V(x_{i2}), V(x_{i3})) \neq (is(x_1), is(x_2), is(x_3)).$$

Together with the definition of the inverse sign is , we obtain that $x_{i1} \vee x_{i2} \vee x_{i3}$ is true under the valuation V . □

Next we prove that (5) holds.

Claim 3 If ϕ is true, then the \oplus -consistent answer of q is true.

Proof of Claim Assume that ϕ is true and let J be a \oplus -repair of I . Suppose for contradiction that $A^J = \emptyset$. As ϕ_1 is true in J , this implies that $Q^J = \emptyset$.

We start by showing that $(P')^J = (P')^I$. Let J_1 be the instance obtained from J by replacing $(P')^J$ by $(P')^I$. Since Σ is satisfied in J and $A^J = \emptyset$, Σ is satisfied in J_1 as well. Moreover, it follows from the definition of J_1 that $I \oplus J_1 \subseteq I \oplus J$. Since J is a \oplus -repair of I , this implies $J = J_1$. Hence, $(P')^J = (P')^I$. Similarly, we can show that $(Q')^J = (Q')^I$.

Next we prove that

$$\text{either } E^J = E^I \text{ or } E^J = \{(0, 0), (1, 1), (1, 0), (0, 1)\}. \tag{8}$$

In order to show (8), we start by proving that

$$E^I \subseteq E^J \subseteq \{(0, 0), (1, 1), (1, 0), (0, 1)\}. \tag{9}$$

We define J_2 as the instance obtained from J by replacing E^J by $(E^J \cup E^I) \cap \{(0, 0), (1, 1), (1, 0), (0, 1)\}$. We can check that Σ is satisfied in J_2 . Moreover, we have that $I \oplus J_2 \subseteq I \oplus J$. Since J is a \oplus -repair of I , this means that $J = J_2$. In particular,

$$E^J = (E^J \cup E^I) \cap \{(0, 0), (1, 1), (1, 0), (0, 1)\}.$$

This implies that (9) holds.

Now in order to derive (8) from (9), it is sufficient to prove that $E^J(0, 1)$ implies $E^J(1, 0)$ and $E^J(1, 0)$ implies $E^J(0, 1)$. We only show that $E^J(0, 1)$ implies $E^J(1, 0)$; the other implication is similar.

Assume that $E^J(0, 1)$. We claim that there exists $1 \leq i \leq n$ such that $P^J(p_i, 0)$ and $P^J(p_i, 1)$. Suppose for contradiction that there is no such a i . Let J_3 be the instance obtained from J by replacing E^J by $E^J \setminus \{(0, 1)\}$. Since there is no i such that $P^J(p_i, 0)$ and $P^J(p_i, 1)$, ϕ_{4i} is true for all $1 \leq i \leq n$. As $Q^J = \emptyset$, ϕ_{3i} is also true for all $1 \leq i \leq k$. It follows that Σ is true in J_3 . As $I \oplus J_3 \subsetneq I \oplus J$, this contradicts the fact that J is a \oplus -repair of I . Hence, there exists $1 \leq i_0 \leq n$ such that $P^J(p_{i_0}, 0)$ and $P^J(p_{i_0}, 1)$.

Since $P^J(p_{i_0}, 1) \wedge P^J(p_{i_0}, 0)$ holds and ϕ_{4i_0} is true in J , we must have $E^J(1, 0)$. This finishes the proof that $E^J(0, 1)$ implies $E^J(1, 0)$; hence the proof of (8).

- Suppose first that $E^J = E^I$. Since for all $1 \leq i \leq n$, ϕ_{4i} is true in J , there is at most one value $v_i \in \{0, 1\}$ such that $P^J(p_i, v_i)$. Hence, we can pick a valuation $V_P : \{p_i \mid 1 \leq i \leq n\} \rightarrow \{0, 1\}$ such that for all $1 \leq i \leq n$,

$$P^J(p_i, v_i) \text{ implies } v_i = V_P(p_i).$$

Since ϕ is true, there exists a valuation $V_Q : \{q_i \mid 1 \leq i \leq k\} \rightarrow \{0, 1\}$ such that ψ is true under the valuation $V_P \cup V_Q$.

Next we define the instance J_1 in the following way:

$$\begin{aligned} A^{J_1} &= \{1\}, \\ (P')^{J_1} &= (P')^I, \\ (Q')^{J_1} &= (Q')^I, \\ E^{J_1} &= E^I, \\ P^{J_1} &= P^J, \\ Q^{J_1} &= \{(q_i, V_q(q_i)) \mid 1 \leq i \leq k\}, \\ T_{a_1 a_2 a_3}^{J_1} &= T_{a_1 a_2 a_3}^I, \\ R_{a_1 a_2 a_3}^{J_1} &= R_{a_1 a_2 a_3}^I. \end{aligned}$$

Using the fact that ψ is true under the valuation $V_P \cup V_Q$, it is possible to check that Σ is true in J_1 . Moreover, as $Q^J = \emptyset$ and $A^J = \emptyset$, we can also show that $I \oplus J_1 \subseteq I \oplus J$. Since J is a \oplus -repair of I , we can conclude that $J = J_1$. This implies that $A^J \neq \emptyset$, which is a contradiction.

- Suppose finally that $E^J = \{(0, 0), (1, 1), (1, 0), (0, 1)\}$. We define an instance J_1 in the following way:

$$\begin{aligned} A^{J_1} &= \{1\}, \\ (P')^{J_1} &= (P')^I, \\ (Q')^{J_1} &= (Q')^I, \\ E^{J_1} &= E^J = \{(0, 0), (1, 1), (1, 0), (0, 1)\}, \\ P^{J_1} &= P^I, \\ Q^{J_1} &= Q^I, \\ T_{a_1 a_2 a_3}^{J_1} &= T_{a_1 a_2 a_3}^I, \\ R_{a_1 a_2 a_3}^{J_1} &= R_{a_1 a_2 a_3}^I. \end{aligned}$$

Using the fact that $E^{J_1} = \{(0, 0), (1, 1), (1, 0), (0, 1)\}$, we can check that Σ is true in J_1 . Next, as $E^J = \{(0, 0), (1, 1), (1, 0), (0, 1)\}$, we also have $I \oplus J_1 \subseteq I \oplus J$. Since J is a \oplus -repair of I , this implies that $J = J_1$. Hence, $A^J \neq \emptyset$, which is a contradiction. □

7 Arbitrary tgds (and egds)

In this section, we obtain results about arbitrary sets of tgds that contrast sharply with our earlier results about weakly acyclic sets of tgds. We begin by pointing out that our previous techniques are not of help in the study of arbitrary sets of tgds. First, as seen in Example 3.4, there is a (non-weakly acyclic) set of LAV tgds Σ and an instance I such that I has superset-repairs w.r.t. Σ (hence also \oplus -repairs w.r.t. Σ) of arbitrarily large sizes that cannot be bounded by any function in the size of the original instance. This contrasts with the state of affairs for weakly acyclic sets of tgds and egds in

Theorem 6.2. Furthermore, recall Theorem 4.6, which states that if Σ is any set of LAV tgds, then given any instance we compute in polynomial time an n -universal superset-repair w.r.t. Σ , for any fixed $n \geq 1$ (even though a universal superset-repair may not exist). But, by Proposition 3.8, in cases where Σ consists of arbitrary tgds, an n -universal superset-repairs may not be effectively computable even for $n = 1$.

We present now the results concerning the consistent query answering problem. We begin with the subset-consistent query answering and the superset-consistent query answering problems.

Theorem 7.1 *The following statements are true.*

1. *If Σ is a set of tgds and q is a conjunctive query, then the subset-consistent query answering problem for q w.r.t. Σ is in Π_2^P .*
2. *There is a set of tgds Σ and a conjunctive query q such that the subset-consistent query answering problem for q w.r.t. Σ is Π_2^P -complete.*
3. *There is a set of tgds Σ and a conjunctive query q such that the superset-consistent query answering problem for q w.r.t. Σ is undecidable.*

Proof We already mentioned the Π_2^P upper bound of subset-consistent query answering in Section 2; the lower bound was shown in Theorem 6.3, even for a weakly acyclic set of tgds.

As regards superset-repairs, it was shown in [28] that there is a (non-weakly acyclic) set Σ of tgds and egds such that the following problem is undecidable: given an instance I , is there an instance $J \models \Sigma$ such that $I \subseteq J$? Since every such instance J contains a superset-repair for I , we have that checking whether a given instance has a superset-repair with respect to Σ is an undecidable problem as well. By taking $q = \exists x P(x)$, where P is a fresh relation, it follows that the superset-consistent query answering problem is also undecidable. The egds of Σ can be eliminated by a standard construction: we add another binary relation E to the schema and extend Σ with GAV tgds stating that E is a congruence, i.e., an equivalence relation such that every other relation in the schema is closed under substitution of equals by equals with respect to the equivalence relation E . Alternatively, a proof of this undecidability result can be obtained by adapting the proof of Theorem 7.2 given later on in this section. \square

We now come to the main result, which asserts that there is a set Σ of tgds and a conjunctive query q such that computing the \oplus -consistent answers of q w.r.t. Σ is an undecidable problem. As mentioned in the Introduction, this improves a result in [3] asserting that there is a set of universal first-order sentences and a universal query (in fact, the negation of a conjunctive query) such that computing the \oplus -consistent answers is an undecidable problem.

Theorem 7.2 *There is a set Σ of tgds and a conjunctive query q such that the \oplus -consistent query answering problem for q w.r.t. Σ is undecidable.*

The remainder of this section is devoted to the proof of Theorem 7.2. We start with the following lemmas.

Lemma 7.3 *Let Σ be a set of tgds and egds. Suppose that R is a relational symbol that does not appear on the right-hand side of any tgd in Σ . Then for all \oplus -repairs J of an instance I , $R^J \subseteq R^I$.*

Proof Let J be a \oplus -repair of I . We define the instance K such that $R^K = R^J \cap R^I$ and for all relational symbols $S \neq R$, $S^K = S^J$. Since $J \models \Sigma$ and R never appears on the right-hand side of a tgd in Σ , Σ remains true in K . It also follows from the definition of K that $I \oplus K \subseteq I \oplus J$. Since J is a \oplus -repair, this implies that $K = J$. In particular, R^J is equal to $R^K (= R^J \cap R^I)$. That is, $R^J \subseteq R^I$. \square

Lemma 7.4 *Let Σ be a set of tgds and let q be a query of the form*

$$q_1 \vee \dots \vee q_k \vee \neg r_1 \vee \dots \vee \neg r_l, \tag{10}$$

where $q_1, \dots, q_k, r_1, \dots, r_l$ are conjunctive queries over a schema \mathbf{R} . There is a set Σ' of tgds and a conjunctive query q' , both over a schema $\mathbf{R}' \supseteq \mathbf{R}$, such that for every instance I over \mathbf{R} , we can compute in polynomial time an instance I' over \mathbf{R}' satisfying

$$\oplus\text{-Con}(q, I, \Sigma) = \oplus\text{-Con}(q', I', \Sigma').$$

Proof Suppose that q is a query of the form

$$q_1 \vee \dots \vee q_k \vee \neg r_1 \vee \dots \vee \neg r_l,$$

where $q_1, \dots, q_k, r_1, \dots, r_l$ are conjunctive queries.

Let A and B be fresh relation symbols. We define \mathbf{R}' as the schema $\mathbf{R} \cup \{A, B\}$. For all $1 \leq i \leq k$, if $q_i = \exists \mathbf{x}_i \phi_i(\mathbf{x}_i)$, we let t_i be the tgd given by

$$\phi_i(\mathbf{x}_i) \rightarrow \exists s B(s).$$

Next, if $r_i = \exists \mathbf{y}_i \psi_i(\mathbf{y}_i)$ for $1 \leq i \leq l$, we define t as the tgd

$$A(s) \wedge \psi_1(\mathbf{y}_1) \wedge \dots \wedge \psi_l(\mathbf{y}_l) \rightarrow \exists u B(u),$$

with s not occurring in $\mathbf{y}_1, \dots, \mathbf{y}_l$. We define Σ' as the set

$$\Sigma \cup \{t, t_i \mid 1 \leq i \leq k\}$$

and we let q' be the conjunctive query $\exists s A(s)$. Let I be an instance over \mathbf{R} . We define I' such that $A^{I'} = \{1\}$, $B^{I'} = \emptyset$ and for all $R \in \mathbf{R}$, $R^{I'} = R^I$. We have to show that

$$\oplus\text{-Con}(q, I, \Sigma) = \top \text{ iff } \oplus\text{-Con}(q', I', \Sigma') = \top. \tag{11}$$

We start by proving the direction from left to right of (11). Assume that $\oplus\text{-Con}(q, I, \Sigma)$ is true. Let J' be a \oplus -repair of I' with respect to Σ' . We have to show that $\exists s A(s)$ is true in J' . We define J'_0 as the instance obtained from J' by adding the fact $A(1)$. It suffices to show that $J' = J'_0$.

Let J be the instance over \mathbf{R} such that for all $R \in \mathbf{R}$, $R^J = R^{J'}$. J is a \oplus -repair of I with respect to Σ . Since $\oplus\text{-Con}(q, I, \Sigma)$ is true, q is true in J . Therefore, q is also true in J' .

We show that this implies that t is true in J'_0 . We make the following case distinction:

- Suppose that for some $1 \leq i \leq k$, q_i is true in J' . Hence, there exists \mathbf{a} such that $\phi_i(\mathbf{a})$ holds in J' . Since Σ' is true in J' , t_i given by

$$\phi_i(\mathbf{x}_i) \rightarrow \exists s B(s)$$

is true in J' . It follows that $\exists s B(s)$ is true in J' ; hence, in J'_0 . This implies that the tgd t is true in J'_0 .

- Next suppose that for some $1 \leq i \leq n$, $\neg r_i$ is true in J' . Since r_i is false in J' , r_i is also false in J'_0 . So there are no \mathbf{y}_i such $\psi_i(\mathbf{y}_i)$ holds in J'_0 . Hence, the left-hand side of t is never satisfied in J'_0 . In particular, t is true in J'_0 .

Since A does not occur in $\Sigma \cup \{t_i \mid 1 \leq i \leq k\}$ and that set of tgds is true in J' , $\Sigma \cup \{t_i \mid 1 \leq i \leq k\}$ is also true in J'_0 . Together with $J'_0 \models t$, we obtain $J'_0 \models \Sigma'$. It is clear that $I' \oplus J'_0 \subseteq I' \oplus J'$. Since J' is a \oplus -repair of I' with respect to Σ' , this implies $J'_0 = J'$ and this finishes the proof that q' is true in J' .

Next we prove the implication from right to left of (11). Assume that $\oplus\text{-Con}(q', I', \Sigma')$ is true. Suppose for contradiction that q is false in a \oplus -repair J of I with respect to Σ . Let J' be the instance such that $A^{J'} = \emptyset$, $B^{J'} = \emptyset$ and for all $R \in \mathbf{R}$, $R^{J'} = R^J$. In order to finish the proof, it suffices to show that J' is a \oplus -repair of I' with respect to Σ' . Indeed, q' is clearly false in J' , which contradicts the fact that $\oplus\text{-Con}(q', I', \Sigma')$ is true.

Now we prove that J' is a \oplus -repair of I' with respect to Σ' . There exists a \oplus -repair K' of I' with respect to Σ' such that $I \oplus K' \subseteq I \oplus J'$. We prove that $J' = K'$. Let K be the instance over \mathbf{R} such that for all $R \in \mathbf{R}$, $R^K = R^{K'}$. K is a \oplus -repair of I with respect to Σ and $I \oplus K \subseteq I \oplus J$. Since J is a \oplus -repair of I with respect to Σ , $K = J$. It follows that for all $R \in \mathbf{R}$, $R^{J'} = R^{K'}$.

In order to prove that $J' = K'$, it remains to show that $B^{J'} = B^{K'}$ and $A^{J'} = A^{K'}$. Since $B^{J'} = B^{I'}$ and $I' \oplus K' \subseteq I' \oplus J'$, we have $B^{K'} = B^{I'} = B^{J'}$.

Suppose for contradiction that $A^{J'} \neq A^{K'}$. Since $A^{J'} = \emptyset$, $A^{I'} = \{1\}$ and $I \oplus K' \subseteq I \oplus J'$, this can only happen if $A^{K'} = \{1\}$. Recall that q is false in J . In particular, for all $1 \leq i \leq l$, r_i is true in J . That is, for all $1 \leq i \leq l$, there exists a tuple \mathbf{b}_i such that $\psi_i(\mathbf{b}_i)$ is true in J ; hence, in J' and in K' . This implies that

$$A(1) \wedge \psi_1(\mathbf{b}_1) \wedge \dots \wedge \psi_l(\mathbf{b}_l)$$

is true in K' . Since K' is a \oplus -repair with respect to Σ' , t given by

$$A(s) \wedge \psi_1(\mathbf{y}_1) \wedge \dots \wedge \psi_l(\mathbf{y}_l) \rightarrow \exists u B(u),$$

is true in K' . Therefore, $\exists u B(u)$ is true in K' . This contradicts the fact that $B^{K'} = \emptyset$. □

Note that Lemma 7.4 remains true if we consider subset-repairs, instead of \oplus -repairs. In the case of superset-repairs, we do not know whether the lemma holds.

Theorem 7.2 will be proved via a reduction from the halting problem for two-register machines. In describing two-register machines, we follow the definition used in [6].

A *two-register machine* (2RM) is similar to Turing machines, except that, instead of a tape, it has two registers r_1, r_2 . Each register contains a natural number. A 2RM is programmed by a numbered sequence $\alpha_0, \dots, \alpha_l$ of instructions. Each instruction α_i is either an addition or a subtraction. An addition has the form $+(rg, j)$, with rg a register number and $j \leq l$ an instruction number. Its semantics is: add one to register rg and move to instruction α_j . A subtraction has the form $-(rg, j, k)$ with rg a register number and $j, k \leq l$ instruction numbers. Its semantics is: if content of register rg is zero, then move to instruction α_j ; otherwise, subtract one from register rg and move to instruction α_k .

An *instantaneous description* (ID) of a 2RM M is a triple (s, t, i) with $i \leq l$ an instruction number and $m, n \geq 0$ natural numbers representing the content of the registers r_1 and r_2 . The unique *successor of an ID* (s, t, i) is the ID (s', t', i') such that

- if $\alpha_i = +(1, j)$, then $s' = s + 1, t' = t, i' = j$;
- if $\alpha_i = +(2, j)$, then $s' = s, t' = t + 1, i' = j$;
- if $\alpha_i = -(1, j, k)$ and $s \neq 0$, then $s' = s - 1, t' = t, i' = j$;
- if $\alpha_i = -(1, j, k)$ and $s = 0$, then $s' = s, t' = t, i' = k$;
- if $\alpha_i = -(2, j, k)$ and $t \neq 0$, then $s' = s, t' = t - 1, i' = j$;
- if $\alpha_i = -(2, j, k)$ and $t = 0$, then $s' = s, t' = t, i' = k$.

The ID $(0, 0, l)$ is called *final*. If $M = (\alpha_0, \dots, \alpha_l)$ is a 2RM, then the *run* of M is the sequence $(D_i)_{i \geq 1}$ of ID's such that $D_0 = (0, 0, 0)$ and D_{i+1} is the successor of D_i , for all $i \geq 1$. The run is *halting* if it contains the ID $(0, 0, l)$. The *halting problem* for 2RM is to determine whether the run of a given 2RM is halting.

Theorem 7.5 [8] *The halting problem for two-register machines is undecidable.*

In the proof of Theorem 7.2, we also make use of the following property of 2RMs.

Lemma 7.6 *Given a 2RM $M = (\alpha_0, \dots, \alpha_l)$, we can compute a 2RM M' such that the successor of the ID $(0, 0, l)$ in M' is itself and*

the run of M is halting iff the run of M' is halting.

Proof Let $M = (\alpha_0, \dots, \alpha_l)$ be a 2RM. We define $M' = (\beta_0, \dots, \beta_{l+4})$ as the following 2RM:

$$\begin{aligned} \beta_i &= \alpha_i, \\ \beta_l &= -(1, l + 2, l + 4), \\ \beta_{l+1} &= \alpha_l, \end{aligned}$$

$$\begin{aligned} \beta_{l+2} &= +(1, l + 1), \\ \beta_{l+3} &= +(2, l + 1), \\ \beta_{l+4} &= -(2, l + 3, l + 4), \end{aligned}$$

where $0 \leq i < l$. By definition of M' , the successor of ID $(0, 0, l + 4)$ is $(0, 0, l + 4)$. Hence, it remains to show that the run of M is halting iff the run of M' is halting. This is implied by the following property: for all IDs a and b such that $a \neq (0, 0, l)$ and b is the successor of a in M , there exists a sequence $a_0 \dots a_n$ of ID such that

$$a_0 = a, a_n = b \text{ and } a_{i+1} \text{ is the successor of the ID } a_i \text{ in } M', \tag{12}$$

where $1 \leq i < n$.

Let b be the successor in M of an ID $a = (s, t, i)$ such that $a \neq (l, 0, 0)$. We prove the existence of a sequence $a_0 \dots a_n$ satisfying (12). If $i \neq l$, then the sequence $a_0 a_1$ with $a_0 = a$ and $a_1 = b$ satisfies (12).

Next suppose that $i = l$. We make the following case distinction. Assume first that $s \neq 0$. Then the successor ID of $a = (s, t, l)$ in M' is $(s - 1, t, l + 2)$. It also follows from the definition of M' that the successor ID of $(s - 1, t, l + 2)$ in M' is $(s, t, l + 1)$. Since $\beta_{l+1} = \alpha_l$, the successor ID of $(s, t, l + 1)$ in M' is b . Hence, the sequence $a_0 \dots a_3$ with $a_0 = a, a_1 = (s - 1, t, l + 2), a_2 = (s, t, l + 1)$ and $a_3 = b$ satisfies (12).

Assume now that $s = 0$. Since $a \neq (0, 0, l)$, this means that $t \neq 0$. By definition of M' , the successor ID of $a = (0, t, l)$ in M' is $a_1 = (0, t, l + 4)$. Since $t \neq 0$, the successor ID of a_1 in M' is $a_2 = (0, t - 1, l + 3)$. Next, the successor ID of a_2 in M' is $a_3 = (0, t, l + 1)$. Since $\beta_{l+1} = \alpha_l$, the successor ID of $(0, t, l + 1)$ in M' is b . The sequence $a_0 \dots a_4$ with $a_0 = a$ and $a_4 = b$ satisfies (12). □

We now proceed with the proof of Theorem 7.2.

Proof of Theorem 7.2 We shall give a reduction from the halting problem for 2RMs. For this, we define a set Σ' of tgds and a conjunctive query q' such that for every 2RM M , we can compute an instance I' such that

$$\text{the run of } M \text{ is halting} \iff \oplus\text{-Con}(q', I', \Sigma') \neq \top. \tag{13}$$

The idea is that each \oplus -repair of I^M encodes the run of a subset of the instructions in M (due to the symmetric difference semantics, we cannot prevent instructions from being dropped). The run of M will be halting if and only if some \oplus -repair of I^M contains the final ID of M . Hence, if q is a conjunctive query expressing that the final ID occurs, then the run of M is halting if and only if q is true in some \oplus -repair of I . That is, the run of M is halting if and only if $\oplus\text{-Con}(\neg q, I^M, \Sigma) \neq \top$. The problem is that $\neg q$ is not a conjunctive query, and this is where we use Lemma 7.4.

By Lemma 7.4, in order to find Σ' and q' satisfying (1), it suffices to find a set Σ of tgds and a query q of the form 10 (as stated in Lemma 7.4) such that for all 2RM M , we can compute an instance I satisfying

$$\text{the run of } M \text{ is halting} \iff \oplus\text{-Con}(q, I, \Sigma) \neq \top. \tag{14}$$

Let Σ be the set of the following tgds:

$$\begin{aligned}
 \phi_1^+ &= S(s, t, i) \wedge R_{1,+}(i, j) \rightarrow \exists s' Succ(s, s') \\
 \psi_1^+ &= S(s, t, i) \wedge R_{1,+}(i, j) \wedge Succ(s, s') \rightarrow S(s', t, j) \\
 \phi_2^+ &= S(s, t, i) \wedge R_{2,+}(i, j) \rightarrow \exists t' Succ(t, t') \\
 \psi_2^+ &= S(s, t, i) \wedge R_{2,+}(i, j) \wedge Succ(t, t') \rightarrow S(s, t', j) \\
 \phi_1^- &= S(s, t, i) \wedge R_{1,-}(i, j, k) \wedge Succ(s', s) \rightarrow S(s', t, j) \\
 \phi_1^- &= S(s, t, i) \wedge R_{1,-}(i, j, k) \wedge Zero(s) \rightarrow S(s, t, k) \\
 \phi_2^- &= S(s, t, i) \wedge R_{2,-}(i, j, k) \wedge Succ(t', t) \rightarrow S(s, t', j) \\
 \phi_2^- &= S(s, t, i) \wedge R_{2,-}(i, j, k) \wedge Zero(t) \rightarrow S(s, t, k) \\
 \phi_1^{Anc} &= Anc(u, s) \wedge Succ(s, s') \rightarrow Anc(u, s') \\
 \phi_2^{Anc} &= Succ(s, s') \rightarrow Anc(s, s') \\
 \phi_f &= Final(j) \wedge S(s, t, j) \wedge Zero(s) \wedge Zero(t) \rightarrow \exists s F(s)
 \end{aligned}$$

Let q be the query $\exists s Anc(s, s) \vee \neg(\exists s F(s))$. The intuition behind the relation symbols is as follows. We use S to encode IDs: the fact $S(s, t, i)$ corresponds to the ID where s is the content of the first register, t is the content of the second register and i is an instruction segment. We use the relation $Succ$ to encode the successor relation on an initial segment of the natural numbers, while Anc encodes the strict natural ordering on that initial segment. The relation Anc is the transitive closure of $Succ$, as expressed by the tgds ϕ_1^{Anc} and ϕ_2^{Anc} . We represent the initial and the final instructions of the machine with the relations $Zero$ and $Final$. The relation $Zero$ will consist of the singleton 0, while $Final$ will consist of the singleton l (where $l + 1$ is the number of instructions of the machine).

The relations $R_{1,+}$, $R_{1,-}$, $R_{2,+}$ and $R_{2,-}$ encode the instructions of the machines. A fact $R_{n,+}(i, j)$ expresses that the instruction α_i is equal to $+(n, j)$. Similarly, a fact $R_{n,-}(i, j, k)$ expresses that the instruction α_i is the subtraction $-(n, j, k)$.

The tgd ϕ_1^+ expresses that if the ID is (s, t, i) and instruction α_i needs to add one to the first register, then the initial segment of the natural numbers that we consider should at least contain the successor of s . The tgd ψ_1^+ expresses that if the ID is (s, t, i) , the instruction α_i is $+(1, j)$, and s' is equal to $s + 1$, then we move to the ID (s', t, j) . The tgds ϕ_2^+ and ψ_2^+ have similar meanings.

The tgd ϕ_1^- expresses that if the ID is (s, t, i) , the instruction α_i is $-(1, j, k)$, and s' is equal to $s - 1$ (in particular, $s \neq 0$), then we move to the ID (s', t, j) . The tgd ϕ_1^- expresses that if the ID is (s, t, i) , the instruction α_i is $+(1, j)$ and s is equal to 0, then we move to the ID (s, t, k) .

Finally, F acts as an indicator for termination. It becomes non-empty, when the instance contains an ID of the form $(0, 0, l)$. This is expressed by the tgd ϕ_f . The query q is false in an instance if and only if the relation F is non-empty and the ancestor relation does not contain any tuple of the form (s, s) . Intuitively, this means that the instance contains an ID of the form $(0, 0, l)$ and that the successor relation does not contain any cycle.

We associate an instance I^M with each 2RM. Let $M = (\alpha_0, \dots, \alpha_l)$ be a 2RM. By Lemma 7.6, we may assume that the successor of the ID $(0, 0, l)$ is itself. We define an instance $I := I^M$ as follows:

$$\begin{aligned} S^I &= \{(0, 0, 0)\}, \\ R_{1,+}^I &= \{(i, j) \mid \alpha_i = +(1, j)\}, \\ R_{1,-}^I &= \{(i, j, k) \mid \alpha_i = -(1, j, k)\}, \\ R_{2,+}^I &= \{(i, j) \mid \alpha_i = +(2, j)\}, \\ R_{2,-}^I &= \{(i, j, k) \mid \alpha_i = -(2, j, k)\}, \\ Zero^I &= \{0\}, \\ Final^I &= \{l\}, \\ Succ^I &= \emptyset, \\ Anc^I &= \emptyset, \\ F^I &= \emptyset, \\ B^I &= \emptyset. \end{aligned}$$

We start by proving the implication from left to right of (14). Assume that the run of M is halting. Let $(s_n, t_n, i_n)_{n \leq \omega}$ be the run of M . There exists m such that $(s_m, t_m, i_m) = (0, 0, l)$. Let n_0 be the maximum of the set $\{s_n, t_n \mid 0 \leq n \leq m\}$. We define an instance J_0 as follows

$$\begin{aligned} S^{J_0} &= \{(s_n, t_n, i_n) \mid 0 \leq n \leq m\}, \\ R_{1,+}^{J_0} &= R_{1,+}^I, \\ R_{1,-}^{J_0} &= R_{1,-}^I, \\ R_{2,+}^{J_0} &= R_{2,+}^I, \\ R_{2,-}^{J_0} &= R_{2,-}^I, \\ Zero^{J_0} &= \{0\}, \\ Final^{J_0} &= \{l\}, \\ Succ^{J_0} &= \{(s, s+1) \mid 0 \leq s < n_0\}, \\ Anc^{J_0} &= \{(s, t) \mid 0 \leq s < t \leq n_0\}, \\ F^{J_0} &= \{1\}, \\ B^{J_0} &= \emptyset. \end{aligned}$$

The query q is false in J_0 . Hence, if J_0 is a \oplus -repair, this implies that $\oplus\text{-Con}(\neg q, I, \Sigma) \neq \top$.

We prove that J_0 is a \oplus -repair. It is a simple exercise to check that Σ is true in J_0 . Next let K_0 be an instance such that $K_0 \models \Sigma$ and $I \oplus K_0 \subseteq I \oplus J_0$. We have to prove that $J_0 = K_0$.

First, for all relational symbols $T \in \{R_{1,+}, R_{1,-}, R_{2,+}, R_{2,-}, Zero, Final\}$, we have $T^{J_0} = T^I$. Since $I \oplus K_0 \subseteq I \oplus J_0$, this implies that for all $T \in \{R_{1,+}, R_{1,-}, R_{2,+}, R_{2,-}, Zero, Final\}$,

$$T^{K_0} = T^I = T^{J_0}.$$

Next, we show that $S^{J_0} = S^{K_0}$ and $Succ^{J_0} = Succ^{K_0}$. Since $S^I \subseteq S^{J_0}$, $Succ^I \subseteq S^{J_0}$ and $I \oplus K_0 \subseteq I \oplus J_0$, we have

$$S^{K_0} \subseteq S^{J_0} \text{ and } Succ^{K_0} \subseteq Succ^{J_0}. \tag{15}$$

Next we prove by induction on n that for all $0 \leq n \leq m$,

$$(s_n, t_n, i_n) \in S^{K_0} \text{ and } \{(s, s + 1) \mid 0 \leq s + 1 \leq f(n)\} \subseteq Succ^{K_0}, \tag{16}$$

where $f(n)$ is the maximum of the set $\{s_j, t_j \mid 0 \leq j \leq n\}$. Suppose first that $n = 0$. Recall that $(s_0, t_0, i_0) = (0, 0, 0)$. Since $(0, 0, 0)$ belongs to $S^I \cap S^{J_0}$ and $I \oplus K_0 \subseteq I \oplus J_0$, $(0, 0, 0)$ also belongs to S^{K_0} .

For the induction step, suppose that (16) holds. The ID $(s_{n+1}, t_{n+1}, i_{n+1})$ is the unique successor of (s_n, t_n, i_n) .

- Suppose that i_n is an addition. Without loss of generality, we may assume that i_n is of the form $+(1, j)$. Hence, $(s_{n+1}, t_{n+1}, i_{n+1})$ is equal to $(s_n + 1, t_n, j)$. First, we show that

$$(s, s + 1) \mid 0 \leq s + 1 \leq f(n + 1)\} \subseteq Succ^{K_0}. \tag{17}$$

Since $t_{n+1} = t_n$ and $s_{n+1} = s_n + 1$, it follows from the definition of $f(n)$ and $f(n + 1)$ that either $f(n + 1) = f(n)$ or $f(n + 1) = s_n + 1$. If $f(n + 1) = f(n)$, then (17) follows immediately from the induction hypothesis.

Suppose next that $f(n + 1) = s_n + 1$. Let s be such that $0 \leq s + 1 \leq s_n + 1$. We have to show that $(s, s + 1)$ belongs to $Succ^{K_0}$. If $s + 1 \leq s_n$, then by definition of s_n , $s + 1 \leq f(n)$. By induction hypothesis, $(s, s + 1)$ belongs to $Succ^{K_0}$. Next suppose that $s + 1 = s_n + 1$. We have to show that $(s_n, s_n + 1)$ belongs to $Succ^{K_0}$.

By induction hypothesis, $S(s_n, t_n, i_n)$ belongs to K_0 . Moreover, by definition of $R_{1,+}^{J_0}$, $R_{1,+}(i_n, j)$ belongs to J_0 . We observed earlier that $R_{1,+}^{J_0} = R_{1,+}^{K_0}$. Hence, $R_{1,+}(i_n, j)$ belongs to K_0 . Hence,

$$S(s_n, t_n, i_n) \wedge R_{1,+}(i_n, j)$$

is true in K_0 . The formula ϕ_1^+ given by

$$S(s, t, i) \wedge R_{1,+}(i, j) \rightarrow \exists s' Succ(s, s'),$$

is true in K_0 . Hence, there is s' such that $Succ(s_n, s')$ belongs to K_0 . Recall that $Succ^{K_0} \subseteq Succ^{J_0}$. By definition of $Succ^{J_0}$, $Succ^{K_0}(s_n, s')$ implies $s' = s_n + 1$. This means that $(s_n, s_n + 1)$ belongs to $Succ^{K_0}$. This finishes the proof of (17).

Next we prove that $S(s_{n+1}, t_{n+1}, i_{n+1})$ belongs to K_0 . We showed in the previous paragraph that

$$S(s_n, t_n, i_n) \wedge R_{1,+}(i_n, j) \wedge Succ(s_n, s_n + 1)$$

is true in K_0 . Recall that ψ_1^+ given by

$$S(s, t, i) \wedge R_{1,+}(i, j) \wedge Succ(s, s') \rightarrow S(s', t, j)$$

is true in K_0 . It follows that $S(s_n + 1, t_n, j)$ belongs to K_0 . That is, $S(s_{n+1}, t_{n+1}, i_{n+1})$ belongs to K_0 .

- Suppose that i_n is a subtraction. Without loss of generality, we may assume that i_n is of the form $-(1, j, k)$.

Since i_n is a subtraction, $f(n + 1) = f(n)$. Hence, using the induction hypothesis, we obtain

$$\{(s, s + 1) \mid 0 \leq s + 1 \leq f(n + 1)\} \subseteq Succ^{K_0}.$$

Next we prove that $(s_{n+1}, t_{n+1}, i_{n+1})$ belongs to K_0 . By definition of $R_{1,-}^{J_0}$, $R_{1,-}(i_n, j, k)$ belongs to J_0 . Recall that $R_{1,-}^{J_0} = R_{1,-}^{K_0}$. Hence, $R_{1,-}(i_n, j, k)$ belongs to K_0 .

Suppose first that $s_n \neq 0$. Therefore, $(s_{n+1}, t_{n+1}, i_{n+1})$ is equal to $(s_n - 1, t_n, j)$. By induction hypothesis, $S(s_n, t_n, i_n)$ belongs to K_0 and $(s_n - 1, s_n)$ belongs to $Succ^{K_0}$. It follows that

$$S(s_n, t_n, i_n) \wedge R_{1,-}(i_n, j, k) \wedge Succ(s_n - 1, s_n)$$

holds in K_0 . Since ϕ_1^- given by

$$S(s, t, i) \wedge R_{1,-}(i, j, k) \wedge Succ(s', s) \rightarrow S(s', t, j)$$

is true in K_0 , $S(s_n - 1, t_n, j)$ belongs to K_0 . That is, $(s_{n+1}, t_{n+1}, i_{n+1})$ belongs to K_0 .

Next assume that $s_n = 0$. Hence, $(s_{n+1}, t_{n+1}, i_{n+1})$ is equal to $(0, t_n, k)$. By definition of $Zero^{J_0}$, $Zero(0)$ holds in J_0 . Together with $Zero^{J_0} = Zero^{K_0}$, we obtain $Zero^{K_0}(0)$. Therefore,

$$S(0, t_n, i_n) \wedge R_{1,-}(i_n, j, k) \wedge Zero(0)$$

holds in K_0 . Since ϕ_1^- given by

$$S(s, t, i) \wedge R_{1,-}(i, j, k) \wedge Q_0(s) \rightarrow S(s, t, k)$$

is true in K_0 , $S(0, t_n, k)$ belongs to K_0 . That is, $(s_{n+1}, t_{n+1}, i_{n+1})$ belongs to K_0 .

This finishes the proof that (16) holds for all $0 \leq n \leq m$.

It follows from (16) that $S^{J_0} \subseteq S^{K_0}$ and $Succ^{J_0} \subseteq Succ^{K_0}$. Together with (15), we obtain $S^{J_0} = S^{K_0}$ and $Succ^{J_0} = Succ^{K_0}$.

Next we show that $Anc^{K_0} = Anc^{J_0}$. Recall that ϕ_1^{Anc} and ϕ_2^{Anc} are true in K_0 . Together with $Succ^{J_0} = Succ^{K_0}$, we obtain that $Anc^{J_0} \subseteq Anc^{K_0}$. Since $Anc^I = \emptyset$ and $I \oplus K_0 \subseteq I \oplus J_0$, this implies $Anc^{J_0} = Anc^{K_0}$.

In order to finish the proof that $K_0 = J_0$, it remains to show that $F^{K_0} = F^{J_0}$. We established earlier that $S^{K_0} = S^{J_0}$. In particular, $(s_m, t_m, i_m) = (0, 0, l)$ belongs to S^{K_0} . We also proved that $Zero^{K_0} = Zero^{J_0} = \{0\}$ and $Final^{K_0} = Final^{J_0} = \{l\}$. Hence,

$$Final(l) \wedge S(0, 0, l) \wedge Zero(0) \wedge Zero(0)$$

holds in K_0 . The formula ϕ_f given by

$$Final(j) \wedge S(s, t, j) \wedge Zero(s) \wedge Zero(t) \rightarrow \exists s F(s)$$

is true in K_0 . This implies that $F^{K_0} \neq \emptyset$. Moreover, since $F^I \subseteq F^{J_0}$ and $I \oplus K_0 \subseteq I \oplus J_0$, we have $F^{K_0} \subseteq F^{J_0}$. Recall that F^{J_0} is a singleton. Together with $F^{K_0} \neq \emptyset$, we obtain $F^{K_0} = F^{J_0}$.

We prove the implication from right to left of (14). Suppose that the consistent answers of q with respect to I and Σ is false. Hence, there is an instance J that is a \oplus -repair of I and such that q is false in J . That is,

$$J \models \neg \exists s Anc(s, s) \quad \text{and} \quad F^J \neq \emptyset. \tag{18}$$

We start by proving

$$R_{i,j}^J \subseteq R_{i,j}^I, \tag{19}$$

$$Zero^J = Zero^I = \{0\}, \tag{20}$$

$$Final^J = Final^I = \{l\} \tag{21}$$

$$(0, 0, 0) \in S^J. \tag{22}$$

Inclusion (19) follows from Lemma 7.3. Next we prove (20); that is, $Zero^J = Zero^I$. The inclusion $Zero^J \subseteq Zero^I$ follows from Lemma 7.3. Since $Zero^I = \{0\}$, $Zero^J = \emptyset$ or $Zero^J = Zero^I$. Suppose for contradiction that $Zero^J = \emptyset$. Let J_3 be the instance such that $F^{J_3} = \emptyset$, and, for all relational symbols R distinct from F , $R^{J_3} = R^J$. Since $Zero^J = \emptyset$, the set Σ remains true in R^{J_3} . Moreover, $I \oplus J_3 \subseteq I \oplus J$. Since J is a \oplus -repair, $J_3 = J$. In particular, $F^J = \emptyset$. This contradicts (18) and finishes the proof of (20).

The proof that (21) holds is similar to the proof of (20). In order to prove (22), suppose for contradiction that $(0, 0, 0)$ does not belong to S^J . Let J_2 be the instance such that $S^{J_2} = \emptyset$, $F^{J_2} = \emptyset$ and for all relational symbols R distinct from S and F , $R^{J_2} = R^J$. We can check that Σ is true in J_2 and $I \oplus J_2 \subseteq I \oplus J$. Hence, $J = J_2$. In particular, $F^J = \emptyset$. This contradicts (18). Therefore, $(0, 0, 0)$ belongs to S^J . This finishes the proof of (19), (20), (21) and (22).

The proof of the implication from right to left of (14) is basically divided in two parts. First we prove that we may identify the relation $Succ$ with the successor relation on a well-chosen initial segment of the natural numbers. Second, we show how we may associate the run of M with the facts of the form $S(s, t, i)$.

We start by associating a natural number with each value occurring in S . Let N be the set $\{s \mid \exists s', S^J(s, s') \text{ or } S^J(s', s)\}$. We pick a map $f : N \rightarrow N \cup \{\perp\}$ such that for all $s \in N$,

- if there is s' such that $Succ^J(s, s')$, then $Succ^J(s, f(s))$,
- otherwise, $f(s) = \perp$.

We show that there is a smallest natural number n_0 such that $f^{n_0+1}(0) = \perp$ where $f^{n_0+1}(0)$ is obtained by applying f ($n_0 + 1$) times to 0. Suppose for contradiction that there is no n such that $f^n(0) = \perp$. Since J is finite, there is a sequence s_1, \dots, s_n such that for all $1 \leq i \leq n - 1$, $Succ^I(s_i, s_{i+1})$ and $Succ(s_n, s_1)$. Since ϕ_1^{Anc} and

ϕ_2^{Anc} are true in J , we obtain that $Anc^J(s_1, s_1)$. This contradicts (18). Hence, there is a smallest natural number n_0 such that $f^{n_0+1}(0) = \perp$.

We define U as the set $\{f^n(0) \mid 0 \leq n \leq n_0\}$. Without loss of generality, we may identify the element $f^n(0)$ (where $0 \leq n \leq n_0$) with the natural number n . By definition of f and since ϕ_1^{Anc} and ϕ_2^{Anc} are true in J , we have

$$\{(s, s + 1) \mid 0 \leq s < n_0\} \subseteq Succ^J \tag{23}$$

$$\{(s, t) \mid 0 \leq s < t \leq n_0\} \subseteq Anc^J \tag{24}$$

Claim 4 (a) $S^J = \{(s, t, i) \in S^J \mid 0 \leq s, t \leq n_0\}$,

(b) $Succ^J = \{(s, s + 1) \mid 0 \leq s < n_0\}$,

(c) $Anc^J = \{(s, t) \mid 0 \leq s < t \leq n_0\}$. □

Proof of Claim Let J_5 be the instance such that S^{J_5} is equal to $\{(s, t, i) \in S^J \mid 0 \leq s, t \leq n_0\}$, $Succ^{J_5}$ is equal to $\{(s, s + 1) \mid 0 \leq s < n_0\}$, Anc^{J_5} is equal to $\{(s, t) \mid 0 \leq s < t \leq n_0\}$ and for all relational symbols $R \notin \{S, Succ, Anc\}$, $R^{J_5} = R^J$. In order to prove (a), (b) and (c), we have to show that $J = J_5$.

First, we show that Σ is true in J_5 .

By definition of $Succ^{J_5}$ and Anc^{J_5} , the formulas ϕ_1^{Anc} and ϕ_2^{Anc} are true in J_5 .

Next, as $F^{J_5} = F^J$ and $F^J \neq \emptyset$ (see (18)), ϕ_f is also true in J_5 . Now we prove that ϕ_1^- is true in J_5 . Suppose that $S^{J_5}(s, t, i)$, $R_{1,-}^{J_5}(i, j, k)$ and $Succ^{J_5}(s', s)$ hold. We have to show that $S^{J_5}(s', t, j)$ holds. By definition of J_5 and by inclusion (23), this implies $S^J(s, t, i)$, $R_{1,-}^J(i, j, k)$ and $Succ^J(s', s)$. Since ϕ_1^- is true in J , we have $S^J(s', t, j)$. Now, as $S^{J_5}(s, t, i)$ and $Succ^{J_5}(s', s)$, s' and t are natural numbers between 0 and n_0 . Putting that together with the fact that $S^J(s', t, j)$, we obtain $S^{J_5}(s', t, j)$. This finishes the proof that ϕ_1^- is true in J_5 . The proofs that $\phi_1^-, \phi_2^-, \phi_2^+, \psi_1^+$ and ψ_2^+ are true in J_5 are similar.

Next we prove that ϕ_1^+ is true in J_5 . Suppose that $S^{J_5}(s, t, i)$ and $R_{1,+}^{J_5}(i, j)$ hold. In particular, s and t belong to $\{n \mid 0 \leq n \leq n_0\}$. We have to find s' such that $Succ^{J_5}(s, s')$. By definition of J_5 , $S^{J_5}(s, t, i)$ and $R_{1,+}^{J_5}(i, j)$ imply $S^J(s, t, i)$ and $R_{1,+}^J(i, j)$. Since ϕ_1^+ is true in J , there exists s'' such that $Succ^J(s, s'')$. The problem is that there is no guarantee that s'' is a natural number between 0 and n_0 .

Since s has a successor s'' in J , we know by definition of f and n_0 that $s < n_0$. Hence, $Succ^{J_5}(s, s + 1)$. By (23), this implies $Succ^J(s, s + 1)$. We may conclude that ϕ_1^+ is true in J_5 . The proof that ϕ_2^+ is true in J_5 is similar. This finishes the proof of the Claim.

It follows from the definition of J_5 and inclusions (23) and (24) that $I \oplus J_5 \subseteq I \oplus J$. Since J is a \oplus -repair, this means that $J = J_5$. Hence, (a), (b) and (c) hold. □

The next step is to show how the relation S encodes the run of M . For that purpose, we introduce the notion of reachability. We say that a tuple $S(s', t', j)$ in J is *1-step reachable* from a tuple $S(s, t, i)$ in J if

- either $R_{1,+}^J(i, j)$, $s' = s + 1$ and $t = t'$,

- or $R_{2,+}^J(i, j), t' = t + 1$ and $s = s'$,
- or $R_{1,-}^J(i, j, k)$ for some $k, s' = s - 1$ and $t = t'$,
- or $R_{1,-}^J(i, k, j)$ for some $k, Zero(s), s = s'$ and $t = t'$,
- or $R_{2,-}^J(i, j, k)$ for some $k, t' = t - 1$ and $s = s'$,
- or $R_{2,-}^J(i, k, j)$ for some $k, Zero(s), s = s'$ and $t = t'$.

Since $Zero^J = \{0\}$ and $R_{i,j}^J \subseteq R_{i,j}^I$ (where $i \in \{1, 2\}$ and $j \in \{+, -\}$), we obtain

$$\begin{aligned} &\text{if } S(s', t', j) \text{ is 1-step reachable from } S(s, t, i), \\ &\text{the successor of the ID } (s, t, i) \text{ is } (s', t', i'). \end{aligned} \tag{25}$$

A tuple $S(j, s', t')$ is *reachable* from a tuple $S(s, t, i)$ if the pair $(S(s, t, i), S(j, s', t'))$ belongs to the transitive closure of the 1-step reachability relation.

Now let m_0 be the maximum of the set

$$\{s, t \mid S(s, t, i) \text{ is reachable from } S(0, 0, 0)\}.$$

We show that

- (i) $S^J = \{S(s, t, i) \mid S(s, t, i) \text{ is reachable from } S(0, 0, 0)\}$,
- (ii) $Succ^J = \{(s, s + 1) \mid 0 \leq s < m_0\}$,
- (iii) $Anc^J = \{(s, t) \mid 0 \leq s < t \leq m_0\}$.

Let J_6 be the database such that S^{J_6} is equal to $\{S(s, t, i) \mid S(s, t, i) \text{ is reachable from } S(0, 0, 0)\}$, $Succ^{J_6}$ is equal to $\{(s, s + 1) \mid 0 \leq s < m_0\}$, Anc^{J_6} is equal to $\{(s, t) \mid 0 \leq s < t \leq m_0\}$ and for all relational symbols $R \notin \{S, Succ, Anc\}, R^{J_6} = R^J$.

We can check that Σ remains true in the instance J_6 . By definition of J_6 , we also have $I \oplus J_6 \subseteq I \oplus J$. Hence, $J = J_6$; that is, (i), (ii) and (iii) hold.

Next we prove that there exists s_m, t_m and j_m such that

$$S^J(s_m, t_m, j_m), Zero(s_m), Zero(t_m) \text{ and } Final(j_m). \tag{26}$$

Suppose for contradiction that there are not such s_m, t_m and j_m . Consider the instance J_7 obtained from J by removing the tuples of the form $F^J(s)$. The tgds of Σ remains true in J_7 . Moreover, $I \oplus J_7 \subseteq I \oplus J$. Since J is a \oplus -repair, $J = J_7$. That is, $F^J = \emptyset$, which contradicts (18). Therefore, there exists s_m, t_m and j_m satisfying (26).

We proved earlier that $Final^J = Final^I$ and $Zero^J = Zero^I$. Hence, $s_m = t_m = 0$ and $j_m = l$. Together with (26), this means that the fact $S(0, 0, l)$ belongs to J . Since J satisfies (i), the tuple $S(0, 0, l)$ is reachable from $(0, 0, 0)$. That is, there is a sequence $(s_0, t_0, i_0), \dots, (s_m, t_m, i_m)$ such that

- $(s_0, t_0, i_0) = (0, 0, 0)$,
- for all $0 \leq n < m, S(s_{n+1}, t_{n+1}, i_{n+1})$ is 1-step reachable from (s_n, t_n, i_n) ,
- $(s_m, t_m, i_m) = (0, 0, l)$.

It follows from (25) and the hypothesis that the ID $(0, 0, l)$ is its unique successor, that the sequence

$$(s_0, t_0, i_0), \dots, (s_{m-1}, t_{m-1}, i_{m-1}), (0, 0, l), (0, 0, l), \dots$$

is the run of the machine M . Hence, the run of M is halting.

8 Concluding Remarks

In this paper, we carried out a fairly comprehensive investigation of the data complexity of the consistent query answering problem for conjunctive queries with respect to classes of constraints that have played a major role in data exchange and data integration. In the process, we brought into front stage and used extensively the notions of universal repairs and n -universal repairs, $n \geq 1$.

One problem left open by our investigation is the data complexity of the superset-consistent answers and of the \oplus -consistent answers for conjunctive queries w.r.t. sets of inclusion dependencies and equality-generating dependencies. As mentioned earlier, this problem has been shown to be undecidable in combined complexity [31]. It would also be interesting to investigate algorithmic aspects concerning the existence of universal repairs. Specifically, what can one say about the complexity of the following decision problem: given a set Σ of dependencies and an instance I , does I have a \star -universal repair w.r.t. Σ , where $\star \in \{\oplus, \text{subset}, \text{superset}\}$? And similarly for n -universal repairs, $n \geq 1$. Note that results in [26, 27] already imply that there is a set Σ of tgds for which the problem of checking whether a given instance has a universal superset-repair is undecidable.

Finally, the results presented give rise to challenging complete classification problems that, if resolved, may take the form of a dichotomy or a trichotomy theorems. For example, is it true that for every set Σ of GAV tgds and every conjunctive query q , the \oplus -consistent answers of q are either in PTIME or coNP-complete? Even for just key constraints only partial results towards such a dichotomy theorem have been obtained so far [19, 21, 29, 36, 37].

Acknowledgments We thank the anonymous reviewers for their comments and pointers to the literature. This research was partially supported by NSF grants IIS-0905276 and IIS-1217869. Fontaine was visiting UC Santa Cruz while this research was carried out.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
2. Afrati, F.N., Kolaitis, P.G.: Repair checking in inconsistent databases: Algorithms and complexity. In: ICDT, pp. 31–41 (2009)
3. Arenas, M., Bertossi, L.E.: On the decidability of consistent query answering. In: AMW (2010)
4. Arenas, M., Bertossi, L.E., Chomicki, J.: Consistent query answers in inconsistent databases. In: PODS, pp. 68–79 (1999)
5. Bárány, V., Gottlob, G., Otto, M.: Querying the guarded fragment. In: Proc. of LICS, pp. 1–10 (2010)
6. Benedikt, M., Fan, W., Geerts, F.: XPath satisfiability in the presence of dtds. J. ACM **55**(2) (2008)
7. Bertossi, L.E.: Consistent query answering in databases. SIGMOD Rec **35**(2), 68–76 (2006)

8. Börger, E., Grädel, E., Gurevich, Y.: *The Classical Decision Problem. Perspectives in Mathematical Logic*. Springer (1997)
9. Bravo, L., Bertossi, L.E.: Consistent query answering under inclusion dependencies. In: *CASCON*, pp. 202–216 (2004)
10. Cali, A., Gottlob, G., Kifer, M.: Taming the infinite chase: Query answering under expressive relational constraints. In: Brevka, G., Lang, J. (eds.) *KR*, pp. 70–80. AAAI Press (2008)
11. Cali, A., Lembo, D., Rosati, R.: On the decidability and complexity of query answering over inconsistent and incomplete databases. In: *PODS*, pp. 260–271 (2003)
12. ten Cate, B., Kolaitis, P.G.: Structural characterizations of schema-mapping languages. *Commun. ACM* **53**(1), 101–110 (2010)
13. Chomicki, J.: Consistent query answering: Five easy pieces. In: *ICDT*, pp. 1–17 (2007)
14. Chomicki, J., Marcinkowski, J.: Minimal-change integrity maintenance using tuple deletions. *Inf. Comput.* **197**(1–2), 90–121 (2005)
15. Deutsch, A., Tannen, V.: Reformulation of XML queries and constraints. In: *ICDT*, pp. 225–241 (2003)
16. Fagin, R.: Horn clauses and database dependencies. *J. ACM* **29**(4), 952–985 (1982)
17. Fagin, R., Kolaitis, P.G., Miller, R.J., Popa, L.: Data exchange: Semantics and query answering. *Theor. Comput. Sci.* **336**(1), 89–124 (2005)
18. Fagin, R., Kolaitis, P.G., Popa, L.: Data exchange: Getting to the core. *TODS* **30**(1), 174–210 (2005)
19. Fontaine, G.: Why is it hard to obtain a dichotomy for consistent query answering? In: *LICS*. To appear (2013)
20. Fuxman, A., Kolaitis, P.G., Miller, R.J., Chiew Tan, W.: Peer data exchange. *ACM Trans. Database Syst.* **31**(4), 1454–1498 (2006)
21. Fuxman, A., Miller, R.J.: First-order query rewriting for inconsistent databases. *J. Comput. Syst. Sci.* **73**(4), 610–635 (2007)
22. Garey, M.R., Johnson, D.S.: *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York (1990)
23. Gottlob, G., Nash, A.: Efficient core computation in data exchange. *J. ACM* **55**(2), 1–49 (2008)
24. Grahne, G., Onet, A.: Data correspondence, exchange and repair. In: *ICDT*, pp. 219–230 (2010)
25. Hell, P., Nešetřil, J.: The core of a graph. *Discret. Math* **109**, 117–126 (1992)
26. Hernich, A.: *Foundations of Query Answering in Relational Data Exchange*. PhD thesis, Goethe-Universität Frankfurt am Main (2010)
27. Hernich, A., Schweikardt, N.: CWA-solutions for data exchange settings with target dependencies. In: *PODS*, pp. 113–122 (2007)
28. Kolaitis, P.G., Panttaja, J., Tan, W.-C.: The complexity of data exchange. In: *Proceedings of PODS'06*, pp. 30–39 (2006)
29. Kolaitis, P.G., Pema, E.: A dichotomy in the complexity of consistent query answering for queries with two atoms. *Inf. Process. Lett.* **112**(3), 77–85 (2012)
30. Papadimitriou, C.M.: *Computational Complexity*. Addison-Wesley, Reading (1994)
31. Rosati, R.: On the finite controllability of conjunctive query answering in databases under open-world assumption. *J. Comput. Syst. Sci.* **77**(3), 572–594 (2011)
32. Staworko, S., Chomicki, J.: Consistent query answers in the presence of universal constraints. *Inf. Syst* **35**(1), 1–22 (2010)
33. ten Cate, B., Fontaine, G., Kolaitis, P.G.: On the data complexity of consistent query answering. In: *Proceedings of the 15th International Conference on Database Theory, ICDT '12*, pp. 22–33. ACM, New York (2012)
34. Vardi, M.Y.: The complexity of relational query languages (extended abstract). In: *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pp. 137–146 (1982)
35. Wijsen, J.: Database repairing using updates. *ACM Trans. Database Syst.* **30**, 722–768 (2005)
36. Wijsen, J.: On the first-order expressibility of computing certain answers to conjunctive queries over uncertain databases. In: *PODS*, pp. 179–190 (2010)
37. Wijsen, J.: A remark on the complexity of consistent conjunctive query answering under primary key violations. *Inf. Process. Lett.* **110**(21), 950–955 (2010)