

Lossless Selection Views under Conditional Domain Constraints

Ingo Feinerer, Enrico Franconi, and Paolo Guagliardo

Abstract—A set of views defined by selection queries splits a database relation into sub-relations, each containing a subset of the original rows. This *decomposition* into horizontal fragments is *lossless* when the initial relation can be reconstructed from the fragments by union. In this paper, we consider horizontal decomposition in a setting where some of the attributes in the database schema are *interpreted* over a specific domain, on which a set of special predicates and functions is defined. We study losslessness in the presence of integrity constraints on the database schema. We consider the class of *conditional domain constraints* (CDCs), which restrict the values that the interpreted attributes may take whenever a certain condition holds on the non-interpreted ones, and investigate lossless horizontal decomposition under CDCs in isolation, as well as in combination with functional and unary inclusion dependencies.

Index Terms—Selection, views, losslessness, constraints, CDC, consistency, separability

1 INTRODUCTION

THE problem of updating a database through a set of views consists in propagating updates issued on the views to the underlying base relations over which the view relations are defined, so that the changes to the database reflect *exactly* those to the views. This is a classical problem in database research, known as the *view update problem* ([1], [2], [3]), which in recent years has received renewed and increasing attention ([4], [5], [6], [7], [8], [9]).

View updates can be consistently propagated in an unambiguous way under the condition that the mapping between database and view relations is *lossless*, which means that not only do the view relations depend on the database relations, but also the converse is true. However, just knowing that such an “inverse” dependency exists is not yet sufficient to effectively propagate the changes from the views to the database. What is essential to know is *how*, in some constructive way, the database relations depend on the view relations. This amounts to being able to define each database relation in terms of the views by means of a query, in much the same way the latter are defined from the former [10]. In such a context, database decompositions [11] play an important role, because their losslessness is associated with the existence of an explicit *reconstruction operator* that, as the name suggests, prescribes how a database relation can be rebuilt from the pieces, called *fragments*, into which it has been decomposed.

Horizontal decomposition is the process of splitting a given relation into sub-relations on the same attributes and of the

same arity, each containing a subset of the rows of the original relation. For example, consider the relation R shown in Fig. 1, recording data about the employees of a company: their name (**EMP**), the department (**DEP**) and the position (**POS**) in which they are employed, and their income (e.g., euros per month) consisting of a fixed salary (**SAL**) plus a variable bonus (**BON**). In Fig. 1, the relation R is decomposed into three fragments: $V1$ selects the rows of R with employees working as managers in departments other than ICT, $V2$ selects the rows of R with employees who get strictly less than 4,000 as bonus, and $V3$ selects the rows of R with employees who do not work as managers. This kind of decomposition is lossless when the original relation can be reconstructed from the fragments by union; in other words, the reconstruction operator for horizontal decomposition is the union. In the example of Fig. 1, the set of views $\{V1, V2, V3\}$ constitutes a lossless horizontal decomposition of R , as the union of $V1$, $V2$ and $V3$ contains all (and only) the rows of R . Each proper subset of $\{V1, V2, V3\}$ constitutes a lossy decomposition of R , because each view selects at least one row that is not selected by any of the others; e.g., the union of $V1$ and $V2$ does not contain the third row of R , $\langle \text{Linda, Finance, Consultant, 5,000, 4,000} \rangle$, which is selected only by $V3$.

Observe that the horizontal decomposition specified by the definitions of views $V1$, $V2$ and $V3$ in Fig. 1 is lossless for the given relation R , but this is not the case for every relation (over the same attributes). For instance, the tuple $\langle \text{Sam, ICT, Manager, 6,000, 5,000} \rangle$ is not selected by any of these views; indeed, every relation containing a row for an employee who works as a manager in the ICT department and receives a bonus greater than 4,000 would not be losslessly decomposed by $V1$, $V2$ and $V3$. In the presence of integrity constraints, however, things may be different, because some tuples, such as the one above, might not be allowed in the input relation.

The study of horizontal decomposition ([12], [13], [14], [15]) has mostly focused on settings where data values can only be compared for equality. However, most real-world

- I. Feinerer is with the Database and Artificial Intelligence Group, Institute of Information Systems, Vienna University of Technology, Austria.
E-mail: ingo.feinerer@tuwien.ac.at.
- E. Franconi and P. Guagliardo are with the KRDB Research Centre, Free University of Bozen-Bolzano, Italy.
E-mail: {franconi, guagliardo}@inf.unibz.it.

Manuscript received 23 Nov. 2013; revised 15 Apr. 2014; accepted 19 June 2014. Date of publication 30 June 2014; date of current version 23 Dec. 2014.

Recommended for acceptance by Z. Ives.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TKDE.2014.2334327

		R				
		EMP	DEP	POS	SAL	BON
V1	{	Mary	Finance	Manager	6000	5000
		Sarah	Sales	Manager	7000	4000
		Linda	Finance	Consultant	5000	4000
V2	{	John	ICT	Analyst	2000	1000
		Tom	ICT	Manager	3000	2000

} V3

```

create view V1 as select * from R where DEP<>"ICT"
                                and POS="Manager"
create view V2 as select * from R where BON<4000
create view V3 as select * from R where POS<>"Manager"

```

Fig. 1. Selection views over a company database.

applications make use of data values coming from domains with a richer structure (e.g., ordering) on which a variety of other restrictions besides equality can be expressed (e.g., that of being within a range or above a threshold). Examples are the attributes **SAL** and **BON** of the relation in Fig. 1, the dimensions, weights and prices in the database of a shipping company, or the various amounts (credits, debits, interest and exchange rates, etc.) recorded in a banking application. It is therefore of practical interest to consider a scenario where some of the attributes in the database schema are *interpreted* over a specific domain, such as the reals or the integers, on which a set of predicates (e.g., smaller/greater than) and functions (e.g., addition and subtraction) are defined, according to a first-order language \mathcal{C} .

In the present paper, we consider horizontal decomposition in a setting with interpreted attributes, in which fragments are defined by selection queries consisting of a condition on the non-interpreted attributes, expressed by a Boolean combination of equalities, and a condition on the interpreted attributes, expressed by a formula in \mathcal{C} . In particular, we study the *losslessness* (w.r.t. every input relation) of horizontal decompositions specified in this way, in the presence of integrity constraints on the database schema. We work under the pure universal relation assumption (URA) [11], that is, we restrict ourselves to a database schema consisting of only one relation symbol, as customary in the study of database decomposition.

1.1 Contribution and Outline

In Section 2, we introduce a class of integrity constraints called *conditional domain constraints* (CDCs). By means of a formula in \mathcal{C} , a CDC restricts the values that the interpreted attributes can take whenever a certain condition is satisfied by the non-interpreted ones. Depending on the expressive power of \mathcal{C} , CDCs can capture constraints that naturally arise in practise; for example, in the scenario of Fig. 1, it may be required that employees in the ICT department have a total income (i.e., salary plus bonus) of at most 5,000, that employees working as managers get a bonus of at least 2,000, and that employees never receive a bonus greater than their salary. These constraints can be expressed as:

$$\text{DEP} = \text{"ICT"} \Rightarrow \text{SAL} + \text{BON} \leq 5,000 \quad (1a)$$

$$\text{POS} = \text{"Manager"} \Rightarrow \text{BON} \geq 2,000 \quad (1b)$$

$$\text{SAL} - \text{BON} \geq 0. \quad (1c)$$

As we shall see, the views of Fig. 1 losslessly decompose every relation satisfying the above CDCs.

In our investigation, we do not commit to any specific language \mathcal{C} and we simply assume that \mathcal{C} is closed under negation.

In Section 3, we characterise consistent sets of CDCs in terms of satisfiability in \mathcal{C} . Whenever the satisfiability of sets of formulae in \mathcal{C} is decidable, our characterisation directly gives a decision procedure for checking whether a set of CDCs is consistent. This is the case, e.g., for the so-called *Unit Two Variable Per Inequality* fragment of linear arithmetic over the integers, whose formulae (referred to as UTVPIs) consist of at most two variables and variables have unit coefficients, as well as for Boolean combinations of such formulae. We prove that deciding consistency is NP-complete for both of these languages.

In Section 4, we characterise lossless horizontal decomposition under CDCs in terms of unsatisfiability in \mathcal{C} . Whenever the satisfiability of sets of formulae in \mathcal{C} is decidable, this characterisation gives a decision procedure for checking whether a horizontal decomposition is lossless under CDCs. We show that this problem is co-NP-complete when \mathcal{C} is the language of either UTVPIs or Boolean combinations of UTVPIs.

In Section 5, we study lossless horizontal decomposition under CDCs in combination with traditional integrity constraints. We show that functional dependencies (FDs) do not interact with CDCs and can thus be allowed without any restriction, whereas this is not the case for unary inclusion dependencies (UINDs). We provide a *domain propagation rule* to derive a set of CDCs that fully captures the interaction between a given set of UINDs and opportunely restricted CDCs w.r.t. lossless horizontal decomposition, which makes possible to employ the general technique for deciding losslessness also in the presence of UINDs. In addition, we consider restricted combinations of CDCs with both FDs and UINDs.

We conclude in Section 6 with a discussion of the results, relevant related work and future research directions.

2 PRELIMINARIES

We start by introducing the necessary notation and notions that will be used throughout the paper. We assume some familiarity with formal logic and its application to database theory.

Basics. An *n-tuple* is an ordered list of n elements, where n is a positive integer. We denote tuples by overlined lowercase letters (e.g., \bar{t}) and we write them as comma-separated sequences enclosed in parentheses; the k th element of a tuple \bar{t} is denoted by $\bar{t}[k]$. For example, if \bar{t} is the 4-tuple (a, b, c, a) , then $\bar{t}[3] = c$. An *n-ary relation* on a set A , where n is called the *arity* of the relation, is a set of n -tuples of elements of A .

A *schema* is a finite set \mathbf{S} of relation symbols, also called a *relational signature*. Each relation symbol S has a positive arity $|S|$ indicating the total number of *positions* in S , which are partitioned into *interpreted* and *non-interpreted* ones.

Relation symbols of arity n are called n -ary; we indicate that $|S| = n$ by writing S/n .

Let \mathbf{dom} be a possibly infinite set of arbitrary values, and let \mathbf{idom} be a set of values from a specific domain (e.g., the integers \mathbb{Z}) on which a set of predicates (e.g., \leq) and functions (e.g., $+$) are defined, according to a first-order language \mathcal{C} closed under negation. An *instance* over a schema \mathbf{S} associates each $S \in \mathbf{S}$ with a relation S^I of appropriate arity on $\mathbf{dom} \cup \mathbf{idom}$, called the *extension* of S under I , such that the values for the interpreted and non-interpreted positions of S are taken from \mathbf{idom} and \mathbf{dom} , respectively. The set of elements of $\mathbf{dom} \cup \mathbf{idom}$ occurring in an instance I is called the *active domain* of I , denoted by $\mathbf{adom}(I)$. An instance is finite if its active domain is, and all instances in this paper are assumed to be finite. A *fact* is given by the association, denoted by $R(\bar{t})$, between a relation symbol R and a tuple \bar{t} of values of appropriate arity; an instance can be represented as a set of facts.

Constraints. A *language* over a relational signature \mathbf{S} is a set of first-order logic (FOL) formulae over \mathbf{S} with constants $\mathbf{dom} \cup \mathbf{idom}$ under the standard name assumption (i.e., the interpretation of each constant is the constant's name itself). A formula in some language \mathcal{L} is called an \mathcal{L} -formula. The sets of constants and relation symbols that occur in a formula φ are denoted by $\text{const}(\varphi)$ and $\text{sig}(\varphi)$, respectively; we extend $\text{const}(\cdot)$ and $\text{sig}(\cdot)$ to sets of formulae in the natural way.

A *constraint* is a closed formula (that is, without free variables) in some language. For a set Γ of constraints, we say that an instance I over $\text{sig}(\Gamma)$ is a *model* of (or *satisfies*) Γ , and write $I \models \Gamma$, to indicate that the relational structure $\langle \mathbf{adom}(I) \cup \text{const}(\Gamma), I \rangle$ makes every formula in Γ true under the standard FOL semantics. We write $I \models \varphi$ as short for $I \models \{\varphi\}$, and say that I satisfies φ . A set of constraints Γ *entails* (or *logically implies*) a constraint φ , written $\Gamma \models \varphi$, if every finite model of Γ also satisfies φ . All sets of constraints in this paper are finite.

Propositional theories. A *propositional variable* is a variable whose value can be either T (true) or F (false). A *propositional formula* is a Boolean combination of *propositional variables*, including the two special propositional variables \top and \perp , whose values are always T and F, respectively. A *propositional theory* is a set of propositional formulae. We denote the set of propositional variables occurring in a propositional formula P by $\text{var}(P)$ and we extend $\text{var}(\cdot)$ to propositional theories in the natural way. A *valuation* of a set of propositional variables (also called a *truth-value assignment*) assigns a *truth-value* (i.e., either T or F) to each propositional variable in the set. The truth-value $\alpha(P)$ of a propositional formula P under a valuation α of its propositional variables is determined by the standard semantics of the Boolean connectives. We say that α *satisfies* (or *makes true*) P , and write $\alpha \models P$, if $\alpha(P) = \text{T}$. Given a propositional theory Π , a valuation of $\text{var}(\Pi)$ *satisfies* Π , written $\alpha \models \Pi$, if α satisfies every propositional formula in Π .

2.1 Horizontal Decomposition

We consider a *source schema* \mathbf{R} , consisting of a single relation symbol R , and a *decomposed schema* \mathbf{V} , disjoint with \mathbf{R} , of *view symbols* with the same arity as R . We formally define horizontal decomposition as follows.

Definition 1. Let $\mathbf{R} = \{R\}$ and $\mathbf{V} = \{V_1, \dots, V_n\}$. Let Δ be a set of constraints over \mathbf{R} and let Σ be a set of exact view definitions, one for each $V_i \in \mathbf{V}$, of the form $\forall \bar{x}. V_i(\bar{x}) \leftrightarrow \varphi(\bar{x})$, where φ is a safe-range¹ formula over \mathbf{R} . Then, Σ is a horizontal decomposition of \mathbf{R} into \mathbf{V} under Δ if $\Delta \cup \Sigma \models \forall \bar{x}. V_i(\bar{x}) \rightarrow R(\bar{x})$ for every $V_i \in \mathbf{V}$. We say that Σ is lossless if $\Delta \cup \Sigma \models \forall \bar{x}. R(\bar{x}) \leftrightarrow V_1(\bar{x}) \vee \dots \vee V_n(\bar{x})$.

For the sake of simplicity, w.l.o.g. we assume that the first $\|R\|$ positions of R and of every $V \in \mathbf{V}$ are non-interpreted, while the remaining $|R| - \|R\|$ positions are interpreted. Under this assumption, instances over $\mathbf{R} \cup \mathbf{V}$ associate each relation symbol with a subset of the Cartesian product $\mathbf{dom}^k \times \mathbf{idom}^{n-k}$, where $n = |R|$ and $k = \|R\|$. Unless otherwise specified, when we speak of a tuple \bar{t} we implicitly assume that \bar{t} is of arity n and that the first k values of \bar{t} are from \mathbf{dom} while the rest are from \mathbf{idom} . W.l.o.g. we also assume that a variable associated with the i th position of R is named x_i if $i \leq k$, and y_{i-k} otherwise. By default, \bar{x} and \bar{y} denote the tuples (x_1, \dots, x_k) and (y_1, \dots, y_{n-k}) , respectively.

Since every \mathcal{C} -formula is over variables associated with interpreted positions, we write $\phi(\bar{y})$ to indicate that ϕ is a \mathcal{C} -formula whose free variables are among the variables in \bar{y} . For a tuple \bar{t}_y of $n - k$ values from \mathbf{idom} , we denote by $\phi(\bar{t}_y)$ the result of replacing every occurrence in ϕ of the free variable y_i with the value $\bar{t}_y[i]$. We say that \bar{t}_y is a *solution* to ϕ if $\phi(\bar{t}_y)$ is true under the semantics of \mathcal{C} . In such a case, we also say that the assignment β associating each y_i with $\bar{t}_y[i]$ satisfies ϕ , and we write $\beta \models \phi$.²

Source constraints. The class of integrity constraints we consider on the source schema \mathbf{R} is that of *conditional domain constraints*, which restrict the admissible values at interpreted positions by means of formulae in \mathcal{C} , when a certain condition holds on the non-interpreted ones. Formally, a CDC is a formula of the form

$$\forall \bar{x}, \bar{y}. (R(\bar{x}, \bar{y}) \wedge \lambda(\bar{x})) \rightarrow \delta(\bar{y}), \quad (2)$$

where $\lambda(\bar{x})$ is a Boolean combination of equalities $x = a$, with x from \bar{x} and a from \mathbf{dom} , and $\delta(\bar{y}) \in \mathcal{C}$. We use $x \neq a$ as short for $\neg(x = a)$ and, for ease of notation, we write (2) simply as $\lambda(\bar{x}) \rightarrow \delta(\bar{y})$. Here, we make use of a more general variant of the CDCs introduced in [16], where the condition $\lambda(\bar{x})$ was limited to a conjunction of possibly negated equalities. In general, (2) is more expressive than a CDC of the form used in [16], as in the latter negation is allowed only atomically, that is, in front of equalities, and so disjunction cannot be expressed. However, there is no difference in expressivity between the two variants when considering sets of CDCs, because the antecedent of (2) can always be rewritten in disjunctive normal form (DNF) and the CDC split into a set of CDCs having the same consequent and each disjunct as antecedent.

View definitions. The view symbols in \mathbf{V} are defined by selection queries with conditions on both interpreted and non-interpreted positions. Formally, each $V \in \mathbf{V}$ is defined by a formula of the form

1. For details on the syntactic notion of range restriction, corresponding to the semantic notion of domain independence, refer to [11].
2. Sometimes, by abuse of terminology, we say that an assignment β is a solution to a \mathcal{C} -formula ϕ , with the obvious meaning.

$$\forall \bar{x}, \bar{y}. V(\bar{x}, \bar{y}) \leftrightarrow (R(\bar{x}, \bar{y}) \wedge \lambda(\bar{x}) \wedge \sigma(\bar{y})), \quad (3)$$

where $\lambda(\bar{x})$ is as in (2) and $\sigma(\bar{y}) \in \mathcal{C}$. In the following, we write (3) simply as $V : \lambda(\bar{x}) \wedge \sigma(\bar{y})$. View definitions of this form clearly generalise those in [16], where $\lambda(\bar{x})$ is limited to a conjunction of possibly negated equalities and disjunction cannot be expressed. While this has an impact on end-users, who can define more expressive views, there is no difference between the two formalisms w.r.t. losslessness, in that any view symbol V defined by (3) can be split into a set of views, defined by formulae of the form used in [16], that together select exactly the same tuples as V ; given $\lambda(\bar{x})$ in DNF, each of these view definitions has the same selection condition as V on the interpreted attributes and a disjunct of $\lambda(\bar{x})$ as selection condition on the non-interpreted ones.

The technique we will present in Section 4 for checking whether a set of selections of the form (3) is lossless can also be applied when (some of) the selections have the form $V : \lambda(\bar{x}) \vee \sigma(\bar{y})$ by considering, in place of each such selection, the two selections $V' : \lambda(\bar{x})$ and $V'' : \sigma(\bar{y})$.

Running example. To clarify the notation and illustrate the concepts introduced so far, we now give an example that will be used also in the rest of the paper. It is based on the source schema of Fig. 1, the CDCs (1a)-(1c) informally described in Section 1 and the views previously specified in Fig. 1 by means of SQL statements.

Example 1. Let R be a relation symbol of arity 5, whose positions are associated with attributes **EMP**, **DEP**, **POS**, **SAL**, **BON**, in this order, with the last two interpreted over the integers. Differently from the example of Fig. 1, for simplicity we assume that salaries and bonuses are given in thousands of euros/month. Let $a = \text{“ICT”}$ and $b = \text{“Manager”}$, and consider the following set Δ of CDCs: $\{x_2 = a \rightarrow y_1 + y_2 \leq 5 ; x_3 = b \rightarrow y_2 \geq 2 ; \top \rightarrow y_1 - y_2 \geq 0\}$. Let $\mathbf{V} = \{V_1, V_2, V_3\}$, and let Σ be the horizontal decomposition given by $V_1 : x_2 \neq a \wedge x_3 = b$, $V_2 : y_2 < 4$, and $V_3 : x_3 \neq b$.

Specific languages. The techniques we will present for deciding whether a set of CDCs is consistent (Section 3) and whether a horizontal decomposition is lossless under CDCs (Section 4) give actual algorithms when satisfiability in \mathcal{C} is decidable; in the case of losslessness, \mathcal{C} is additionally required to be closed under negation. Thus, even though our investigation is in general independent of the choice of \mathcal{C} , from a practical point of view it makes sense to consider concrete languages that enjoy both of the above properties. Two prominent such languages are *Unit Two-Variable Per Inequality* formulae (UTVPIs) and Boolean combinations thereof. UTVPIs, a.k.a. *Generalised 2 SAT* (G2SAT) formulae [17], are a fragment of linear arithmetic over the integers. Formally, a UTVPI formula has the form $ax + by \leq d$, where x and y are integer variables, $a, b \in \{-1, 0, 1\}$ and $d \in \mathbb{Z}$.

UTVPIs can express comparisons between two variables and between a variable and an integer, as well as compare the sum or difference of two variables with an integer. As integers allow to represent also real numbers with fixed precision, UTVPIs may be sufficient for most applications. The CDCs and the view definitions of Example 1 can be expressed when \mathcal{C} is the language of UTVPIs.

Whether a set of UTVPIs is satisfiable can be checked in polynomial time ([18], [19], [20]). We refer to a Boolean combination of UTVPIs as BUTVPI; deciding the satisfiability of a set of BUTVPIs is NP-complete [21].

3 CONSISTENT SETS OF CDCs

Before turning our attention to horizontal decomposition, we first deal with the relevant problem of determining whether a set of CDCs is *consistent*, that is, whether it has a non-empty model.³ It is important to make sure that the integrity constraints over the source schema are consistent, as every horizontal decomposition is meaninglessly lossless when there are in fact no legal relations to decompose.

In this section, we will characterise the consistency of a set of CDCs in terms of satisfiability in \mathcal{C} , where \mathcal{C} is *not* required to be closed under negation. The *consistency problem* for CDCs is the decision problem that takes as input a set Δ of CDCs and answers the question: “Is Δ consistent?” We will show that when \mathcal{C} is the language of either UTVPIs or BUTVPIs this problem is NP-complete. The technique employed here provides the basis for the approach we follow in Section 4 in the study of lossless horizontal decomposition.

Observe that, given their form, CDCs affect only one tuple at a time, and so whether an instance satisfies a set of CDCs depends on each tuple of the instance in isolation from the others. Indeed, a set of CDCs is consistent precisely if it is satisfiable on an instance consisting of only one tuple, therefore we can restrict our attention to single tuples. Moreover, we are not really interested in the actual values of a tuple at non-interpreted positions; what we need to know is simply whether such values satisfy the conditions in the antecedent of each CDC or not. To this end, with each equality between a variable x_i and a constant a we associate a propositional variable p_i^a , whose truth-value indicates whether the value in the i th position is a . To each valuation of such propositional variables corresponds the (possibly infinite) set of tuples satisfying the equalities associated with the names of the propositional variables. For example, a valuation assigning true to p_1^a and false to p_2^b identifies all the tuples in which the value of the first element is a and the value of the second is different from b . A bit more care is needed with valuations of propositional variables that refer to the same position (i.e., have the same subscript) but to different constants (i.e., have different superscripts). For example, p_1^a and p_1^b (with $a \neq b$) should never be both evaluated to true.

As we shall see, checking whether a set Δ of CDCs is consistent amounts to first building a propositional theory by replacing the equalities with the corresponding propositional variables, and then looking for a valuation α such that:

- any two propositional variables referring to the same position but to different constants are not evaluated both to true; and
- the set of \mathcal{C} -formulae that “apply” under α is satisfiable.

Definition 2. Let $\Delta = \{\phi_1, \dots, \phi_n\}$ be a set of CDCs over R . For each $\phi_i \in \Delta$, recalling it has the form (2), we construct

3. Since CDCs are universally-quantified closed implicational formulae, any set thereof is always trivially satisfied by the empty instance.

$$\text{prop}(\phi_i) = P \rightarrow v_i, \quad (4)$$

where P is a propositional formula (possibly \top) obtained from the condition $\lambda(\bar{x})$ in the antecedent of ϕ by replacing each equality $x_i = a$ between a variable x_i and $a \in \text{idom}$ with the propositional variable p_i^a , and v_i is a fresh propositional variable associated with the \mathcal{C} -formula $\delta(\bar{y})$, denoted by $\text{idf}(v_i)$,⁴ in the consequent of ϕ . We denote $\{\text{prop}(\phi) \mid \phi \in \Delta\}$ by Π_Δ and we call it the propositional theory associated with Δ .

We consider the set $\text{var}(\Pi_\Delta)$ of propositional variables occurring in Π_Δ partitioned into $\text{var}_p(\Pi_\Delta) = \{\text{var}(P) \mid (P \rightarrow v_i) \in \Pi_\Delta\}$ and $\text{var}_v(\Pi_\Delta) = \text{var}(\Pi_\Delta) \setminus \text{var}_p(\Pi_\Delta)$.

For a pair of distinct propositional variables p_i^a and p_i^b associated with the same position i but distinct **dom** constants a and b , we consider the propositional formula $p_i^a \wedge p_i^b \rightarrow \perp$, called the *axiom of unique value* for p_i^a and p_i^b , intuitively stating that two distinct constants are not allowed in the same position. The *axioms of unique value* for a set of propositional variables consist of the axiom of unique value for each pair of distinct propositional variables p_i^a and p_i^b in the set. A tuple \bar{t} is *consistent with a valuation* α if, for every propositional variable p_i^a , it holds that $\bar{t}[i] = a$ precisely if $\alpha(p_i^a) = \text{T}$. In general, given a valuation α of a set of propositional variables, by construction there exists a tuple consistent with α if and only if α satisfies the corresponding axioms of unique value for that set.

Definition 3. Let Δ be a set of CDCs over R . The auxiliary theory Π_{aux} for Π_Δ consists of the axioms of unique value for $\text{var}_p(\Pi_\Delta)$.

Example 2. The propositional theory associated with Δ of Example 1 is $\Pi_\Delta = \{p_2^a \rightarrow v_1, p_3^b \rightarrow v_2, \top \rightarrow v_3\}$ for which $\text{var}_p(\Pi_\Delta) = \{p_2^a, p_3^b\}$ and $\text{var}_v(\Pi_\Delta) = \{v_1, v_2, v_3\}$. The auxiliary theory for Π_Δ is $\Pi_{\text{aux}} = \emptyset$. The association between the propositional variables in $\text{var}_v(\Pi_\Delta)$ and the set of UTVPIs from the CDCs in Δ is $\text{idf} = \{v_1 \mapsto y_1 + y_2 \leq 5, v_2 \mapsto y_2 \geq 2, v_3 \mapsto y_1 - y_2 \geq 0\}$.

Given a set Δ of CDCs and a valuation α of $\text{var}_p(\Pi_\Delta)$, we say that a CDC $\phi \in \Delta$ is *applicable* under α if α makes the l.h.s. of $\text{prop}(\phi)$ true. We can use α to “filter” Π_Δ and construct a set consisting of the consequent of each CDC in Δ that is applicable under α . This set contains the \mathcal{C} -formulae that must be necessarily satisfied by the values at interpreted positions of every tuple consistent with α , that is, whose values at non-interpreted positions satisfy the antecedents of the CDCs applicable under α .

Definition 4. Let Δ consist of CDCs over R , and let α be a valuation of $\text{var}_p(\Pi_\Delta)$. The α -filtering of Π_Δ is the set

$$\Pi_\Delta^\alpha = \{\text{idf}(v) \mid (P \rightarrow v) \in \Pi_\Delta, \alpha(P) = \text{T}\}, \quad (5)$$

consisting of \mathcal{C} -formulae associated with propositional variables that occur in some propositional formula of Π_Δ whose l.h.s. holds true under α .

The main result of this section characterises the consistency of a set of CDCs in terms of satisfiability in \mathcal{C} . We remark again that the result holds in general for any

language \mathcal{C} , not necessarily closed under negation. This requirement will become essential only in the upcoming Sections 4 and 5.

Theorem 1. Let Δ be a set of CDCs over R , and let Π_{aux} be the auxiliary theory for Π_Δ . Then, Δ is consistent if and only if there exists a valuation α of $\text{var}_p(\Pi_\Delta)$ satisfying Π_{aux} and such that Π_Δ^α is satisfiable.

Whenever the satisfiability of sets of \mathcal{C} -formulae is decidable, Theorem 1 gives an algorithm to check whether a set of CDCs is consistent, as we illustrate below in our running example, where \mathcal{C} is the language of UTVPIs.

Example 3. With respect to Π_Δ of Example 2, consider the valuation $\alpha = \{p_2^a \mapsto \text{T}, p_3^b \mapsto \text{F}\}$, for which we have $\Pi_\Delta^\alpha = \{y_1 + y_2 \leq 5, y_1 - y_2 \geq 0\}$. Obviously, α satisfies the (empty) auxiliary theory Π_{aux} for Π_Δ . In addition, Π_Δ^α is satisfiable as, e.g., $\{y_1 \mapsto 3, y_2 \mapsto 2\}$ is a solution to every UTVPI in it.

We will now give the proof of Theorem 1, for which we first need to prove a technical lemma. Let $n = |R|$ and $k = \|\text{ID}\|$; with each tuple \bar{t} is associated the assignment $\beta: \{y_1, \dots, y_{n-k}\} \rightarrow \text{idom}$, which we refer to as the *assignment induced by the interpreted positions of \bar{t}* , such that $\beta(y_{i-k}) = \bar{t}[i]$ for every $i \in \{k+1, \dots, n\}$. Intuitively, the following lemma shows that any tuple that is consistent with a valuation α satisfies a set of CDCs precisely if the assignment induced by its interpreted positions satisfies the α -filtering.

Lemma 1. Let Δ be a set of CDCs over R , and let α be a valuation of $\text{var}_p(\Pi_\Delta)$. Let \bar{t} be consistent with α , and let β be the assignment induced by the interpreted positions of \bar{t} . Then, $\{R(\bar{t})\} \models \Delta$ if and only if β satisfies Π_Δ^α .

Proof. Let $n = |R|$ and $k = \|\text{ID}\|$.

Claim 1. Let $\phi \in \Delta$ and $\text{prop}(\phi) = P \rightarrow v$. Then, $\alpha(P) = \text{T}$ iff $\lambda(\bar{x})$ is true under $\{x_1 \mapsto \bar{t}[1], \dots, x_k \mapsto \bar{t}[k]\}$.

Proof. Since \bar{t} is consistent with α , for $i \in \{1, \dots, k\}$ we have that $\bar{t}[i] = a$ if and only if $\alpha(p_i^a) = \text{T}$.

Claim 2. For each $\text{prop}(\phi) = P \rightarrow v$ with $\phi \in \Delta$, it is the case that $I \not\models \phi$ if and only if $\alpha(P) = \text{T}$ and $\beta \not\models \text{idf}(v)$.

Proof. As ϕ is a CDC, $I \not\models \phi$ if and only if the antecedent $\lambda(\bar{x})$ of ϕ holds true under $\{x_1 \mapsto \bar{t}[1], \dots, x_k \mapsto \bar{t}[k]\}$ and the consequent $\delta(\bar{y})$ of ϕ is not true under $\{y_1 \mapsto \bar{t}[k+1], \dots, y_{n-k} \mapsto \bar{t}[n]\}$. In turn, this is the case if and only if both $\alpha(P) = \text{T}$ (by Claim 1) and β does not satisfy $\text{idf}(v) = \delta(\bar{y})$ (by construction).

We prove Lemma 1 by showing that $I \not\models \Delta$ if and only if β does not satisfy Π_Δ^α .

“if” Assume $\beta \not\models \Pi_\Delta^\alpha$, that is, there is some \mathcal{C} -formula $\psi \in \Pi_\Delta^\alpha$ not satisfied by β . By construction of Π_Δ^α , ψ is the consequent of a CDC $\phi \in \Delta$ such that $\text{prop}(\phi) = P \rightarrow v$, with $\psi = \text{idf}(v)$ and $\alpha(P) = \text{T}$. Thus, as β does not satisfy ψ , by Claim 2 $I \not\models \phi$, and therefore $I \not\models \Delta$.

“Only if” Assume $I \not\models \Delta$. Then, there exists some $\phi \in \Delta$ which is not satisfied by I . Since $\text{prop}(\phi) = P \rightarrow v$, Claim by 2 $\alpha(P) = \text{T}$ and $\beta \not\models \text{idf}(v)$. Hence, $\text{idf}(v) \in \Pi_\Delta^\alpha$. Therefore, $\beta \not\models \Pi_\Delta^\alpha$. \square

4. idf stands for “interpreted domain formula”.

Proof of Theorem 1. Let $n = |R|$ and $k = \|R\|$.

“*if*” Let α and β be such that $\alpha \models \Pi_{\text{aux}}$ and $\beta \models \Pi_{\Delta}^{\alpha}$. Then, as $\alpha \models \Pi_{\text{aux}}$, it is never the case that two distinct propositional variables in $\text{var}_p(\Pi_{\Delta})$ associated with the same position are both true under α . Thus, there exists a tuple \bar{t} consistent with α and such that β is the assignment induced by its interpreted positions. Therefore, as $\beta \models \Pi_{\Delta}^{\alpha}$, the instance $\{R(\bar{t})\}$ is a model of Δ by Lemma 1.

“*only if*” Assume that Δ is consistent, that is, it has a non-empty model. In particular, as every formula in Δ is in one tuple, there is a tuple \bar{t} such that the instance $I = \{R(\bar{t})\}$ is a model of Δ . Take α as follows: for every propositional variable $p \in \text{var}_p(\Pi_{\Delta})$, $\alpha(p) = \text{T}$ if $p = p_i^a$ and $\bar{t}[i] = a$, otherwise $\alpha(p) = \text{F}$. By construction, $\alpha \models \Pi_{\text{aux}}$ and \bar{t} is consistent with α . Therefore, as $I \models \Delta$, the assignment β induced by the interpreted positions of \bar{t} satisfies Π_{Δ}^{α} by Lemma 1. \square

The satisfiability problem for \mathcal{C} takes as input a set Γ of \mathcal{C} -formulae and answers the question: “Is Γ satisfiable?”

Lemma 2. *The satisfiability problem for \mathcal{C} linearly reduces to the consistency problem for CDCs.*

Proof. Let $\Gamma = \{\phi_1, \dots, \phi_n\}$ be a set \mathcal{C} -formulae. Then, take $\Delta = \{\top \rightarrow \phi_i \mid \phi_i \in \Gamma\}$, let $\Pi_{\Delta} = \{\top \rightarrow v_i \mid i = 1, \dots, n\}$ and $\text{idf} = \{v_i \mapsto \phi_i \mid i = 1, \dots, n\}$. The auxiliary theory for Π_{Δ} is $\Pi_{\text{aux}} = \emptyset$. As $\text{var}_p(\Pi_{\Delta}) = \emptyset$, the only valuation of $\text{var}_p(\Pi_{\Delta})$ is $\alpha = \emptyset$, which satisfies Π_{aux} and for which $\Pi_{\Delta}^{\alpha} = \{\text{idf}(v_i) \mid v_i \in \text{var}_v(\Pi_{\Delta})\} = \Gamma$. Thus, by Theorem 1, the set Δ of CDCs is consistent iff Γ is satisfiable. The reduction is linear in the size of Γ . \square

With regard to the consistency problem for CDCs whose consequents are either UTVPIs or BUTVPIs, we have the following complexity results.

Theorem 2. *When \mathcal{C} is the language of either BUTVPIs or UTVPIs, the consistency problem for CDCs is NP-complete.*

Proof sketch. Constructing the propositional theories Π_{Δ} and Π_{aux} requires linear time, checking that a valuation α of $\text{var}_p(\Pi_{\Delta})$ satisfies Π_{aux} takes polynomial time, and checking that an assignment from the variables in \bar{v} to integers satisfies Π_{Δ}^{α} (whose construction takes linear time) can be done in polynomial time, whether Π_{Δ}^{α} consists of either UTVPIs or BUTVPIs. Hence, in light of Theorem 1, we can verify a given solution to the consistency problem, when \mathcal{C} is the language of either UTVPIs or BUTVPIs, in polynomial time.

The NP-hardness of the consistency problem when \mathcal{C} is the language of BUTVPIs follows by Lemma 2 from the fact that the satisfiability problem for BUTVPIs is NP-hard.

The NP-hardness of the consistency problem when \mathcal{C} is the language of UTVPIs can be shown by a reduction from SAT. \square

4 LOSSLESS SELECTIONS UNDER CDCS

The technique described in the previous section can be opportunely extended and applied for checking whether

a set of selection views of the form (3) is lossless under CDCs, that is, whether every source relation satisfying the given CDCs can be reconstructed by union from the fragments into which it is decomposed by the given view definitions.

In this section, we will characterise lossless horizontal decomposition in terms of unsatisfiability in \mathcal{C} , where \mathcal{C} is closed under negation. The *losslessness problem* in \mathcal{C} is the decision problem that takes as input a horizontal decomposition Σ specified by selections of the form (3) and a set Δ of CDCs and answers the question: “Is Σ lossless under Δ ?” We will show that this problem is co-NP-complete when \mathcal{C} is the language of either UTVPIs or BUTVPIs. For these languages, our characterisation provides an exponential-time algorithm for deciding the losslessness of Σ under Δ , by means of a number of unsatisfiability checks in \mathcal{C} which is exponentially bounded by the size of Δ .

By definition, a horizontal decomposition Σ of R into V_1, \dots, V_n is lossless under a set Δ of CDCs over R if $R^I = V_1^I \cup \dots \cup V_n^I$ for every model I of $\Delta \cup \Sigma$. As the extension of each view symbol is always included in the extension of R , the problem is equivalent to checking that there is no model I of $\Delta \cup \Sigma$ where a tuple $\bar{t} \in R^I$ does not belong to any V_i^I . In turn, this means that for each definition in Σ , which has the form (3), the values in \bar{t} at non-interpreted positions do not satisfy λ , or the values in \bar{t} at interpreted positions do not satisfy the \mathcal{C} -formula σ .

The formulae in Σ apply to one tuple at a time and, as already observed in Section 3, so do CDCs; therefore we can again focus on single tuples. With each equality we associate, as before, a propositional variable whose truth-value determines whether the equality is satisfied. Given a valuation α , we consider the set consisting of \mathcal{C} -formulae in the r.h.s. of all the CDCs that are applicable under α and the *negation* of the selection condition $\delta(\bar{v})$ of each view definition in Σ whose selection condition $\lambda(\bar{x})$ is satisfied by α . Then, checking losslessness is equivalent to checking that there exists no valuation α for which the above set of \mathcal{C} -formulae is satisfiable. Indeed, from such a valuation and the corresponding assignment of values from idom satisfying the relevant \mathcal{C} -formulae, we can obtain a tuple that provides a counterexample to losslessness.

Similarly to what we did in Section 3 for sets of CDCs, we build a propositional theory associated with a given horizontal decomposition.

Definition 5. *Let $\Sigma = \{\phi_1, \dots, \phi_n\}$ be a horizontal decomposition. For each $\phi_i \in \Sigma$, which has the form (3), we build*

$$\text{prop}(\phi_i) = P \rightarrow v'_i, \quad (6)$$

in which v'_i is either a fresh propositional variable associated (by means of idf) with the \mathcal{C} -formula $\sigma(\bar{v})$, if any, occurring in ϕ_i , or \perp otherwise.⁵ We denote $\{\text{prop}(\phi) \mid \phi \in \Sigma\}$ by Π_{Σ} and we call it the propositional theory associated with Σ .

We consider the set $\text{var}(\Pi_{\Sigma})$ of propositional variables occurring in Π_{Σ} partitioned into $\text{var}_p(\Pi_{\Sigma}) = \{\text{var}(P) \mid (P \rightarrow v_i) \in \Pi_{\Sigma}\}$ and $\text{var}_v(\Pi_{\Sigma}) = \text{var}(\Pi_{\Sigma}) \setminus \text{var}_p(\Pi_{\Sigma})$.

5. This is because the constraints in Σ may not specify a \mathcal{C} -formula.

Given a set Δ of CDCs over R and a horizontal decomposition Σ of R , the propositional theory associated with $\Delta \cup \Sigma$ is $\Pi = \Pi_\Delta \cup \Pi_\Sigma$, where Π_Δ and Π_Σ are the propositional theories of Definitions 2 and 5 associated with Δ and Σ , respectively. The set $\text{var}(\Pi) = \text{var}(\Pi_\Delta) \cup \text{var}(\Pi_\Sigma)$ of propositional variables occurring in Π is partitioned into $\text{var}_p(\Pi) = \text{var}_p(\Pi_\Delta) \cup \text{var}_p(\Pi_\Sigma)$ and $\text{var}_v(\Pi) = \text{var}_v(\Pi_\Delta) \cup \text{var}_v(\Pi_\Sigma)$.

Definition 6. Let Δ be a set of CDCs over R and let Σ be a horizontal decomposition of R . The auxiliary theory Π_{aux} for $\Pi = \Pi_\Delta \cup \Pi_\Sigma$ consists of the propositional formulae in Π_Σ whose r.h.s. is \perp and the axioms of unique value for $\text{var}_p(\Pi)$.

Observe that the above is a proper extension of Definition 3: whenever Σ is empty, the auxiliary theory for Π coincides with the auxiliary theory for Π_Δ .

Example 4. The propositional theory associated with Σ of Example 1 is $\Pi_\Sigma = \{-p_2^a \wedge p_3^b \rightarrow \perp, \top \rightarrow v_2', \neg p_3^b \rightarrow \perp\}$. Let $\Pi = \Pi_\Delta \cup \Pi_\Sigma$, where Π_Δ is the propositional theory already given in Example 2. The association between the propositional variables in $\text{var}_v(\Pi)$ and UTVPIs is idf of Example 2 extended with $v_2' \mapsto y_2 < 4$, and the auxiliary theory for Π is $\Pi_{\text{aux}} = \{-p_2^a \wedge p_3^b \rightarrow \perp, \neg p_3^b \rightarrow \perp\}$.

Definition 7. Let Σ be a horizontal decomposition, and let α be a valuation of $\text{var}_p(\Pi_\Sigma)$. The α -filtering of Π_Σ is the set

$$\Pi_\Sigma^\alpha = \{\neg \text{idf}(v') \mid (P \rightarrow v') \in \Pi_\Sigma, \alpha(P) = \top, v' \neq \perp\} \quad (7)$$

consisting of the negation of \mathcal{C} -formulae associated with propositional variables that occur in some propositional formula of Π_Σ whose l.h.s. holds true under α .

Observe that in (7), differently from (5), \mathcal{C} -formulae are negated. This is because a counter-instance I to losslessness is such that $V_1^I \cup \dots \cup V_n^I = \emptyset$ and R^I has only one tuple; therefore, whenever the formula $\lambda(\bar{x})$ in the selection that defines a view symbol is satisfied by I , the \mathcal{C} -formula $\delta(\bar{y})$, if any, is not. On the other hand, the \mathcal{C} -formula in the consequent of a CDC must hold whenever the condition in the antecedent is satisfied.

For a valuation α of $\text{var}_p(\Pi)$, the α -filtering of Π is the set $\Pi^\alpha = \Pi_\Delta^\alpha \cup \Pi_\Sigma^\alpha$, which, as \mathcal{C} is closed under negation, consists of \mathcal{C} -formulae.

The main result of this section is the following characterisation of lossless horizontal decomposition in terms of unsatisfiability in \mathcal{C} .

Theorem 3. Let Σ be a horizontal decomposition of R , let Δ be a set of CDCs over R , and let Π_{aux} be the auxiliary theory for $\Pi = \Pi_\Delta \cup \Pi_\Sigma$. Then, Σ is lossless under Δ if and only if the α -filtering $\Pi^\alpha = \Pi_\Delta^\alpha \cup \Pi_\Sigma^\alpha$ of Π is unsatisfiable for every valuation α of $\text{var}_p(\Pi)$ satisfying Π_{aux} .

Whenever the satisfiability of \mathcal{C} -formulae is decidable, Theorem 3 provides an algorithm for deciding whether a given horizontal decomposition is lossless. We illustrate this in our running example with UTVPIs.

Example 5. Consider Π and Π_{aux} from Example 4. The only valuation of $\text{var}_p(\Pi)$ satisfying Π_{aux} is $\alpha = \{p_2^a \mapsto \top, p_3^b \mapsto \top\}$, for which the α -filtering of Π is

$$\Pi^\alpha = \underbrace{\{y_1 + y_2 \leq 5, y_2 \geq 2, y_1 - y_2 \geq 0\}}_{\Pi_\Delta^\alpha} \cup \underbrace{\{y_2 \geq 4\}}_{\Pi_\Sigma^\alpha}.$$

Note that $y_2 \geq 4$ in Π_Σ^α is $\neg \text{idf}(v_2')$, that is, the negation of $y_2 < 4$. The set $\Pi^\alpha = \Pi_\Delta^\alpha \cup \Pi_\Sigma^\alpha$ is unsatisfiable because from $y_1 + y_2 \leq 5$ and $y_2 \geq 4$ we get $y_1 \leq 1$, which together with $y_1 - y_2 \geq 0$ yields $y_2 \leq 1$, in conflict with $y_2 \geq 2$. So, the horizontal decomposition Σ is lossless under Δ .⁶

We will now give the proof of Theorem 3, for which we first need to prove two additional lemmas. In the following, and in the rest of the paper, let $\tilde{\varphi}$ denote the formula $\forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \leftrightarrow \bigvee_{V \in \mathcal{V}} V(\bar{x}, \bar{y})$, and recall that a horizontal decomposition Σ is lossless under Δ if and only if $\Delta \cup \Sigma \models \tilde{\varphi}$. We start by showing that, when Δ consists of CDCs, $\Delta \cup \Sigma$ does not entail $\tilde{\varphi}$ precisely if there is a counterexample to it with only one tuple.

Lemma 3. Let Σ be a horizontal decomposition of R and let Δ be a set of CDCs over R . Then, $\Delta \cup \Sigma \not\models \tilde{\varphi}$ if and only if there exists a tuple \bar{t} such that the instance $I = \{R(\bar{t})\}$ is a model of $\Delta \cup \Sigma$.

Proof. The “if” is trivial. For the “only if”, assume that $\Delta \cup \Sigma \not\models \tilde{\varphi}$. Then, there exists a model J of $\Delta \cup \Sigma$ such that $J \not\models \tilde{\varphi}$, that is, $R^J \neq V_1^J \cup \dots \cup V_n^J$. The extension of each V_i always contains a subset of the tuples in the extension of R under every instance, hence there must be $\bar{t} \in R^J$ such that $\bar{t} \notin V_i^J$ for every $i \in \{1, \dots, n\}$. Let $I = \{R(\bar{t})\}$; as every constraint in $\Delta \cup \Sigma$ is in one tuple and $J \models \Delta \cup \Sigma$, we have that $I \models \Delta \cup \Sigma$. \square

The next lemma is more technical: intuitively, it shows that any tuple that is consistent with a valuation α satisfying the auxiliary theory provides a counterexample to losslessness if and only if the assignment induced by its interpreted positions satisfies the α -filtering.

Lemma 4. Let Σ be a horizontal decomposition of R , let Δ be a set of CDCs over R , and let Π_{aux} be the auxiliary theory for $\Pi = \Pi_\Delta \cup \Pi_\Sigma$. Let α be a valuation of $\text{var}_p(\Pi)$, let \bar{t} be a tuple consistent with α , and let β be the assignment induced by the interpreted positions of \bar{t} . Whenever $\alpha \models \Pi_{\text{aux}}$, we have that $\{R(\bar{t})\} \models \Delta \cup \Sigma$ if and only if β satisfies Π^α .

Proof. Let $n = |R|$ and $k = \|R\|$.

Claim 1. Let $\text{prop}(\phi) = P \rightarrow v'$, with $\phi \in \Sigma$. Then, $I \not\models \phi$ iff $\alpha(P) = \top$ and, whenever $v' \neq \perp$, $\beta \models \text{idf}(v')$.

Proof. Since $\phi \in \Sigma$ has the form (3), $I \not\models \phi$ iff $\lambda(\bar{x})$ is true under $\{x_1 \mapsto \bar{t}[1], \dots, x_k \mapsto \bar{t}[k]\}$ and $\sigma(\bar{y})$, if any, holds true under $\{y_1 \mapsto \bar{t}[k+1], \dots, y_{n-k} \mapsto \bar{t}[n]\}$. As $v' \neq \perp$ iff ϕ contains a \mathcal{C} -formula $\sigma(\bar{y}) = \text{idf}(v')$, the claim follows by construction of α and β .

Assume $\alpha \models \Pi_{\text{aux}}$. We will show that $I \not\models \Delta \cup \Sigma$ if and only if β does not satisfy Π^α .

“if”. Assume $\beta \not\models \Pi^\alpha$. Then, there is a \mathcal{C} -formula $\psi \in \Pi^\alpha$ that is not satisfied by β . By construction of Π^α , either ψ or its negation appear in some $\phi \in \Delta \cup \Sigma$, depending on whether $\phi \in \Delta$ or $\phi \in \Sigma$, respectively. If

6. In the scenario of our running example it would make sense to require salaries and bonuses to be non-negative quantities, which can be done by consistently adding the CDCs $\top \rightarrow y_1 \geq 0$ and $\top \rightarrow y_2 \geq 0$ without affecting the losslessness of the decomposition.

$\phi \in \Delta$, then $\text{prop}(\phi) = P \rightarrow v$ with $\alpha(P) = \top$, so $I \not\models \phi$ (by Claim 2 in the proof of Lemma 1). If $\phi \in \Sigma$, then $\text{prop}(\phi) = P \rightarrow v'$ with $v' \neq \perp$ and $\alpha(P) = \top$, hence $I \not\models \phi$ by Claim 1. In either case $I \not\models \Delta \cup \Sigma$.

“only if”. Assume $I \not\models \Delta \cup \Sigma$. Then, there is some $\phi \in \Delta \cup \Sigma$ that is not satisfied by I . If ϕ is in Δ , by Lemma 1 $\beta \not\models \Pi_\Delta^\alpha$, hence $\beta \not\models \Pi^\alpha$. If ϕ is in Σ , $\text{prop}(\phi) = P \rightarrow v'$; as $I \not\models \phi$, by Claim 1 $\alpha(P) = \top$ and $v' \neq \perp$ implies $\beta \models \text{idf}(v)$. Suppose $v' = \perp$, then $\text{prop}(\phi)$ is in Π_{aux} and, since $\alpha \models \Pi_{\text{aux}}$, we obtain $\alpha(P) = \top$, which is a contradiction. So, $v' \neq \perp$ and $\beta \models \text{idf}(v')$. In turn, we have that $\beta \not\models \neg \text{idf}(v')$ and $\neg \text{idf}(v') \in \Pi^\alpha$. Therefore, $\beta \not\models \Pi^\alpha$. \square

Proof of Theorem 3. Let $n = |R|$ and $k = \|R\|$. We will show that $\Delta \cup \Sigma \not\models \tilde{\varphi}$ if and only if there exist α and β satisfying Π_{aux} and Π^α , respectively.

“if”. Let α and β be such that $\alpha \models \Pi_{\text{aux}}$ and $\beta \models \Pi^\alpha$. Since $\alpha \models \Pi_{\text{aux}}$, no two distinct propositional variables in $\text{var}_p(\Pi)$ associated with the same position are both true under α . Hence, there is a tuple \bar{t} consistent with α and such that β is the assignment induced by its interpreted positions. So, the instance $I = \{R(\bar{t})\}$ is a model of $\Delta \cup \Sigma$ by Lemma 4. Thus, as $I \not\models \tilde{\varphi}$, $\Delta \cup \Sigma \not\models \tilde{\varphi}$ by Lemma 3.

“only if”. Assume that $\Delta \cup \Sigma \not\models \tilde{\varphi}$. By Lemma 3, there exists a tuple \bar{t} such that $I = \{R(\bar{t})\}$ is a model of $\Delta \cup \Sigma$. Let β be the assignment induced by the interpreted positions of \bar{t} , and let α be the valuation such that, for each $p \in \text{var}_p(\Pi)$, $\alpha(p) = \top$ if $p = p_i^a$ and $\bar{t}[i] = a$, and $\alpha(p) = \text{F}$ otherwise. We will show that α satisfies Π_{aux} and, in turn, $\beta \models \Pi^\alpha$ by Lemma 4, since $I \models \Delta \cup \Sigma$. By construction, α satisfies every propositional formula in Π_{aux} of the form $p_i^a \wedge p_i^b \rightarrow \perp$ with $p_i^a, p_i^b \in \text{var}_p(\Pi)$ and $p_i^a \neq p_i^b$. All other propositional formulae in Π_{aux} have the form $\text{prop}(\phi) = P \rightarrow \perp$, where ϕ is a constraint in Σ that does not contain a \mathcal{C} -formula $\sigma(\bar{y})$. As $I \models \Delta \cup \Sigma$, the condition $\lambda(\bar{x})$ in each such $\phi \in \Sigma$ is not true under $\{x_1 \mapsto \bar{t}[1], \dots, x_k \mapsto \bar{t}[k]\}$. Therefore, $\text{prop}(\phi) = P \rightarrow \perp$ is true under α as $\alpha(P) = \text{F}$ by construction of α . \square

The unsatisfiability problem for \mathcal{C} is the complement of the satisfiability problem for \mathcal{C} .

Lemma 5. *The unsatisfiability problem for \mathcal{C} linearly reduces to the losslessness problem in \mathcal{C} .*

Proof. Let $\Gamma = \{\phi_1, \dots, \phi_n\}$ be a set \mathcal{C} -formulae. We will show how to construct a horizontal decomposition that is lossless under $\Delta = \emptyset$ precisely if Γ is unsatisfiable. To this end, take $\Sigma = \{V_i : \neg \phi_i \mid \phi_i \in \Gamma\}$ and observe that, as \mathcal{C} is closed under negation, $\neg \phi_i \in \mathcal{C}$. Thus, Σ consists of selections of the form (3), where $\sigma = \neg \phi_i$ and $\lambda = \top$. Therefore, Σ is indeed a horizontal decomposition.

Let $\Pi = \Pi_\Delta \cup \Pi_\Sigma = \emptyset \cup \{v'_i \mid i = 1, \dots, n\}$ for which $\text{idf} = \{v'_i \mapsto \neg \phi_i \mid \phi_i \in \Gamma\}$. Then, the auxiliary theory for Π is $\Pi_{\text{aux}} = \emptyset$. Since $\text{var}_p(\Pi) = \emptyset$, the only valuation of $\text{var}_p(\Pi)$ is $\alpha = \emptyset$, which satisfies Π_{aux} and for which $\Pi^\alpha = \{\neg \text{idf}(v'_i) \mid v'_i \in \text{var}_v(\Pi_\Sigma)\} = \Gamma$. Therefore, by Theorem 3, Σ is lossless under $\Delta = \emptyset$ if and only if Γ is unsatisfiable. The reduction is linear in the size of Γ . \square

With regard to the losslessness problem in the languages of UTVPs and BUTVPs, we have the following complexity results.

Theorem 4. *When \mathcal{C} is the language of either BUTVPs or UTVPIs, the losslessness problem in \mathcal{C} is co-NP-complete.*

Proof sketch. Constructing the propositional theories Π and Π_{aux} takes linear time, checking whether a valuation α of $\text{var}_p(\Pi)$ satisfies Π_{aux} requires polynomial time, and checking that an assignment of integers to the variables in \bar{y} satisfies Π^α (whose construction takes linear time) can be done in polynomial time, whether Π^α consists of UTVPs or BUTVPs. Hence, in light of Theorem 3, we can verify a given solution to the complement of the losslessness problem, in either language, in polynomial time. Therefore, the losslessness problem is in co-NP in both cases.

The co-NP-hardness in the case of BUTVPs follows by Lemma 5 from the fact that the satisfiability problem for BUTVPI-formulae is NP-hard and so its complement is, in turn, co-NP-hard. The co-NP-hardness in the case of UTVPIs can be shown by a reduction from UNSAT. \square

5 ADDING FDS AND UINDS

So far, we have considered lossless horizontal decomposition under CDCs in isolation; in this section, we extend our study to the case in which the integrity constraints over the source schema are combinations of CDCs with more traditional database constraints. This investigation is vital to understand whether, how and to what extent the techniques we described in Section 4 can be applied to existing database schemas on which a set of integrity constraints other than CDCs is already defined.

Here, we focus on two well-known classes of integrity constraints, namely functional dependencies and unary inclusion dependencies [11]. Under certain restrictions—as we shall see—their interaction with CDCs can be fully captured, w.r.t. lossless horizontal decomposition, in terms of CDCs. It is important to remark that we consider restrictions solely on the CDCs, so that existing integrity constraints need not be modified in any way in order to allow for CDCs.

Let us recall that an instance I satisfies a UIND $R[i] \subseteq R[j]$ if every value in the i th column of R^I appears in the j th column of R^I . The following example shows that, if we allow CDCs together with constraints from another class, such as UINDs, their interaction may influence the losslessness of horizontal decomposition.

Example 6. Let R and V be relation symbols of arity 2, whose positions are interpreted over the integers. Let Σ be the horizontal decomposition defined by $V : y_1 > 3$, and let Δ be a set of integrity constraints on R consisting of the CDC $\top \rightarrow y_2 > 3$ and the UIND $R[1] \subseteq R[2]$. It is easy to see that Δ entails $\top \rightarrow y_1 > 3$. Therefore, Σ is lossless under Δ because V selects all of the tuples in R , which is clearly not the case without the UIND.

We now introduce a general property, *separability*, that will constitute the main technical tool for the subsequent analysis of combinations of CDCs with FDs and UINDs.

Informally, a class of constraints is separable from CDCs if, after making explicit the result of their interaction, which is captured by a suitable set of inference rules, we can disregard constraints from that class and focus solely on CDCs, as far as lossless horizontal decomposition is concerned. In what follows, for a set Δ of constraints we denote by $\text{cdc}(\Delta)$ the maximal subset of Δ consisting solely of CDCs.

Definition 8 (Separability). Let \mathbf{C} be a class of integrity constraints, let \mathcal{S} be a finite set of sound inference rules⁷ for \mathbf{C} extended with CDCs, and let Δ consist of CDCs and \mathbf{C} -constraints. We say that the \mathbf{C} -constraints are \mathcal{S} -separable in Δ from the CDCs if every horizontal decomposition is lossless under Δ exactly when it is lossless under $\text{cdc}(\Delta^*)$, where Δ^* denotes the \mathcal{S} -closure of Δ .⁸ We say that the \mathbf{C} -constraints are separable if there is some \mathcal{S} for which they are \mathcal{S} -separable.

Thus, to check whether a horizontal decomposition Σ is lossless under an \mathcal{S} -separable combination Δ of CDCs and other constraints, one can proceed as follows:

- 1) compute the deductive closure Δ^* of Δ w.r.t. \mathcal{S} , which makes explicit the interaction between CDCs and the other constraints in Δ by adding entailed constraints;
- 2) by using the technique of Section 4, check whether Σ is lossless under $\text{cdc}(\Delta^*)$, that is, the set obtained by discarding from Δ^* all of the constraints that are not CDCs.

Observe that \mathcal{S} -separability implies \mathcal{S}' -separability for every sound $\mathcal{S}' \supseteq \mathcal{S}$.

5.1 Functional Dependencies

We begin our investigation of separability by showing that FDs do not interact with CDCs and so, as far as the losslessness of horizontal decomposition is concerned, they can be freely allowed in combination with them.

Theorem 5. Let Δ be a set of CDCs and FDs. Then, the FDs are \emptyset -separable in Δ from the CDCs.

Proof. We will prove that a horizontal decomposition is lossless under Δ if and only if it is lossless under $\text{cdc}(\Delta)$.

“if”. We have that $\text{cdc}(\Delta) \subseteq \Delta$ and, in turn, Δ entails $\text{cdc}(\Delta)$; therefore $\text{cdc}(\Delta) \models \tilde{\varphi}$ implies $\Delta \models \tilde{\varphi}$.⁹

“only if”. Whenever a horizontal decomposition is not lossless, by Lemma 3 there is a witness instance I with only one tuple. Since the violation of an FD involves at least two tuples, I satisfies all of the FDs in Δ .¹⁰ \square

5.2 Unary Inclusion Dependencies

Since in general it is not possible to compare values from dom with values from idom , we consider only UINDs of the form $R[i] \subseteq R[j]$ where positions i and j are either both non-interpreted or both interpreted. We refer to the UINDs

in the former case as X -UINDs and in the latter as Y -UINDs. Let $n = |R|$ and $k = \|R\|$; we write $R[x_i] \subseteq R[x_j]$ with $i, j \in \{1, \dots, k\}$ to denote the X -UIND $R[i] \subseteq R[j]$ and we write $R[y_i] \subseteq R[y_j]$ with $i, j \in \{1, \dots, n - k\}$ to denote the Y -UIND $R[i + k] \subseteq R[j + k]$.

5.2.1 UINDs on Interpreted Attributes

First, we study the interaction between Y -UINDs (that is, UINDs at interpreted positions) and a restricted form of CDCs, which we shall introduce shortly. This interaction is captured by the following *domain propagation* rule:

$$\frac{\top \rightarrow \delta(y_i) \quad R[y_j] \subseteq R[y_i]}{\top \rightarrow \delta(y_j)}, \quad (\text{dp})$$

whose soundness is easily shown below.

Theorem 6. Let Δ be a set of CDCs and UINDs. If $\Delta \models \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_i)$ and $\Delta \models R[y_j] \subseteq R[y_i]$, then $\Delta \models \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_j)$.

Proof. If Δ is inconsistent, the claim follows trivially. Thus, let I be a model of Δ ; hence I satisfies the CDC $\forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_i)$ and the UIND $R[y_j] \subseteq R[y_i]$. If $R^I = \emptyset$, then trivially $I \models \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_j)$. So, let $R^I \neq \emptyset$ and suppose $I \not\models \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_j)$. Then, there exists $\bar{t} \in R^I$ for which $\delta(\bar{t}[j + k])$ holds true, with $k = \|R\|$. By the UIND, there must be $\bar{t}' \in R^I$ such that $\bar{t}'[i + k] = \bar{t}[j + k]$. Hence $\delta(\bar{t}'[i + k])$ is not true, in contradiction of $I \models \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \rightarrow \delta(y_i)$. \square

It turns out that when all of the CDCs that mention a variable y corresponding to an interpreted position affected by some Y -UIND have the form $\top \rightarrow \delta(y)$, the domain propagation rule fully captures the interaction between such CDCs and Y -UINDs w.r.t. losslessness.

Definition 9. We say that a set Δ of CDCs and Y -UINDs is *dp-controllable* if, for every Y -UIND $R[y_i] \subseteq R[y_j]$ in Δ with $i \neq j$, all of the CDCs in Δ mentioning the variable y , where y is y_i or y_j , are of the form $\top \rightarrow \delta(y)$.

Theorem 7. Let Δ be a *dp-controllable* set of CDCs and Y -UINDs. Then, the Y -UINDs are $\{(\text{dp})\}$ -separable in Δ from the CDCs.

The above theorem is a special case of a more general result (Theorem 10) given later on.

Even though in general *dp-controllability* is not a necessary condition for the $\{(\text{dp})\}$ -separability of Y -UINDs from CDCs, the following examples show two different situations where, in the absence of *dp-controllability*, the Y -UINDs are not $\{(\text{dp})\}$ -separable from the CDCs.

Example 7. Let R be a ternary relation symbol, whose last two positions are interpreted over the integers. Let Δ consist of the Y -UIND $R[y_1] \subseteq R[y_2]$ and of the CDCs $x_1 = a \rightarrow y_2 > 2$, $x_1 \neq a \rightarrow y_1 > 0$, $x_1 \neq a \rightarrow y_1 < 0$, and consider the view symbol $V : y_1 > 1$. For $x_1 \neq a$ there is no suitable value for y_1 to satisfy the above CDCs, thus every model I of Δ is such that, for every $\bar{t} \in R^I$, $\bar{t}[1] = a$ and $\bar{t}[3] > 2$. Moreover, by the Y -UIND $R[y_1] \subseteq R[y_2]$, we also have that $\bar{t}[2] > 2$, and therefore every tuple in R^I is

7. We assume the reader to be familiar with the standard notions (from proof theory) of inference rule, soundness, deductive closure.

8. As the constraints that are not CDCs are in any case filtered out from Δ^* , it does not matter whether \mathbf{C} extended with CDCs is closed under \mathcal{S} or not.

9. Recall that $\tilde{\varphi} = \forall \bar{x}, \bar{y}. R(\bar{x}, \bar{y}) \leftrightarrow \bigvee_{i=1}^n V_i(\bar{x}, \bar{y})$.

10. As a matter of fact, it satisfies any FD.

also in V^I , which means that V is lossless under Δ . Clearly, this is not the case in the absence of the Y-UIIND, that is, under $\text{cdc}(\Delta)$. Let Δ^* be the $\{(dp)\}$ -closure of Δ . Then, as $\Delta^* = \Delta$, we have that V is lossy under $\text{cdc}(\Delta^*)$ and, therefore, the Y-UIIND is not $\{(dp)\}$ -separable in Δ from the CDCs.

Example 8. Let R be a relation symbol of arity 4 and with all of its positions interpreted over the integers. Consider the view symbol $V : y_3 < 3 \wedge y_4 > 4$, and let Δ consist of the Y-UIIND $R[y_1] \subseteq R[y_2]$ and the CDCs $\top \rightarrow y_1 + y_3 > 0$, $\top \rightarrow y_2 + y_4 \leq 0$, $\top \rightarrow y_3 - y_4 \leq 0$. The above CDCs entail $\top \rightarrow y_1 - y_2 \geq 1$, thus in every model I of Δ each tuple $\bar{t} \in R^I$ must be such that $\bar{t}[1] - \bar{t}[2] \geq 1$.

By the Y-UIIND $R[y_1] \subseteq R[y_2]$, for each d in $\pi_2(R^I)$ ¹¹ there exists $d' \in \pi_2(R^I)$ with $d' \geq d + 1$. Then, as $d' \neq d$, the instance I is either infinite or empty. Hence, every horizontal decomposition is lossless under Δ .

On the other hand, let Δ^* be the $\{(dp)\}$ -closure of Δ and observe that $\Delta^* = \Delta$. Let $J = \{R(1, 0, 0, 0)\}$; then, since $J \models \text{cdc}(\Delta^*)$, V is lossy under $\text{cdc}(\Delta^*)$. Therefore, the Y-UIIND is not $\{(dp)\}$ -separable in Δ from the CDCs.

5.2.2 UIINDs on Non-Interpreted Attributes

We now turn our attention to combinations of CDCs and X-UIINDs (i.e., UIINDs at non-interpreted positions). First, we show that the syntactic restrictions introduced in [16] on the CDCs are *not* sufficient for the \emptyset -separability of X-UIINDs. Indeed, the following is a counterexample to Theorem of [16].

Example 9. Let R be a ternary relation symbol, with the third position interpreted over the integers. Let Δ consist of the CDC $x_2 = a \rightarrow y_1 \leq 0 \wedge y_1 > 0$ and the X-UIIND $R[1] \subseteq R[2]$. The CDCs in Δ are trivially *non-overlapping* with the UIINDs [16] and *partition-free* [16]. Consider the horizontal decomposition Σ specified by the selections $V_1 : x_1 \neq a$, $V_2 : x_2 \neq b$ and $V_3 : y_1 \neq 0$. Observe that every tuple other than $(a, b, 0)$ is captured by at least one of the above selections. Let $I = \{R(a, b, 0)\}$; clearly, $I \models \text{cdc}(\Delta) \cup \Sigma$ but $I \not\models \tilde{\varphi}$, hence Σ is not lossless under $\text{cdc}(\Delta)$. However, Σ is lossless under Δ as every model of $\Delta \cup \Sigma$ also satisfies $\tilde{\varphi}$. This is due to the fact that there exists no instance J such that $J \models \Delta$ and $R(a, b, 0) \in J$. Indeed, to satisfy the UIIND $R[1] \subseteq R[2]$, such an instance J must also contain a tuple $\bar{t} \in R^J$ with $\bar{t}[2] = a$ which, on the other hand, does not satisfy the CDC $x_2 = a \rightarrow y_1 \leq 0 \wedge y_1 > 0$. Hence, the X-UIINDs are not \emptyset -separable in Δ from the CDCs.

Below, we introduce a restriction on the CDCs, which ensures the \emptyset -separability of the X-UIINDs.

Definition 10. A set Δ of CDCs is globally consistent if, for every $\|R\|$ -tuple \bar{t}_x of **dom** constants, there is a tuple \bar{t}_y of $|R| - \|R\|$ values from **idom** such that the instance $\{R(\bar{t})\}$, with $\bar{t} = \langle \bar{t}_x, \bar{t}_y \rangle$, is a model of Δ .

Note that $\text{cdc}(\Delta)$ in Example 9 is not globally consistent.

11. π_i denotes projection on the i th position.

Theorem 8. Let Δ consist of CDCs and X-UIINDs such that $\text{cdc}(\Delta)$ is globally consistent. Then, the X-UIINDs are \emptyset -separable in Δ from the CDCs.

The above theorem, like Theorem 7, is a special case of a more general result (Theorem 10) given later on.

It is possible to check for the global consistency of a set Δ of CDCs in a way similar to the one described in Theorem 1 for consistency, by building the propositional theory Π_Δ associated with Δ , along with the auxiliary theory Π_{aux} for Π_Δ , and then checking that the α -filtering Π_Δ^α of Π_Δ is satisfiable for every (rather than just for one) valuation α of $\text{var}_p(\Pi_\Delta)$ that satisfies Π_{aux} . Indeed, under the assumptions of Theorem 1, Δ is globally consistent if and only if Π_Δ^α is satisfiable for every valuation α of $\text{var}_p(\Pi_\Delta)$ satisfying Π_{aux} . Checking for global consistency is expensive, because it requires an exponential number of satisfiability checks in \mathcal{C} ; the associated decision problem is in PSPACE (the space used for one satisfiability check can be reused for the next) for UTVPIs as well as for BUTVPIs.

Devising purely syntactic restrictions that guarantee the global consistency of CDCs on the specific constraint language \mathcal{C} in use, which is indeed what we overlooked in [16]. As it turns out, the non-overlapping and partition-free restrictions of [16] ensure global consistency (and so also the \emptyset -separability of X-UIINDs) only for sets of CDCs whose consequents are UTVPIs. This is not the case anymore for CDCs whose consequents are BUTVPIs, which indeed allow to express Example 9.

We provide a condition that, although not guaranteeing global consistency, ensures the \emptyset -separability of the X-UIINDs. Moreover, this restriction can be checked more efficiently than global consistency, as it requires only a polynomial number of \mathcal{C} -satisfiability checks.

Definition 11. Let Δ be a set of CDCs. We say that the CDCs in Δ are disjoint w.r.t. an X-UIIND $R[x_i] \subseteq R[x_j]$ if for any two CDCs $\phi_1(\bar{x}_1, \bar{y}_1)$ and $\phi_2(\bar{x}_2, \bar{y}_2)$ in Δ , with x_j in \bar{x}_1 , the consequent of ϕ_1 is satisfiable and has no variables in common with the consequent of ϕ_2 .

Intuitively, the above requires that all of the variables appearing in the consequent of any CDC ϕ whose antecedent mentions the variable x_j affected by an X-UIIND $R[x_i] \subseteq R[x_j]$ do not occur in the consequent of any other CDC; moreover, the consequent of each such ϕ must be satisfiable.

Theorem 9. Let Δ be a set of CDCs and X-UIINDs, where the CDCs are disjoint w.r.t. each X-UIIND in Δ . Then, the X-UIINDs are \emptyset -separable in Δ from the CDCs.

The above theorem is a special case of a more general result (Theorem 11) given later on in Section 5.3.

Clearly, as global consistency is a property of the CDCs in isolation, whereas disjointness is relative to a X-UIIND, these two notions are incomparable, in the sense that one does not imply the other and vice versa, as shown below.

Example 10. Let R be a ternary relation symbol, whose third position is interpreted over the integers, and let ψ be the X-UIIND $R[2] \subseteq R[1]$. The set Δ_1 consisting of the CDCs $x_1 = a \rightarrow y_1 < 0$ and $x_1 = a \rightarrow y_1 > 0$ is not globally consistent, as there is no suitable value for the

third position (associated with y_1) whenever the value of the first (associated with x_1) is a , but the CDCs in Δ_1 are disjoint w.r.t. ψ . On the other hand, the set Δ_2 consisting of $x_1 = a \rightarrow y_1 > 0$ and $x_2 = a \rightarrow y_1 > 1$ is not disjoint w.r.t. ψ (because both CDCs mention a variable affected by ψ and their consequents have y_1 in common) but Δ_2 is globally consistent.

5.2.3 UINDs on All Attributes

We now study the separability of UINDs (i.e., X-UINDs and Y-UINDs together)¹² from CDCs. The following is a generalisation of both Theorem 7 and Theorem 8.

Theorem 10. *Let Δ be a set of globally consistent CDCs, X-UINDs and Y-UINDs, such that the CDCs and Y-UINDs are dp-controllable. Then, the UINDs are $\{(dp)\}$ -separable in Δ from the CDCs.*

To give the proof of the above theorem, we will need to prove several lemmas, showing how any given model of the (saturated set of) CDCs can be extended in order to satisfy the UINDs as well.

Lemma 6. *Let Δ be a dp-controllable set of CDCs and Y-UINDs, and let \bar{t} be a tuple such that $\{R(\bar{t})\} \models \text{cdc}(\Delta^*)$, where Δ^* is the $\{(dp)\}$ -closure of Δ . Let $\psi = R[i] \subseteq R[j]$ be a Y-UIND in Δ , and let \bar{t}' be identical to \bar{t} except for $\bar{t}'[j] = \bar{t}[i]$. Then, $\{R(\bar{t}')\} \models \text{cdc}(\Delta^*) \cup \{\psi\}$.*

Proof. Let $\Delta' = \text{cdc}(\Delta^*)$. Since all of the UINDs in Δ are Y-UINDs, $i, j > k$ with $k = \|R\|$. As \bar{t} satisfies Δ' and \bar{t}' differs from \bar{t} only on the j th element, \bar{t}' satisfies every CDC in Δ' not mentioning the variable y_{j-k} . The only CDCs in Δ' which are allowed to mention y_{j-k} have the form $\top \rightarrow \delta(y_{j-k})$. For each such CDC, since $R[i] \subseteq R[j]$ is in Δ' , by (dp) also $\top \rightarrow \delta(y_{i-k})$ is in Δ' . Hence $\delta(\bar{t}[i])$ holds true, and in turn $\delta(\bar{t}'[j])$ is true as well, because $\bar{t}'[j] = \bar{t}[i]$. Therefore, \bar{t}' satisfies all the CDCs of the form $\top \rightarrow \delta(y_{j-k})$. Moreover, \bar{t}' trivially satisfies the UIND ψ , as $\bar{t}'[i] = \bar{t}'[j] = \bar{t}[i]$. \square

Lemma 7. *Let Δ be a dp-controllable set of CDCs and Y-UINDs, and let I be a model of $\text{cdc}(\Delta^*)$, where Δ^* is the $\{(dp)\}$ -closure of Δ . Then, there exists an instance $J \supseteq I$ such that $J \models \Delta^*$.*

Proof sketch. Let $J_0 = I$. We will iteratively add tuples to J_0 so as to obtain a model of Δ^* . At each iteration k , proceed as follows:

- 1) Find a violation of some Y-UIND $R[i] \subseteq R[j]$ in Δ^* , that is, a value $d \in \pi_i(R^{J_k})$ which is not in $\pi_j(R^{J_k})$.
- 2) Take $\bar{t} \in R^{J_k}$ such that $\bar{t}[i] = d$ and $\bar{t}[j] \neq d$.
- 3) Let $J_{k+1} = J_k \cup \{R(\bar{t}')\}$, with \bar{t}' identical to \bar{t} except for $\bar{t}'[j] = d$.

By construction, J satisfies all of the Y-UINDs in Δ and, since at each step $\{R(\bar{t}')\} \models \text{cdc}(\Delta)$ by Lemma 6, it can be shown that $J \models \text{cdc}(\Delta)$ by an easy induction. \square

12. Recall that UINDs between non-interpreted and interpreted positions are not allowed, as they make little sense.

Lemma 8. *Let Δ consist of X-UINDs and globally consistent CDCs, and let I be a model of $\text{cdc}(\Delta)$. Then, there exists an instance J such that $J \supseteq I$ and $J \models \Delta$.*

Proof sketch. Let $\Delta' = \text{cdc}(\Delta)$ and $J_0 = I$. We will show how to build a model $J \supseteq I$ of Δ by iteratively adding tuples to J_0 . At each iteration k , proceed as follows:

- 1) Find a violation of some X-UIND $R[i] \subseteq R[j]$ in Δ , i.e., a **dom** constant $a \in \pi_i(R^{J_k})$ which is not in $\pi_j(R^{J_k})$.
- 2) Take $\bar{t} \in R^{J_k}$ such that $\bar{t}[i] = a$ and $\bar{t}[j] \neq a$.
- 3) Let $J_{k+1} = J_k \cup \{R(\bar{t}')\}$, where \bar{t}' agrees with \bar{t} at non-interpreted positions except for $\bar{t}'[j] = a$. Suitable values of \bar{t}' at interpreted positions exist by Definition 10 as Δ' is globally consistent.

By construction, J satisfies all of the X-UINDs in Δ and by an easy induction we get that $J \models \text{cdc}(\Delta)$, as at each step $\{R(\bar{t}')\} \models \text{cdc}(\Delta)$ by the global consistency of Δ' . \square

Proof of Theorem 10. Let Δ^* be the closure of Δ under $\{(dp)\}$, and let $\Delta' = \text{cdc}(\Delta^*)$. Observe that $\Delta^* \equiv \Delta$, as (dp) is sound by Theorem 6. According to Definition 8, we will show that $\Delta \cup \Sigma \models \tilde{\varphi}$ if and only if $\Delta' \cup \Sigma \models \tilde{\varphi}$.

“if”. As $\Delta' \cup \Sigma \subseteq \Delta^* \cup \Sigma$, every model of $\Delta^* \cup \Sigma$ is also a model of $\Delta' \cup \Sigma$. Hence, $\Delta^* \cup \Sigma \models \tilde{\varphi}$ and, since $\Delta^* \equiv \Delta$, in turn $\Delta \cup \Sigma \models \tilde{\varphi}$. Therefore, $\Delta \cup \Sigma \models \tilde{\varphi}$ whenever $\Delta' \cup \Sigma \models \tilde{\varphi}$.

“only if”. By contraposition. Assume that $\Delta' \cup \Sigma \not\models \tilde{\varphi}$. Then, as Δ' consists solely of CDCs, by Lemma 3 there is a tuple \bar{t} such that $I = \{R(\bar{t})\}$ satisfies $\Delta' \cup \Sigma$. In turn, as Δ' is over R , I is also a model of Δ' . By Lemma 8, there exists an instance $J' \supseteq I$ satisfying all of the X-UINDs in Δ^* and, by Lemma 7, there exists $J'' \supseteq J'$ satisfying all of the Y-UINDs in Δ^* . Moreover, by construction, for each tuple in J'' there is a tuple in J' having the same values at non-interpreted positions, thus J'' also satisfies all of the X-UINDs in Δ^* . Therefore, J'' is model of Δ^* and, as $\Delta^* \equiv \Delta$, of Δ as well. Let J be the instance over $\mathbf{R} \cup \mathbf{V}$ with $R^J = R^{J''}$ (the extension of each V_i under J is unambiguously determined by R^J). Clearly, $J \models \Delta \cup \Sigma$ but $J \not\models \tilde{\varphi}$, because $\bar{t} \in R^J$ while $\bar{t} \notin V_1^J \cup \dots \cup V_n^J$. \square

Next, we show that replacing global consistency of the CDCs in the assumptions of Theorem 10 by disjointness w.r.t. the X-UINDs yields another sufficient condition for the $\{(dp)\}$ -separability of the UINDs from the CDCs.

Theorem 11. *Let Δ be a set of CDCs and UINDs such that the CDCs are disjoint w.r.t. each X-UIND in Δ , and the CDCs and Y-UINDs are dp-controllable. Then, the UINDs are $\{(dp)\}$ -separable in Δ from the CDCs.*

The proof of the above theorem is analogous to that of Theorem 10, with the difference that in the “only if” direction the existence of the instance J' is guaranteed by the following lemma rather than Lemma 8.

Lemma 9. *Let Δ be a set of X-UINDs and CDCs such that the CDCs are disjoint w.r.t. each X-UIND in Δ , and let I be a*

model of $\text{cdc}(\Delta)$. Then, there exists an instance J such that $J \supseteq I$ and $J \models \Delta$.

Proof sketch. The construction of J is the same as in Lemma 8, with the difference that in step 3 suitable values for \vec{t}' at interpreted positions exist by the disjointness of the CDCs w.r.t. the X-UIINDs. \square

5.3 FDs and UIINDs Together

Unfortunately, the separability results presented above for combinations of CDCs and UIINDs do not automatically carry over to the case in which FDs are also present. In fact, although FDs do not directly interact with CDCs, they do in general interact with UIINDs,¹³ which in turn interact with CDCs.

We write FDs over R as implications between sets of positions of R (e.g., $\{1, 3\} \rightarrow \{4\}$). We call *X-FD* (resp., *Y-FD*) an FD whose l.h.s. and r.h.s. both consist of non-interpreted (resp., interpreted) positions; and we call *XY-FD* (resp., *YX-FD*) an FD where the l.h.s. consists of non-interpreted (resp., interpreted) positions and the r.h.s. of interpreted (resp., non-interpreted) ones.

The following generalises Theorem 7 in the presence of X-FDs and YX-FDs.

Theorem 12. *Let Δ be a set of CDCs, Y-UIINDs, X-FDs and YX-FDs, where the CDCs and Y-UIINDs are dp-controllable. Then, the X-FDs, YX-FDs and Y-UIINDs are $\{(\text{dp})\}$ -separable in Δ from the CDCs.*

Proof sketch. The proof of Theorem 10 can be modified as follows: in the “only if” direction take $J' = I$, which contains only the tuple \vec{t} , and construct J'' as in Lemma 7 (with $\Delta^* = \Delta$), by extending J' with tuples that have the same values as \vec{t} at non-interpreted positions. Therefore, J'' satisfies any FD whose r.h.s. is a set of non-interpreted positions. \square

Theorem 8 does not hold anymore in the presence of Y-FDs, that is, X-UIINDs and Y-FDs are not \emptyset -separable in general from globally consistent CDCs, as shown below.

Theorem 13. *There is a set of X-UIINDs, Y-FDs and globally consistent CDCs, in which the X-UIINDs and Y-FDs are not \emptyset -separable from the CDCs.*

Proof. Let R be a relation symbol of arity 4, whose last two positions are interpreted over the integers. Let Δ consist of the X-UIIND $R[1] \subseteq R[2]$, the Y-FD $R : \{3\} \rightarrow \{4\}$, and the CDCs $x_1 = a \wedge x_2 = b \rightarrow y_1 = 0 \wedge y_2 > 1$ and $x_2 = a \rightarrow y_1 = 0 \wedge y_2 < 1$. These CDCs are globally consistent, because their consequents are satisfiable and their antecedents are never true at the same time (as x_2 cannot be simultaneously equal to b and a). Let Σ be the horizontal decomposition specified by $V_1 : x_1 \neq a$ and $V_2 : x_2 \neq b$. Clearly, Σ is lossy under $\text{cdc}(\Delta)$ as the instance $I = \{R(a, b, 0, 2)\}$ satisfies $\text{cdc}(\Delta)$ and Σ ; indeed, the tuple $(a, b, 0, 2)$ is in R^I but it is not selected by either V_1 or V_2 . Suppose that Σ is lossy under Δ . Then, there exists a model

J of $\Delta \cup \Sigma$ and a tuple $\vec{t} \in R^J$ such that $\vec{t} \notin V_1^J \cup V_2^J$. By definition of V_1 and V_2 , we have that $\vec{t}[1] = a$ and $\vec{t}[2] = b$ and, in turn, $\vec{t}[3] = 0$ and $\vec{t}[4] > 1$ by the first CDC. By the X-UIIND, there must be $\vec{t}' \in R^J$ such that $\vec{t}'[2] = a$ and, in turn, $\vec{t}'[3] = 0$ and $\vec{t}'[4] < 1$ by the second CDC. But then, \vec{t} and \vec{t}' violate the Y-FD, since they agree on the third position but must differ on the fourth. So $J \not\models \Delta$, which is a contradiction. Hence, Σ is lossless under Δ , and we conclude that the X-UIIND and Y-FD are not \emptyset -separable in Δ from the CDCs. \square

The CDCs in the above proof are globally consistent, but not disjoint w.r.t. the X-UIIND. However, Theorem 9 does not hold either in the presence of Y-FDs, that is, not even disjointness is enough to ensure the \emptyset -separability of X-UIINDs and Y-FDs from CDCs.

Theorem 14. *There exists a set of CDCs, X-UIINDs and Y-FDs, in which the CDCs are disjoint w.r.t. each X-UIIND, but the X-UIINDs and Y-FDs are not \emptyset -separable from the CDCs.*

Proof. Let R be a relation symbol of arity 4 with its last two positions interpreted over the integers. Let Δ consist of the X-UIIND $R[1] \subseteq R[2]$, the Y-FD $R : \{3\} \rightarrow \{4\}$, and the CDC $x_2 = a \rightarrow y_1 = 0 \wedge y_2 = 2$, trivially disjoint with the X-UIIND. Consider the horizontal decomposition Σ specified by $V : x_1 \neq a \wedge x_2 \neq b \wedge y_1 \neq 0 \wedge y_2 \neq 1$. Clearly, Σ is lossy under $\text{cdc}(\Delta)$ because the instance $I = \{R(\vec{t})\}$, where $\vec{t} = (a, b, 0, 1)$, satisfies $\text{cdc}(\Delta)$ and Σ ; indeed, \vec{t} is not selected by V . Suppose that Σ is lossy under Δ ; since V selects any tuple other than \vec{t} , there is a model J of $\Delta \cup \Sigma$ such that $\vec{t} \in R^J$ but $\vec{t} \notin V^J$. By the X-UIIND, there must be $\vec{t}' \in R^J$ such that $\vec{t}'[2] = a$ and, in turn, $\vec{t}'[3] = 0$ and $\vec{t}'[4] = 2$ by the CDC. But then, \vec{t} and \vec{t}' violate the Y-FD, because they agree on the third position but differ on the fourth. So $J \not\models \Delta$, which is a contradiction. Hence, Σ is lossless under Δ , and we conclude that the X-UIIND and Y-FD are not \emptyset -separable in Δ from the CDCs. \square

6 DISCUSSION AND OUTLOOK

In this paper, we studied lossless horizontal decomposition under constraints in a setting where the values for some of the attributes in the schema are taken from an *interpreted* domain. Data values in such a domain can be compared in ways beyond equality, according to a first-order language \mathcal{C} . We did not make any assumption on \mathcal{C} , other than requiring it to be closed under negation.

In the above setting, we considered a class of integrity constraints, CDCs, based on those introduced in [16]. We have characterised the consistency of a set of CDCs in terms of satisfiability in \mathcal{C} and we have shown that the problem of deciding consistency is NP-complete when \mathcal{C} is the language of either UTVPIs or BUTVPIs.

We considered a more general form of selections than in [16] and characterised, in terms of unsatisfiability in \mathcal{C} , whether a horizontal decomposition specified by such selections is lossless under CDCs. We have shown that the problem of deciding losslessness is co-NP-complete when \mathcal{C} is the language of either UTVPIs or BUTVPIs.

We also considered losslessness under CDCs in combination with FDs and UIINDs. We introduced and

13. The interaction between FDs and UIINDs can be fully captured, as there is a sound and complete axiomatization for the finite implication of FDs and UIINDs [11].

TABLE 1

Summary of S -Separability Results (unr = Unrestricted, dpc = dp-Controllable, gc = Globally Consistent, dis = Disjoint)

Constraints	CDCs	S	Theorem
FDs	unr	\emptyset	5
Y-UINDs (+ X-FDs + YX-FDs)	dpc	{(dp)}	7 12
X-UINDs	gc dis	\emptyset	8 9
X-UINDs + Y-UINDs	dpc + gc dpc + dis	{(dp)}	10 11

studied the important notion of separability, which indicates whether constraints other than CDCs can be disregarded w.r.t. losslessness, after incorporating the effect of their interaction in terms of entailed CDCs. A summary of all the separability results presented in this paper is given in Table 1.

A promising direction for future research we are currently investigating is the generalisation of the separability results for UINDs to arbitrary inclusion dependencies (INDs). Observe that INDs, differently from UINDs, can affect both interpreted and non-interpreted attributes at the same time, e.g., in $R[x_1, y_1] \subseteq R[x_2, y_2]$. Some care is needed in allowing FDs in this setting as well, because logical implication for unrestricted combinations of FDs and INDs is undecidable and has no axiomatization [11].

Another interesting direction is that of allowing equalities between two variables in the antecedents of CDCs as well as in the selection conditions on non-interpreted attributes of view definitions. We believe our approach could be extended in this direction by representing such equalities by propositional variables and by adding suitable axioms to the auxiliary theory to handle transitivity and symmetry.

The main motivation for our study of lossless horizontal decomposition is that it provides the groundwork for the consistent and unambiguous propagation of updates in the context of selection views. By applying the general criterion of [6], given a lossless horizontal decomposition it is possible to determine whether an update issued on some (possibly all) of the fragments can be propagated to the underlying database without affecting the other fragments. Similarly, it is possible to partition the source relation by adding suitable conditions in the selections that define the fragments, so that each is disjoint with the others. In general, a lossy horizontal decomposition can always be turned into a lossless one by defining an additional fragment, called a *complement*, which selects the missing tuples. In particular, there is a *unique minimal complement* selecting all and only the rows of the source relation that are not selected by any of the other fragments. In follow-up work, we will show how to compute the definition of such a complement, in the scope of an in-depth study of partitioning and update propagation in the setting studied in this paper.

Most of the work in the field of horizontal decomposition has been carried out in the context of distributed databases systems, where one is mainly concerned with finding an optimal decomposition w.r.t. some parameters (e.g., workload,

query-execution time, storage quotas), rather than determining whether a *given* horizontal decomposition is lossless.

De Bra ([12], [13]) developed a theory of horizontal decomposition to partition a relation into two sub-relations such that one satisfies certain FDs that the other does not. The approach is based on constraints that capture partial implications between sets of FDs and exceptions to sets of FDs, for which a sound and complete set of inference rules is provided. These constraints are \emptyset -separable from our CDCs (for the same reason FDs are).

Maier and Ullman [15] consider horizontal decomposition involving physical and virtual fragments over the same attributes. Fragments are defined in an arbitrary (first-order) language closed under Booleans, where entailment is decidable and consisting of formulae that, as in our case, can be evaluated by examining one tuple at a time, in isolation from the others. Differently from our case, the language allows to express equalities between variables associated with non-interpreted attributes. But, if such equalities are forbidden, the setting of [15] can be recast into ours: the union of the physical fragments is the single source relation R^I we consider here, the definitions of the physical fragments can be taken as integrity constraints over R , and the definition of each virtual fragment (given in terms of the physical fragments and other virtual ones) can be expressed only in terms of R by query unfolding. Then, the problem of determining whether the virtual fragments constitute a lossless horizontal decomposition of the physical fragments, which is not addressed in [15], can be solved by applying the techniques we described in this paper. Virtual fragments in [15] are defined by selection and union, that is, in our notation, by formulae of either the form $\lambda(\bar{x}) \wedge \sigma(\bar{y})$ or $\lambda(\bar{x}) \vee \sigma(\bar{y})$. As we remarked in Section 2, in such a case losslessness can be checked by considering two views $\lambda(\bar{x})$ and $\sigma(\bar{y})$ in place of each view of the latter form.

ACKNOWLEDGMENTS

A preliminary version of this paper appeared in *Big Data: Proceedings of the 29th British National Conference on Databases*, vol. 7968 of LNCS, pp. 77–91, Springer, 2013, under the title “Lossless horizontal decomposition with domain constraints on interpreted attributes”. This work was supported by the Austrian Science Fund (FWF), project P25207-N23. The authors would like to thank the anonymous reviewers for their valuable comments and suggestions, which contributed to the improvement of the paper’s quality.

REFERENCES

- [1] F. Bancilhon and N. Spyratos, “Update semantics of relational views,” *ACM Trans. Database Syst.*, vol. 6, no. 4, pp. 557–575, Dec. 1981.
- [2] U. Dayal and P. A. Bernstein, “On the correct translation of update operations on relational views,” *ACM Trans. Database Syst.*, vol. 7, no. 3, pp. 381–416, Sept. 1982.
- [3] A. M. Keller, “Updating relational databases through views,” Ph.D. dissertation, Dept. Comput. Sci., Stanford Univ., Stanford, CA, USA, Feb. 1985.
- [4] Y. Peng, B. Choi, J. Xu, H. Hu, and S. S. Bhowmick, “Side-effect estimation: A filtering approach to the view update problem,” *IEEE Trans. Knowl. Data Eng.*, vol. 26, no. 9, pp. 2307–2322, Sep. 2013.
- [5] L. Caroprese, I. Trubitsyna, M. Truszczynski, and E. Zumpano, “The view-update problem for indefinite databases,” in *Proc. 13th Eur. Conf. Logics Artif. Intell.*, 2012, pp. 134–146.

- [6] E. Franconi and P. Guagliardo, "On the translatability of view updates," in *Proc. Mendelzon Workshop Found. Data Manage.*, 2012, pp. 154–167.
- [7] M. Johnson and R. Rosebrugh, "Constant complements, reversibility and universal view updates," in *Proc. 12th Int. Conf. Algebraic Methodol. Softw. Technol.*, 2008, pp. 238–252.
- [8] A. Bohannon, B. C. Pierce, and J. A. Vaughan, "Relational lenses: A language for updatable views," in *Proc. 25th ACM SIGMOD-SIGACT-SIGART Symp. Principles Database Syst.*, 2006, pp. 338–347.
- [9] S. J. Hegner, "An order-based theory of updates for closed database views," *Ann. Math. Artif. Intell.*, vol. 40, no. 1–2, pp. 63–125, Jan. 2004.
- [10] E. Franconi and P. Guagliardo, "Effectively updatable conjunctive views," presented at the 7th Alberto Mendelzon Int. Workshop Foundations of Data Manage., Puebla/Cholula, Mexico, 2013.
- [11] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*. Reading, MA, USA: Addison-Wesley, 1995.
- [12] P. D. Bra, "Horizontal decompositions in the relational database model," Ph.D. dissertation, Univ. Antwerp, Antwerpen, Belgium, 1987.
- [13] P. De Bra, "Horizontal decompositions based on functional-dependency-set-implications," in *Proc. Int. Conf. Database Theory*, 1986, pp. 157–170.
- [14] S. Ceri, M. Negri, and G. Pelagatti, "Horizontal data partitioning in database design," in *Proc. SIGMOD Conf.*, 1982, pp. 128–136.
- [15] D. Maier and J. D. Ullman, "Fragments of relations," in *Proc. SIGMOD Conf.*, 1983, pp. 15–22.
- [16] I. Feinerer, E. Franconi, and P. Guagliardo, "Lossless horizontal decomposition with domain constraints on interpreted attributes," in *Proc. 29th Brit. Nat. Conf. Big Data*, 2013, pp. 77–91.
- [17] C. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli, "Satisfiability modulo theories," in *Handbook of Satisfiability*. Amsterdam, The Netherlands: IOS Press, 2009, ch. 26, pp. 825–885.
- [18] A. Schutt and P. J. Stuckey, "Incremental satisfiability and implication for UTVPI constraints," *INFORMS J. Comput.*, vol. 22, no. 4, pp. 514–527, 2010.
- [19] S. K. Lahiri and M. Musuvathi, "An efficient decision procedure for UTVPI constraints," in *Proc. 5th Int. Workshop Frontiers Combining Syst.*, 2005, pp. 168–183.
- [20] J. Jaffar, M. J. Maher, P. J. Stuckey, and R. H. C. Yap, "Beyond finite domains," in *Proc. 2nd Int. Workshop Principles Practice Constraint Programm.*, 1994, pp. 86–94.
- [21] S. A. Seshia, K. Subramani, and R. E. Bryant, "On solving boolean combinations of UTVPI constraints," *J. Satisfiability, Boolean Model. Comput.*, vol. 3, no. 1–2, pp. 67–90, 2007.



Ingo Feinerer received the PhD degree in computer science from the Vienna University of Technology and the PhD degree in business administration from the Vienna University of Economics and Business. He has been an assistant professor at the Vienna University of Technology since 2008. In 2007, he was a postdoctoral visiting scholar in the Computer Science Department, Carnegie Mellon University. His expertise and research interests comprise configuration, formal methods, text mining, and databases.



Enrico Franconi is an associate professor at the Free University of Bozen-Bolzano, Italy, where he is the director of the KRDB Research Centre for Knowledge and Data. His main research interest is on knowledge representation and reasoning technologies applied to databases, and in particular description logics, temporal representations, conceptual modeling, intelligent access to information, and natural language interfaces. Recently, he has been involved as a principal investigator in many European projects (SEWASIE, DWQ), networks of excellence (KnowledgeWeb, InterOp), and Marie Curie actions (NET2). He is a member of the Advisory Committee of the World Wide Web Consortium (W3C). He is in the editorial board of the *Journal of Applied Logic (JAL)*, published by Elsevier. He is the author of more than 150 publications in international workshops, conferences, and journals.



Paolo Guagliardo completed the European Master in computational logic in 2009, received the MSc degree in computer science with honors from the Vienna University of Technology and the Free University of Bozen-Bolzano. He is currently working toward the PhD degree at the Free University of Bozen-Bolzano under the supervision of Enrico Franconi. His research interests include database theory, knowledge representation, and Semantic Web.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.