TLP **17** (2): 148–204, 2017. © Cambridge University Press 2016
doi:10.1017/S147106841600017X First published online 29 July 2016

Consistency and trust in peer data exchange systems

LEOPOLDO BERTOSSI
School of Computer Science, Carleton University, Ottawa, Canada
(*e-mail:* bertossi@scs.carleton.ca)

LORETO BRAVO
Facultad de Ingeniería, Universidad del Desarrollo, Santiago, Chile
(*e-mail:* bravo@udd.cl)

submitted 4 May, 2014; revised 4 September, 2015; accepted 4 May, 2016

Abstract

We propose and investigate a semantics for *peer data exchange systems* where different peers are related by data exchange constraints and trust relationships. These two elements plus the data at the peers' sites and their local integrity constraints are made compatible via a semantics that characterizes sets of *solution instances* for the peers. They are the intended – possibly virtual – instances for a peer that are obtained through a data repair semantics that we introduce and investigate. The semantically correct answers from a peer to a query, the so-called *peer consistent answers*, are defined as those answers that are invariant under all its different solution instances. We show that solution instances can be specified as the models of logic programs with a stable model semantics. The repair semantics is based on null values as used in SQL databases, and is also of independent interest for repairs of single databases with respect to integrity constraints.

KEYWORDS: Peer data exchange, answer set programs, disjunctive stable model semantics, metadata management, schema mappings, relational databases, integrity constraints, database repairs, consistency

1 Introduction

A *peer data exchange system* (PDES) (also known as a peer data management system) can be conceived as a finite set $\mathscr{P} = \{\mathtt{P}_1, \ldots, \mathtt{P}_n\}$ of peers, each of them owning a local relational schema and a database instance. When a peer P receives a query, in order to answer it, P's data is completed or modified according to relevant data that other peers may have. For this to be possible, P has to be directly related to some neighbor peers Q through sets, $\Sigma(\mathtt{P}, \mathtt{Q})$, of *logical mappings* or *data exchange constraints* (DECs), from P to Q. They are first-order sentences expressed in terms of the elements of the database schemas of P and Q. In their turn, P's neighbors may have their own neighbors, who through their data may, transitively and eventually, have an impact on P's data as well.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

DECs between two peers are expected to be satisfied by the combination of the two local instances. In this sense, DECs act as integrity constraints (ICs) on the combination of two schemas and their instances. However, they are not forced to be satisfied, and there is no mechanism in place for maintaining them. Actually, *it is the inconsistency of a DEC that will enable the movement of data between peers, at query-answering time*: When a peer P receives a query, it examines its DECs, on that basis poses queries to its neighbors, who return *consistent* data to P. P also queries its own database. If its DECs are not satisfied by the collected data, P, after a – possibly virtual – consistency repair process, composes a *consistent* answer to the original query. A peer's answering of a query received from an external user or another peer triggers an iterative process of interleaved consistency checking of DECs with respect to the data at hand, and the repair of the latter, if necessary.

This movement of data among peers is used for query answering, but not for updating the local physical instances with the purpose of having the DECs satisfied. Of course, this could be done by a peer who receives data from other peers, if it decides to do so. However, it is not the goal of data exchange in such a system to have all the instances synchronized and globally satisfying the DECs. Peers are autonomous, their instances are possibly being independently updated, and other peers are not expected to be aware of that. DECs constrain and trigger data exchange through their inconsistency, which is detected and resolved locally, by a single peer, namely the one who owns those DECs. *We do not assume any kind of central or global coordination; nor that any two neighbors interact for DEC evaluation.*

A peer P that is answering a query may, at query-answering time, import data from its neighbors, to complement its data and/or ignore part of its own data. The way a peer uses the imported data of course depends on its DECs, but also on the *trust relationships* to its neighbors: A peer P may trust its data the same as or less than a neighbor's data. As a consequence, *in our PDESs, data exchange is driven by inconsistency and constrained by trust*.

Example 1.1

Peers P1 and P2 have relational schemas $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot,\cdot), S^1(\cdot)\}$ and $\mathscr{S}(\mathtt{P2}) = \{R^2(\cdot,\cdot), S^2(\cdot,\cdot)\}$. Here, P1 is connected to peer P2 by P1's set of DECs:

$$\Sigma(\mathtt{P1},\mathtt{P2}) = \{\forall x\forall y(R^2(x,y) \wedge S^2(y,z) \rightarrow R^1(x,y,z)),\ \forall x(S^1(x) \rightarrow S^2(5,x))\}.$$

These DECs belong to P1 and P2 may not even be aware of them. Let's assume that P1 trusts P2 more than itself. The existence of DECs from P1 to P2, and the trust relationship are indicated in Figure 1 by the labeled arrow from one instance to the other.

We can see that the DECs are not satisfied by the combined instance, which is perfectly acceptable; we do not have to do anything.

Now, assume that the query $\mathscr{Q}(x) : \exists yz R^1(x,y,z)$ is posed to P1. In order to answer it, P1 notices through the DECs that P2's schema is related to its predicate R^1. Accordingly, P1 has to adjust its own data, so that the DECs with P2 are satisfied. In order to do this, P1 requests P2's data, by issuing to P2 atomic queries in

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

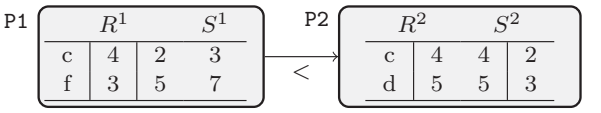


Fig. 1. A simple PDES.

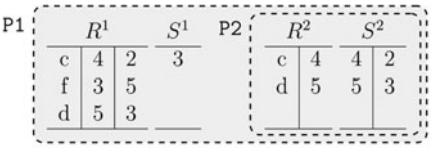


Fig. 2. A neighborhood solution instance for P1.

P2's language, namely $\mathcal{Q}_1(x, y) : R^2(x, y)$ and $\mathcal{Q}_2(x, y) : S^2(x, y)$.[1] Notice that, at this stage and from the point of view of P2, the queries $\mathcal{Q}_1$ and $\mathcal{Q}_2$ can be considered as queries posed by an external user, so as query $\mathcal{Q}$ was posed to P1.

Since P2 has neither DECs to other peers nor local ICs, it will return to P1 its physical data, without any modification, as query answers. Now P1, confronted with a combined instance, can detect that the first DEC is not satisfied, because the tuple $R^1(d, 5, 3)$ is not in R^1. Nor the second DEC, because the tuple $S^2(5, 7)$ is not in S^2.

Since P1 trusts P2's data more than its own, the inconsistencies are locally solved by P1 by inserting $R^1(d, 5, 3)$ into R^1, and deleting tuple $S^1(7)$ from S^1. These changes could be only virtual, and for the purpose of answering the query at hand. In this case, P1 has a single *neighborhood solution* (instance), that is, an instance for the combined schema that satisfies the DECs, and, in some sense, *minimally departs* from the (inconsistent) instance that had been formed at P1' neighborhood. This neighborhood solution for P1 is shown in Figure 2. The restriction of this instance to P1's schema is *a solution* (instance) for P1, and is used to answer the initial query posed to P1. Thus, the answers to $\mathcal{Q}$ are $\langle c \rangle, \langle f \rangle, \langle d \rangle$.

In the example, for illustration purposes, we used a particular and well-known *repair semantics*, i.e. a particular way of restoring consistency on a database that does not satisfy a given set of ICs[2], while still staying "as close as possible" to the given, inconsistent instance. In this case, repairs minimize, under set inclusion, the set of tuples that are deleted or inserted (Arenas *et al.* 1999). Repair semantics come with their minimality criteria. (Cf. Bertossi (2011) for a general discussion of repair semantics, and for references to different repair semantics.) In this work, we will introduce, investigate, and use a variation of the repair semantics used in the example.

In some cases, a peer P may have several neighborhood solutions, leading to also possibly several (local) solution instances. The answers returned by P (to an external user or another peer) will be the *peer consistent answers* (PCAs), i.e. those that are shared (or returned) by all the different solution instances for P. That is, a *cautious* (also known as *skeptical* or *certain*) semantics is applied to query answering

[1] Actually, P1 does not need the whole extensions of P2's tables. However, asking for P2's whole tables allow us to better illustrate our general approach.

[2] We emphasize that DECs can be seen as ICs on combined schemas and instances.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

(Leone *et al.* 2006; Brewka *et al.* 2011). The same do P's neighbors, who may have solutions of their own, when they hand over their data to P: They give away their own peer-consistent data.

The notion of solution instance for a peer is used as an auxiliary notion, to characterize the semantically correct answers from P's point of view. When trying to correctly answer a query, we try to avoid, as much as possible, the generation of material solutions instances. Ideally, a peer would compute its PCAs to a user query just by querying the already available local instances, and dealing with the involved DECs on-the-fly, at query-answering time.

In this work, we formalize these ideas, by first providing a general semantic framework for PDESs and their solution-semantics. It captures in abstract terms, and in this order, the notions of: (a) PDES, including a broad and useful syntactic class of DECs; (b) neighborhood solution for a peer, appealing to an *abstract repair semantics* that involves trust relationships; (c) solution for a peer, a recursive concept due to a possible (always finite) chain of peers building each, one after another, neighborhood solutions; and (d) PCAs.

The result is a formal semantics for a system of peers who exchange data for query answering. It is a model-theoretic *possible-world semantics* that characterizes the *class of intended instances for a peer*, the above mentioned *solutions*. The expected answers from a peer are *certain* with respect to that class, i.e. true in all of them.

Next, as a second main subject, we instantiate the general semantic framework, and introduce and investigate a specific repair (or solution) semantics for PDESs. We do so by making some specific commitments and assumptions: (a) Database instances may have null values as those used in SQL relational databases. (b) Those null values behave and are handled as in DBMSs that (at least partially) comply with the SQL Standard, in particular with respect to IC satisfaction and query answering. (c) The null-based solution semantics that we introduce heavily depends on the presence of nulls à la SQL, and on the use of nulls to restore consistency. (d) This null-based repair semantics comes with its own minimality criterion. (e) Query answering, and DEC and IC satisfaction follow a *a logic* in databases with null *à la* SQL that we formalize for conjunctive queries and a broad class of DECs and ICs.

For the gist, the null-based repair semantics we adopt does not allow, for example, to introduce a null value to satisfy a join, such as that with the existential z in the DEC $\forall x \forall y(R(x, y) \rightarrow \exists z(S(x, z) \wedge T(z, y)))$. In this case, a tuple deletion from R will be preferred. However, a null can be used for variable z in a tuple inserted to satisfy the DEC $\forall x \forall y(R(x, y) \rightarrow \exists z S(x, z))$ (a deletion from R is also possible).

Our decision to consider null values, as used in SQL relational database systems, for single database repairs and solutions in PDESs is based upon and motivated by the facts that: (a) in this way, our research becomes close to database practice; (b) it can be made compatible with and applicable in that practice; (c) SQL DBMSs implicitly implement a sort of null-based repair semantics (when maintaining ICs); (d) that "repair semantics" deserves a theoretical investigation; and (e) at least a sizeable part of the SQL standard and practice can be put on solid logical grounds. Actually, we are able to bring into this research essentially all the classes of ICs and logical mappings that are used in both database practice and research, and to treat them according to the just mentioned null-based SQL semantics.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

We emphasize, however, that our abstract framework is sufficiently modular and flexible to adopt alternative repair semantics for dealing with inconsistency and incompleteness of data.

We see our work also as a contribution to the subjects of repairs and consistent query answering (CQA) for single relational databases. Repairs based on the use of nulls; actually a single one with a semantics in the style of SQL – which we also logically formalize in this work – had been presented in preliminary form in Bravo and Bertossi (2006) and Bravo (2007). Here, we provide an extended and definitive formulation, and we apply it to peer data exchange.

We also investigate complexity-theoretic and algorithmic issues related to the solution semantics. More specifically, we show that deciding if a neighborhood instance is a neighborhood solution for a peer is *coNP*-complete in data, and deciding if a tuple is a PCA to a local query is Π_2^P-complete in data. We investigate a particular, important, and common case of the PDES semantics, where peers just import data from other peers to complete their own data set, without giving semantic priority to the latter. This leads to a reduction in data complexity for the solution checking problem.

Finally, we show that the null-based, model-theoretic semantics for PDESs can be captured by means of disjunctive logic programs with stable model semantics (Gelfond and Lifschitz 1991), also known as *answer set programs* (Brewka *et al.* 2011). More specifically, we establish an explicit correspondence between the solutions for a peer and the stable models of the program, obtaining, in this way, a declarative semantics for PDESs that can also be made executable.

In relation to most immediately related work, the idea of peer CQA resembles that of CQA in databases (Arenas *et al.* 1999): Consistent answers to a query posed to a database that may be inconsistent with respect to certain ICs are invariant under all *repairs*, i.e. the minimally repaired and consistent versions of the original instance. Furthermore, there are mechanisms for computing consistent answers that avoid or minimize the physical generation of repairs. Logic programs for specifying database repairs as their stable models have been proposed and investigated in Arenas *et al.* (2003), Greco *et al.* (2003), Barcelo and Bertossi (2003), Caniupan and Bertossi (2010). Cf. Bertossi (2011, 2006), Chomicki (2007) for recent surveys of CQA.

Our work can be classified among those on *semantic- and schema-based approaches to peer-to-peer data management systems* (Halevy *et al.* 2003, 2004; Calvanese *et al.* 2004; Franconi *et al.* 2004a; Bertossi and Bravo 2004b; Calvanese *et al.* 2005; Fuxman *et al.* 2006). The emphasis is on specifying the intended and legal database instances behind a system of peers who are connected with each other by means of schema mappings.

Different notions and forms of trust have been considered in P2P information sharing systems (Sabater and Sierra 2005; Marti and Garcia-Molina 2006; Artz and Gil 2007; Demolombe 2011) (cf. the special issue (Boutaba and Marshall 2006)), but not much in the context of semantic PDESs, where the trust relationships between peers are used to essentially modulate or qualify the use of the DECs between peers. Our emphasis is on the integration of trust and DECs, with a notion of trust that is closer to the notion of relative reputation or reliability of a peer as a source of

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

information, in relation to the quality of its data (Marti and Garcia-Molina 2006). In this work, trust relationships are given, and not computed or updated (Jøsang *et al.* 2006, 2012).

This paper is structured as follows. Section 2 provides technical preliminaries and basic elements of PDESs. Section 3 introduces a general semantic framework for PDESs. Section 4, in preparation for the introduction of what will be the official, concrete peer-solution semantics discusses and formalizes database instances with null values as used in SQL. It also formulates query answering and integrity satisfaction in those databases. Section 5 instantiates the abstract semantic framework proposed in Section 3 by introducing a particular consistency restoration semantics based on the use of null values. This section also investigates the complexity of some computational problems. Section 6 captures the semantics of the previous section in terms of logic programs with stable model semantics. In Section 7, we discuss related work. In Section 8, some final conclusions are drawn. The Online Appendix A discusses several additional and alternative approaches and issues with respect to the previously introduced general and specific semantics. The Online Appendix B gives proofs of our main results. This work considerably extends and develops the semantics for PDESs first suggested in Bertossi and Bravo (2004b) and further developed in Bertossi and Bravo (2007).

2 The basic PDES scenario

Every peer in a PDES will have a local relational schema with a local relational instance. For this reason, we recall first some basic notions from relational databases. A relational schema $\mathscr{S}$ consists of a set of relational predicates. A relational predicate $R \in \mathscr{S}$ with arity n and attributes $A_1, \ldots, A_n$, is commonly denoted with $R(A_1, \ldots, A_n)$, or sometimes simply $R(\cdot, \ldots, \cdot)$. Each attribute A of (a predicate in) $\mathscr{S}$ has a possibly infinite data domain, $Dom(A)$. In general, we will denote with $\mathscr{U}$ the union of the attribute domains, obtaining a single, possibly infinite *data domain*.

In addition to the *database predicates* in a relational schema $\mathscr{S}$, we have a set $\mathscr{B}$ of built-in predicates (that have a fixed semantics), e.g. comparison predicates, such as $=, \neq, <$, etc. We assume that $\mathscr{B}$ contains the propositional predicate **false** that is always false. (Later on, we will introduce some additional, specific, built-in predicates.)

The predicates of the relational schema $\mathscr{S}$ plus those in $\mathscr{B}$ (that we leave implicit), and the constants in $\mathscr{U}$ determine a language $\mathscr{L}(\mathscr{S})$ of first-order predicate logic. The schema may also contain *ICs*, that are sentences of $\mathscr{L}(\mathscr{S})$.

An instance D for a schema $\mathscr{S}$ is a finite set of a ground atoms of the form $R(\bar{c})$, where $R(A_1, \ldots, A_n) \in \mathscr{S}$ has some arity n, and $\bar{c} = \langle c_1, \ldots, c_n \rangle$, with $c_i \in Dom(A_i)$. For each n-ary predicate $R \in \mathscr{S}$, an instance D for $\mathscr{S}$ determines an extension for R that is a finite n-ary relation over the data domain. If $t \in D$, we denote by $t[\bar{A}]$ the sequence of values in $\bar{t}$ for attributes $A \in \bar{A}$. Given an instance D, the *active domain* of D, denoted, $Adom(D)$, is the finite subset of $\bigcup_A Dom(A)$ that contains all the constants that appear in relations in D.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

If $\mathcal{S}'$ is a subschema of $\mathcal{S}$, i.e. contains some of the relational predicates in $\mathcal{S}$, and D is an instance for $\mathcal{S}$, then $D \restriction \mathcal{S}'$ denotes the restriction of D to $\mathcal{S}'$, i.e. $D \restriction \mathcal{S}' = \{R(\bar{t}) \mid R \in \mathcal{S}' \text{ and } R(\bar{t}) \in D\}$.

A database instance D for the schema $\mathcal{S}$ serves then as an interpretation for the language $\mathcal{L}(\mathcal{S})$. If D is an instance for $\mathcal{S}$, and Ψ is a set of sentences of $\mathcal{L}(\mathcal{S})$, then $D \models \Psi$ denotes that D satisfies (makes true) all the sentences in Ψ. If Ψ is the set of ICs that comes with the schema $\mathcal{S}$ and $D \models \Psi$, we say that D is *consistent*. Otherwise, D is *inconsistent*[3].

A query $\mathcal{Q}(\bar{x})$ is a formula of a language $\mathcal{L}(\mathcal{S})$, where $\bar{x}$ is the list of free variables. A sequence of constants $\bar{c}$ is an answer to $\mathcal{Q}(\bar{x})$ in instance D for $\mathcal{S}$ if $D \models \mathcal{Q}[\bar{c}]$, i.e. the formula becomes true in D when the variables in $\bar{x}$ are replaced by the corresponding constants in $\bar{c}$. When $\bar{x}$ is empty, the query $\mathcal{Q}$ is a *Boolean query*, i.e. a sentence. In this case, the answer in D can be *yes* or *no* depending on whether it is true or not in D, denoted $D \models \mathcal{Q}$, resp. $D \not\models \mathcal{Q}$.

Now, we introduce some of the elements that will form a PDES. After introducing them, we will give the formal definition of a PDES (cf. Definitions 2.1 and 2.2 below).

A PDES contains a finite set $\mathcal{P} = \{\mathtt{P1}, \ldots, \mathtt{Pn}\}$ of peers. Each peer $\mathtt{P}$ owns a relational database schema $\mathcal{S}(\mathtt{P})$, and a database instance $D(\mathtt{P})$ for the schema $\mathcal{S}(\mathtt{P})$. We denote with $\mathfrak{S}$ the set of all schemas of a PDES, that is, $\mathfrak{S} = \{\mathcal{S}(\mathtt{P1}), \ldots, \mathcal{S}(\mathtt{Pn})\}$. We assume, to simplify the presentation and, without loss of generality, that the peers' schemas are mutually disjoint, but share a common, possibly infinite database domain $\mathcal{U}$. Each $D(\mathtt{P})$ can be seen as a finite set of ground atoms over $\mathcal{U}$, with predicates in $\mathcal{S}(\mathtt{P})$. It holds, $Adom(D(\mathtt{P})) \subseteq \mathcal{U}$.

The peers' schemas, or unions thereof, determine first-order languages, e.g. $\mathcal{L}(\mathtt{P})$, $\mathcal{L}(\mathtt{P}, \mathtt{Q})$, which are $\mathcal{L}(\mathcal{S}(\mathtt{P}))$ and $\mathcal{L}(\mathcal{S}(\mathtt{P}) \cup \mathcal{S}(\mathtt{Q}))$, resp. A DEC between peers $\mathtt{P}$, $\mathtt{Q}$ is an $\mathcal{L}(\mathtt{P}, \mathtt{Q})$-sentence. We will consider the following two syntactic classes of DECs:

(a) A *universal data exchange constraint* (UDEC) between peers $\mathtt{P}$, $\mathtt{Q}$ is an $\mathcal{L}(\mathtt{P}, \mathtt{Q})$-sentence of the form:

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow \bigwedge_{k=1}^{l} \left(\bigvee_{j=1}^{m} Q_{kj}(\bar{y}_{kj}) \right) \right) \tag{2.1}$$

where the R_i are predicates in $\mathcal{S}(\mathtt{P}) \cup \mathcal{S}(\mathtt{Q})$, the Q_{kj} are atomic formulas with predicates in $\mathcal{S}(\mathtt{P}) \cup \mathcal{S}(\mathtt{Q})$ or atoms with predicates in $\mathscr{B}$, $\bar{x} = \bigcup_{i=1}^{n} \bar{x}_i$, and $\bar{y}_{kj} \subseteq \bar{x}$.

(b) A *referential data exchange constraint* (RDEC) between peers $\mathtt{P}$, $\mathtt{Q}$ is an $\mathcal{L}(\mathtt{P}, \mathtt{Q})$-sentence of the form:

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow \exists \bar{y} \left(\bigwedge_{k=1}^{l} Q_k(\bar{x}_k, \bar{y}_k) \wedge \varphi_k\left(\bar{x}'_k, \bar{y}'_k\right) \right) \vee \varphi(\bar{x}) \right) \tag{2.2}$$

[3] In this work, whenever we consider sets Σ of ICs, we assume that Σ is *logically consistent*.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

where the R_i, Q_k are predicates in $\mathscr{S}(\mathtt{P}) \cup \mathscr{S}(\mathtt{Q})$, φ_k, and φ are a conjunction, resp. a disjunction, of atoms with predicates in $\mathscr{B}$, and $\bar{x}_i, \bar{x}_k \subseteq \bar{x}$, $\bar{x}'_k \subseteq \bar{x}_k$, $\bar{y}'_k \subseteq \bar{y}_k$, and $\bar{y} = \bigcup_{k=1}^{l} \bar{y}_k \neq \emptyset$.
The formulas $\varphi_k(\bar{x}'_k, \bar{y}'_k))$ are used to impose conditions on the existential values, and φ for conditions on the values for $\bar{x}$ in the antecedent.

The classes of exchange constraints that we are considering are broad enough to capture all the relevant ICs and logical mappings found in data exchange and virtual data integration that are usually considered in the theoretical, technical, and applied literature on data management. In particular, UDECs can be used to express *equality-generating dependencies* and RDECs can express general *tuple-generating dependencies* (Abiteboul *et al.* 1995).

Each peer $\mathtt{P}$ of $\mathscr{P}$ has a possibly empty collection of sets of DECs, $\Sigma(\mathtt{P}, \mathtt{Q})$, between $\mathtt{P}$ and peers $\mathtt{Q} \in \mathscr{P}$, with at most one $\Sigma(\mathtt{P}, \mathtt{Q})$ for each peer $\mathtt{Q}$. Each $\Sigma(\mathtt{P}, \mathtt{Q})$ is finite and logically consistent. Due to the local nature of PDESs systems, it is possible for $\Sigma(\mathtt{P}, \mathtt{Q})$ (which is owned by $\mathtt{P}$) to be different from $\Sigma(\mathtt{Q}, \mathtt{P})$ (which is owned by $\mathtt{Q}$). The DECs in $\Sigma(\mathtt{P}, \mathtt{P})$ are the ICs for (instances of) peer $\mathtt{P}$. We denote with Σ the class formed by of all non-empty sets $\Sigma(\mathtt{P}, \mathtt{Q})$ of a PDES.

A PDES also has a relation $\mathit{Trust} \subseteq \mathscr{P} \times \{\mathit{less}, \mathit{same}\} \times \mathscr{P}$, with exactly one triple of the form $(\mathtt{P}, \cdot, \mathtt{Q})$ for each $\Sigma(\mathtt{P}, \mathtt{Q}) \in \Sigma$. The intended semantics of $(\mathtt{P}, \mathit{less}, \mathtt{Q}) \in \mathit{Trust}$ is that peer $\mathtt{P}$ trusts itself less than $\mathtt{Q}$; while $(\mathtt{P}, \mathit{same}, \mathtt{Q}) \in \mathit{Trust}$ indicates that $\mathtt{P}$ trusts itself the same as $\mathtt{Q}$.[4] The trust relation is not necessarily symmetric, e.g. it could hold $(\mathtt{P}, \mathit{less}, \mathtt{Q}), (\mathtt{Q}, \mathit{same}, \mathtt{P}) \in \mathit{Trust}$. For a peer $\mathtt{P}$, when $\Sigma(\mathtt{P}, \mathtt{P}) \neq \emptyset$, we assume $(\mathtt{P}, \mathit{same}, \mathtt{P}) \in \mathit{Trust}$.

Definition 2.1

(a) The *schema of a PDES* $\mathfrak{P}$ is a sequence $\langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, where $\mathscr{P}$ is a set of peers, $\mathfrak{S}$ is a corresponding set of peer database schemas, Σ the set of DECs, and *Trust* the set of trust relationships.

(b) An instance $\mathfrak{D}$ of a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ is a set containing one database instance $D(\mathtt{P})$ for the schema $S(\mathtt{P}) \in \mathfrak{S}$, for each peer $\mathtt{P} \in \mathscr{P}$.

Intuitively, and due to the locality of peers, a peer $\mathtt{P}$ will be aware only of the sets $\Sigma(\mathtt{P}, \mathtt{Q})$ and elements of *Trust* whose first argument is $\mathtt{P}$. More precisely, a peer stores its own database schema and instance, its DECs to other peers, and its trust relationships. Furthermore, for an instance $\mathfrak{D}$ of schema $\mathfrak{P}$, the instances $D(\mathtt{P}) \in \mathfrak{D}$ are not required to (jointly) satisfy the DECs in Σ. Consistency and consistency restoration is a concern only when queries are posed to peers.

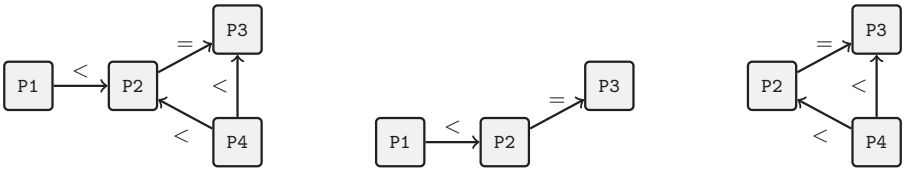
Definition 2.2

Given a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$:

(a) The *schema of a peer* $P \in \mathscr{P}$ is $\mathfrak{P}(\mathtt{P}) = \langle \mathscr{S}(\mathtt{P}), \Sigma(\mathtt{P}), \mathit{Trust}(\mathtt{P}) \rangle$, where $\mathscr{S}(\mathtt{P}) \in \mathfrak{S}$, $\Sigma(\mathtt{P}) = \{\varphi \mid \varphi \in \Sigma(\mathtt{P}, \mathtt{Q}), \Sigma(\mathtt{P}, \mathtt{Q}) \in \Sigma, \text{ and } \mathtt{Q} \in \mathscr{P}\}$; and $\mathit{Trust}(\mathtt{P}) = \{(\mathtt{P}, t, \mathtt{Q}) \mid (\mathtt{P}, t, \mathtt{Q}) \in \mathit{Trust}, \text{ and } \mathtt{Q} \in \mathscr{P}\}$. (We can safely identify a peer $\mathtt{P}$ with its schema $\mathfrak{P}(\mathtt{P})$.)

[4] We do not consider the case when a peer $\mathtt{P}$ trusts itself more than another peer, since the information of the latter should be irrelevant to $\mathtt{P}$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press



(a) $\mathcal{G}(\mathfrak{P})$ (b) $\mathcal{G}(\mathtt{P1})$ (c) $\mathcal{G}(\mathtt{P4})$

Fig. 3. Graphs for Example 2.1.

(b) The *accessibility graph* $\mathcal{G}(\mathfrak{P})$ contains a vertex for each $\mathtt{P} \in \mathscr{P}$, and a directed edge from P to Q, if $\mathtt{P} \neq \mathtt{Q}$ and $\Sigma(\mathtt{P}, \mathtt{Q}) \neq \emptyset$. An edge from P to Q is labeled with "<" when $(\mathtt{P}, \mathit{less}, \mathtt{Q}) \in \mathit{Trust}$, and with "=" when $(\mathtt{P}, \mathit{same}, \mathtt{Q}) \in \mathit{Trust}$.

(c) Peer P' is *accessible* from P if there is a path in $\mathcal{G}(\mathfrak{P})$ from P to P' or if P'=P. $\mathscr{AC}(\mathtt{P})$ denotes the set of peers accessible from P. $\mathcal{G}(\mathtt{P})$ denotes the restriction of $\mathcal{G}(\mathfrak{P})$ to the peers in $\mathscr{AC}(\mathtt{P})$ (i.e. it contains as vertices only the peers in $\mathscr{AC}(\mathtt{P})$ and the edges between them are those inherited from $\mathcal{G}(\mathfrak{P})$).

(d) Peer P' is a *neighbor* of P if there is an edge from P to P' in $\mathcal{G}(\mathfrak{P})$, or if P' = P. We denote with $\mathscr{N}(\mathtt{P})$ the set of neighbors of P; and with $\mathscr{S}(\mathscr{N}(\mathtt{P}))$ the union of their schemas, i.e. $\bigcup_{\mathtt{Q} \in \mathscr{N}(\mathtt{P})} \mathscr{S}(\mathtt{Q})$. Furthermore, $\mathscr{N}^{\circ}(\mathtt{P}) := \mathscr{N}(\mathtt{P}) \setminus \{\mathtt{P}\}$.

Example 2.1

Consider $\mathfrak{P} = \langle \mathscr{P}, \mathscr{S}, \Sigma, \mathit{Trust} \rangle$ with $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}, \mathtt{P4}\}$, $\mathfrak{S} = \{\mathscr{S}(\mathtt{P1}), \mathscr{S}(\mathtt{P2}), \mathscr{S}(\mathtt{P3}), \mathscr{S}(\mathtt{P4})\}$, $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P2}) = \{R^2(\cdot,\cdot), S^2(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P3}) = \{R^3(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P4}) = \{R^4(\cdot,\cdot,\cdot)\}$. Furthermore, $\mathit{Trust} = \{(\mathtt{P1}, \mathit{less}, \mathtt{P2}), (\mathtt{P2}, \mathit{same}, \mathtt{P3}), (\mathtt{P4}, \mathit{less}, \mathtt{P2}), (\mathtt{P4}, \mathit{less}, \mathtt{P3})\}$, and Σ contains:

$$\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xy(R^2(x,y) \rightarrow R^1(x,y))\} \tag{2.3}$$

$$\Sigma(\mathtt{P2}, \mathtt{P2}) = \{\forall x \forall y(R^2(x,y) \wedge S^2(x,y) \rightarrow \mathbf{false}) \tag{2.4}$$

$$\forall x \forall y \forall z(R^2(x,y) \wedge R^2(x,z) \rightarrow y = z)\}$$

$$\Sigma(\mathtt{P2}, \mathtt{P3}) = \{\forall xy(R^2(x,y) \wedge R^3(x,y) \rightarrow \mathbf{false})\} \tag{2.5}$$

$$\Sigma(\mathtt{P4}, \mathtt{P2}) = \{\forall xyz(R^2(x,y) \wedge S^2(y,z) \rightarrow R^4(x,y,z))\} \tag{2.6}$$

$$\Sigma(\mathtt{P4}, \mathtt{P3}) = \{\forall xy(R^3(x,y) \rightarrow \exists z R^4(x,y,z))\} \tag{2.7}$$

The built-in atom **false** in (2.4) and (2.5) is false in every instance. The DECs in $\Sigma(\mathtt{P2}, \mathtt{P2})$ are the local ICs for P2, in this case, a denial constraint (i.e. a prohibited join of positive atoms) and a functional dependency, requiring that (the values for) the first attribute functionally determines (determine the values for) the second one. Here, $\Sigma(\mathtt{P4}) = \{\forall xyz(R^2(x,y) \wedge S^2(y,z) \rightarrow R^4(x,y,z)), \forall xy(R^3(x,y) \rightarrow \exists z R^4(x,y,z))\}$.

The DEC in (2.7) is an RDEC, and all the others are UDECs. Some of the accessibility graphs are shown in Figure 3. Notice that $\mathscr{AC}(\mathtt{P1}) = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}\}$, $\mathscr{AC}(\mathtt{P2}) = \{\mathtt{P2}, \mathtt{P3}\}$, $\mathscr{AC}(\mathtt{P3}) = \{\mathtt{P3}\}$, $\mathscr{AC}(\mathtt{P4}) = \{\mathtt{P2}, \mathtt{P3}, \mathtt{P4}\}$, $\mathscr{N}(\mathtt{P1}) = \{\mathtt{P1}, \mathtt{P2}\}$, $\mathscr{N}(\mathtt{P2}) = \{\mathtt{P2}, \mathtt{P3}\}$, $\mathscr{N}(\mathtt{P3}) = \{\mathtt{P3}\}$, $\mathscr{N}(\mathtt{P4}) = \{\mathtt{P2}, \mathtt{P3}, \mathtt{P4}\}$, $\mathscr{N}^{\circ}(\mathtt{P1}) = \{\mathtt{P2}\}$, $\mathscr{N}^{\circ}(\mathtt{P2}) = \{\mathtt{P3}\}$, $\mathscr{N}^{\circ}(\mathtt{P3}) = \emptyset$, and $\mathscr{N}^{\circ}(\mathtt{P4}) = \{\mathtt{P2}, \mathtt{P3}\}$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

As an example, peer P4 only knows its schema and database, the DECs from it to peers P2 and P3, and how much it trusts them. More precisely, P4 knows its schema $\mathfrak{P}(\mathtt{P4}) = \langle \mathscr{S}(\mathtt{P4}), \Sigma(\mathtt{P4}), \mathit{Trust}(\mathtt{P4})\rangle$, with $\Sigma(\mathtt{P4}) = \{\forall xyz(R^2(x,y) \wedge S^2(y,z) \rightarrow R^4(x,y,z)), \forall xy(R^3(x,y) \rightarrow \exists z R^4(x,y,z))\}$, and $\mathit{Trust}(\mathtt{P4})$=$\{(\mathtt{P4}, \mathit{less}, \mathtt{P2}), (\mathtt{P4}, \mathit{less}, \mathtt{P3})\}$. This peer also has its database instance $D(\mathtt{P4})$.

The schema of P1 is $\mathfrak{P}(\mathtt{P1}) = \langle\{R^1(\cdot,\cdot)\}, \{\forall xy(R^2(x,y) \rightarrow R^1(x,y))\}, \{(\mathtt{P1}, \mathit{less}, \mathtt{P2})\}\rangle$. As will be determined by the semantics later on, the combination of P1's DEC in (2.3) and the trust relationship to P2, will make P1 import all the missing data from the extension of its neighbor's R^2 into the extension of its own R^1.

Example 2.2
(example 1.1 cont.) The PDES can be formalized through the schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust}\rangle$, where $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}\}$, $\mathfrak{S} = \{\mathscr{S}(\mathtt{P1}), \mathscr{S}(\mathtt{P2})\}$, $\Sigma = \{\Sigma(\mathtt{P1}, \mathtt{P2})\}$, $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xy(R^2(x,y) \wedge S^2(y,z) \rightarrow R^1(x,y,z)), \forall x(S^1(x) \rightarrow S^2(5,x))\}$, and $\mathit{Trust} = \{(\mathtt{P1}, \mathit{less}, \mathtt{P2})\}$.

The schema of P1 is $\mathfrak{P}(\mathtt{P1}) = \langle \mathscr{S}(\mathtt{P1}), \Sigma(\mathtt{P1}, \mathtt{P2}), \{(\mathtt{P1}, \mathit{less}, \mathtt{P2})\}\rangle$, and that of P2 is $\mathfrak{P}(\mathtt{P2}) = \langle \mathscr{S}(\mathtt{P2}), \emptyset, \emptyset\rangle$. The instance for the PDES, $\mathfrak{D} = \{D(\mathtt{P1}), D(\mathtt{P2})\}$, contains the two instances shown in Figure 1.

Remark 2.1
From now on, we assume that the graph $\mathscr{G}(\mathfrak{P})$ is acyclic. As a consequence, for each particular peer P, $\mathscr{G}(\mathtt{P})$ is also acyclic. This is a global assumption, not local to any particular peer, that will be used to define the semantics of a PDES. (A discussion around this and related assumptions can be found in the Online Appendix A.)

Notice that according to its definition, the accessibility graph has no self-loops. The finiteness of the set of peers and the acyclicity of the accessibility graph imply that there must be "sink" (or terminal) peers, without any outgoing edges. Also, notice that we are not making the assumption that for a peer P its initial instance $D(\mathtt{P})$ has to satisfy the local constraints in $\Sigma(\mathtt{P}, \mathtt{P})$. This applies in particular to sink peers in the accessibility graph. They can have local ICs, which may be violated.

3 A general semantic framework for a PDES

Before presenting the formal semantics of PDESs, we describe the intended semantics in intuitive and operational terms.

A query $\mathscr{Q} \in \mathscr{L}(\mathtt{P})$ is posed to a peer P by a certain user U, who may be an external user or another peer in $\mathscr{P}$. Now P, depending on the query predicates, inspects the DECs in $\Sigma(\mathtt{P})$ to identify data owned by other neighboring peers Q that may be related to its own data. If predicate $R^{\mathtt{Q}} \in \mathscr{S}(\mathtt{Q})$ appears in $\Sigma(\mathtt{P})$, then P requests from Q the (contents of the) relation $R^{\mathtt{Q}}$. Q returns to P a possible modified version of its relation, because Q may have to take into account data from its own neighbors (in the same way P did when it received the user query), which could produce local inconsistencies at Q's level, and they have to be repaired; and so on.

For the purpose of presenting the semantics of the system, that should be query-independent, we will assume for the moment that P receives whole relation instances

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

from its neighbors. We emphasize that, in their turn, P's neighboring peers Q consider data from their own neighbors, and so on; data that may also be eventually used by P, by transitivity (cf. details below).

The data P receives from another peer Q is defined in a recursive manner, because Q may have connections to other peers who may contribute with data of their own. Eventually, after receiving the requested relation instances from its neighbors, P has now a database instance D for the expanded database schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, which extends the initial instance $D(\mathtt{P})$ by adding predicates and data from its neighbors. D can be used to interpret the DECs in $\Sigma(\mathtt{P})$. Most likely, D will not satisfy $\Sigma(\mathtt{P})$.

When this extended instance D is inconsistent with respect to $\Sigma(\mathtt{P})$, different alternatives to restore consistency with respect to $\Sigma(\mathtt{P})$ can be considered, but the (possibly virtual) updates performed on the extended instance D will have to both respect the trust relationships and make sure that the consistent alternative instances stay *as close as possible* to instance D (for which a form of distance has to be introduced, as we do later in this section). In this way, a collection of consistent, and possibly virtual, extended instances for P's neighborhood emerges, the *neighborhood instances*. By restricting those instances to the schema of peer P, the *solution instances* for P are obtained.

Now, we give the precise definition of neighborhood solution. In order to do so, we will assume that it is possible to compare arbitrary instances D_1, D_2 with respect to a fixed instance D by means of a pre-order relation $\preceq_D$ (i.e. a reflexive and transitive binary relation). If D, D_1, D_2 are database instances for the same schema, the intuition behind the relationship $D_1 \preceq_D D_2$ is that D_1 *is at least as close to* D *as* D_2 *(is to* D*)*. We can define $D_1 \prec_D D_2$ iff $D_1 \preceq_D D_2$, but not $D_2 \preceq_D D_1$, with the intuition that D_1 is closer to D than D_2.

In our semantics, the participating instances will always be for a neighborhood schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$ of a peer P. Even more, the pre-order relation will also depend on $\Sigma(\mathtt{P})$ (which has to be satisfied).

As a consequence, we will assume, for D *a fixed instance for* $\mathscr{S}(\mathscr{N}(P))$*, the existence of a pre-order relation* $\preceq_D^{\Sigma(P)}$ *on instances for the schema* $\mathscr{S}(\mathscr{N}(P))$.

Definition 3.1
Given a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, a peer $\mathtt{P} \in \mathscr{P}$, and an instance $\bar{D}$ for the schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$:

(a) An instance D' for the schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$ is a *neighborhood solution* for P and $\bar{D}$ if:
  (i) $D' \models \Sigma(\mathtt{P})$.
  (ii) There is no instance D'' satisfying (i), and $D'' \prec_{\bar{D}}^{\Sigma(\mathtt{P})} D'$.
  (iii) $D' \downharpoonright \{R\} = \bar{D} \downharpoonright \{R\}$ for every predicate $R \in \mathscr{S}(\mathtt{Q})$ with $(\mathtt{P}, \mathit{less}, \mathtt{Q}) \in \mathit{Trust}$.

(b) $NS^{\mathfrak{P}}(\mathtt{P}, \bar{D})$ denotes the set of neighborhood solutions for P and $\bar{D}$. When clear from the context, we will simply use $NS(\mathtt{P}, \bar{D})$.

A neighborhood solution for P is a database for its whole neighborhood that satisfies P's DECs (including its local constraints), and respects its trust relationships. A neighborhood solution also stays as close as possible to the original neighborhood instance, while staying the same for trustable peers. In operational and intuitive

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

R^1			S^1	R^2		S^2	
c	4	2	3	c	4	4	2
f	3	5	7	d	5	5	3

Fig. 4. A database instance D for schema $\mathscr{S}(\mathscr{N}(\mathtt{P1}))$.

terms, a minimal data set with respect to Definition 3.1 is imported or given up to satisfy the DECs.

Notice that Definition 3.1 introduces an abstract *repair semantics*, similar to those used to handle inconsistency in single databases (Bertossi 2011). For example, a well-studied repair semantics for single databases was introduced in Arenas *et al.* (1999). It is based on insertions or deletions of whole tuples into/from the original inconsistent instance D, and the "distances" of two instances D_1, D_2 from D are compared using the symmetric differences: $D_1 \preceq_D^{\Delta,\Sigma(\mathtt{P})} D_2 :\Leftrightarrow \Delta(D, D_1) \subseteq \Delta(D, D_2)$.[5]

For a peer P in isolation, i.e. with $\mathscr{N}^\circ(\mathtt{P}) = \emptyset$, a neighborhood solution of its original instance $D(\mathtt{P})$ will simply be an instance D' for its schema $\mathscr{S}(\mathtt{P})$, such that: (a) $D' \models \Sigma(\mathtt{P}, \mathtt{P})$ (i.e. P's own ICs); and (b) there is no instance D'' satisfying (a) and $D'' \prec_{D(\mathtt{P})}^{\Sigma(\mathtt{P},\mathtt{P})} D'$. This defines a class of *repairs* of instance $D(\mathtt{P})$ with respect to $\Sigma(\mathtt{P}, \mathtt{P})$.[6] Thus, this neighborhood-solution semantics for PDESs naturally extends the notion of repair to neighborhood solutions.

Example 3.1 (*examples 1.1 and 2.2 cont.*)
The elements of the PDES schema are all as before, in particular,

$$\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xy(R^2(x, y) \wedge S^2(y, z) \rightarrow R^1(x, y, z)),\ \forall x(S^1(x) \rightarrow S^2(5, x))\} \tag{3.8}$$

but this time consider $\{(\mathtt{P1}, \mathit{same}, \mathtt{P2})\} \in \mathit{Trust}$.

Now assume, in this example, for illustration purposes, that the pre-order relations are defined by $D_1 \preceq_D^{\Delta,\Sigma(\mathtt{P})} D_2$ iff $\Delta(D, D_1) \subseteq \Delta(D, D_2)$.

If a query is posed to P1, it will have to compute (possibly virtually) its neighborhood solutions. Peer P1's knowledge of the schema is limited to its own schema $\mathfrak{P}(\mathtt{P1}) = \langle \mathscr{S}(\mathtt{P1}), \Sigma(\mathtt{P1}), \mathit{Trust}(\mathtt{P1})\rangle$, where $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot,\cdot), S^1(\cdot)\}$, $\Sigma(\mathtt{P1}) = \Sigma(\mathtt{P1}, \mathtt{P2})$, and $\mathit{Trust}(\mathtt{P1}) = \{(\mathtt{P1}, \mathit{same}, \mathtt{P2})\}$.

In order to enforce $\Sigma(\mathtt{P1})$, P1 poses two atomic queries to peer P2: $\mathscr{Q}_1(x, y)$: $R^2(x, y)$ and $\mathscr{Q}_2(x, y)$: $S^2(x, y)$. Since P2 is not related to any other peer, and has no local ICs, it will provide as answers the content of those relations in $D(\mathtt{P2})$. In this way, P1 has an extended instance D, shown in Figure 4, that corresponds to the union of the two instances in Figure 1, that is, $D = D(\mathtt{P1}) \cup D(\mathtt{P2})$. Since D does not satisfy $\Sigma(\mathtt{P1})$, D has to be repaired.

Since $(\mathtt{P1}, \mathit{same}, \mathtt{P2})$, there are several neighborhood solutions for P1, which are obtained by virtually modifying both peers' data. For example, the inconsistencies with respect to the first UDEC in (3.8) can be (minimally) solved by either removing

[5] For sets S_1 and S_2, $\Delta(S_1, S_2) := (S_1 \setminus S_2) \cup (S_2 \setminus S_1)$.
[6] We are using abstract pre-order relations to define repair semantics. It should also be possible to use abstract "distance measures" to define repairs and repair semantics. Cf. Arieli *et al.* (2007).

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

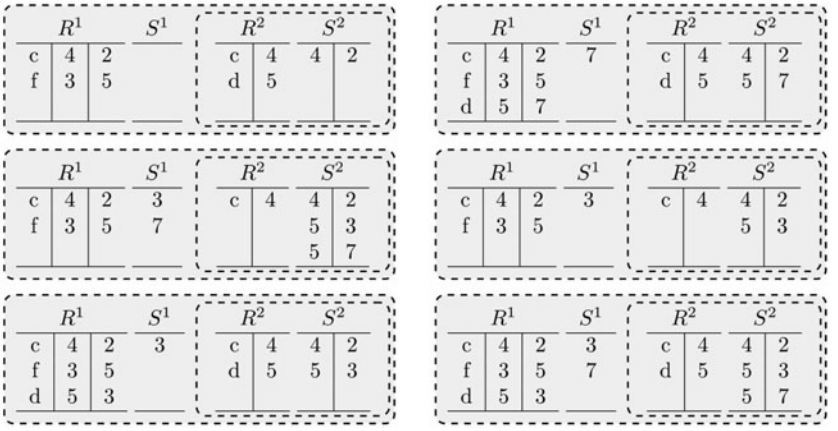


Fig. 5. Several neighborhood solution instances for peer P1.

$R^2(d, 5)$ from R^2, or removing $S^2(5, 3)$ from S^2, or inserting $R^1(d, 5, 3)$ into R^1. If $S^2(5, 3)$ were removed from S^2, a new inconsistency is created with respect to the second UDEC in (3.8). This one can be solved by removing $S^1(3)$ from S^1. The inconsistencies with respect to the second DEC can be solved by either removing $S^1(7)$ from S^1 or inserting $S^2(5, 7)$ into S^2.

Figure 5 shows the six neighborhood solutions in $NS^{\mathfrak{P}}(\mathtt{P1}, D)$. All these neighborhood instances are *repairs* in the sense of Arenas *et al.* (1999) of instance D with respect to P1's DECs. If we are locally interested only in P1, we consider their restrictions to P1' schema $\mathscr{S}(\mathtt{P1})$.

As the following example shows, trust relationships in combination with inconsistent DECs may cause that neighborhood solutions do not exist.

Example 3.2
Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ with $\mathscr{P} = \{\mathtt{P}, \mathtt{Q}, \mathtt{R}\}$, and

(1) $\mathfrak{S} = \{\mathscr{S}(\mathtt{P}), \mathscr{S}(\mathtt{Q}), \mathscr{S}(\mathtt{R})\}$, $\mathscr{S}(\mathtt{P}) = \{P\}$, $\mathscr{S}(\mathtt{Q}) = \{Q\}$, $\mathscr{S}(\mathtt{R}) = \{R\}$;
(2) $\Sigma = \{\Sigma(\mathtt{P}, \mathtt{Q}), \Sigma(\mathtt{P}, \mathtt{R}))\}$, $\Sigma(\mathtt{P}, \mathtt{Q}) = \{\forall xy(Q(x, y) \rightarrow P(x, y))\}$,
$\Sigma(\mathtt{P}, \mathtt{R}) = \{\forall xy(P(x, y) \rightarrow R(x, y))\}$;
(3) $\mathit{Trust} = \{(\mathtt{P}, \mathit{less}, \mathtt{Q}), (\mathtt{P}, \mathit{less}, \mathtt{R})\}$;
(4) $\mathfrak{D} = \{D(\mathtt{P}), D(\mathtt{Q}), D(\mathtt{R})\}$, with $D(\mathtt{P}) = D(\mathtt{R}) = \emptyset$, and $D(\mathtt{Q}) = \{Q(a, b)\}$.

For peer P and $D = D(\mathtt{P}) \cup D(\mathtt{Q}) \cup D(\mathtt{R}) = \{Q(a, b)\}$, there is no neighborhood solution, i.e. $NS(\mathtt{P}, D(\mathtt{P}) \cup D(\mathtt{Q}) \cup D(\mathtt{R}))$ is empty. This is because P's DEC with Q forces the insertion of $P(a, b)$, but at the same time the interaction of P with R requires relation P to be empty. Since peer P trusts both peers more than itself, there is no neighborhood solution that satisfies the DECs and respect the trust relations.

Notice that the DEC in $\Sigma(\mathtt{P}, \mathtt{R})$ acts as a restriction on tuples for its owner, P, rather than as a tuple generator for it.

In the previous example, the trust relationships impose unsolvable conditions on neighborhood solutions. However, if all the trust relationships are of the flexible form $(\mathtt{P}, \mathit{same}, \mathtt{Q})$, a peer always has solutions.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Proposition 3.1
Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ and a peer $\mathtt{P} \in \mathscr{P}$ whose trust relationships are all of the form $(\mathtt{P}, \mathit{same}, \mathtt{Q})$. Given an instance $\bar{D}$ for $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, peer P always has a neighborhood solution for $\bar{D}$.

Having defined the notion of neighborhood solution for a peer, now we can define the notion of solution for a peer, which is a local instance but takes all the accessible peers into account. The data distributed across different peers has to be appropriately gathered to build solution instances for the peer. This definition of solution instance is recursive, and appeals to that of neighborhood solution. The intuition and idea are as follows. Assume each neighbor Q of peer P has the class $\mathit{Sol}(\mathtt{Q}, \mathfrak{D})$ of its solution instances (hence the recursion). Each of these Q passes to P the intersection $\bigcap \mathit{Sol}(\mathtt{Q}, \mathfrak{D})$ of its solution set[7]. With them and its own initial instance $D(\mathtt{P})$, P builds an instance D for its neighborhood $\mathscr{N}(\mathtt{P})$. The solutions for P become the restrictions to P's schema of the neighborhood solutions for P with respect to D, $\Sigma(\mathtt{P})$, and P's trust relationships.

Under this recursive definition, the solutions for the neighbors have to be determined, under the same semantics. Base cases of the recursion are peers with no DECs or with only local constraints, that is, when either $\Sigma(\mathtt{P}) = \emptyset$ or $\Sigma(\mathtt{P}) = \{\Sigma(\mathtt{P}, \mathtt{P})\}$. These peers are sinks in the accessibility graph of P. In this regard, we recall that we made the assumption that the accessibility graphs are acyclic.

Definition 3.2
Given an instance $\mathfrak{D}$ for the PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, and a peer P, an instance D for the schema $\mathscr{S}(\mathtt{P})$ is a *solution instance* (or simply *solution*) for P, denoted $D \in \mathit{Sol}^{\mathfrak{P}}(\mathtt{P}, \mathfrak{D})$ (or simply, $\mathit{Sol}(\mathtt{P}, \mathfrak{D})$), iff:

(a) For $\Sigma(\mathtt{P}) = \emptyset$: $D = D(\mathtt{P})$ (P's initial instance, the projection of $\mathfrak{D}$ on $\mathscr{S}(\mathtt{P})$).
(b) For $\Sigma(\mathtt{P}) = \{\Sigma(\mathtt{P}, \mathtt{P})\}$: $D \in \mathit{NS}^{\mathfrak{P}}(\mathtt{P}, D(\mathtt{P}))$.
(c) For $\emptyset \neq \Sigma(\mathtt{P}) \neq \{\Sigma(\mathtt{P}, \mathtt{P})\}$: $D = \overline{D} \restriction \mathscr{S}(\mathtt{P})$, where $\overline{D} \in \mathit{NS}^{\mathfrak{P}}(\mathtt{P}, D(\mathtt{P}) \cup \bigcup_{\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})} \mathit{Core}(\mathtt{Q}, \mathfrak{D}))$.
Here, $\mathit{Core}(\mathtt{Q}, \mathfrak{D}) := \bigcap \mathit{Sol}^{\mathfrak{P}}(\mathtt{Q}, \mathfrak{D})$, when $\mathit{Sol}^{\mathfrak{P}}(\mathtt{Q}, \mathfrak{D}) \neq \emptyset$ and $\mathit{Core}(\mathtt{Q}, \mathfrak{D}) := \{\mathbf{inc}_{\mathtt{Q}}\}$, otherwise. The propositional built-in predicate $\mathbf{inc}_{\mathtt{Q}}$ in $\mathscr{S}(\mathtt{Q})$ is true of an instance iff it is contained in the latter as an atom.

The base cases of the recursion are (a) and (b), where the solutions of a peer can be computed without data from other peers. Case (c) corresponds to the properly recursive case, where, before determining P's solutions, there is an extended instance $\overline{D}$ around P formed by its local instance $D(\mathtt{P})$ plus, for each neighbor Q, the intersection, $\mathit{Core}(\mathtt{Q}, \mathfrak{D})$, of all its solutions. The combined database $\overline{D}$ is for the schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, and starting from $\overline{D}$, neighborhood solutions for P can be determined; and their restrictions to P's schema become P's solutions.

[7] When $\mathscr{C}$ is a class of sets, then we will usually denote with $\bigcap \mathscr{C}$ the intersection of its elements, i.e. $\bigcap \mathscr{C} := \bigcap_{C \in \mathscr{C}} C$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Although, in Definition 3.2, cases (a) and (b) are special cases of (c), we include them explicitly, for clarity. Notice that case (b) amounts to obtaining the repairs of a local instance with respect to a set of local ICs.

The following is an immediate consequence of Proposition 3.1.

Corollary 3.1
Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ with trust relationships all of the form $(\mathtt{R}, \mathit{same}, \mathtt{Q})$, and a peer $\mathtt{P} \in \mathscr{P}$. Given an instance $\mathfrak{D}$ for $\mathfrak{P}$, peer P always has a solution, that is, $\mathit{Sol}^{\mathfrak{P}}(\mathtt{P}, \mathfrak{D}) \neq \emptyset$.

Remark 3.1
As we have seen, even in the absence of cycles in $\mathscr{G}(\mathscr{P})$, a peer may have no neighborhood solutions (cf. Example 3.2), and then no solutions, i.e. $\mathit{Sol}(\mathtt{P}, \mathfrak{D}) = \emptyset$. That is why, in this case, we make the convention in Definition 3.2 that $\mathit{Core}(\mathtt{P}, \mathfrak{D}) = \{\mathbf{inc}_{\mathtt{P}}\}$. When $\mathit{Sol}(\mathtt{P}, \mathfrak{D}) = \emptyset$, if another peer requests from P the intersection of its solutions, P returns the instance $\{\mathbf{inc}_{\mathtt{P}}\}$.

The idea is that when a peer becomes inconsistent, it makes another peer Q who requests its data notice its inconsistency by sending this special instance. In this case, Q is expected to ignore its mappings with P. This can be made precise and formally accommodated by modifying the DECs through the introduction of an extra conjunct in the antecedent. For example, the DEC $\Sigma(\mathtt{Q}, \mathtt{P}) = \{\forall x(P(x) \rightarrow Q(x))\}$ would be replaced by $\{\forall x(P(x) \wedge \neg \mathbf{inc}_{\mathtt{P}} \rightarrow Q(x))\}$, which would become trivially satisfied. Alternatively, if we want to keep the built-ins in the consequent, we could replace $\Sigma(\mathtt{Q}, \mathtt{P})$ by $\{\forall x(P(x) \rightarrow (Q(x) \vee \mathbf{inc}_{\mathtt{P}}))\}$.

In order not to complicate the notation, we will refrain from explicitly introducing the $\mathbf{inc}_{\mathtt{P}}$s in the DECs.

According to this convention about the treatment of peers without solutions, a peer that becomes intrinsically inconsistent, "irreparable", is ignored by its neighbors when they receive the notification of inconsistency. However, this form of ignoring is put here on a solid logical foot, that is compatible and uniform with the treatment of other peers. As long as a peer declares itself as inconsistent, there is no much a neighbor can do. However, an inconsistent peer might decide to relax its own consistency requirements and send to other peers only "partially consistent" data, which would be transparent to those receiving peers[8].

The PCAs from a peer to a query are the semantically correct answers, which means that when answering the query, the peer consistently considers the data of its neighbors and the trust relationships with them.

Definition 3.3
Consider an instance $\mathfrak{D}$ for the PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, and a peer $\mathtt{P} \in \mathscr{P}$. Let $\mathscr{Q}(\bar{x}) \in \mathscr{L}(\mathtt{P})$ be a query, with $\bar{x}$ a possibly empty list of free variables.

[8] An alternative to this design choice could be, in the case a peer P trusts an inconsistent peer Q more than itself, that P becomes or declares itself inconsistent as well. This alternative may be worth exploring, but we do not pursue it here any further.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

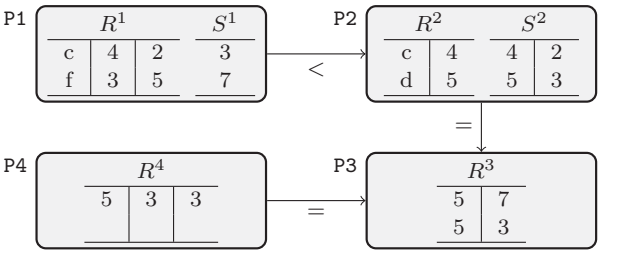


Fig. 6. PDES for Example 3.3.

(1) If $Sol(\mathtt{P}, \mathfrak{D}) \neq \emptyset$:
(a) If $\bar{x} \neq \emptyset$, a finite sequence $\bar{c}$ of constants in $\mathscr{U}$ of the same length as $\bar{x}$ is a PCA to $\mathscr{Q}$ from P iff $D \models \mathscr{Q}[\bar{c}]$ for every $D \in Sol(\mathtt{P}, \mathfrak{D})$.
(b) If $\mathscr{Q}$ is Boolean and $D \models \mathscr{Q}$ for every $D \in Sol(\mathtt{P}, \mathfrak{D})$, then *yes* is the only PCA to $\mathscr{Q}$.
(2) If $Sol(\mathtt{P}, \mathfrak{D}) = \emptyset$, then $\mathbf{inc}_{\mathtt{P}}$ is the only PCA to $\mathscr{Q}(\bar{x})$.

For illustration, in Example 3.1, if P1 has to answer a query $\mathscr{Q} \in \mathscr{L}(\mathtt{P1})$, it returns the answers that are simultaneously true in all its neighborhood solutions. For example, under this semantics, the answers to query $\mathscr{Q}(x) : \exists yz R^1(x, y, z)$ posed to P1 will be $\langle c \rangle, \langle f \rangle$, which are shared by all the six neighborhood solutions for P1 (or better, by their restrictions to P1's schema). The answers to queries $\mathscr{Q}_1$ and $\mathscr{Q}_2$ posed by P1 to P2 are answered in the same way, taking into consideration the P2's DECs.

We can see that the answers from a peer to a query are *certain answers* (Imielinski and Lipski 1984). This condition makes the data moved from one peer to a neighbor always certain. In particular, this allows us to treat external queries and inter-peer queries in a uniform manner. In particular, we will be in position to conceive and implement the passage of data from one peer to a neighbor as a query-answering process.

Example 3.3
Consider the PDES schema $\mathfrak{P}$ and instance $\mathfrak{D}$ represented in graph $\mathscr{G}(\mathfrak{P})$ in Figure 6. Here, $\Sigma = \{\Sigma(\mathtt{P1}, \mathtt{P2}), \Sigma(\mathtt{P2}, \mathtt{P3}), \Sigma(\mathtt{P4}, \mathtt{P3})\}$, and:
– $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xyz\ (R^2(x, y) \wedge S^2(y, z) \rightarrow R^1(x, y, z)), \forall x\ (S^1(x) \rightarrow S^2(5, x))\}$.
– $\Sigma(\mathtt{P2}, \mathtt{P3}) = \{\forall xy\ (S^2(x, y) \rightarrow R^3(x, y))\}$.
– $\Sigma(\mathtt{P4}, \mathtt{P3}) = \{\forall xyz\ (R^3(x, y) \rightarrow R^4(x, y, 3))\}$.

In intuitive and procedural terms, if a query is posed to P1, it will send queries to P2, to check the satisfaction of the DECs in $\Sigma(\mathtt{P1}, \mathtt{P2})$. But, in order for P2 to answer those queries, it will send queries to peer P3 to check the DECs in $\Sigma(\mathtt{P2}, \mathtt{P3})$. Since P3 is not connected to any other peer, it will answer P2's queries using its initial material instance $D(\mathtt{P3})$. Thus, the solutions for P1 and its PCAs will be affected by the peers in $\mathscr{NC}(\mathtt{P1}) = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}\}$. More precisely, we can now apply the definitions of solution instance and PCA.

Since P3 has neither DECs with other peers nor local ICs, its only neighborhood solution is its instance $D(\mathtt{P3}) \in \mathfrak{D}$, which is sent back to P2. Now, P2 has to repair

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

the extended instance $D = \{R^2(c,4), R^2(d,5), S^2(4,2), S^2(5,3), R^3(5,7), R^3(5,3)\}$ with respect to $\Sigma(\mathtt{P2},\mathtt{P3})$, which is not satisfied due to the presence of tuple $S^2(4,2)$. As P2 trusts P3 the same as itself, it can modify its own data or the data it got from P3.

Assuming – for illustration purposes for the moment – that the pre-order relation on instances is given in terms of the symmetric set-difference, P2 has two neighborhood solutions: $\{R^2(c,4), R^2(d,5), S^2(5,3), R^3(5,7), R^3(5,3)\}$ and $\{R^2(c,4), R^2(d,5), S^2(4,2), S^2(5,3), R^3(5,7), R^3(5,3), R^3(4,2)\}$. They minimally depart from D, and their restrictions to P2's schema lead to two solutions for P2: $Sol(\mathtt{P2},\mathfrak{D}) = \{\{R^2(c,4), R^2(d,5), S^2(5,3)\}, \{R^2(c,4), R^2(d,5), S^2(4,2), S^2(5,3)\}\}$.

Peer P2 will send to P1 the intersection of its solutions, namely $Core(\mathtt{P2},\mathfrak{D}) = \bigcap Sol(\mathtt{P2},\mathfrak{D}) = \{R^2(c,4), R^2(d,5), S^2(5,3)\}$. Now, P1 has to repair the extended instance $\{R^1(c,4,2), R^1(f,3,5), S^1(3), S^1(7), R^2(c,4), R^2(d,5), S^2(5,3)\}$ with respect to $\Sigma(\mathtt{P1},\mathtt{P2})$, which is not satisfied.

Since P1 trusts P2 more, it will solve inconsistencies by modifying its own data, in this case inserting tuples into R^1 and deleting tuples from S^1. This produces only one neighborhood solution: $\{R^1(c,4,2), R^1(f,3,5), R^1(d,5,3), S^1(3), R^2(c,4), R^2(d,5), S^2(5,3)\}$. Thus, $Sol(\mathtt{P1},\mathfrak{D}) = \{\{R^1(c,4,2), R^1(f,3,5), R^1(d,5,3), S^1(3)\}\}$.

If P1 had received the query $\mathcal{Q}(x) : \exists yz(R^1(x,y,z) \wedge S^1(y))$, the only PCA would have been $\langle f \rangle$, obtained from its only solution.

Notice that when a peer Q passes its core, $Core(\mathtt{Q},\bar{D})$, with $\bar{D}$ a neighborhood instance, to a peer P, it is delivering the peers consistent answers to the atomic queries from P to Q of the form $\mathcal{Q}^R(\bar{x}) : R(\bar{x})$, where $R \in \mathscr{S}(\mathtt{Q})$. However, when P computes its local PCAs to a query, it does not use its own core, $\bigcap Sol(\mathtt{P},\mathfrak{D})$, but the collection $Sol(\mathtt{P},\mathfrak{D})$ as a whole. The reason is that this core is unnecessarily restrictive for peer CQA. For example, if the query is $\mathcal{Q} : \exists x P(x)$, and P's solutions are $\{P(a)\}$ and $\{P(b)\}$, the PCA to this Boolean query would be *no* if evaluated on the empty core, but *yes* according to our definition.

Remark 3.2

Notice that the definitions of solution instance for a peer and of PCA are parameterized by:

(a) The class of DECs. We have considered certain syntactic classes of FO sentences, and we will also do so in the rest of this work. However, our presentation so far has been general enough to accommodate broader classes of FO sentences in the combined language of any two peers. The restriction on the classes of DECs has not been required or used yet.

(b) A notion of satisfaction, $D \models \varphi$, where D is a database instance and φ is a FO sentence in the language of D's schema. Most commonly, D is a relational database considered as a FO structure, and classical logical satisfaction is used. However, if D contains uncertain information, then we may have to depart from FO logic, as we will see in Section 4, for a particular PDES semantics.

(c) The pre-order relations $\preceq_D^{\Sigma(\mathtt{P})}$ between instances for the schema of an instance D. Different pre-order relations can be considered. In the examples above, we have

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

considered, just to fix ideas, the common pre-order based on the symmetric set difference between instances.

(d) The repair semantics, i.e. by a characterization of the instances that minimally depart from a given one D in order to satisfy local ICs or DECs between two peers. The repair semantics is based on the pre-orders $D'' \preceq_D^{\Sigma(\mathsf{P})} D'$, and the associated minimality conditions (what we did in Definition 3.1).

4 Toward a special PDES semantics with NULL

Considering the rather abstract nature and flexibility of the semantic framework introduced in Section 3, in this section, we make specific commitments about the semantic parameters discussed in Remark 3.2.

In the rest of this section, we consider a classical relational schema $\Sigma = (\mathscr{U}, \mathscr{R}, \mathscr{B})$, consisting of the data domain, a set of database predicates, and a set of built-in predicates[9]. Now, we will introduce in the schema some extra elements related to the special constant, *null*, that we will use in the rest of this work.

Remark 4.1

We assume from now on that every attribute domain, $Adom(A)$, contains the constant *null*. In particular, $null \in \mathscr{U}$. Furthermore, among the built-in predicates in $\mathscr{B}$, we will also find $IsNull(\cdot)$ and $IsNotNull(\cdot)$. The first one is true only with constant *null*, and the second one, on any constant c other than *null*.[10] Accordingly, *null* and these built-in predicates may appear in DECs as defined in (2.1) and (2.2), in particular, in ICs, and also in queries. They all become sentences or formulas of the FO language $\mathscr{L}(\Sigma^{\mathtt{null}})$ associated to $\Sigma^{\mathtt{null}} = (\mathscr{U}, \mathscr{R}, \mathscr{B}^{\mathtt{null}})$ with $\mathscr{B}^{\mathtt{null}} = \mathscr{B} \cup \{IsNull, IsNotNull\}$.

The special constant *null* in the data domain $\mathscr{U}$ is intended to behave and be used as the null value, NULL, in SQL relational databases. The new, unary, built-in predicates correspond to the SQL predicates `IS NULL` and `IS NOT NULL` used to check null values. Constant *null* may appear in database tuples, and will be used to restore consistency with respect to DECs and ICs.

Using a single null, with its SQL semantics, is clearly different from using multiple labeled null values, as is done in data exchange for enforcing the satisfaction of existential quantifications (Kolaitis 2005). It is also different from using arbitrary elements of the underlying data domain for the same purpose (Cali *et al.* 2003). However, to a large extent, the semantics of this section could be developed without the specific restrictions imposed on the representation and use of null values, adopting other forms of handling incomplete and inconsistent data. Cf. Section 7 and the Online Appendix A for additional discussions of these issues.

The semantics of *null* we introduce next captures the way nulls are handled by relational DBMSs that follow the SQL standard. More precisely, our semantics provides a partial logical reconstruction in first-order predicate logic of the way

[9] In the coming sections, where we will apply the semantics of this section, $\mathscr{R}$ will be the set of database predicates for a peer or a pair thereof.

[10] Each is the negation of the other, but we will keep both in order to avoid using explicit negation.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

nulls are handled in SQL databases. It refines and extend previous work presented in Bravo and Bertossi (2006) on database repairs with and in the presence of NULL.

The SQL Standard leaves many issues around NULL unspecified, and different DBMSs depart from the standard in different ways. As a consequence, it is not possible to provide a full logical reconstruction of SQL databases with NULL. For this reason, Accordingly, our semantics concentrates on the notion of satisfaction of DECs and ICs, and query answering for a broad classes of queries. Furthermore, the proposed semantics extends the "classical" notion of DEC and IC satisfaction, and query answering in databases without NULL[11].

The rest of this section continues as follows. In Section 4.1, we illustrate some of the elements of and issues around the notion of query answer as used in SQL databases with NULL. It serves a motivation for Section 4.2, where we formalize the semantics of query answering. In Section 4.3, we introduce a rewriting-based semantics for constraint satisfaction in the presence of NULL, leading to a fully classical semantics. Finally, in Section 4.4, we apply the rewriting methodology to the semantics of query answering under NULL.

4.1 Query answering under NULL: motivation

A tuple $\bar{c}$ of elements of $\mathscr{U}$ is an answer to query $\mathscr{Q}(\bar{x})$, denoted $D \models_N \mathscr{Q}(\bar{c})$, if the formula (that represents) $\mathscr{Q}$ is *classically true* when the quantifiers on its *relevant variables* (or attributes) run over $(\mathscr{U} \setminus \{null\})$; and those on the non-relevant variables run over $\mathscr{U}$. The free relevant variables, i.e. relevant among those in $\bar{x}$, cannot take the value *null* either. (Relevance is made precise in Section 4.2.)

Example 4.1
Consider the instance D_2 and the query below:

R	A	B	C
	1	1	1
	2	*null*	*null*
	null	3	3

S	B
	null
	1
	3

$$\mathscr{Q}_2(x):\ \exists y \exists z (R(x,y,z) \wedge S(y) \wedge y > 2) \quad (4.9)$$

A variable v (quantified or not) in a conjunctive query is *relevant* if it appears (non-trivially) twice in the formula after the quantifier prefix. Occurrences of the form $v = null$ and $v \neq null$ do not count though.

In query (4.9), the only relevant quantified variable is y, because it participates in a join and a built-in in the quantifier-free matrix of (4.9). So, there are two reasons for y to be relevant. The only free variable is x, which is not relevant. As for query answers, the only candidate values for x are: $null, 2, 1$. In this case, *null* is a candidate value because x is a non-relevant variable.

First, $x = null$ is an answer to the query, because the formula $\exists y \exists z (R(x,y,z) \wedge S(y) \wedge y > 2)$ is true in D_2, with a non-null witness value for y and a witness value

[11] We reserve the use of constant NULL to illustrate issues in relation its use in SQL databases; and to emphasize this, also in subsection titles. Otherwise, we keep using only the constant *null* of the underlying data domain.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

for z that combined make the (non-quantified) formula true. Namely, $y = 3, z = 3$. So, it holds $D_2 \models_N \mathcal{Q}_2[null]$.

Next, $x = 2$ is not an answer. For this value of x, because the candidate value for y, namely *null* that accompanies 2 in P, makes the formula $(R(x, y, z) \wedge S(y) \wedge y > 2)$ false. Even if it were true, this value for y would not be allowed.

Finally, $x = 1$ is not an answer, because the only candidate value for y, namely 1, makes the formula false. In consequence, *null* is the only answer.

The next example with SQL queries and `NULL` provides additional intuition and motivation for the formal semantics of Section 4.2. Notice the use in logical queries of the new unary predicates *IsNull* and *IsNotNull*.

Example 4.2

Consider the schema $\mathcal{S} = \{R(A, B), S(B, C)\}$ and the instance in the table below. In it, `NULL` is the SQL null. If this instance is stored in an SQL database, we can observe the behavior of the following queries when they are directly translated into SQL and run on an SQL DB:

R	A	B
	a	b
	a	c
	d	NULL
	d	e
	u	u
	v	NULL
	v	r
	NULL	NULL

S	B	C
	b	h
	NULL	s
	l	m

(a) $\mathcal{Q}_1(x, y) : R(x, y) \wedge y = null$
SQL: `Select * from R where B = NULL;`
Result: No tuple.

(b) $\mathcal{Q}'_1(x, y) : R(x, y) \wedge IsNull(y)$
SQL: Now uses `IS NULL`
Result: ⟨d, NULL⟩, ⟨v, NULL⟩, ⟨NULL, NULL⟩.

(c) $\mathcal{Q}_2(x, y) : R(x, y) \wedge y \neq null$
SQL: `Select * from R where B <> NULL;`
Result: No tuple.

(d) $\mathcal{Q}'_2(x, y) : R(x, y) \wedge IsNotNull(y)$
SQL: Now uses `IS NOT NULL`
Answer: The five expected tuples.

(e) $\mathcal{Q}_3(x, y) : R(x, y) \wedge x = y$
SQL: `Select * from R` where A = B;
Result: ⟨u, u⟩.

(f) $\mathcal{Q}_4(x, y) : R(x, y) \wedge x \neq y$
SQL: `Select * from R where A <>` B; Result: ⟨a, b⟩, ⟨a, c⟩, ⟨d, e⟩, ⟨v, r⟩.

(g) $\mathcal{Q}_5(x, y, x, z) : R(x, y) \wedge R(x, z) \wedge y \neq z$
SQL: `Select * from R r1, R r2 where r1.A = r2.A and r1.B <> r2.B;`
Result: ⟨a, b, a, c⟩, ⟨a, c, a, b⟩.

(h) $\mathcal{Q}_6(x, y, z, t) : R(x, y) \wedge S(z, t) \wedge y = z$
SQL: `Select * from R r1, S s1 where r1.B = s1.B;`
Result: ⟨a, b, b, h⟩.

(i) As in (h), but now SQL: `Select * from R r1 join S s1 on r1.B = s1.B;`
Result[12]: ⟨a, b, b, h⟩.

(j) $\mathcal{Q}_7(x, y, z, t) : R(x, y) \wedge S(z, t) \wedge y \neq z$. In SQL:

[12] The same result is obtained from DBMSs that do not require an explicitly equality together with the join.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

`Select R1.A, R1.B, S1.B, S1.C from R R1, S S1 where R1.B <> S1.B';`
Result: $\langle a, c, b, h\rangle, \langle d, e, b, h\rangle, \langle u, u, b, h\rangle, \langle v, r, b, h\rangle$, $\langle a, b, l, m\rangle, \langle a, c, l, m\rangle, \langle d, e, l, m\rangle$, $\langle u, u, l, m\rangle, \langle v, r, l, m\rangle$.

We need to introduce predicates *IsNull* and *IsNotNull* in our formal treatment of nulls, because, as shown in Example 4.2, in the presence of `NULL`, SQL treats `IS NULL` and `IS NOT NULL` differently from classical $=$ and $\neq$, resp. For example, the queries

$$\mathcal{Q}(x) : \exists y(R(x, y) \wedge \mathit{IsNull}(y)) \quad \text{and} \quad \mathcal{Q}'(x) : \exists y(R(x, y) \wedge y = \mathit{null}) \tag{4.10}$$

are both conjunctive queries, but in SQL databases, they have different semantics.

4.2 *Query answering under* `NULL`*: the semantics*

Here, we introduce the semantics of FO conjunctive query answering in relational databases with null values. More precisely, in SQL relational databases with a single null value, *null*, that is handled like the SQL `NULL`. We will exclude from the "SQL-like" conjunctive queries those such as (a) and (c) in Example 4.2. We will concentrate only on conjunctive queries with built-ins.

Definition 4.1
(a) $\mathit{Conj}(\Sigma^{\mathtt{null}})$ denotes the class of conjunctive queries in $\mathscr{L}(\Sigma^{\mathtt{null}})$ of the form

$$\mathcal{Q}(\bar{x}) : \exists \bar{y}(A_1(\bar{x}_1) \wedge \cdots \wedge A_n(\bar{x}_n)) \tag{4.11}$$

where $\bar{y} \subseteq \bigcup_i \bar{x}_i$, $\bar{x} = (\bigcup_i \bar{x}_i) \smallsetminus \bar{y}$, and the A_i are atoms containing any of the predicates in $\mathscr{R} \cup \mathscr{B}^{\mathtt{null}}$ plus terms, i.e. variables or constants in $\mathscr{U}$.

(b) $\mathit{Conj}^{\mathtt{sql}}(\Sigma^{\mathtt{null}})$ denotes the class of conjunctive queries as in (a) whose conjuncts are never of the form $t = \mathit{null}$, $t \neq \mathit{null}$, with t a term (*null* or not).

For example, for the queries in (4.10), $\mathcal{Q}, \mathcal{Q}' \in \mathit{Conj}(\Sigma^{\mathtt{null}})$, $\mathcal{Q} \in \mathit{Conj}^{\mathtt{sql}}(\Sigma^{\mathtt{null}})$ but $\mathcal{Q}' \notin \mathit{Conj}^{\mathtt{sql}}(\Sigma^{\mathtt{null}})$. Notice that $\mathit{Conj}^{\mathtt{sql}}(\Sigma^{\mathtt{null}}), \mathit{Conj}(\Sigma) \subseteq \mathit{Conj}(\Sigma^{\mathtt{null}})$. The idea is to force conjunctive queries *à la* SQL that explicitly mention the null value in (in)equalities, to use the built-ins *InNull* or *IsNotNull*.

Definition 4.2
Given a query in $\mathit{Conj}(\Sigma^{\mathtt{null}})$ of the form $\mathcal{Q}(\bar{x})$: $\exists \bar{y}\psi(\bar{x}, \bar{y})$, with ψ quantifier-free, a variable v is *relevant* for $\mathcal{Q}$ if it occurs at least twice in ψ, without considering the atoms $\mathit{IsNull}(v)$, $\mathit{IsNotNull}(v)$, $v\ \theta\ \mathit{null}$, or $\mathit{null}\ \theta\ v$, with $\theta \in \mathscr{B}$. $\mathit{RelV}(\mathcal{Q})$ denotes the set of relevant variables for $\mathcal{Q}$.

For example, for the query $\mathcal{Q}(x) : \exists y(P(x, y, z) \wedge Q(y) \wedge \mathit{IsNull}(y))$, $\mathit{RelV}(\mathcal{Q}(x)) = \{y\}$, because y is used twice in the subformula $P(x, y, z) \wedge Q(y)$.

As usual in FO logic, we consider assignments from the set, *Var*, of variables to the underlying database domain $\mathscr{U}$ (that contains constant *null*), i.e. $\sigma : \mathit{Var} \to \mathscr{U}$. Such an assignment can be extended to terms, mapping variables x to $\sigma(x)$, and $c \in \mathscr{U}$ to c. For an assignment σ, a variable y and a constant c, $\sigma\frac{y}{c}$ denotes the assignment that coincides with σ everywhere, except possibly on y, that takes the

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

value c. Given a formula ψ, $\psi[\sigma]$ denotes the formula obtained from ψ by replacing its free variables by their values according to σ.

Now, given a formula (query) χ and an assignment σ, we verify if instance D satisfies $\chi[\sigma]$ by assuming that the quantifiers on relevant variables range over $(\mathscr{U} \setminus \{null\})$, and those on non-relevant variables range over $\mathscr{U}$.

Definition 4.3
Let $\chi \in Conj(\Sigma^{\mathtt{null}})$, and σ be an assignment. Instance D with σ *satisfies* χ *under the null-semantics*, denoted $D \models_N \chi[\sigma]$, exactly in the following cases: (below $t, t_1, \ldots$ are terms; and x, x_1, x_2 variables)

(1) (a) $D \models_N IsNull(t)[\sigma]$, with $\sigma(t) = null$; (b) $D \models_N IsNotNull(t)[\sigma]$, with $\sigma(t) \neq null$.
(2) $D \models_N (t_1 < t_2)[\sigma]$, with $\sigma(t_1) \neq null \neq \sigma(t_2)$, and $\sigma(t_1) < \sigma(t_2)$.[13]
(3) (a) $D \models_N (x = c)[\sigma]$, with $\sigma(x) = c \in (\mathscr{U} \setminus \{null\})$. (or symmetrically)[14]
    (b) $D \models_N (x_1 = x_2)[\sigma]$, with $\sigma(x_1) = \sigma(x_2) \neq null$.
    (c) $D \models_N (c = c)[\sigma]$, with $c \in (\mathscr{U} \setminus \{null\})$.
(4) (a) $D \models_N (x \neq c)[\sigma]$, with $null \neq \sigma(x) \neq c \in (\mathscr{U} \setminus \{null\})$. (or symmetrically)
    (b) $D \models_N (c_1 \neq c_2)[\sigma]$, with $c_1 \neq c_2$, and $c_1, c_2 \in (\mathscr{U} \setminus \{null\})$.
(5) $D \models_N R(t_1, \ldots, t_n)[\sigma]$, with $R \in \mathscr{R}$, and $R(\sigma(t_1), \ldots, \sigma(t_n)) \in D$.
(6) $D \models_N (\alpha \wedge \beta)[\sigma]$, with α, β quantifier-free, $\sigma(y) \neq null$ for every $y \in RelV(\alpha \wedge \beta)$, and $D \models_N \alpha[\sigma]$ and $D \models_N \beta[\sigma]$.
(7) $D \models_N (\exists y\ \alpha)[\sigma]$ when: (a) if $y \in RelV(\alpha)$, there is c in $(\mathscr{U} \setminus \{null\})$ with $D \models_N \alpha[\sigma \frac{y}{c}]$; or (b) if $y \notin RelV(\alpha)$, there is c in $\mathscr{U}$ with $D \models_N \alpha[\sigma \frac{y}{c}]$.

This semantics also applies to $Conj^{\mathtt{sql}}(\Sigma^{\mathtt{null}})$.

Definition 4.4
Let $\mathscr{Q}(\bar{x}) : \exists \bar{y} \psi(\bar{x}, \bar{y})$ be in $Conj(\Sigma^{\mathtt{null}})$, with $\bar{x} = x_1, \ldots, x_n$, and ψ quantifier-free.

(a) A tuple $\langle c_1, \ldots, c_n \rangle \in \mathscr{U}^n$ is an *N-answer from* D *to* $\mathscr{Q}$, denoted $D \models_N \mathscr{Q}[c_1, \ldots, c_n]$, iff there is an assignment σ such that $\sigma(x_i) = c_i$, for $i = 1, \ldots, n$; and $D \models_N (\exists \bar{y} \psi)[\sigma]$.
(b) If $\mathscr{Q}$ is a sentence (a Boolean query), the N-answer is *yes* iff $D \models_N \mathscr{Q}$, and *no*, otherwise.
(c) $\mathscr{Q}^{Nsem}(D)$ denotes the set of N-answers to $\mathscr{Q}$ from instance D.

Notice that $D \models_N (\exists \bar{y} \psi)[\sigma]$ in (a) above requires, according to Definition 4.3, that the relevant variables in the existential prefix $\exists \bar{y}$ do not take the value *null*. The free variables, i.e. in $\bar{x}$, may take the value *null* only when they are not relevant in the query. For illustration, in Example 4.1, since the free variable x is not relevant, $\mathscr{Q}_2^{Nsem}(D_2) = \{\langle null \rangle\}$.

Example 4.3
Consider the instance D and the conjunctive query below:

[13] Of course, when there is an order relation on $\mathscr{U}$. We could introduce ">" similarly.
[14] Notice the use of the symbols $=$ and $\neq$ both at the object and the meta levels.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

R	A	B
	a	b
	c	d
	e	*null*

S	B	C
	b	f
	d	g
	null	j

$\mathcal{Q}(x, z)$: $\exists y(R(x, y) \wedge S(y, z))$.

Here, under the classical semantics, $\mathcal{Q}(D) = \{\langle a, f\rangle, \langle c, g\rangle, \langle e, j\rangle\}$, treating *null* as any other constant. However, $\mathcal{Q}^{Nsem}(D) = \{\langle a, f\rangle, \langle c, g\rangle\} \subseteq \mathcal{Q}(D)$.

It is easy to prove that, for queries in $Conj(\Sigma^{\mathtt{null}})$: $\mathcal{Q}^{Nsem}(D) \subseteq \mathcal{Q}(D)$. Furthermore, the N-query-answering semantics coincides with classical FO query-answering semantics in databases without *null*. More precisely, if $null \notin \mathcal{U}$ (and then it does not appear in D or $\mathcal{Q}$ either): $D \models_N \mathcal{Q}[\bar{t}]$ iff $D \models \mathcal{Q}[\bar{t}]$.

4.3 Constraint satisfaction under NULL via FO rewriting

The notions of relevant attributes (or variables) and formula satisfaction under the null-semantics can be both extended to more complex formulas. In particular, they can be applied to constraint satisfaction under SQL NULL (Bravo and Bertossi 2006; Bravo 2007). As expected, the satisfaction of a constraint ψ by a database that may contain *null* depends upon the presence of *null* in the *relevant attributes* of ψ. The following is a generalization of Definition 4.2 to a larger class of formulas.

Definition 4.5
For $\psi \in \mathcal{L}(\Sigma^{\mathtt{null}})$ in prenex normal form[15], a variable v is *relevant* if it occurs at least twice in ψ, without considering occurrences in quantifiers or atoms of the forms $IsNull(v)$, $IsNotNull(v)$, $v\ \theta\ null$, or $null\ \theta\ v$, where θ a built-in comparison predicate. $RelV(\psi)$ denotes the set of relevant variables of ψ.

The constraints we are considering in this work, particularly those of the form (2.2), may not be in prenex normal form, but can be easily transformed while keeping the same variables and their occurrences, so that relevant variables can be determined.

Example 4.4
Consider the *referential integrity constraint* (RIC) on schema $\mathcal{R} = \{P(A, B, C), R(A, B, E)\}$: ψ : $\forall x \forall y \forall z(P(x, y, z) \rightarrow \exists v R(x, y, v))$; and the instance D:

P	A	B	C
	a	5	d
	b	*null*	a

R	A	B	E
	a	5	3
	a	3	7

DBMSs implement the so-called "simple semantics" of the SQL Standard for satisfaction of ICs. According to it, the database D above satisfies the RIC. This is because, for every tuple t in P, if $t[A]$ and $t[B]$ are different from *null*, there is a tuple t' in R with $t[A, B] = t'[A, B]$.

[15] That is, of the form $\bar{Q}\chi$, where $\bar{Q}$ is a prefix of quantifiers, and χ is a quantifier-free formula. Actually, in this work, all the formulas have a quantifier prefix of the form $\bar{\forall}\bar{\exists}$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

In this case, and informally, the attributes that are relevant for checking the satisfaction of the RIC (i.e. those we attempt to capture through the relevant variables) are A and B, in both P and R. If we try to insert tuple $P(c, d, null)$ into P, the DBMS will reject the insertion, because none of the attributes that are relevant for checking the constraint are *null*, and there is no tuple $R(c, d)$.

More precisely, the set of relevant variables for ψ is $RelV(\psi) = \{x, y\}$, because x and y appear twice in ψ. Accordingly, the values for attributes C and E are not relevant when checking the satisfaction of ψ, which makes sense.

We will now formalize constraint satisfaction under the null-semantics. In this direction, we consider a single class of constraints that includes those in (2.1) and (2.2), and can be handled in a uniform manner. It also includes all the common constraints used in data management. More precisely, we consider constraints that are sentences in $\mathcal{L}(\Sigma^{\mathtt{null}})$ of the form:

$$\forall \bar{x}(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow \bigvee_{j=1}^{m} C_j) \tag{4.12}$$

where $m, n \neq 0$, R_i is a predicate in $\mathscr{R}$, C_j is a conjunctive formula of the form $\exists \bar{y}_j \bigwedge_{k=1}^{l} Q_{jk}(\bar{x}'_{jk}, \bar{y}_{jk})$, where each Q_{jk} is a predicate in $\mathscr{R}$ or a built-in[16], $\bar{x} = \bigcup_{i=1}^{n} \bar{x}_i$, $\bar{x}'_{jk} \subseteq \bar{x}$ and $\bar{y}_j = \bigcup_{k=1}^{l} \bar{y}_{jk}$.

Without loss of generality, we assume that the existentially quantified variables (the $\bar{y}_j$) do not appear in $\bar{x}$ and are different for each conjunctive literal C_j. Notice that (4.12) allows formulas with only built-in atoms in the consequent. We also assume that they do not have any *explicit* occurrence of the constant *null*.[17] Those formulas may contain the *IsNull* or *IsNotNull*, but as built-in predicates.

In order to define N-satisfaction of a constraint $\psi \in \mathcal{L}(\Sigma^{\mathtt{null}})$ of the form (4.12), we first rewrite it into a new FO formula ψ^N, which makes explicit the role played by the relevant attributes in ψ (as in Definition 4.5) and the way nulls are handled in them. Next, satisfaction is defined in terms of the rewriting.

Definition 4.6

Let $\psi \in \mathcal{L}(\Sigma^{\mathtt{null}})$ be a constraint of the form (4.12), i.e. $\forall \bar{x}(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow \bigvee_{j=1}^{m} C_j)$.

(a) *The N-rewriting of* ψ is the $\mathcal{L}(\Sigma^{\mathtt{null}})$-sentence:

$$\psi^N: \quad \forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \rightarrow \left(\bigvee_{v \in RelV(\psi) \cap \bar{x}} IsNull(v) \vee \bigvee_{j=1}^{m} C_j^N \right) \right) \tag{4.13}$$

[16] Occurrences of variables in built-ins have to be *safe*, i.e. they also appear in $\bar{x}$ or in database predicate in the same conjunction.

[17] This is not an essential requirement, but will simplify the presentation. Furthermore, all reasonable DECs and local ICs do not require the explicit use of *null* as a constant. If we need to express a relational NOT-NULL-constraint, we can say, e.g. $\forall x \forall y(R(x, y) \rightarrow IsNotNull(x))$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

with $\bar{x} = \cup_{i=1}^{n} \bar{x}_i$ and

$$C_j^N = \exists \bar{y}_j \left(\bigwedge_{k=1}^{l} Q_{jk}(\bar{x}'_{jk}, \bar{y}_{jk}) \ \wedge \bigwedge_{w \in ((RelV(\psi) \setminus \bar{x}) \cap \bar{y}_j)} IsNotNull(w) \right) \tag{4.14}$$

(b) For an instance D, possibly containing *null*, ψ is *N-satisfied by* D, denoted $D \models_N \psi$, iff $D \models \psi^N$. Here, $D \models \psi^N$ refers to classical first-order satisfaction, with *null* treated as any other constant of the domain.

We can see from Definition 4.6 that there are basically two cases when a ground instantiation of ψ (obtained by assigning constants to the variables in $\bar{x}$) is immediately satisfied due to the presence of *null*: (a) When *null* appears in any of the relevant attributes in the antecedent. (b) At least one of the conjunctive formulas C_j is true, considering that when they are checked according to equation (4.14), *null* is treated as any other constant, but the variables in existential joins do not take the value *null* (hence the condition based on $w \in (RelV(\psi) \setminus \bar{x}) \cap \bar{y}_j$ in (4.14)). The rewriting can be applied in particular to our UDECs and RDECs, as follows.

For a UDEC ψ of the form (2.1), ψ^N is

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \ \rightarrow \ \left(\bigvee_{v \in RelV(\psi)} IsNull(v) \ \vee \ \bigwedge_{k=1}^{l} \left(\bigvee_{j=1}^{m} Q_{kj}(\bar{y}_j) \right) \right) \right) \tag{4.15}$$

For an RDEC ψ of the form (2.2), ψ^N is

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \ \rightarrow \ \left(\bigvee_{v \in (RelV(\psi) \cap \bar{x})} IsNull(v) \ \vee \exists \bar{y} \left(\bigwedge_{k=1}^{l} (Q_k(\bar{x}_k, \bar{y}_k) \wedge \varphi_k(\bar{x}'_k, \bar{y}'_k)) \wedge \right. \right. \right.$$
$$\left. \left. \left. \bigwedge_{w \in (RelV(\psi) \setminus \bar{x})} IsNotNull(w) \right) \right) \right) \tag{4.16}$$

Example 4.5
(example 4.4 cont.) The rewriting ψ^N of the RIC ψ is, according to (4.16):

$$\psi^N : \ \forall x \forall y \forall z (P(x, y, z) \ \rightarrow \ IsNull(x) \vee IsNull(y) \vee \exists w R(x, y, w)).$$

D classically satisfies ψ^N, treating *null* as any other constant. Then, $D \models_N \psi$.

Example 4.6
For $\psi : \ \forall x (R(x) \rightarrow \exists y (T(x, y) \wedge S(y)), \ RelV(\psi) = \{x, y\}$. From (4.16):

$$\psi^N : \ \forall x (R(x) \ \rightarrow \ IsNull(x) \vee \exists y (T(x, y) \wedge S(y) \wedge IsNotNull(y)).$$

(a) For $D = \{R(a)\}$, $D \not\models \psi^N$; and then, $D \not\models_N \psi$.
(b) For $D = \{R(a), T(a, null), S(null)\}$, $D \not\models \psi^N$; and then, $D \not\models_N \psi$.
(c) For $D = \{R(a), T(a, b), S(b)\}$, $D \models \psi^N$; and then, $D \models_N \psi$.
(d) For $D = \{R(null)\}$, $D \models \psi^N$; and then, $D \models_N \psi$.
(e) For $D = \emptyset$, $D \models \psi^N$; and then, $D \models_N \psi$.

Example 4.7

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Consider the schema $\mathcal{S} = \{R(A, B, E)\}$ with the primary *key constraint* (KC) $R : AB \rightarrow E$, expressing that attributes A, B, together, functionally determine attribute E.[18] An SQL database with the instance $D = \{R(a, 5, 3), R(a, 3, 7)\}$, which contains no nulls, would satisfy the KC.

The insertion of $R(null, 4, 5)$ into D would be rejected since the KC would be violated: *null* is not allowed in a key attribute.The insertion of $R(a, 5, null)$ would also be rejected, but now due to the potential occurrence of two R-tuples t_1, t_2 with $t_1[A] = t_2[A] = a$ and $t_1[B] = t_2[B] = 5$, but being unknown whether $t_1[E] = t_2[E]$.

In order to fully capture this KC in predicate logic and the intended semantics we just described, we need the following sentences to represent the KC:

$$\psi_1 : \ \forall x \forall y \forall z_1 \forall z_2 (R(x, y, z_1) \wedge R(x, y, z_2) \rightarrow z_1 = z_2) \tag{4.17}$$

$$\psi_2 : \ \forall x \forall y \forall z_1 \forall z_2 (R(x, y, z_1) \wedge R(x, y, z_2) \wedge IsNull(z_1) \rightarrow IsNull(z_2)) \tag{4.18}$$

$$\psi_3 : \ \forall x \forall y \forall z (R(x, y, z) \wedge IsNull(x) \rightarrow \textbf{false}) \tag{4.19}$$

$$\psi_4 : \ \forall x \forall y \forall z (R(x, y, z) \wedge IsNull(y) \rightarrow \textbf{false}) \tag{4.20}$$

The first two constraints ensure that, if two tuples coincide in attributes A and B, then they have the same value in z or they are both *null*. The last two constraints ensure that the values in attributes A or B in relation R cannot be *null*.

Actually, (4.18) can be written as

$$\forall x \forall y \forall z_1 \forall z_2 (R(x, y, z_1) \wedge R(x, y, z_2) \rightarrow \neg IsNull(z_1) \vee IsNull(z_2)),$$

or equivalently, as a UDEC of the form (2.1):

$$\psi_2' : \ \forall x \forall y \forall z_1 \forall z_2 (R(x, y, z_1) \wedge R(x, y, z_2) \rightarrow IsNotNull(z_1) \vee IsNull(z_2)).$$

Similarly, (4.19) and (4.20) can be written as UDECs of the form (2.1):

$$\psi_3' : \ \forall x \forall y \forall z (R(x, y, z) \ \rightarrow \ IsNotNull(x)),$$
$$\psi_4' : \ \forall x \forall y \forall z (R(x, y, z) \ \rightarrow \ IsNotNull(y)).$$

So, finally, the KC is represented by the UDECs $\psi_1, \psi_2', \psi_3', \psi_4'$.

For illustration, considering that $RelV(\psi_1) = \{x, y, z_1, z_2\}$, and $RelV(\psi_3') = \emptyset$, we obtain from (4.15):

$$\psi_1^N : \ \forall x \forall y \forall z_1 \forall z_2 (R(x, y, z_1) \wedge R(x, y, z_2) \ \rightarrow \ IsNull(x) \vee IsNull(y) \vee IsNull(z_1) \vee IsNull(z_2) \vee z_1 = z_2) \tag{4.21}$$

$$(\psi_3')^N : \ \forall x \forall y \forall z (R(x, y, z) \ \rightarrow \ IsNotNull(x)) \tag{4.22}$$

Finally, observe that when the atom $z_1 = z_2$ in ψ_1^N is evaluated on the domain, the constant *null* is treated as any other constant, i.e. $null = c$ is true only when constant c is *null*.

We have managed to reduce constraint satisfaction under `NULL` as handled by SQL databases to formula satisfaction in FO predicate logic, without assigning any

[18] A KC is a particular kind of functional dependency, where a set of attributes of a relational predicate functionally determines all the attributes of the predicate. Declaring a KC as *primary* in an SQL-based relational DBMS has in particular the effect that `NULL` is not accepted in key attributes.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

special status to *null*, which is treated as an ordinary constant. We considered only constraints of the form (4.12), which include all our DECs and ICs for peers, in which case, the generic set of database predicates $\mathscr{R}$ will contain those for a single peer or a pair thereof. However, the semantics can be extended to more general FO formulas (Bravo 2007).

4.4 *Query answering under* `NULL` *via FO rewriting*

The notion of query answer under `NULL`, that of N-answer given in Section 4.2, being interesting *per se*, allowed us to motivate the treatment of constraints via relevant variables. However, query answering in presence of *null* can also be treated via the rewriting in Definition 4.6. This is achieved by considering conjunctive queries as a special case of (4.12), without the $\forall\bar{x}$ and empty antecedents; and obtaining their relevant variables through Definition 4.5.

More precisely, a conjunctive query $\mathscr{Q}(\bar{x}) \in Conj(\Sigma^{\mathtt{null}})$, i.e. of the form (4.11), is rewritten into a new conjunctive query as follows:

$$\mathscr{Q}^N(\bar{x}):\ \exists\bar{y}(A_1(\bar{x}_1) \wedge \cdots \wedge A_n(\bar{x}_n)\ \wedge \bigwedge_{v \in RelV(\mathscr{Q})} v \neq null) \tag{4.23}$$

It holds: $D \models_N \mathscr{Q}[\bar{c}]$ iff $D \models \mathscr{Q}^N[\bar{c}]$, where the latter is classic FO satisfaction, with *null* treated as an ordinary constant in the domain.

This transformation ensures that relevant variables range over $(\mathscr{U} \setminus \{null\})$. $\mathscr{Q}^N(\bar{x}) \in Conj(\Sigma^{\mathtt{null}})$, and may contain atoms of the form $IsNull(t)$ or $IsNotNull(t)$. However, they can be replaced by $t = null$ or $t \neq null$, resp., leading to a query in $Conj(\Sigma)$, but with the same answers as (4.23).

Example 4.8
Consider the instance D below, and the query $\mathscr{Q}(x):\ \exists y(P(x, y) \wedge y > 5)$, for which $RelV(\mathscr{Q}) = \{y\}$. In this case, the bound variable is the only relevant variable, and then ranges over non-null values when checking query satisfaction.

P	A	B
	f	7
	f	5
	null	8
	b	*null*

Accordingly, $D \models_N \mathscr{Q}[f]$ holds, because $\exists y(P(f, y) \wedge y > 5)$ is true in D, with 7 as a non-null value for y that makes the formula true. This result is confirmed by the rewriting of $\mathscr{Q}(x)$.

For $\mathscr{Q}^N(x):\ \exists y(P(x, y) \wedge y > 5 \wedge IsNotNull(y))$, $D \models \mathscr{Q}^N[f]$ holds, under classical satisfaction, with *null* treated as any other constant. Similarly, $D \models_N \mathscr{Q}[null]$.

Example 4.9
(example 4.1 continued) The query in (4.9) can be rewritten as

$$\mathscr{Q}_2^N(x):\ \exists y\exists z(R(x, y, z) \wedge S(y) \wedge y > 2 \wedge y \neq null).$$

We had $D \not\models_N \mathscr{Q}_2[1]$. Now, $D \not\models \exists y\exists z(R(1, y, z) \wedge S(y) \wedge y > 2 \wedge y \neq null)$, classically, with *null* treated as an ordinary constant. As expected, $D \not\models \mathscr{Q}_2^N[2]$ due to the new conjunct $y \neq null$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Finally, $D \models \mathcal{Q}_2^N[null]$ because $D \models (R(null, 3, 3) \wedge S(3) \wedge 3 > 2 \wedge 3 \neq null)$. Since *null* is treated as any other constant, we can compare it with 3. By the *unique names assumption*, $3 \neq null$ holds.

Our query-answering N-semantics for $Conj(\Sigma^{\mathtt{null}})$ can be applied in particular any SQL-like conjunctive query, but first expressing it as a query $\mathcal{Q}$ in $Conj^{\mathtt{sql}}(\Sigma^{\mathtt{null}})$, and then computing and classically evaluating $\mathcal{Q}^N$.

The notions of constraint satisfaction and query answer in the presence of `NULL` *à la* SQL we introduced in this section coincide with the classic notions in databases without *null*.

In the rest of this work, the notion of formula satisfaction that we have denoted with $\models_N$ will be simply denoted with $\models$. When constraints are not satisfied in this sense, we will apply a particular repair semantics that captures the special role of *null*. It is introduced in the next section.

5 Solution semantics with NULL

In the preceding section, we introduced a notion of formula satisfaction with database instances that may contain *null*. We now use it for defining the semantics of a peer system where the data movement process relies on consistency restoration of peer instances with respect to DECs. The value *null* will be used to replace existentially quantified variables in consequents of referential DECs and local referential ICs as a possibility to consider for DEC and IC enforcement. Another possibility for the same task is tuple deletion.

The repair semantics that supports consistency restoration will have to capture and be sensitive to the (possibly multiple) presence of *null* in the database and its use for consistency enforcement. In particular, it has to give an account of the facts that:

(1) When atoms are inserted into the database, an existential variable that appears in a join or in a built-in atom in an RDEC's consequent (a so-called *problematic existential variable*) is never replaced by (or takes the value) *null*.
(2) Arbitrary non-null constants from the domain are not used for these problematic variables either. Actually, rather than introducing arbitrary values for existential variables of this kind, tuple deletions from antecedents in DECs will be privileged.

5.1 A restricted chase

In order to achieve the just stated goals, we first introduce a restricted, *ad hoc* form of the chase (Abiteboul *et al.* 1995), as an auxiliary construct. It is applied only with the DECs and ICs at hand that do not have problematic existential variables in their consequents. The enforcement of this subset of the constraints introduces tuples with *null* for existential variables (that do not appear in joins though).

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

In the end, this chase will return a finite set of atoms that will be used as a possibly generous upper-bound for the admissible tuple insertions that create repairs and solutions on the basis of the whole sets of DECs at hand.

As any other form of chase, our restricted chase enforces the satisfaction of constraints, but in our case, the notion of satisfaction corresponds to that introduced in Section 4.3.

In the rest of this subsection, we consider a database instance D and a set of DECs $\Sigma = \Sigma_1 \cup \Sigma_2$, with Σ_1 a set of UDECs of the form (2.1) and Σ_2 a set of RDECs of the form (2.2). We consider the subset, Σ_2^- of Σ_2 that contains all the RDECs except for those involving existentially quantified variables in joins or built-in atoms. Accordingly, we consider $\Sigma^- := \Sigma_1 \cup \Sigma_2^-$.

Definition 5.1
The *restricted chase*, *r-Chase*$^{null}(D, \Sigma^-)$, with an instance D and $\Sigma^- = \Sigma_1 \cup \Sigma_2^-$, is the instance D' obtained as a fix-point of the following iterative procedure:

(1) $D'_0 := D$.
(2) Given instance D'_s, if $D'_s \models \Sigma^-$ (relative to $\models_N$), then $D'_{s+1} := D'_s$.
(3) Given instance D'_s, if a ground instantiation, $\varphi\downarrow$, of a constraint $\varphi \in \Sigma_1$ (with constants or *null*) is not satisfied by D'_s (relative to $\models_N$), but its antecedent is, then D'_{s+1} is obtained from D'_s by adding the ground database atoms appearing in every disjunct of every conjunct in the consequent of $\varphi\downarrow$ (but only when the built-ins in the conjunct are satisfied).
Notice that this excludes the generation of atoms with any kind of built-ins, including the atom **false**.
(4) Given instance D'_s, if a ground instantiation (obtained by replacing universal variables with constants or *null*, but keeping the existential quantifiers) of a constraint in Σ_2^- is not satisfied by D'_s , but its antecedent is, then replace the existentially quantified variables by *null*, and build D'_{s+1} by adding to D'_s the corresponding ground atoms in the conjunct in the consequent (but only when the built-ins in it are satisfied).

This procedure finitely terminates as we show below, but first some intuitions, explanations, and examples. The chase procedure propagates and invents values due to the enforcement of RECs with non-problematic existential quantifiers and UDECs. As a result, the constants in the chase instance will be those already appearing in the initial instance D, those that appear explicitly in consequents of DECs, or *null* as a value invented for (non-problematic) existentially quantified variables.

Example 5.1
Consider the following set of constraints, classified according to Definition 5.1:

$$\Sigma_1 = \{\forall x \forall y (T(x,y) \rightarrow R(x,y)) \tag{5.24}$$
$$\forall x \forall y \forall z (R(x,y) \wedge S(y,z) \rightarrow Q(x,y,z) \vee T(x,z)) \tag{5.25}$$
$$\forall x \forall y \forall z (Q(x,y,z) \rightarrow S(x,y) \wedge R(y,z)) \tag{5.26}$$
$$\forall x \forall y \forall z (T(x,y) \wedge T(x,z) \rightarrow y = z) \tag{5.27}$$

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

$$\forall x \forall y (T(x, y) \wedge S(x, y) \rightarrow \textbf{false}) \tag{5.28}$$

$$\Sigma_2^- = \{\forall x \forall y (R(x, y) \rightarrow \exists z (Q(x, y, z) \wedge x \neq y) \tag{5.29}$$

$$\forall x \forall y \forall z (Q(x, y, z) \rightarrow \exists w (R(x, z) \wedge S(x, w))\} \tag{5.30}$$

Independently from the instance at hand (initial or not) at a chase step, constraints (5.27) and (5.28), will never be applied (or enforced).

If D contains $T(a, null)$, then the ground instantiation $T(a, null) \rightarrow R(a, null)$ of (5.24) is satisfied under $\models_N$, even if D does not contain $R(a, null)$. This is because both variables x, y in (5.24) are relevant (cf. Definition 4.5). So, $R(a, null)$ is not generated. However, if $T(a, b) \in D$ and $R(a, b) \notin D$, the latter atom is generated.

Now, the ground instantiation $R(a, a) \rightarrow \exists z (Q(a, a, z) \wedge a \neq a)$ of (5.29) will not create new tuples because of the inequality $a \neq a$. However, if $R(a, b) \in D$, and the instantiation $R(a, b) \rightarrow \exists z (Q(a, b, z) \wedge a \neq b)$ is not satisfied, the tuple $Q(a, b, null)$ will be created.

The ground instantiation $R(a, b) \wedge S(b, c) \rightarrow Q(a, b, c) \vee T(a, c)$ of (5.25), with $R(a, b)$, $S(b, c) \in D$, will generate both $Q(a, b, c)$ and $T(a, c)$ (if it was not already satisfied).

Notice that creation of $Q(a, b, null)$ above, due to (5.29), feeds the antecedent of (5.26). However, its instantiation $Q(a, b, null) \rightarrow S(a, b) \wedge R(b, null))$ is satisfied, because *null* appear for variable z that is relevant. So, $R(b, null)$ is not generated.

Remark 5.1

(a) Notice that in the DEC (5.26) all variables, x, y, z, are relevant. As a consequence, any ground atom with predicate Q with *null* in it will not trigger the rule. We can see that, every time the chase introduces *null* in a predicate position that turns out to be relevant in another DEC, this latter DEC will not be triggered (because it will be automatically satisfied). This possibly very common situation will contribute to a fast termination of the chase.

(b) We can assume that the chase applies in parallel all possible instantiations of DECs in Σ^-. It is clear that this restricted *null*-based chase always terminates, because only the elements of the initial active domain plus possibly *null* are used to fill arguments in database predicates, which leads to a saturation point in polynomial-time in the size of instance D and the schema. In particular, the resulting instance is of polynomial size in the size of D.

(c) Instead of using the chase procedure (in case we wanted a more model-theoretic or declarative approach to the chase instance), we could replace each sentence in Σ_1 by a collection of Datalog rules, one for each disjunct in the consequent; and each sentence in Σ_2^- by a Datalog rule obtained replacing existential variables by *null*. The chase instance would coincide with the minimal model of the resulting Datalog program (Abiteboul *et al.* 1995).

Notice that $r\text{-}Chase^{null}(D, \Sigma^-)$ may not satisfy Σ, for the trivial reason that not all DECs in Σ are considered in the chase. Actually, $r\text{-}Chase^{null}(D, \Sigma^-)$ may not even satisfy Σ^-, because it ignores DECs with only built-ins in consequents. For illustration, in Example 5.1, the chase of $D = \{T(a, b), T(a, c)\}$ will not satisfy (5.27), and the chase of $D = \{T(a, b), S(a, b)\}$ will not satisfy (5.28).

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

The next step will be imposing $r\text{-}Chase^{null}(D, \Sigma^{-})$ as an upper bound on the extensions of possible repairs. With this, we will discard candidate repairs that introduce arbitrary, non-mandatory, non-null constants. In other words, we will use the result of this chase to test candidate repairs (or solutions), demanding their inserted atoms to be contained in the chase applied with the subset Σ^{-} of the set Σ of DECs at hand.

5.2 Repair semantics with NULL and solutions

A repair semantics for single relational databases and sets of ICs that may include referential ICs is proposed in Bravo and Bertossi (2006) and Bravo (2007). It introduces null values of the kind described in Section 4.3. We will adopt and adapt this repair semantics in the context of PDESs, by taking into account the restrictions imposed by the chase we just introduced, and also the trust relationships.

So as repairs of an inconsistent instance with respect to a set of ICs, solutions for a peer in a PDES are expected to stay "close" to the original peer's instance, while satisfying its DECs. In order to capture "closeness" in our null-based semantics, we need to compare instances and their tuples, which may contain constant *null*.

Definition 5.2
(Leven and Loizou 1997) For constants $c, d \in \mathscr{U}$, c provides *less or equal information* than d, denoted $c \sqsubseteq d$, iff $c = null$ or $c = d$. For sequences $\bar{s}_1 = \langle c_1, \ldots, c_n \rangle$ and $\bar{s}_2 = \langle d_1, \ldots, d_n \rangle$ with $c_i, d_i \in \mathscr{U}$, $\bar{s}_1$ provides less or equal information than $\bar{s}_2$, denoted $\bar{s}_1 \sqsubseteq \bar{s}_2$, iff $c_i \sqsubseteq d_i$ for every $i = 1, \ldots, n$. Finally, $\bar{s}_1 \sqsubset \bar{s}_2$ means $\bar{s}_1 \sqsubseteq \bar{s}_2$ and $\bar{s}_1 \neq \bar{s}_2$.

Definition 5.3
Consider D, D', D'' be database instances for the same schema, and a set Σ of DECs of the form (2.1) or (2.2) in the language of the schema. Let Σ^{-} be the subset of Σ that excludes the DECs with problematic existential variables (cf. Section 5.1).

(a) D' *is at least as close to* D *as* D'' (is to D), denoted $D' \leqslant_D^{\Sigma} D''$, iff one of the following holds:
  (1) $D'' \notin r\text{-}Chase^{null}(D, \Sigma^{-})$.
  (2) For every $P(\bar{a}) \in \Delta(D, D')$, there exists $P(\bar{a}') \in \Delta(D, D'')$, such that:
    (i) $\bar{a} \sqsubseteq \bar{a}'$ and (ii) if $\bar{a} \sqsubset \bar{a}'$, then $P(\bar{a}') \notin \Delta(D, D')$.

(b) D' *is closer to* D *than* D'' (is to D), denoted $D' <_D^{\Sigma} D''$, iff $D' \leqslant_D^{\Sigma} D''$, but not $D'' \leqslant_D^{\Sigma} D'$.

Definition 5.4
Given an instance D and a set Σ of DECs of the forms (2.1) or (2.2), a *null-based repair* of D with respect to Σ is an instance D' such that $D' \models_N \Sigma$, and there is no D'', such that $D'' \models_N \Sigma$ and $D'' <_D^{\Sigma} D'$.

This definition ensures, in particular, that a database that, due to the enforcement of an RDEC with a non-problematic existential variable, inserts a tuple with *null* for that variable, is closer to D than another that adds some other, arbitrary constant. At the same time, the definition makes us prefer repairs obtained via tuple deletions

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

instead of tuple insertions (with values null or not), when enforcing RDECs with problematic existential variables.

The null-based repair semantics, in particular Definition 5.4, is also of independent interest in the context of single inconsistent databases with respect to a set of ICs that are like the DECs we are considering in the PDES setting. Actually, this is the official and detailed formalization of the repair semantics first sketched and used in Bravo and Bertossi (2006) for databases with referential ICs. The idea in that case is that inconsistencies are preferably repaired through the insertion of tuples with the special constant *null*, that is used as the null in SQL. When this is not possible, tuple deletions are preferred. This happens either because a constraint cannot be solved via tuple insertions at all (e.g. functional dependencies) or an inclusion dependency would have to be satisfied through the use of *null* for variables in a join. The next two examples show how this repair semantics would be used for repairing a single database.

Example 5.2
(example 4.6 cont.) Consider $\Sigma = \{\forall x(R(x) \rightarrow \exists y(T(x,y) \wedge S(y))\}$ and the instance $D = \{R(a)\}$.

Here, $\Sigma^- = \emptyset$ and $r\text{-}Chase^{null}(D, \Sigma^-) = \{R(a)\}$. Furthermore, $D \not\models \Sigma$ (this is case (a) in Example 4.6).

(i) $D'_1 = \{R(a), T(a, null), S(null)\} \not\models \Sigma$ (this is case (b) in Example 4.6). So, D'_1 is not a repair.

(ii) $D'_2 = \emptyset \models \Sigma$ (this is case (e) in Example 4.6). D'_2 is a repair.
Indeed, for any instance D'' that is not contained in $r\text{-}Chase^{null}(D, \Sigma^-)$, $D'_2 \leqslant^{\Sigma}_{D} D''$ holds. So, the only possible improvement on D'_2 could be $\{R(a)\}$, which does not satisfy Σ.

(iii) $D'_3 = \{R(a), T(a,a), S(a)\} \models \Sigma$ (as case (c) in Example 4.6), but it is not a repair.
Indeed, first notice that $D'_2 \leqslant^{\Sigma}_{D} D'_3$, because $D'_3 \notin r\text{-}Chase^{null}(D, \Sigma^-)$.
Now, $D'_3 \not\leqslant^{\Sigma}_{D} D'_2$, because condition (2) of Definition 5.3 does not hold for $S(a)$ (or $T(a,a)$). We have obtained that $D'_2 <^{\Sigma}_{D} D'_3$, i.e. the former, which is a repair, is strictly closer to D than the latter. The same holds for any instance of the form $D'_4 = \{R(a), T(a,b), S(b)\}$ (and its supersets).

Example 5.3
Consider $\Sigma = \{\forall x \forall y \forall z(T(x,y) \wedge T(x,z) \rightarrow y = z),\ \forall x \forall y(T(x,y) \wedge S(x,y) \rightarrow \textbf{false})\}$, containing a functional dependency and a denial constraint. The instance $D = \{T(a,b), T(a,c), S(a,c)\}$ is inconsistent with respect to Σ.

Here, $\Sigma^- = \Sigma$. However, the restricted chase does not apply any of the ICs, because they contain only built-ins in the consequent. So, $r\text{-}Chase^{null}(D, \Sigma^-) = D$. As a consequence, only repairs based on tuple deletions are acceptable. In this case, $D' = \{T(a,b), S(a,c)\}$ is the only repair.

Since the $\leqslant^{\Sigma}_{D}$ relations can take the place of the generic relation $\leq^{\Sigma(\mathtt{P})}_{D}$ used in Section 3, $\leqslant_D$ determines, through the instantiation of Definition 3.1, a concrete repair semantics that determines null-based neighborhood solutions for a peer P,

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

in which case the definition has to be applied with $\Sigma(\mathtt{P})$. However, Definition 5.4 does not take into account the trust relationships in a neighborhood that are imposed in Definition 3.1. As a consequence, to have a null-based notion of neighborhood solution, we have to impose the additional requirement that certain database relations in a neighborhood instance have to stay untouched (they belong to a peer that is most trusted).

Having a specific notion of neighborhood solution, it is possible to obtain, following the developments in Section 3, also specific definitions of solution and core for a peer, and the notion of PCA. In the rest of this section, we provide some additional examples that show aspects of this specific semantics, and in the following subsections, we investigate computational problems related to this semantics.

Example 5.4

(example 5.2 continued) Consider a PDES instance $\mathfrak{D}$ for the schema $\mathfrak{P}$ with peers P1 and P2, with $\mathscr{S}(\mathtt{P1}) = \{R\}$, $\mathscr{S}(\mathtt{P2}) = \{T, S\}$, and $Trust = \{(\mathtt{P1}, same, \mathtt{P2})\}$. Assume the only DECs are those in $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall x(R(x) \rightarrow \exists y(T(x, y) \wedge S(y))\}$. So, $\Sigma^{-}(\mathtt{P1}, \mathtt{P2}) = \emptyset$.

Assume that the instance of the neighborhood of P1 is $D = \{R(a)\}$. Here, $D \not\models \Sigma(\mathtt{P1}, \mathtt{P2})$. So, in order to obtain a neighborhood solution, we have to repair D with respect to $\Sigma(\mathtt{P1}, \mathtt{P2})$.

Since P1 trusts P2 the same as itself, in principle, we could insert tuples into P2's relations or delete tuples from P1's relation. However, in this case (cf. Example 5.2), since $r\text{-}Chase^{null}(D, \Sigma^{-}(\mathtt{P1}, \mathtt{P2})) = \{R(a)\}$, the only neighborhood solution is the empty instance: $\emptyset \models_N \Sigma(\mathtt{P1}, \mathtt{P2})$ and $\emptyset <_D^{\Sigma(\mathtt{P1})} D'$, for any $\mathscr{S}(\mathtt{P1}, \mathtt{P2})$ – instance D' that inserts tuples.

Similarly, the generic Definition 3.2 of solution and core for peers, can be instantiated with the null-based solution semantics, i.e. on the basis of the $\leqslant_D^{\Sigma}$ relations, obtaining specific versions of those definitions. Accordingly, we recall that, given an instance $\mathfrak{D} = \{D(\mathtt{Q}) \mid \mathtt{Q} \in \mathscr{P}\}$ for the schema $\mathfrak{P}$, the *core* of peer P is the intersection of its solutions: $Core(\mathtt{P}, \mathfrak{D}) := \bigcap Sol(\mathtt{P}, \mathfrak{D})$.

Example 5.5

(example 2.1 continued) Consider the following peers' instances: $D(\mathtt{P1}) = \{R^1(a, 2)\}$, $D(\mathtt{P2}) = \{R^2(c, 4), R^2(d, 5)\}$, $D(\mathtt{P3}) = \{R^3(c, 4)\}$, and $D(\mathtt{P4}) = \{R^4(d, 5, 1)\}$. So, the PDES instance is $\mathfrak{D} = \{D(\mathtt{P1}), D(\mathtt{P2}), D(\mathtt{P3}), D(\mathtt{P4})\}$.

If we want the solutions for P4, we first need the solutions for P3 and P2, who also needs the solution for P3. Since P3 has no DECs with other peers, its only neighborhood solution is its local instance $D(\mathtt{P3})$, which is sent back to P2 or P4 if request.

Peer P2 needs to find the neighborhood solutions for $\{R^2(c, 4), R^2(d, 5), R^3(c, 4)\}$ with respect to $\Sigma(\mathtt{P2}, \mathtt{P3})$. Since P2 trusts P3 the same as itself, it can modify its own data or the data it got from P3. There are two neighborhood solutions for P2: $\{R^2(c, 4), R^2(d, 5)\}$ and $\{R^2(d, 5), R^3(c, 4)\}$, that restricted to P2's schema lead to two solutions for P2: $Sol(\mathtt{P2}, \mathfrak{D}) = \{\{R^2(c, 4), R^2(d, 5)\}, \{R^2(d, 5)\}\}$. Peer P2 sends to P4 the intersection of its solutions: $Core(\mathtt{P2}, \mathfrak{D}) = \{R^2(d, 5)\}$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Neighborhood solutions for P4 are obtained by repairing $\{R^4(d, 5, 1), R^2(d, 5), R^3(c, 4)\}$ with respect to $\Sigma(\mathtt{P4}, \mathtt{P2})$ and $\Sigma(\mathtt{P4}, \mathtt{P3})$. The DECs in $\Sigma(\mathtt{P4}, \mathtt{P2})$ are already satisfied, but not those in $\Sigma(\mathtt{P4}, \mathtt{P3})$. Since P4 trusts the data in P3 more than its own, the only neighborhood solution for P4 is obtained by inserting a tuple with *null* into P4's: $\{R^4(d, 5, 1), R^2(d, 5), R^3(c, 4), R^4(c, 4, null)\}$. Consequently, $Sol(\mathtt{P4}, \mathfrak{D}) = \{\{R^4(d, 5, 1), R^4(c, 4, null)\}\}$.

Due to a particular combination of DECs, trust relationships, and local instances, it is possible that a peer has not solution, which may happen with or without the null-based repair semantics. For illustration, Example 3.2, now with the null-based solution semantics, still does not have solutions.

5.3 Complexity of neighborhood solutions

We now investigate the complexity of decision problems related to the general case of DECs of the form (4.12) (cf. Section 4), in combination with the null-based repair semantics we introduced in this section.

We concentrate mostly on the case that is directly relevant to our forthcoming answer-set programming approach to specifying solutions for individual peers when they have already gathered their neighbors solutions. That is, we consider specifications and reasoning at the neighborhood level (cf. Section 6). Accordingly, we start by analyzing the complexity of deciding if an instance is a neighborhood solution (cf. Definition 3.1).

Definition 5.5

Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, Trust \rangle$ and a peer $\mathtt{P} \in \mathscr{P}$. Given an instance $\bar{D}$ for the neighborhood schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$ around P, and an instance J for the schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, the NeighborhoodSol decision problem is about determining if J is a neighborhood solution for P and $\bar{D}$, i.e. about membership of the set:

$$\mathsf{NeighborhoodSol}(\mathfrak{P}, \mathtt{P}) = \{(J, \bar{D}) \mid J \in NS(\mathtt{P}, \bar{D})\}.$$

This decision problem is parameterized by peer schemas and selected peers. The inputs are database instances. Then, we are considering data complexity.

Proposition 5.1

NeighborhoodSol is *coNP*-complete in the size of $J \cup \bar{D}$. □

Notice that Proposition 5.1 has to be interpreted as follows (and similarly those that follow in this section): For every peer schema $\mathfrak{P}$, $\mathsf{NeighborhoodSol}(\mathfrak{P}, \mathtt{P})$ is in *coNP*; and there is a peer schema $\mathfrak{P}_0$ and a peer $\mathtt{P}_0$ in it, such that the associated problem $\mathsf{NeighborhoodSol}(\mathfrak{P}_0, \mathtt{P}_0)$ is *coNP*-hard.

Proposition 5.1 holds already for PDES with a single peer with local ICs; actually with (cyclic) sets of RDECs of the simple form $\forall \bar{x}(R(\bar{x}) \longrightarrow \exists \bar{y}\ Q(\bar{x}', \bar{y}))$. The proof in this case is a bit more involved since tuple deletions and insertions are in principle possible, but the latter were blocked with the trust relationships. A proof

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

can be found in Bravo (2007)[19]. The proposition also holds in the case where $\bar{D} = \bigcup_{\mathtt{Q} \in \mathscr{N}(\mathtt{P})} J_{\mathtt{Q}}$, but every $J_{\mathtt{Q}}$ with $\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})$ is fixed, i.e. only J and $J_{\mathtt{P}}$ vary.

5.4 Complexity of the core and peer-consistent answers

Definition 5.6
Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ and a peer $\mathtt{P}$ in it. Given a neighborhood instance $\bar{D} = \bigcup_{\mathtt{Q} \in \mathscr{N}(\mathtt{P})} J_{\mathtt{Q}}$ for $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, the *local core* of $\mathtt{P}$ is the intersection of the neighborhood solutions for $\mathtt{P}$ and $\bar{D}$, but restricted to $\mathtt{P}$'s schema $\mathscr{S}(\mathtt{P})$: $\mathit{localCore}(\mathtt{P}, \bar{D}) := (\bigcap NS(\mathtt{P}, \bar{D})) \downharpoonright \mathscr{S}(\mathtt{P})$.

In this definition, $J_{\mathtt{Q}}$ denotes an arbitrary instance for peer $\mathtt{Q}$, which may be different from what we have called the initial instance $D(\mathtt{Q})$ for $\mathtt{Q}$ (cf. Definition 3.2). Actually, $J_{\mathtt{Q}}$ could be $D(\mathtt{Q})$, but also a neighborhood solution for $\mathtt{Q}$ restricted to its schema, or the intersection of the latter, etc. In this regard, we recall that a neighborhood solution for $\mathtt{P}$ is defined in terms of its neighborhood schema, its local instance $D(\mathtt{P})$ in $\mathfrak{D}$, and instances $J_{\mathtt{Q}}$ for other peers $\mathtt{Q}$ in its neighborhood (cf. Definition 3.1). Each $J_{\mathtt{Q}}$ for a neighboring peer $\mathtt{Q}$ is typically the restriction to $\mathscr{S}(\mathtt{Q})$ of the intersection of $\mathtt{Q}$'s local neighborhood solutions, which may not necessarily be the same as the initial instance $D(\mathtt{Q}) \in \mathfrak{D}$.

We can see that the problems of defining and computing neighborhood solutions for a peer can be formulated with arbitrary instances for the neighbors, and do not require taking into consideration the recursive relations between peers. However, the computation of the solutions of a peer is a global problem that does include recursion. In this case, we fix all the instances to be the initial ones for each peer.

Notice from Definition 3.2, that when $\bar{D} = D(\mathtt{P}) \cup \bigcup_{\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})} \mathit{Core}(\mathtt{Q}, \mathfrak{D})$ in Definition 5.6, it holds $\mathit{localCore}(\mathtt{P}, \bar{D}) = \mathit{Core}(\mathtt{P}, \mathfrak{D})$. In particular, Proposition 5.1 remains true if $J_{\mathtt{Q}} = \mathit{Core}(\mathtt{Q}, \mathfrak{D})$, for $\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})$ and $J_{\mathtt{P}} = D(\mathtt{P})$.

Definition 5.7
Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ and a peer $\mathtt{P} \in \mathscr{P}$, the InLocalCore decision problem is about membership of the set:

$$\mathsf{InLocalCore}(\mathfrak{P}, \mathtt{P}) = \{(R(\bar{t}), \bar{D}) \mid R(\bar{t}) \in \mathit{localCore}(\mathtt{P}, \bar{D}) \text{ and } \bar{D} \text{ is an instance for } \mathscr{S}(\mathscr{N}(\mathtt{P}))\}.$$

Proposition 5.2
For a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ and a peer $\mathtt{P} \in \mathscr{P}$, $\mathsf{InLocalCore}(\mathfrak{P}, \mathtt{P})$ is Π_2^P-complete in the size of $\bar{D}$, the input neighborhood instance.

Despite the assumption that the peer graphs $\mathscr{G}(\mathfrak{P})$ are acyclic (cf. Remark 2.1), sets Σ of DECs may contain cycles through inclusion dependencies, in particular with existential quantifiers, i.e. RDECs.

[19] It is also possible to modify the proof just given by introducing auxiliary, dummy joins in the consequent of the RDEC in $\Sigma(\mathtt{P1}, \mathtt{P1})$, to make attributes relevant and avoid so the insertion of tuples with *null*, which has the effect of enforcing deletions.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

More precisely, assume Σ is (or contains) a set of sentences φ of the form $\forall\bar{x}(\psi(\bar{x}) \rightarrow \exists\bar{y}\chi(\bar{x},\bar{y}))$, of the form (4.12). When $\bar{y}$ is not empty, φ is called *existential*. A directed graph can be associated to Σ. The nodes are the database predicates, and there is an edge from predicate P to predicate Q when P appears in the antecedent, and Q in the consequent of a $\varphi \in \Sigma$. The edge is *marked* if φ is existential. Σ is *ref-acyclic* if there no cycles with marked edges in the graph[20].

For example, $IC_1 = \{\forall x(S(x) \rightarrow Q(x)),\ \forall x(Q(x) \rightarrow \exists y\ T(x,y))\}$ and $IC_2 = \{\forall x(S(x) \rightarrow Q(x)), \forall x(Q(x) \rightarrow S(x))\}$ are ref-acyclic sets, whereas $IC_3 = IC_1 \cup \{\forall xy\ (T(x,y) \rightarrow Q(y))\}$ is not. This notion can be applied without changes to the sets $\Sigma(\mathtt{P})$ in PDESs.

Example 5.6
Given a PDES $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust}\rangle$ with $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}\}$, $\mathfrak{S} = \{\mathscr{S}(\mathtt{P1}), \mathscr{S}(\mathtt{P1})\}$, $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P2}) = \{R^2(\cdot,\cdot)\}$, $\Sigma(\mathtt{P1},\mathtt{P2}) = \{\forall x\forall z(R^1(x,z) \rightarrow \exists y R^2(x,y))$, $\forall x\forall z(R^2(x,z) \rightarrow \exists y R^1(x,y)\}$, $\Sigma(\mathtt{P2},\mathtt{P1}) = \emptyset$, and $\mathit{Trust} = \{(\mathtt{P1}, \mathit{less}, \mathtt{P2})\}$. In this case, the peer graph $\mathscr{G}(\mathfrak{P})$ is acyclic. However, schema $\mathfrak{P}$ is not ref-acyclic, because $\Sigma(\mathtt{P1}) = \Sigma(\mathtt{P1},\mathtt{P2})$ has a cycle through RDECs.

In some cases, we will make the assumption that the sets of DECs at hand are *ref-acyclic*. For this reason, we make notice that the proof of Proposition 5.2 uses a ref-acyclic set of DECs. So, the proposition still holds for this class of DECs, which will become relevant in Section 6. (Cf. the Online Appendix A for an additional discussion.)

Corollary 5.1
For a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust}\rangle$ and a peer $\mathtt{P} \in \mathscr{P}$, with ref-acyclic sets of DECs, the $\mathsf{InLocalCore}(\mathfrak{P},\mathtt{P})$ decision problem is Π_2^P-complete in the size of $\bar{D}$, the input neighborhood instance.

We can also apply the null-based solution semantics to Definition 3.3, obtaining a specific definition of PCA. More precisely, given a PDES schema $\mathfrak{P}$, an instance $\mathfrak{D}$ for it, a peer $\mathtt{P}$ and a query $\mathscr{Q}(\bar{x}) \in \mathscr{L}(\mathtt{P})$, the set of PCAs to $\mathscr{Q}$ from $\mathtt{P}$ is

$$PCA^{\mathfrak{D}}_{\mathfrak{P},\mathtt{P}}(\mathscr{Q}) = \{t \quad | \quad D \models_N \mathscr{Q}[t],\ \text{for all } D \in \mathit{Sol}(\mathtt{P},\mathfrak{D})\ \} \tag{5.31}$$

with $\mathit{Sol}(\mathtt{P},\mathfrak{D})$ as defined in this section, on the basis of the $\leqslant_D^{\Sigma(\mathtt{P})}$ relations.

Proposition 5.2 and its Corollary 5.1 still hold when, for each $\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})$, $J_{\mathtt{Q}} = \mathit{Core}(\mathtt{Q},\mathfrak{D})$, and $J_{\mathtt{P}} = D(\mathtt{P})$. Therefore, the local computation of PCAs to a conjunctive query – for which the cores of the neighboring peers are used – is also Π_2^P-complete.

Corollary 5.2
For a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust}\rangle$, an instance $\mathfrak{D}$ for $\mathfrak{P}$, and a peer $\mathtt{P} \in \mathscr{P}$, deciding answers to a conjunctive query posed to $\mathtt{P}$ that are true in $\bigcap NS(\mathtt{P},\bar{D})$, where $\bar{D} = D(\mathtt{P}) \cup \bigcup_{\mathtt{Q}\in\mathscr{N}^{\circ}(\mathtt{P})} \mathit{Core}(\mathtt{Q},\mathfrak{D})$, is Π_2^P-complete in the size of $\bar{D}$. This is also true for $\mathfrak{P}$ with ref-acyclic sets of DECs.

[20] The condition of ref-acyclicity was already used in Bravo and Bertossi (2006), as *RIC-acyclicity*, for sets of ICs on a single database schema.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

The complexity of the decision and computational problems at the neighborhood level we have obtained so far are the most interesting, due to the modular and local manner solutions and PCAs are computed. These results will be useful in Section 6, where specification and computation of neighborhood solutions and PCAs are addressed.

We make only some final remarks in relation to global solutions and PCAs on their basis. Corollary 5.2 already tells us that deciding PCAs (which involves global solutions as opposed to neighborhood solutions) will be at least Π_2^P-hard. Actually, using Proposition 5.2, it is possible to provide a non-deterministic polynomial time algorithm (in data) with a Π_2^P-oracle to decide if an instance for a peer is a (global) solution for the peer. On this basis, one can obtain that *PCA*, as a decision problem, belongs to Π_3^P, in data (that includes those of all peers in the system).

5.5 The import case

In this section, we consider a common situation, namely *the import case*, where we find *Trust* relationships are only of the form $(\mathtt{P}, \mathit{less}, \mathtt{Q})$, when $\mathtt{P} \neq \mathtt{Q}$; and the DECs are used for importing data from other peers.

Definition 5.8

Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, and two different peers P and Q:

(a) An *import UDEC* (IUDEC) *to* P *from* Q is a UDEC in $\mathscr{L}(\mathtt{P}, \mathtt{Q})$ of the form:

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow (Q(\bar{x}') \vee \varphi(\bar{x}'')) \right) \tag{5.32}$$

where the R_i are predicates in $\mathscr{S}(\mathtt{Q})$, Q is a predicate in $\mathscr{S}(\mathtt{P})$, $\bar{x}', \bar{x}'' \subseteq \cup \bar{x}_i = \bar{x}$, and $\varphi(\bar{x}'')$ is a disjunction of built-ins (representing a conjunction of built-ins on the variables in the antecedent).

(b) An *import RDEC* (IRDEC) *to* P *from* Q is an RDEC in $\mathscr{L}(\mathtt{P}, \mathtt{Q})$ of the form:

$$\forall \bar{x} \left(\bigwedge_{i=1}^{n} R_i(\bar{x}_i) \longrightarrow \exists \bar{z} (Q(\bar{y}, \bar{z}) \wedge \varphi_1(\bar{x}', \bar{z}')) \vee \varphi_2(\bar{x}'') \right) \tag{5.33}$$

where the R_i are predicates in $\mathscr{S}(\mathtt{Q})$, Q is a predicate in $\mathscr{S}(\mathtt{P})$, $\bar{x}', \bar{x}'', \bar{y} \subseteq \cup_i \bar{x}_i = \bar{x}$, $\bar{z}' \subseteq \bar{z}$, and φ_1 (φ_2) is a conjunction (disjunction) of built-ins. Notice that φ_1 basically imposes conditions with built-ins on the values for $\bar{z}$ that, under the *null*-based semantics, will be all *null*. ($\varphi_2(\bar{x}'')$ plays the same role on φ in (5.32).)

(c) A PDES is of the *import kind* if, for every peer P, the DECs in $\Sigma(\mathtt{P}, \mathtt{Q})$ with any peer Q different from P, are import DECs to P from Q. Furthermore, if this set of DECs is non-empty, $(\mathtt{P}, \mathit{less}, \mathtt{Q}) \in \mathit{Trust}$.

Notice that this definition does not make any assumptions on the possible sets of local constraints, $\Sigma(\mathtt{P}, \mathtt{P})$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Example 5.7
Consider a PDES with $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}, \mathtt{P4}\}$, $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot)\}, \mathscr{S}(\mathtt{P2}) = \{R^2(\cdot,\cdot), S^2(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P3}) = \{R^3(\cdot,\cdot)\}, \mathscr{S}(\mathtt{P4}) = \{R^4(\cdot,\cdot,\cdot)\}$, and the following sets of DECs:

$$\Sigma(\mathtt{P1},\mathtt{P2}) = \{\forall x \forall y (R^2(x,y) \rightarrow (R^1(x,y) \vee x \leqslant y))\},$$
$$\Sigma(\mathtt{P4},\mathtt{P2}) = \{\forall x \forall y \forall z (R^2(x,y) \wedge S^2(y,z) \rightarrow R^4(x,y,z))\},$$
$$\Sigma(\mathtt{P4},\mathtt{P3}) = \{\forall x \forall y (R^3(x,y) \rightarrow \exists z R^4(x,y,z))\}.$$
$$\Sigma(\mathtt{P3},\mathtt{P2}) = \{\forall x \forall y (R^3(x,y) \rightarrow R^2(x,y))\}.$$

The first three sets of DECs are formed by import DECs, but not the last one.

5.5.1 The unrestricted case

The *unrestricted import case* of PDES $\mathfrak{P}$ as in Definition 5.8 occurs when, for every peer $\mathtt{P} \in \mathscr{P}$, $\Sigma(\mathtt{P},\mathtt{P}) = \emptyset$. That is, in this case, a peer may have DECs of the forms (5.32) and (5.33) to neighboring peers, in whom it has more trust than in itself, but no local ICs.

As a consequence of an unrestricted local repair process, a peer will get data from its neighbors and will integrate them, at least virtually, into its own neighborhood instance. The data from a neighboring peer will be obtained by just posing it a conjunctive query, the one corresponding to the antecedent of the DEC.

For illustration, in Example 5.7, if P4 uses its import DEC from P2 (that in $\Sigma(\mathtt{P4},\mathtt{P2})$) to retrieve data from P2, then P4 sends to P2 the query: $\mathscr{Q}(x,y): R^2(x,y) \wedge S^2(y,z)$.

In the import case, we can define a Datalog program for which its minimal model is the neighborhood solution of the peer.

Definition 5.9
Given an unrestricted import PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, a peer $\mathtt{P} \in \mathscr{P}$ and an instance $\bar{D}$ for the neighborhood schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$ around P, *the import program* $\mathfrak{I}(\mathtt{P}, \bar{D})$ is a Datalog program containing:

(1) The facts: $R(\bar{a})$, for every atom $R(\bar{a}) \in \bar{D}$.
(2) For every IUDEC in $\Sigma(\mathtt{P})$ of the form (5.32), the rule:

$$Q(\bar{y}) \leftarrow R_1(\bar{x}_1), \ldots R_n(\bar{x}_n), \neg\varphi(\bar{x}'').$$

(3) For every import RDEC in $\Sigma(\mathtt{P})$ of the form (5.33), the rule:

$$Q(\bar{y}, \overline{null}) \leftarrow R_1(\bar{x}_1), \ldots R_n(\bar{x}_n), \ \neg\varphi_2(\bar{x}'').$$

Here, $\neg\varphi(\bar{x}'')$ and $\neg\varphi_2(\bar{x}'')$ become conjunctions of built-in literals, i.e. atomic formulas with a built-in or negations thereof; and $\overline{null}$ is a sequence of nulls for variables $\bar{z}$. According to the discussion of (5.33), conditions associated to φ_1 become conditions on $\bar{x}'$, which can all be made part of $\bar{x}''$.

Proposition 5.3
Given an unrestricted import PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ and an instance $\mathfrak{D}$ over it, every peer $\mathtt{P} \in \mathscr{P}$ has a unique solution instance. Furthermore, there is

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

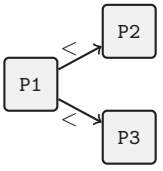


Fig. 7. Accessibility graph for Example 5.8.

an algorithm, that uses the import program $\mathfrak{I}(\mathtt{P}, D')$, and computes the solution for a peer P in polynomial time in the size of $\mathfrak{D} \lfloor \mathscr{S}(\mathscr{AC}(\mathtt{P}))$.

This uniqueness result relies on the fact that we are repairing IRDECs through the insertion of null values, as sanctioned by the official repair semantics. Otherwise, uniqueness would hold in general only for IUDECs.

5.5.2 *The restricted case*

This case appears when, under all the import assumptions above, peers are allowed to have local constraints, that is, it may be that $\Sigma(\mathtt{P}, \mathtt{P}) \neq \emptyset$. In this case, a peer will import data without restrictions from its neighbors, but when building neighborhood solutions, also the local ICs will be taken into account. In this case, P may have none or several solutions.

Example 5.8

Consider the PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$ corresponding to Figure 7. Here, $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}\}$, $\mathfrak{S} = \{\mathscr{S}(\mathtt{P1}), \mathscr{S}(\mathtt{P2}), \mathscr{S}(\mathtt{P3})\}$, $\mathscr{S}(\mathtt{P1}) = \{R^1(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P2}) = \{R^2(\cdot,\cdot)\}$, $\mathscr{S}(\mathtt{P3}) = \{R^3(\cdot,\cdot)\}$.

The local instances are: $D(\mathtt{P1}) = \{\}$, $D(\mathtt{P2}) = \{R^2(a,b)\}$, $D(\mathtt{P3}) = \{R^3(a,c)\}$. The DECs are:

(1) $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall x \forall y (R^2(x,y) \rightarrow R^1(x,y))\}$,
(2) $\Sigma(\mathtt{P1}, \mathtt{P3}) = \{\forall x \forall y (R^3(x,y) \rightarrow R^1(x,y))\}$, and
(3) $\Sigma(\mathtt{P1}, \mathtt{P1}) = \{\forall x \forall y \forall z (R^1(x,y) \wedge R^1(x,z) \rightarrow y = z)\}$.

$\Sigma(\mathtt{P1}, \mathtt{P2})$ and $\Sigma(\mathtt{P1}, \mathtt{P3})$ contain import DECs. However, due to the local constraints in $\Sigma(\mathtt{P1}, \mathtt{P1})$, there are no solutions for P1.

Example 5.9

Consider $D(\mathtt{P}) = \{P(a,b), P(a,c)\}, D(\mathtt{Q}) = \{Q(a,d)\}$, and $\mathit{Trust} = \{(\mathtt{P}, \mathit{less}, \mathtt{Q})\}$. $\Sigma(\mathtt{P}, \mathtt{Q}) = \{\forall x \forall y (Q(x,y) \rightarrow P(x,y)\}$, and $\Sigma(\mathtt{P}, \mathtt{P}) = \{\forall x \forall y \forall z \forall v (P(x,y) \wedge P(x,z) \wedge P(x,v) \rightarrow y = z \vee z = v \vee v = y)\}$. The local constraints in $\Sigma(\mathtt{P}, \mathtt{P})$ ensures that there are at most two tuples in P with the same first attribute. In this case, $\Sigma(\mathtt{P}, \mathtt{Q})$ is of the import kind and therefore $P(a,d)$ will belong to all solutions. The local constraint in $\Sigma(\mathtt{P}, \mathtt{P})$ will force the removal of either $P(a,b)$ or $P(a,c)$, and therefore P has two solutions: $\{P(a,d), P(a,b)\}$ and $\{P(a,d), P(a,c)\}$.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

6 Answer set programs and the solutions for a peer

Continuing with the special PDES semantics based on the use of *null* introduced in Section 4, in this section, we show how to specify neighborhood solutions for a peer (which is the central notion in this paper, upon which the recursive notion of global solution is defined) by means of disjunctive logic programs with stable model semantics (Gelfond and Lifschitz 1991; Eiter *et al.* 1997). As above, we assume that DECs and ICs have the forms (2.1) or (2.2).

Our aim in this section is to show the gist of the approach, by considering in precise terms a particular but still interesting case, one where the RDECs have a limited syntactic form, without joins in the consequents. The general case, i.e. with arbitrary UDECs and RDECs as considered in this work so far, is possible and not difficult from the conceptual or technical point of view, but more difficult or lengthy to present.

Remark 6.1
In this section, we make the additional assumption that given a schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, for each $\mathtt{P} \in \mathscr{P}$, its set of DECs $\Sigma(\mathtt{P})$ is *ref-acyclic*. In this case, we say that $\mathfrak{P}$ is ref-acyclic. The formulation of the programs is independent from this assumption, but – as we will see later in this section – it ensures desirable properties of the program.

Now, we show how to specify the neighborhood solutions for a peer P as the stable models of a disjunctive logic program $\Pi(\mathtt{P})$. *The program can be used for this purpose under the assumption that* P *has already gathered, for each of its neighbors* Q, *the intersection of its solutions*, i.e. $\mathit{Core}(\mathtt{Q}, \mathfrak{D})$ (cf. Definition 5.6). However, the program works with arbitrary instances for the neighbors. Accordingly, the facts of the program come from the union $\mathscr{D}$ of the given instances for the peers in $\mathscr{N}(\mathtt{P})$, the neighborhood centered around peer P. As a consequence, $\Pi(\mathtt{P})$ can be used to compute the restrictions to P'schema of P's neighborhood solutions, on the basis of a neighborhood instance $\mathscr{D}$.

In $\Pi(\mathtt{P})$, we find each predicate $R \in \mathscr{S}(\mathscr{N}(\mathtt{P}))$ and also its copy $R_$ that has an extra argument (an augmented *nickname*). This argument is the last one and is used to place an annotation constant. Thus, if R has arity n, then $R_$ has arity $n + 1$. The possible annotation constants are: $\mathbf{t}, \mathbf{f}, \mathbf{t}^\star, \mathbf{f}^\star, \mathbf{t}^{\star\star}$. Their occurrences in database tuples have the following intended semantics:

Annotation	The tuple $R(\bar{a})$ ...
$R_(\bar{a}, \mathbf{t})$	is made true in (inserted into) the database
$R_(\bar{a}, \mathbf{f})$	is made false in (deleted from) the database
$R_(\bar{a}, \mathbf{t}^\star)$	was true or is made true
$R_(\bar{a}, \mathbf{f}^\star)$	was false or is made false
$R_(\bar{a}, \mathbf{t}^{\star\star})$	was true or made true, and is not deleted

This intended semantics is formally captured in the program by means of appropriate rules. The (possibly virtual) insertions and deletions are made in order to satisfy P's DECs. Actually, for each DEC ψ, a rule captures through its disjunctive head the alternative virtual updates that can be performed to satisfy ψ (cf. rules (2)

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

and (3) in Definition 6.1 for UDECs, and (4) and (5) for RDECs). That is why we use annotations $\mathbf{t}$ or $\mathbf{f}$ in rule heads.

The annotations $\mathbf{t}^\star, \mathbf{f}^\star$ are used to execute sequences of virtual updates, which may be necessary when there are interacting DECs. Finally, atoms annotated with $\mathbf{t}^{\star\star}$ are those that become true in a solution. They are the relevant atoms, and are used to read off the database atoms in the solutions (rules (8) below).

Before presenting the logic program, let us recall that $RelV(\psi)$ denotes the set of relevant variables of a constraint ψ, those where the occurrence of *null* is relevant for its satisfaction (see Definition 4.5). Those attributes/variables receive a special treatment in the program.

In the following, for a DEC ψ, $A(\psi)$, and $C(\psi)$ denote the set of database atoms (without built-ins) in the antecedent of, resp. the consequent, of ψ.

A consequence of repairing using *null*, is that any RDEC with joins between existential quantifiers can only be repaired by deleting tuples and not by inserting tuples with *null*. Thus, the most relevant RDECs in this setting are of the form:

$$\forall \bar{x}(R(\bar{x}) \longrightarrow \exists \bar{y}\ Q(\bar{x}', \bar{y})) \tag{6.34}$$

For this reason, and to simplify the presentation, the logic program that follows considers RDECs of this form only.

Definition 6.1

Consider a PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, Trust \rangle$ and a peer $\mathtt{P} \in \mathscr{P}$ with local instance $D(\mathtt{P})$ (for the schema $\mathscr{S}(\mathtt{P})$). Let $\bar{D}$ be an instance for the schema $\mathscr{S}(\mathscr{N}(\mathtt{P}))$, i.e. for P's neighborhood $\mathscr{N}(\mathtt{P})$, with $\bar{D} \restriction \mathscr{S}(\mathtt{P}) = D(\mathtt{P})$. The *solution program* $\Pi(\mathtt{P}; \bar{D})$ for P contains:

(1) The following facts: $dom(a)$, for every $a \in Adom(\bar{D})$; $dom(null)$, and $R(\bar{a})$, for every atom $R(\bar{a}) \in \bar{D}$.[21]

(2) For every UDEC $\psi \in \Sigma(\mathtt{P}, \mathtt{Q})$ of the form (2.1), with $\mathtt{Q} \in \mathscr{N}(\mathtt{P})$ and $(\mathtt{P}, less, \mathtt{Q}) \in Trust$, the rule:

$$\bigvee\{R__(\bar{x}_i, \mathbf{f}) \mid R(\bar{x}_i) \in A(\psi), R \in \mathscr{S}(\mathtt{P})\} \vee \bigvee\{Q__(\bar{y}_j, \mathbf{t}) \mid Q(\bar{y}_j) \in C(\psi), Q \in \mathscr{S}(\mathtt{P})\} \leftarrow \bigwedge_{i=1}^{n} R_i_(\bar{x}_i, \mathbf{t}^\star),\ \bigwedge_{j=1}^{m} Q_j_(\bar{y}_j, \mathbf{f}^\star),\ \bigwedge_{x_l \in RelV(\psi)} x_l \neq null.$$

Here, when $Q_j \in \mathscr{B}$ (i.e. a built-in), $Q_j_(\bar{y}_j, \mathbf{f}^\star)$ simply denotes the negated atom $Q_j_(\bar{y}_j)$, e.g. $=$ becomes $\neq$, *IsNull* becomes, *IsNotNull*, etc., without annotations. Notice that the condition $Q \in \mathscr{S}(\mathtt{P})$ in the head of the rule discards built-in atoms (cf. Example 6.1 below).

(3) For every UDEC $\psi \in \Sigma(\mathtt{P}, \mathtt{Q})$ of the form (2.1), with $\mathtt{Q} \in \mathscr{N}(\mathtt{P})$ and $(\mathtt{P}, same, \mathtt{Q}) \in Trust$, the rule:

$$\bigvee\{R__(\bar{x}_i, \mathbf{f}) \mid R(\bar{x}_i) \in A(\psi)\} \vee \bigvee\{Q__(\bar{y}_j, \mathbf{t}) \mid Q(\bar{y}_j) \in C(\psi)\} \leftarrow \bigwedge_{i=1}^{n} R_i_(\bar{x}_i, \mathbf{t}^\star), \bigwedge_{j=1}^{m} Q_j_(\bar{y}_j, \mathbf{f}^\star), \bigwedge_{x_l \in RelV(\psi)} x_l \neq null.$$

We apply the same convention as in (2) about the rule body. This case also covers UDECs in $\Sigma(\mathtt{P}, \mathtt{P})$, i.e. local universal ICs for P.

[21] We can also add $IsNull(null)$, and $IsNotNull(c)$, for every non-null constant in the active domain, but this is not necessary if these two predicates do not appear in the program.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

(4) For every RDEC $\psi \in \Sigma(\mathtt{P},\mathtt{Q})$ of the form (6.34), with $\mathtt{Q} \in \mathscr{N}(\mathtt{P})$ and $(\mathtt{P}, \mathit{same}, \mathtt{Q}) \in \mathit{Trust}$, the rule:

$$R_(\bar{x},\mathbf{f}) \vee Q_(\bar{x}',\overline{null},\mathbf{t}) \leftarrow R_(\bar{x},\mathbf{t}^\star),\ \mathit{not}\ \ \mathit{aux}_\psi(\bar{x}'),\ \bar{x}' \neq \mathit{null}.$$

Here, aux_ψ is an auxiliary predicate used to obtain a *safe* rule (i.e. its variables in negated atoms also appear in a positive atom), whose body captures the violations of (6.34). In FO logic, this would be done through: $R(\bar{x}) \wedge \neg\exists\bar{y}Q(\bar{x}',\bar{y})$, which basically leads to the following rules that define the auxiliary predicate (taking into account the special semantics of *null*):
For every $y_i \in \bar{y}$:

$$\mathit{aux}_\psi(\bar{x}') \leftarrow Q_(\bar{x}',\bar{y},\mathbf{t}^\star),\ \mathit{not}\ Q_(\bar{x}',\bar{y},\mathbf{f}),\bar{x}' \neq \mathit{null}, y_i \neq \mathit{null} \tag{6.35}$$

$$\mathit{aux}_\psi(\bar{x}') \leftarrow Q(\bar{x}',\overline{null}),\ \mathit{not}\ Q_(\bar{x}',\overline{null},\mathbf{f}),\bar{x}' \neq \mathit{null} \tag{6.36}$$

This case covers the RDECs in $\Sigma(\mathtt{P},\mathtt{P})$, i.e. local referential ICs for P. (Cf. Example 6.1 below.) Here, and in the rest of this work, atoms of the form $Q_(\bar{x}',\overline{null},\ldots)$ associated to an atom in a DEC, have all the existential variables appearing in the latter and predicate Q replaced by *null*. Furthermore, for $\bar{x} = x_1,\ldots,x_n$, $\bar{x} \neq \mathit{null}$ abbreviates $x_1 \neq \mathit{null},\ldots,x_n \neq \mathit{null}$.

(5) For every RDEC $\psi \in \Sigma(\mathtt{P},\mathtt{Q})$ of the form (6.34), with $\mathtt{Q} \in \mathscr{N}(\mathtt{P})$ and $(\mathtt{P}, \mathit{less}, \mathtt{Q}) \in \mathit{Trust}$, the rules:

(a) If $R \in \mathscr{S}(\mathtt{P})$: $R_(\bar{x},\mathbf{f}) \leftarrow R_(\bar{x},\mathbf{t}^\star),\ \mathit{not}\ \ \mathit{aux}_\psi(\bar{x}'),\ \bar{x}' \neq \mathit{null}.$

(b) If $Q \in \mathscr{S}(\mathtt{P})$: $Q_(\bar{x}',\overline{null},\mathbf{t}) \leftarrow R_(\bar{x},\mathbf{t}^\star),\ \mathit{not}\ \ \mathit{aux}_\psi(\bar{x}'),\ \bar{x}' \neq \mathit{null}.$

Plus rules (6.36), (6.35).

(6) For each predicate $R \in \mathscr{S}(\mathscr{N}(\mathtt{P}))$, the annotation rules:

$R_(\bar{x},\mathbf{f}^\star) \leftarrow \mathit{dom}(\bar{x}),\ \mathit{not}\ R(\bar{x}).$ $\qquad$ $R_(\bar{x},\mathbf{f}^\star) \leftarrow R_(\bar{x},\mathbf{f}).$

$R_(\bar{x},\mathbf{t}^\star) \leftarrow R(\bar{x}).$ $\qquad$ $R_(\bar{x},\mathbf{t}^\star) \leftarrow R_(\bar{x},\mathbf{t}).$

(7) For each predicate $R \in \mathscr{S}(\mathscr{N}(\mathtt{P}))$, the program constraint:
$\leftarrow R_(\bar{x},\mathbf{t}),\ R_(\bar{x},\mathbf{f}).$

(8) For each predicate $R \in \mathscr{S}(\mathtt{P})$, the interpretation rule:
$R_(\bar{x},\mathbf{t}^{\star\star}) \leftarrow R_(\bar{x},\mathbf{t}^\star),\ \mathit{not}\ R_(\bar{x},\mathbf{f}).$

Example 6.1
If the UDEC ψ is $\forall x\forall y(R(x) \rightarrow S(x,y) \vee \mathit{IsNull}(y))$, then the corresponding rule according to (2) above is: $R(x,\mathbf{f}) \leftarrow R(x,\mathbf{t}^\star), S(x,y,\mathbf{f}^\star), \mathit{IsNotNull}(y), x \neq \mathit{null}$.

Notice that $\mathit{RelV}(\psi) = \{x\}$. Since the program interprets *null* as any other constant in the domain, this rule can be replaced by $R(x,\mathbf{f}) \leftarrow R(x,\mathbf{t}^\star), S(x,y,\mathbf{f}^\star), y \neq \mathit{null}, x \neq \mathit{null}$.

If the RDEC ψ is $\forall x(R(x) \rightarrow \exists y S(x,y))$, then the corresponding rules according to (4) above are:

$$R_(x,\mathbf{f}) \vee S_(x,\mathit{null},\mathbf{t}) \leftarrow R(x,\mathbf{t}^\star),\ \mathit{not}\ \ \mathit{aux}(x), x \neq \mathit{null}.$$
$$\mathit{aux}(x) \leftarrow S(x,\mathit{null}),\ \mathit{not}\ S_(x,\mathit{null},\mathbf{f}), x \neq \mathit{null}.$$
$$\mathit{aux}(x) \leftarrow S_(x,y,\mathbf{t}^\star),\ \mathit{not}\ S_(x,y,\mathbf{f}), x \neq \mathit{null}, y \neq \mathit{null}.$$

In this case, the RDEC is "repaired", as usual with referential ICs in databases, by either deleting the tuple from table R or inserting a tuple with a null value in

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

table S, both cases equally acceptable. The two cases for item (5) in the program correspond each to only one of these possible choices.

Since instance $\bar{D}$ and its active domain are finite, the program has a finite number of facts in item (1). The most relevant part of the program corresponds to the rules (2)–(5). They capture through their bodies the violations of DECs; and through their heads, the (possibly alternative) virtual updates that are necessary on the peers' instances to restore the satisfaction of DECs and local ICs. In the bodies of these rules associated to DECs or ICs ψ, the conditions of the form $x \neq null$, with x, a variable appearing in a relevant attribute of ψ, are used to capture the special semantics of null values introduced in Section 4.

The fixed semantics of annotations is captured by the rules (6)–(8). They appear in every solution program. The program constraints in (7) discard models where an atom is both inserted and deleted.

The instances $\bar{D} \upharpoonright \mathscr{S}(\mathtt{Q}))$ used in the program for P's neighbors may not coincide with the initial, physical instances $D(\mathtt{Q})$ in an instance $\mathfrak{D}$ for the PDES $\mathfrak{P}$ (except when Q is P). However, the idea is that P, given a global instance $\mathfrak{D}$, uses its program with $D(\mathtt{P})$ and $\bar{D} \upharpoonright \mathscr{S}(\mathtt{Q})) = Core(\mathtt{Q}, \mathfrak{D})$ for each $\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})$. In this way, the program computes through its stable models the neighborhood solutions for P, whose restrictions to P's schema will be the solutions for P. Notice that atoms annotated with $\mathbf{t}^{\star\star}$ in a stable model of P's program have predicates in $\mathscr{S}(\mathtt{P})$ only; and they define a database instance for P.

The adoption of the stable model semantics for the solution programs (Gelfond and Lifschitz 1991), plus the appropriate use of conditions with *null* in rules' bodies, guarantee the minimal discrepancy between the generated and the original instances according to Definitions 3.1 and 5.3.

Example 6.2

Consider $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}\}$ with $D(\mathtt{P1}) = \{R^1(a, 2)\}$, $D(\mathtt{P2}) = \{R^2(d, 5)\}$, and neighborhood instance $\bar{D} = D(\mathtt{P1}) \cup D(\mathtt{P2})$. Assume $(\mathtt{P1}, less, \mathtt{P2}) \in Trust$, and the only set of DECs in the system is $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xy(R^2(x, y) \rightarrow R^1(x, y))\}$.

The solution program $\Pi(\mathtt{P1}; \bar{D})$ contains (itemized in correspondence to Definition 6.1)

(1) the facts $dom(a), dom(d), dom(2), dom(5), dom(null), R^1(a, 2), R^2(d, 5)$;

(2) $R^1_(x, y, \mathbf{t}) \leftarrow R^2_(x, y, \mathbf{t}^{\star}), R^1_(x, y, \mathbf{f}^{\star}), x \neq null, y \neq null$;

(6) $R^1_(x, y, \mathbf{f}^{\star}) \leftarrow R^1_(x, y, \mathbf{f})$, $\quad R^1_(x, y, \mathbf{f}^{\star}) \leftarrow dom(x), dom(y), \; not \; R^1(x, y)$,
$R^1_(x, y, \mathbf{t}^{\star}) \leftarrow R^1_(x, y, \mathbf{t})$, $\quad R^1_(x, y, \mathbf{t}^{\star}) \leftarrow R^1(x, y)$;

(7) $\leftarrow R^1_(x, y, \mathbf{t}), R^1_(x, y, \mathbf{f})$;

(8) $R^1_(x, y, \mathbf{t}^{\star\star}) \leftarrow R^1_(x, y, \mathbf{t}^{\star}), \; not \; R^1_(x, y, \mathbf{f})$.

The rule in (2) makes sure that a *null*-free R^2-tuple that is not in R^1, is also virtually inserted into R^1 (if it had *null*, the DEC would be satisfied since all attributes are relevant). Since P1 trusts P2 more than itself, virtual changes affect only peer P1. According to Definition 6.1, we should also include rules similar to those in (6) and (7) for R^2. However, they are not necessary because R^2 does not change. If we

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

decide to omit them, we have to replace the atom $R^2_(x, y, \mathbf{t}^\star)$ in the body of rule in (2) by $R^2(x, y)$.

Notice that since P2, which is a sink peer, does not have local ICs, its instance $D(\mathtt{P2})$ is also its only solution, and the one passed over to P1, who uses it to build the facts of its program.

Example 6.3
Consider $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}\}$ with $D(\mathtt{P1}) = \{R^1(s,t),\ R^1(a, null)\}$, $D(\mathtt{P2}) = \{R^2(c,d), R^2(a,e)\}$, and $Trust = \{(\mathtt{P1}, same, \mathtt{P2})\}$.

Assume the only set of DECs are $\Sigma(\mathtt{P1}, \mathtt{P1}) = \{\forall xyz(R^1(x,y) \wedge R^1(x,z) \rightarrow y = z)\}$ and $\Sigma(\mathtt{P1}, \mathtt{P2}) = \{\forall xy(R^2(x,y) \rightarrow \exists z R^1(x,z))\}$. The neighborhood instance is $\bar{D} = D(\mathtt{P1}) \cup D(\mathtt{P2})$ (the second disjunct is also P2's only solution).

The solution program $\Pi(\mathtt{P1}; \bar{D})$ for P1 is as follows: (omitting rules 6 and 7):

(1) $dom(a), \ldots, dom(null), R^1(a, null), R^1(s,t), R^2(c,d),\ R^2(a,e).$

(3) $R^1(x,y,\mathbf{f}) \vee R^1(x,z,\mathbf{f}) \leftarrow R^1(x,y,\mathbf{t}^\star), R^1(x,z,\mathbf{t}^\star),\ x \neq null,\ y \neq null,$
$z \neq null,\ \ y \neq z.$

(4) $R^2_(x,y,\mathbf{f}) \vee R^1_(x, null, \mathbf{t}) \leftarrow R^2_(x,y,\mathbf{t}^\star),\ not\ \ aux(x), x \neq null.$
$aux(x) \leftarrow R^1(x, null),\ not\ R^1_(x, null, \mathbf{f}).$
$aux(x) \leftarrow R^1(x,y,\mathbf{t}^\star),\ not\ R^1_(x,y,\mathbf{f}), x \neq null, y \neq null.$

(8) $R^1_(x,y,\mathbf{t}^{\star\star}) \leftarrow R^1_(x,y,\mathbf{t}^\star),\ not\ R^1_(x,y,\mathbf{f}).$

Rule (3) takes care of P1's the local functional dependency. In case of a violation, one of the two tuples in conflict has to be deleted. Rule (4) has the role of satisfying the RDEC. Since P1 trusts P2 the same as itself, consistency is restored by either deleting the tuple from R^2 or introducing a *null* into R^1.

Here, only $R^1_$ -atoms become annotated with $\mathbf{t}^{\star\star}$. These annotations are used to build solutions for P. In this example, we have two stable models, whose restrictions to $\mathbf{t}^{\star\star}$-annotated atoms are: $\{R^1_(a, null, \mathbf{t}^{\star\star}), R^1_(s,t,\mathbf{t}^{\star\star}), R^1_(c, null, \mathbf{t}^{\star\star})\}$, corresponding to the insertion of tuple $R^1(c, null)$ and $\{R^1_(a, null, \mathbf{t}^{\star\star}),\ R^1_(s,t,\ \mathbf{t}^{\star\star})\}$, corresponding to the deletion of tuple $R^2(c,d)$. They correspond to the two solutions for P: $\{R^1(a, null), R^1(s,t), R^1(c, null)\}$ and $\{R^1(a, null), R^1(s,t)\}$.

Notice that each peer P in a PDES has a solution program that, except for the facts in it, is fixed. It depends only on its DECS in $\Sigma(\mathtt{P})$ and the trust relationships to its neighbors. The same program can be used with different initial neighborhood instances $\bar{D}$. On their basis, the program will compute stable models, whose double-starred atoms will determine a local instance.

Definition 6.2
Consider a PDES $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, Trust \rangle$, a peer $\mathtt{P} \in \mathscr{P}$, and a neighborhood instance $\bar{D}$ for $\mathscr{S}(\mathscr{N}(\mathtt{P}))$. Let M be a stable model of $\Pi(\mathtt{P}; \bar{D})$. The database instance for peer P (and schema $\mathscr{S}(\mathtt{P})$) associated to M is $D_M = \{R(\bar{a}) \mid R \in \mathscr{S}(\mathtt{P})$ and $R_(\bar{a}, \mathbf{t}^{\star\star}) \in M\}$.

The following result tells us that the instances for a peer obtained via the stable models of its program are all and only the solutions for the peer, under the assumption that its neighbors contribute with their own cores.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Proposition 6.1
Consider a PDES $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, Trust \rangle$, with an instance $\mathfrak{D}$ for $\mathfrak{P}$, and $\mathtt{P} \in \mathscr{P}$ whose set of DECs $\Sigma(\mathtt{P})$ is ref-acyclic. Given the neighborhood instance for $\mathscr{S}(\mathscr{N}(\mathtt{P}))$: $\bar{D} := D(\mathtt{P}) \cup \bigcup_{\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})} Core(\mathtt{Q}, \mathfrak{D})$, it holds:

$$Sol(\mathtt{P}, \mathfrak{D}) = \{D_M \mid M \text{ is a stable model of program } \Pi(\mathtt{P}, \bar{D})\}.$$

This result generalizes one about the correctness of similar programs for null-based repairs of the kind considered in this work, for single databases with respect to denial constraints and referential constraints and null values (Bravo and Bertossi 2006). The hypothesis of ref-acyclicity is necessary for the soundness of the program (every instance D_M is a solution). Otherwise, only the completeness of the program can be guaranteed (every solution is a D_M) (Bravo and Bertossi 2006).

Under the assumption that we have already computed the (intersection of the) solution instances for P's neighbors, the program for P allows us to compute its solution instances. This generates a recursive process that can be applied because $\mathscr{G}(\mathtt{P})$ is acyclic. Terminal peers P' in $\mathscr{G}(\mathtt{P})$, i.e. without outgoing edges, become the base cases for the recursion. If their local instances $D(\mathtt{P}')$ are (locally) consistent, they pass those instances to their neighbors. Otherwise, they first, using program $\Pi(\mathtt{P}'; D(\mathtt{P}'))$ compute local repairs for $D(\mathtt{P}')$, and pass the intersection of them to their neighbors.

Several optimizations can be applied to solution programs (Caniupan and Bertossi 2010). An important one has to do with the materialization of the *closed-world-assumption*, which results through the rule $R^1_(x, y, \mathbf{f}^\star) \leftarrow dom(x), dom(y)$, *not* $R^1(x, y)$ in (6) in Example 6.2. This is clearly undesirable and can be avoided. We do not provide optimized versions here, because they are more difficult to read.

Proposition 6.1 still holds if P, instead of collecting the intersection of the solutions of a neighbor Q, uses the intersection of the solutions for Q restricted to the subschema of Q that contains Q's relations that appear in $\Sigma(\mathtt{P}, \mathtt{Q})$, which are those P needs to run its program.

As we expressed at the beginning of this section, our solution programs can be much more general. In particular, they can be modified to include rules for REDCs with existential quantifiers and joins in the consequents. For example, if the DEC is, say $\forall xy(R(x, y) \wedge S(y, z) \rightarrow \exists w(P(x, w) \wedge Q(w, z)))$, and the antecedent is true with *null*-free tuples[22], but not the consequent, one of the (satisfying) tuples in the antecedent has to be deleted. All we need to handle this DEC is a disjunctive, "deleting" rule. Built-in comparisons (with *null*) are used by the rule as well.

Remark 6.2
For simplicity, the programs we introduced above do not consider inconsistent peers, i.e. a peer P whose instance is $D^{\mathtt{P}} = \{\mathbf{inc}_{\mathtt{P}}\}$. As the following example shows, it is easy to modify the rules to capture this situation. This change is compatible with the corresponding general DECs (cf. Remark 3.1).

[22] If it is true via tuples with nulls, then due to the relevance of attributes, the DEC is immediately satisfied.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Example 6.4 (*example 6.2 continued*)
The system is exactly as before, except that now $D^{\mathtt{P2}} = \bigcap Sol(\mathtt{P2}, \mathfrak{D}) = \{\mathbf{inc}_{\mathtt{P2}}\}$, reflecting the fact that P2 has no solutions (due to its mappings and trust relationships to other peers, which we are not showing here). The general program that allows for inconsistent peers would now have, instead of 1 and 2:

(1′) $dom(a), dom(2), R^1(a, 2), \mathbf{inc}_{\mathtt{P2}}$.

(2′) $R^1_(x, y, \mathbf{t}) \leftarrow R^2_(x, y, \mathbf{t}^\star), R^1_(x, y, \mathbf{f}^\star), \mathit{not} \quad \mathbf{inc}_{\mathtt{P2}}, \ x \neq \mathit{null}, y \neq \mathit{null}$. In this case, and as expected, the rule has no effect on P1.

6.1 ASPs and PCAs

With a solution program for P, PCAs to a query $\mathcal{Q}$ posed to P can be obtained by running a query program in combination with the solution program. First, a query program $\Pi(\mathcal{Q})$ has to be produced, which is rather standard, and next, $\Pi(\mathtt{P}) \cup \Pi(\mathcal{Q})$ is run under the *cautious stable model semantics*, the one that declares as true what is simultaneously true in all the stable models. Of course, the same program $\Pi(\mathtt{P})$ can be used with different queries.

Example 6.5 (*example 6.2 continued*)
In order to obtain P1's PCAs to the query $\mathcal{Q}_1(x, y) : R^1(x, y)$, the rule $Ans_1(x, y) \leftarrow R^1(x, y, \mathbf{t}^{\star\star})$ has to be added to $\Pi(\mathtt{P}; D(\mathtt{P1}) \cup D(\mathtt{P2}))$. The PCAs are the ground Ans_1-atoms in the intersection of all stable models of the combined program. For the query $\mathcal{Q}_2(x) : \exists y R^1(x, y)$, the query rule is $Ans_2(x) \leftarrow R^1(x, y, \mathbf{t}^{\star\star})$.

In general, given a conjunctive query $\mathcal{Q}$ (for which we want PCAs), first, its rewriting $\mathcal{Q}^N$ is produced (as in Definition 4.6). Next, $\mathcal{Q}^N$ is rewritten in its turn as a query program with annotation $\mathbf{t}^{\star\star}$. The query program is added to the peer's solution program, and the combination is run under the skeptical semantics.

Deciding the truth of ground atomic queries by means of the solution program for a peer under the skeptical semantics amounts to deciding membership of the local core, which, by Corollary 5.1, is Π_2^P-complete in data. This is exactly the data complexity of deciding skeptical entailment for atomic queries for disjunctive logic programs under the stable model semantics (Dantsin *et al.* 2001). So, the solution programs have the right expressive power for the problem at hand.

Using solution programs, our semantics could be naively implemented as follows. When P is posed a query, P has to run its program, for which it needs as facts those in the intersections of the solutions of its neighbors. So, P sends to each neighbor Q queries of the form $\mathcal{Q} : R(\bar{x})$, where $R \in \mathscr{S}(\mathtt{Q})$ and appears in $\Sigma(\mathtt{P}, \mathtt{Q})$. Peer P expects to receive from Q the PCAs to $\mathcal{Q}$, because they corresponds to the extension of R in the intersection of solutions for Q.

In order to return to P the PCAs to its queries, each of its neighbors Q has to run its own program $\Pi(\mathtt{Q})$. As before, they need PCAs from their own neighbors, etc. This recursion eventually reaches peers that have no DECs, and assuming its local ICs are satisfied, they return query answers to its neighbors directly from their

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

original instances. This is the start of the backward propagation of PCAs process, which goes on until reaching P. Eventually, P gets all the facts to run its program and obtain the PCAs to the original query. If the local instance of a terminal peer is not consistent with respect to its local ICs, the local solution program can still be used to restore consistency with respect to the local ICs, as in CQA. The data to be sent back to the neighbors comes from the intersection of the repairs, in the form of consistent query answers (Arenas *et al.* 1999).

Example 6.6 (*example 2.1 continued*)
Consider an instance $\mathfrak{D} = \{D(\mathtt{P1}), D(\mathtt{P2}), D(\mathtt{P3})\}$ with $D(\mathtt{P1}) = \{R^1(a,2)\}$, $D(\mathtt{P2}) = \{R^2(c,4), R^2(d,5)\}$, and $D(\mathtt{P3}) = \{R^3(c,4)\}$. A user poses the query $\mathcal{Q}_0 : R^1(x,y)$ to P1, expecting its PCAs.

To run its program, P1 needs the intersection of the solutions of peer P2. So, P1 sends to P2 the queries $\mathcal{Q}_1^1 : R^2(x,y)$ and $\mathcal{Q}_2^1 : S^2(x,y)$ (actually, P1 does not need the latter, S^2 is not relevant to P1).

In order to peer-consistently answer these queries, P2 needs from P3 the PCAs to $\mathcal{Q}_3^2 : R^3(x,y)$. Since P3 has no neighbors, it returns to P2 the entire extension in its local database, i.e. $D(\mathtt{P3}) = \{R^3(c,4)\}$. Now, P2 runs its program $\Pi(\mathtt{P2}; D(\mathtt{P2}) \cup D(\mathtt{3}))$. It contains, among others, the facts $R^2(c,4), R^2(d,5), R^3(c,4)$; and (assuming $\Sigma(\mathtt{P2},\mathtt{P2}) = \emptyset$) the main rule:

$$R^2_(x,y,\mathbf{f}) \vee R^3_(x,y,\mathbf{f}) \leftarrow R^2_(x,y,\mathbf{t}^\star), R^3_(x,y,\mathbf{t}^\star), x \neq null, y \neq null.$$

P2 has two solutions, $\{R^2(d,5)\}$ and $\{R^2(c,4), R^2(d,5)\}$, whose intersection, $Core(\mathtt{P2}, \mathfrak{D}) = \{R^2(d,5)\}$ is given to P1.

Finally, the program $\Pi(\mathtt{P1}; D(\mathtt{P1}) \cup Core(\mathtt{P2}, \mathfrak{D})$ (cf. Example 6.2) is run. It has only one solution, namely $\{R^1(a,2), R^1(d,5)\}\}$; and the PCAs to $\mathcal{Q}_0$ are $\langle a,2\rangle$ and $\langle d,5\rangle$.

6.2 The import case revisited

We first consider the unrestricted case introduced in Section 5.5.1. Here, each peer P has a solution program $\Pi^-(\mathtt{P})$ as in Definition 6.1, with rules of the form

(2) $\quad S_(\bar{y}_j, \mathbf{t}) \leftarrow \bigwedge_{i=1}^n R_i_(\bar{x}_i, \mathbf{t}^\star),\ S_(\bar{y}_j, \mathbf{f}^\star),\ \bigwedge_{x_l \in RelV(\psi)} x_l \neq null,\ \bar{\varphi},$

where $S \in \mathscr{S}(\mathtt{P})$ and $R_i \in \mathscr{S}(\mathtt{Q})$, for $\mathtt{Q} \in \mathscr{N}^\circ(\mathtt{P})$;

(5) (b) $\quad Q_(\bar{x}', \overline{null}, \mathbf{t}) \leftarrow R_(\bar{x}, \mathbf{t}^\star),\ not\ \ aux_\psi(\bar{x}'),\ \bar{x}' \neq null,$

where $Q \in \mathscr{S}(\mathtt{P})$ and $R \in \mathscr{S}(\mathtt{Q})$, for $\mathtt{Q} \in \mathscr{N}^\circ(\mathtt{P})$.

We also need the auxiliary rules (6.36), (6.35), and only for predicates in $\mathscr{S}(\mathtt{P})$ that appear in consequents of RDECs.

Now, for a predicate $Q \in \mathscr{S}(\mathtt{P})$, a solution program like this never generates a tuple of the form $Q(\bar{a}, \mathbf{t})$. In consequence, the negative literals can be eliminated from the bodies of both (6.36) and (6.35). The resulting program still contains negation, in rules (6) and (8). However, the solution program becomes a (non-disjunctive) stratified normal program (Abiteboul *et al.* 1995). In consequence, for every set of facts, it has only one stable model that can be computed in polynomial time in data. It corresponds to the only solution for peer P mentioned in Section 6.1.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

If the query $\mathcal{Q}$ (or, better, its rewriting $\mathcal{Q}^N$) posed to P can also be represented as a stratified normal program $\Pi(\mathcal{Q})$ (which is the case when $\mathcal{Q}$ is first-order, for example), then the combined program $\Pi(\mathtt{P}) \cup \Pi(\mathcal{Q})$ is also normal and stratified. In consequence, computing PCAs can also be done in polynomial time. This is under the assumption that the peer has collected, for each of its neighbors, the intersection of its (local) solutions. We have obtained the following result.

Proposition 6.2
For an unrestricted import PDES schema $\mathfrak{P} = \langle \mathscr{P}, \mathfrak{S}, \Sigma, \mathit{Trust} \rangle$, an instance $\mathfrak{D}$ for $\mathfrak{P}$, and a peer $\mathtt{P} \in \mathscr{P}$, deciding answers to a conjunctive query posed to P that are true in $\bigcap NS(\mathtt{P}, \bar{D})$, where $\bar{D} = D(\mathtt{P}) \cup \bigcup_{\mathtt{Q} \in \mathscr{N}^{\circ}(\mathtt{P})} \mathit{Core}(\mathtt{Q}, \mathfrak{D})$, can be done in polynomial time in the size of $\bar{D}$.

In the restricted import case of Section 5.5.2, we allowed peers to have local ICs. In this case, as illustrated in Examples 5.8 and 5.9, the situation may change: Due to the local ICs, the solution program may be disjunctive, and have none or several stable models, which is also reflected in the peers' solutions.

Example 6.7 (*example 5.8 continued*)
In this case, we have no solution. This is captured by the corresponding program through the program constraint $\leftarrow R^1_(x, y, \mathbf{t}), R^1_(x, y, \mathbf{f})$. In each of the two models of the program without the constraint, there will be the atoms $R^1(a, b, \mathbf{t}), R^1(a, c, \mathbf{t})$, $R^1(a, b, \mathbf{f})$ or the atoms $R^1(a, b, \mathbf{t}), R^1(a, c, \mathbf{t}), R^1(a, c, \mathbf{f})$. Both models will be discarded for violating the constraint.

Example 6.8
(example 5.9 continued) The solution program for P has as main rules, the following:

$$P_(x, y, \mathbf{t}) \leftarrow Q_(x, y, \mathbf{t}^\star), P_(x, y, \mathbf{f}^\star), x \neq \mathit{null}, y \neq \mathit{null}.$$

$$P_(x, w_1, \mathbf{f}) \vee P_(x, w_2, \mathbf{f}) \vee P_(x, w_3, \mathbf{f}) \leftarrow P_(x, w_1, \mathbf{t}^\star), P_(x, w_2, \mathbf{t}^\star), P_(x, w_3, \mathbf{t}^\star), \bigwedge_{i \neq j} w_i \neq w_j, x \neq \mathit{null}, w_1 \neq \mathit{null}, w_2 \neq \mathit{null}, w_3 \neq \mathit{null}.$$

We have a disjunctive solution program; actually one with two stable models, as expected according to Example 5.9.

In Barcelo *et al.* (2003) and Bravo and Bertossi (2006), in the context of CQA, several classes of ICs have been identified for which the logic program becomes *head-cycle free*, in which case, it can be replaced by and equivalent non-disjunctive, i.e. normal, program (Dantsin *et al.* 2001). Cautious reasoning from normal logic programs is *coNP*-complete (Dantsin *et al.* 2001). In consequence, for those classes of ICs, peer CQA is in *coNP* in data complexity[23]. Among those constraints, we find denial constraints, such as constraint (2.5) in Example 2.1. Even in the restricted import case with denial constraints as local constraints, peer CQA is *coNP*-complete. (This can be obtained from CQA under denial constraints (Barcelo *et al.* 2003). (Cf. Bertossi (2011) for a survey of complexity results and references.)

[23] Again, under the assumption that the peer has collected, for each of its neighbors, the intersection of its (local) solutions.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

7 Related work

Our work can be placed among those on *semantic peer data management systems*, a research direction that was started, to the best of our knowledge, with Halevy *et al.* (2003) (cf. also Halevy *et al.* (2004)). Mappings relate two conjunctive queries that are expressed in terms of the schemas of two disjoint sets of peers. Already in those papers problematic cases of cyclic dependencies, which implicitly involve a cyclic accessibility graph, were identified. For example, we may have $\mathscr{P} = \{\mathtt{P1}, \mathtt{P2}, \mathtt{P3}\}$ with relations $R^1(\cdot), R^2(\cdot), R^3(\cdot)$, respectively, and DECs $\Sigma(\mathtt{P1}) = \{\forall x(R^2(x) \rightarrow R^1(x))\}$, $\Sigma(\mathtt{P2}) = \{\forall x(R^3(x) \rightarrow R^2(x))\}$, $\Sigma(\mathtt{P3}) = \{\forall x(R^1(x) \rightarrow R^3(x))\}$, each of them satisfied only by importing data into the peer who owns the DEC, i.e. there is the implicit trust relation $\{(\mathtt{P1}, \mathit{less}, \mathtt{P2}), (\mathtt{P2}, \mathit{less}, \mathtt{P3}), (\mathtt{P3}, \mathit{less}, \mathtt{P1})\}$. This is an unrestricted and cyclic import case. In the presence of cycles like this, peer CQA becomes undecidable.

In Halevy *et al.* (2003), mappings are represented and given a semantics using classical first-order logic. This choice requires consistency of peers (i.e. their data, local ICs, and local mappings) with respect to the system as a whole. More precisely, certain answers from a peer are defined as those true in every global instance that is consistent with the local data and metadata, which is rigid. Some criticisms for the use of first-order logic as a basis for the semantics of PDESs were raised in Calvanese *et al.* (2004), motivating a new research direction (cf. below).

For comparison, our approach is inconsistency tolerant, and does not provide a first-order semantics to the global system. Instead, we appeal to a non-monotonic semantics for only neighborhoods of peers. Our mappings are between two peers only, but in that case, more expressive than those considered in Halevy *et al.* (2003).

It is worth noticing that all approaches adopting a first-order logic interpretation for PDESs, e.g. Halevy *et al.* (2003, 2004), do not consider peer mappings as constraints, but as logical implications from a peer to another. This in fact coincides with having trust relationships of the form $(\mathtt{P}, \mathit{less}, \mathtt{Q})$ for any mapping specified from P to Q.

Among the approaches to PDESs that are the closest to ours, we should mention Calvanese *et al.* (2004), Franconi *et al.* (2004b), Calvanese *et al.* (2005, 2008). The DECs are of the "import kind", and there are no explicit trust relationships. For example, in Calvanese *et al.* (2008), DECs are of the form $CQ_i \rightarrow CQ_j$, where CQ_i, CQ_j are conjunctive queries over Pi and Pj's schemas, respectively. This kind of DECs keep the schemas separate, on different sides of the implication. The latter has to be interpreted appropriately. Actually, the semantics of the system, in particular of the DECs, is given by an *epistemic logic*. Local ICs violations are avoided by ignoring a peer that is inconsistent with respect to its local ICs. New atoms are added into a peer by interaction with other peers only if this does not produce a local IC violation. In consequence, a peer trust other peers more than itself, as long as no local inconsistencies are produced. There is not consistency restoration process involved.

As shown in Figure 8, in Calvanese *et al.* (2008), each peer is considered as a local virtual data integration system that follows the *global-as-view* approach (Lenzerini 2002). In this example, the three peers are locally defined by the GAV

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

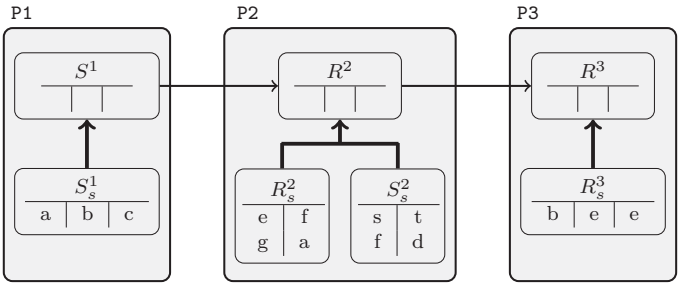


Fig. 8. Integration systems as peers.

mappings $\forall x\forall y\forall z(S^1_s(x,y,z) \rightarrow S^1(x,y,z))$, $\forall x\forall y\forall z(R^2_s(x,y) \wedge S^2_s(y,z) \rightarrow R^2(x,y,z))$ and $\forall x\forall y\forall z(R^3_s(x,y,z) \rightarrow R^3(x,y,z))$, respectively. The predicates of the form P_s correspond to local sources, and those of the form P^i correspond to the mediated schema provided by peer `Pi`. In addition, the DECs establish mappings between the peers. In this case, $\forall x\forall y(R^2(x,y,z) \rightarrow \exists w S^1(x,y,w))$, $\forall x\forall y(R^3(x,y,y) \rightarrow \exists u\exists v R^2(u,x,v))$. These DECs are used to pass data from the peer mentioned in the antecedent to the one mentioned in the consequent. This activity is based on a theory written in epistemic logic which is generated from the local and inter-peer mappings. More specifically, the following epistemic theory is used for query answering:

$$\left.\begin{array}{l}\mathbf{K_1}(\forall x\forall y\forall z(S^1_s(x,y,z) \rightarrow S^1(x,y,z))) \\ \forall x\forall y(\mathbf{K_2}(R^2(x,y,z)) \rightarrow \mathbf{K_1}(\exists w S^1(x,y,w)))\end{array}\right\} \text{Specification of } \texttt{P1}$$

$$\left.\begin{array}{l}\mathbf{K_2}(\forall x\forall y\forall z(R^2_s(x,y) \wedge S^2_s(y,z) \rightarrow R^2(x,y,z))) \\ \forall x\forall y(\mathbf{K_3}(R^3(x,y,y)) \rightarrow \mathbf{K_2}(\exists u\exists v R^2(u,x,v)))\}\end{array}\right\} \text{Specification of } \texttt{P2}$$

$$\left.\mathbf{K_3}(\forall x\forall y\forall z(R^3_s(x,y,z) \rightarrow R^3(x,y,z)))\right\} \text{Specification of } \texttt{P3}$$

$\mathbf{K_i}\varphi$ is interpreted as φ is known by peer `Pi`. The idea behind using the epistemic theory is that data that is known (or certain) is passed from local sources to mediated schemas and from peers to other peers. A tuple $\bar{t}$ is a PCA to a query $\mathcal{Q}$ posed to peer `Pi`, if $\mathbf{K_i}\mathcal{Q}(\bar{t})$ is a logical consequence of the epistemic theory.

An advantage of this approach is that the semantics can be applied in the presence of cycles. However, possibly the whole epistemic theory has to be used by a peer `Pi` to do query answering, which requires not only data, but also the mappings and DECs; and this not only of its neighbors, but of all accessible peers. (This semantics is similar in spirit to the one called *"send all"* in the discussion in the Online Appendix A.)

Our approach can be easily and uniformly adapted in order to make each peer a local data integration system. For this, the answer set programs that specify the legal instances of a virtual data integration introduced in Bravo and Bertossi (2003, 2005), Bertossi and Bravo (2004a) can be used in combination with those presented here.

The epistemic theory has also been extended in Calvanese *et al.* (2008) in order to make the PDES "inconsistency tolerant". This is done by using additional modal operators and extending the epistemic theory, achieving that: (a) A peer whose local data is inconsistent with respect to its ICs, is not considered for data exchange, that

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

is, the PDES behaves as if that peer didn't exist. In our case, however, to such a peer, whenever possible, a repair semantics is applied. (b) Inconsistencies due to data imported from other peers, referred as P2P inconsistencies, are solved by removing a minimal amount of data imported from other peers in order to preserve consistency. In other words, atoms are imported into a peer by interaction with other peers only if this does not produce a violation of an IC of a peer. In our case, under the *same* or *less* trust conditions, the data is accepted and the peer applies again a local repair semantics.

In order to answer a query (Calvanese *et al.* 2005), a peer traverses the network eventually collecting at its site all DECs, ICs, and data of other logically related peers. With these elements, the peer can construct its epistemic theory, that is used for query answering. An accessibility cycle can be detected by using request identifiers. The use of epistemic logic makes sure that *certain data*, the one a peer really knows, is passed to another peer. In our case, a peer collects only data from its neighbors; and certainty is achieved by using the PCAs of a peer, or more generally, the intersection of its solutions. A more detailed comparison can be found in Bravo (2007).

The semantics in Franconi *et al.* (2004a,b) coincides with the epistemic semantics in Calvanese *et al.* (2004). The former also provide a distributed algorithm, where peers' data is updated by instruction of a super peer. When a query is posed to a peer, it can answer the query right away with its data, because the PDES is already updated.

The data exchange problem among distributed independent sources has been investigated in Caroprese *et al.* (2006), Caroprese and Zumpano (2008, 2011). In Caroprese *et al.* (2006), the authors define a declarative semantics for P2P systems that allows one to import, into each peer, maximal subsets of atoms that do not violate the local ICs. The framework has been extended in Caroprese and Zumpano (2008) with a mechanism to set different degrees of reliability for neighbor peers, and in Caroprese and Zumpano (2011), where "dynamic" preferences can be used to import data in the case of conflicting sets of atoms, depending on the properties of the peers' data. This extended framework allows one to model preferences like "in the case of conflicting information, it is preferable to import data from the neighbor peer that can provide the maximum number of tuples" or "in the case of conflicting information, it is preferable to import data from the neighbor peer such that the sum of the values of an attribute is minimum", without having to select preferred peers *a-priori*. To enforce this preference mechanisms, the P2P framework has been enriched with aggregate functions.

In Fuxman *et al.* (2006), the case of two peers that belong to a PDES is analyzed. From a source peer, `P`, to a target peer, `Q`, there may be both source-to-target dependencies, as in data exchange (Kolaitis 2005), and also target-to-source dependencies. The former are used to transport data from `P` to `Q`, and the second, to filter the received data in `Q`. In addition, the target may have an instance and local constraints. The existence of solutions for the target peer is investigated. The semantics we have introduced, but using labeled nulls instead of *null*, could be adapted to a situation like this (cf. the discussion in the Online Appendix A). It is

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

also possible to use the solution programs in the case of Fuxman *et al.* (2006). In this scenario, the source-to-target dependencies would give rise to program rules that do the data exchange, and both target-to-source dependencies and target dependencies would become hard program constraints.

More material and references on PDES can be found in Aberer (2011). Consistency issues that appear in our PDES scenario may also appear in the context of ontologies, when they have to be aligned in the presence of *bridges* between them. When ontologies are merged, it becomes necessary to solve inconsistencies between them (Serafini *et al.* 2005). However, in relation to mappings, we go beyond the cases considered in that area. The mappings, or bridges, between ontologies are in general much simpler than the general case considered in our work.

In this work, we have concentrated on null values that are handled as in SQL relational databases. Actually, in that case, there is a single null, namely the constant `NULL`. It can be interpreted in different ways in a database, e.g. as a representative for an existing but unknown value, as a missing or non-existent value, as a withheld or hidden value, etc. In our case, we do not make any ontological assumption about this constant. Our purpose consists in logically capturing the way it is handled in data management operations by an SQL database. More specifically, with respect to query answering, integrity checking, and interaction with other data elements.

Most of the work in database theory bypasses the issue of SQL null, mainly because it is considered to be an undesirable oddity that could be replaced by a logically cleaner solution. However, SQL `NULL` exists, is used, and will most like stay with us. It would be better to accept it, but trying to clarify its logical and operational status. This is our approach in this work. From this point of view, we accept that SQL `NULL` appears, not only in databases, but could also appear explicitly in DECs and ICs, and in their semantics. For example, a referential constraint or DEC could be of the form $\forall xy(S(x, y) \rightarrow T(x, \texttt{NULL}))$, i.e. requiring that T's second attribute should be filled with `NULL`[24].

A similar approach is adopted in Franconi and Tessaris (2012), where an alternative reconstruction of SQL `NULL` is attempted. They propose a corresponding relational algebra that takes care of the special status of `NULL`, but in the end, after the reconstruction, this constant disappears from the data domain. We keep it, but rewrite queries and constraints in such a way that it can be treated as any other constant.

There is a large literature on nulls in relational databases (see, e.g. Levene and Loizou (1999)). However, to the best of our knowledge those nulls are "ideal" nulls, and not the real SQL `NULL`. There are also some logical approaches, starting with Reiter's treatment of nulls (Reiter 1984), which can be multiple, existential values not subject to the *unique names assumption*. They are also different from the SQL `NULL`. More recent work extending Reiter's approach can be found in Traylor and Gelfond (1994), Lifschitz *et al.* (2012). See also Libkin (2014, 2016) for a recent work on incomplete information in relational databases.

[24] In our case, given our repair semantics of Section 5, even a constraint of the form $\forall xy(S(x, y) \rightarrow \exists z\, T(x, z))$ will be enforced by giving to z the value `NULL` (unless some other constraint prevents this).

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

8 Conclusions

We have introduced a general framework for peer data exchange with trust relationships. The methodology is flexible and *inconsistency tolerant* in that each peer solves its data and semantic conflicts at query time, when querying its own and other peers' data.

The general semantic framework can be specialized in several ways, and we presented some possibilities. In particular, we developed a specific semantics based on UDECs and RDECs, and on the use on SQL null values to deal with incomplete information. In this scenario, logic programs can be used to specify the solution instances for a peer and to obtain PCAs.

Techniques for optimizing the logic programs and their execution could be applied for query answering. Among them, we find the partial evaluation of programs and the solution instances they specify, since we are not interested in the solution *per se*, but in the PCAs. More specifically, techniques used in CQA, such as *magic sets* for stable model semantics (Faber *et al.* 2007), and identification of predicates that are relevant to queries and constraints, could also be used in this setting. In this way, the number of rules and the amount of data that are needed to run the program are reduced (Eiter *et al.* 2008; Caniupan and Bertossi 2010).

The problem of query evaluation from disjunctive programs is Π_2^P-complete in data (Dantsin *et al.* 2001), which matches the complexity of PCA, as we established here. In spite of this, we have also identified syntactic classes of PDESs for which peer CQA has a lower complexity. For them, specifically tailored mechanisms to solve this problem could be developed, as for CQA.

We should emphasize that the DECs we can handle are more general that those found in related work on peer data exchange.

Our semantics allows for inconsistent peers and inconsistencies between peers, without unraveling logical reasoning or having to exclude peers whose data participate in inconsistencies. Actually, our semantics for PDES smoothly extend the repair semantics for CQA from inconsistent databases.

Along the way, and interesting on its own merits, we developed a semantics of conjunctive query answering and constraint satisfaction in terms of classical FO predicate logic. It puts on a solid ground the mechanisms implemented in SQL databases and specifications in the SQL Standard for handling null values.

Acknowledgements

Research supported by NSERC DG (250279-2011) and a CITO/IBM-CAS Student Internship. Loreto Bravo is funded by CONICYT grant PSD-57. We appreciate valuable and detailed comments from the anonymous reviewers.

Supplementary materials

For supplementary material for this article, please visit http://dx.doi.org/10.1017/S147106841600017X.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

References

Aberer, K. 2011. *Peer-to-Peer Data Management*. Morgan & Claypool Publishers, San Rafael, CA.

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

Arenas, M., Bertossi, L. and Chomicki, J. 1999. Consistent query answers in inconsistent databases. In *Proc. ACM Symposium on Principles of Database Systems*, V. Vianu and C. Papadimitriou, Eds. ACM, New York, USA, 68–79.

Arenas, M., Bertossi, L. and Chomicki, J. 2003. Answer sets for consistent query answers. *Theory and Practice of Logic Programming 3* (4–5), 393–424.

Arieli, O., Denecker, M. and Bruynooghe, M. 2007. Distance semantics for database repair. *Annals of Mathematics and Artificial Intelligence 50* (3–4), 389–415.

Artz, D. and Gil, Y. 2007. A survey of trust in computer science and the semantic web. *Journal of Web Semantics 5* (2), 58–71.

Barcelo, P. and Bertossi, L. 2003. Logic programs for querying inconsistent databases. In *Proc. Practical Aspects of Declarative Languages*, V. Dahl and P. Wadler, Eds. LNCS, vol. 2562. Springer-Verlag, Berlin-Heidelberg-New York, 208–222.

Barcelo, P., Bertossi, L. and Bravo, L. 2003. Characterizing and computing semantically correct answers from databases with annotated logic and answer Sets. In '*Semantics in Databases', Springer-Verlag, LNCS 2582*, L. Bertossi, G O. H. Katona, K.-D. Schewe and B. Thalheim, Eds. Springer-Verlag, Berlin-Heidelberg-New York, 7–33.

Bertossi, L. 2006. Consistent query answering in databases. *ACM Sigmod Record 35* (2), 68–76.

Bertossi, L. 2011. *Database Repairing and Consistent Query Answering*. Morgan & Claypool Publishers, San Rafael, CA.

Bertossi, L. and Bravo, L. 2004a. Consistent query answers in virtual data integration systems. In *Inconsistency Tolerance*, L. Bertossi, A. Hunter and T. Schaub, Eds. LNCS, vol. 3300. Springer-Verlag, Berlin-Heidelberg-New York, 42–83.

Bertossi, L. and Bravo, L. 2004b. Query answering in peer-to-peer data exchange systems. In *Proc. EDBT Workshop on Peer-to-Peer Computing and Databases*, W. Lindner, M. Mesiti, C. Türker, Y. Tzitzikas and A. Vakali, Eds. LNCS, vol. 3268. Springer-Verlag, Berlin-Heidelberg-New York, 476–485.

Bertossi, L. and Bravo, L. 2007. The semantics of consistency and trust in peer data exchange systems. In *Proc. International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, N. Dershowitz and A. Voronkov, Eds. LNCS, vol. 4790. Springer-Verlag, Berlin-Heidelberg-New York, 107–122.

Boutaba, R. and Marshall, A. (eds.) 2006. Special issue on management in peer-to-peer systems. *Computer Networks 50* (4), 472–484.

Bravo, L. 2007. *Handling Inconsistency in Databases and Data Integration Systems*. Ph.D. thesis, Carleton University, Department of Computer Science. URL: http://people.scs.carleton.ca/~bertossi/papers/Thesis36.pdf

Bravo, L. and Bertossi, L. 2003. Logic programs for consistently querying data sources. In *Proc. of the 18th International Joint Conference on Artificial Intelligence*, G. Gottlob and T. Walsh, Eds. Morgan Kaufmann, Burlington, MA, 10–15.

Bravo, L. and Bertossi, L. 2005. Disjunctive deductive databases for computing certain and consistent answers to queries from mediated data integration systems. *Journal of Applied Logic 3* (2), 329–367.

Bravo, L. and Bertossi, L. 2006. Semantically correct query answers in the presence of null values. In *Current Trends in Database Technology – EDBT 2006*, T. Grust, H. Höpfner,

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

A. Illarramendi, S. Jablonski, M. Mesiti, S. Müller, P.-L. Patranjan, K.-U. Sattler, M. Spiliopoulou and J. Wijsen, Eds. LNCS, vol. 4254. Springer-Verlag, Berlin-Heidelberg-New York, 336–357.

Brewka, G., Eiter, T. and Truszczynski, M. 2011. Answer set programming at a glance. *Communications of the ACM 54* (12), 92–103.

Cali, A., Lembo, D. and Rosati, R. 2003. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *Proc. ACM Symposium on Principles of Database Systems*, F. Neven, C. Beeri and T. Milo, Eds. ACM Press, New York, USA, 260–271.

Calvanese, D., Damaggio, E., De Giacomo, G., Lenzerini, M. and Rosati, R. 2004. Semantic data integration in P2P systems. In *Proc. VLDB International Workshop on Databases, Information Systems and Peer-to-Peer Computing*, K. Aberer, M. Koubarakis and V. Kalogeraki, Eds. LNCS, vol. 2944. Springer-Verlag, Berlin-Heidelberg-New York, 77–90.

Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. 2005. Inconsistency tolerance in P2P Data integration: An epistemic logic approach. In *Proc. International Symposium on Database Programming Languages*, G. Bierman and C. Koch, Eds. LNCS, vol. 3774. Springer-Verlag, Berlin-Heidelberg-New York, 90–105.

Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. 2008. Inconsistency tolerance in P2P data integration: An epistemic logic approach. *Information Systems 33* (4–5), 360–384.

Calvanese, D., De Giacomo, G., Lenzerini, M. and Rosati, R. 2004. Logical foundations of peer-to-peer data integration. In *Proc. ACM Symposium on Principles of Database Systems*, C. Beeri and A. Deutsch, Eds. ACM Press, New York, 241–251.

Caniupan, M. and Bertossi, L. 2010. The consistency extractor system: Answer set programs for consistent query answering in databases. *Data & Knowledge Engineering 69* (6), 545–572.

Caroprese, L., Greco, S. and Zumpano, E. 2006. A logic programming approach to querying and integrating P2P deductive databases. In *Proc. FLAIRS*, G. Sutcliffe and R. Goebel, Eds. AAAI Press, New York, 31–36.

Caroprese, L. and Zumpano, E. 2008. Modeling cooperation in P2P data management systems. In *Foundations of Intelligent Systems*, A. An, S. Matwin, Z. W. Raś and D. Ślęzak, Eds. LNCS, Vol. 4994, Springer-Verlag, Berlin-Heidelberg-New York, 225–235.

Caroprese, L. and Zumpano, E. 2011. Aggregates and priorities in P2P data management systems. In *Proc. International Database Engineering and Applications Symposium (IDEAS 2011)*, B. C. Desai, I. F. Cruz and J. Bernardino, Eds. ACM Press, New York, 1–7.

Chomicki, J. 2007. Consistent query answering: Five easy pieces. In *Proc. International Conference on Database Theory*, T. Schwentick and D. Suciu, Eds. LNCS, vol. 4353. Springer-Verlag, Berlin-Heidelberg-New York, 1–17.

Dantsin, E., Eiter, T., Gottlob, G. and Voronkov, A. 2001. Complexity and expressive power of logic programming. *ACM Computing Surveys 33* (3), 374–425.

Demolombe, R. 2011. Transitivity and propagation of trust in information sources: An analysis in modal logic. In *Proc. CLIMA XII*, J. Leite, P. Torroni, T. Ågotnes, G. Boella and L. van der Torre, Eds. LNAI, vol. 6814. Springer-Verlag, Berlin-Heidelberg-New York, 13–28.

Eiter, T., Fink, M., Greco, G. and Lembo, D. 2008. Repair localization for query answering from inconsistent databases. *ACM Transactions on Database Systems 33,* 2.

Eiter, T., Gottlob, G. and Mannila, H. 1997. Disjunctive datalog. *ACM Transactions on Database Systems 22* (3), 364–418.

Faber, W., Greco, G., and Leone, N. 2007. Magic Sets and their Application to Data Integration. *Journal of Computer and System Sciences 73* (4), 584–609.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Franconi, E., Kuper, G., Lopatenko, A. and Serafini, L. 2004a. A robust logical and computational characterisation of peer-to-peer database systems. In *Proc. VLDB Workshop on Databases, Information Systems and P2P Computing*, K. Aberer, M. Koubarakis and V. Kalogeraki, Eds. LNCS, vol. 2944. Springer-Verlag, Berlin-Heidelberg-New York, 64–76.

Franconi, E., Kuper, G., Lopatenko, A. and Zaihrayeu, I. 2004b. A distributed algorithm for robust data sharing and updates in P2P database networks. In *Proc. EDBT Workshop on Peer-to-peer Computing and Databases*, W. Lindner, M. Mesiti, C. Türker, Y. Tzitzikas and A. Vakali, Eds. LNCS, vol. 3268. Springer-Verlag, Berlin-Heidelberg-New York, 446–455.

Franconi, E. and Tessaris, S. 2012. On the logic of SQL nulls. In *Proc. Alberto Mendelzon WS on Foundations of Data Management*, J. Freire and D. Suciu, Eds. Vol. 866. CEUR WS Proceedings. Available at: http://ceur-ws.org/

Fuxman, A., Kolaitis, P., Miller, R. and Tan, W. 2006. Peer data exchange. *ACM Transactions on Database Systems 31* (4), 1454–1498.

Gelfond, M. and Lifschitz, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing 9*, 365–385.

Greco, G., Greco, S. and Zumpano, E. 2003. A logical framework for querying and repairing inconsistent databases. *IEEE Transactions on Knowledge and Data Engineering 15* (6), 1389–1408.

Halevy, A., Ives, Z., Madhavan, J., Mork, P., Suciu, D. and Tatarinov, I. 2004. The piazza peer data management system. *IEEE Transactions on Knowledge and Data Engineering 16* (7), 787–798.

Halevy, A., Ives, Z., Suciu, D. and Tatarinov, I. 2003. Schema mediation in peer data management systems. In *Proc. International Conference on Data Engineering*, U. Dayal, K. Ramamritham and T. M. Vijayaraman, Eds. IEEE Computer Society, Washington, DC, 505–518.

Imielinski, Th. and Lipski, W. 1984. Incomplete information in relational databases. *Journal of the ACM 31* (4), 761–791.

Jøsang, A., Azderska, T. and Marsh, S. 2012. Trust transitivity and conditional belief reasoning. In *Trust Management VI, IFIP Advances in Information and Communication Technology, Vol. 374*, T. Dimitrakos, R. Moona, D. Patel and D. H. McKnight, Eds. Springer, Heidelberg-Dordrecht-London-New York, 68–83

Jøsang, A., Marsh, S. and Pope, S. 2006. Exploring different types of trust propagation. In *Proc. iTrust*, K. Stølen, W. H. Winsborough, F. Martinelli and F. Massacci, Eds. LNCS, vol. 3986. Springer-Verlag, Berlin-Heidelberg-New York, 179–192.

Kolaitis, P. 2005. Schema mappings, data exchange, and metadata management. In *Proc. ACM Symposium on Principles of Database Systems*, C. Li, Ed. ACM Press, New York, 61–75.

Lenzerini, M. 2002. Data integration: A theoretica perspective. In *Proc. ACM Symposium on Principles of Database Systems*, L. Popa, S. Abiteboul and P. G. Kolaitis, Eds. ACM Press, New York, 233–246.

Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S. and Scarcello, F. 2006. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic 7* (3), 499–562.

Levene, M. and Loizou, G. 1997. Null inclusion dependencies in relational databases. *Information and Computation 136* (2), 67–108.

Levene, M. and Loizou, G. 1999. *A Guided Tour of Relational Databases and Beyond*. Springer, London.

Libkin, L. 2014. Incomplete data: What went wrong, and how to fix it. In *Proc. ACM Symposium on Principles of Database Systems*, R. Hull and M. Grohe, Eds. ACM Press, New York, 1–13.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press

Libkin, L. 2016. Certain answers as objects and knowledge. *Artificial Intelligence 232* (1), 1–19.

Lifschitz, V., Pichotta, K. and Yang, F. 2012. Relational theories with null values and non-Herbrand stable models. *Theory and Practice of Logic Programming 12* (4–5), 565–582.

Marti, S. and Garcia-Molina, H. 2006. Taxonomy of trust: Categorizing P2P reputation systems. *Computer Networks 50* (4), 472–484.

Papadimitriou, Ch. 1994. *Computational Complexity*. Addison-Wesley, Reading, MA.

Reiter, R. 1984. Towards a logical reconstruction of relational database theory. In *On Conceptual Modelling*, M. Brodie, J. Mylopoulos, and J. Schmidt, Eds. Springer, New York, 191–233.

Sabater, J. and Sierra, C. 2005. Review on computational trust and reputation models. *Artificial Intelligence Review 24* (1), 33–60.

Serafini, L., Borgida, A. and Tamilin, A. 2005. Aspects of distributed and modular ontology reasoning. In *Proc. International Joint Conference on Artificial Intelligence*, L. P. Kaelbling and A. Saffiotti, Eds. Morgan Kaufmann, Burlington, MA, 570–575.

Traylor, B. and Gelfond, M. 1994. Representing null values in logic programming. In *Logical Foundations of Computer Science*, A. Nerode and Y. V. Matiyasevich, Eds. LNCS, vol. 813. Springer-Verlag, Berlin-Heidelberg-New York, 341–352.

https://doi.org/10.1017/S147106841600017X Published online by Cambridge University Press