

# The Semantics of Consistency and Trust in Peer Data Exchange Systems

Leopoldo Bertossi<sup>1</sup> and Loreto Bravo<sup>2</sup>

<sup>1</sup> Carleton University, School of Computer Science, Ottawa, Canada.  
bertossi@scs.carleton.ca

<sup>2</sup> University of Edinburgh, School of Informatics, Edinburgh, UK.  
lbravo@inf.ed.ac.uk

**Abstract.** We propose and investigate a semantics for *peer data exchange systems* (or peer data management systems) where different peers are pairwise related to each other by means of data exchange constraints and trust relationships. These two elements plus the data at the peers' sites and the local integrity constraints for a peer are made compatible via the proposed semantics by determining a set of *solution instances*, which are the intended virtual instances for the peer. The semantically correct answers from a peer to a query, called its *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.

## 1 Introduction

A peer data exchange system (PDES) consists of a finite set of peers  $\{P_1, \dots, P_n\}$ , each of them with a local database instance. Peers may be pairwise related by means of logical sentences, called *data exchange constraints* (DECs), which are expressed in terms of the participating schemas and are expected to be satisfied by the combined data. Furthermore, a peer  $P$  may trust its data the same as or less than other peers' data, i.e. there may be trust relationships between pairs of peers. We may also have integrity constraints (ICs) that are local to each peer.

The DECs could be seen as ICs on a global database obtained by conceptually putting together all the peers' schemas and data. Most likely, these DECs will not be satisfied in this global database, but virtually enforcing their satisfaction at query time has the effect of shipping (sub)queries and data between peers. Actually, in such a PDES, a query  $Q$  is posed to a peer  $P$ , who, in order to answer it, may need to consider both its own data and the data stored at other peers' sites that are related to  $P$  by DECs. Keeping  $P$ 's DECs satisfied at query time may imply getting data from other peers to complement  $P$ 's data, but also not using part of its own data.

The decision by a peer  $P$  on what other data to consider does not depend only on its DECs, but also on the *trust relationships* that  $P$  has with other peers. For example, if peer  $P$  trusts peer  $Q$ 's data more than its own,  $P$  will accommodate its data to  $Q$ 's in order to keep the DECs between them satisfied.

In order for  $P$  to return meaningful answers, its local semantic constraints have to be taken into account. The consistency of its instance should be preserved when it is virtually updated due to the interaction with other peers. In consequence, still at query time,  $P$  may also need to "virtually repair" its data.

Our semantics makes all these elements compatible, by defining a set of virtual global instances called *solution instances* (or simply solutions) for a peer. In consequence, the “data” for a peer, the one in its solution instances, depends upon its instance, the local instances of the related peers, the satisfaction of the DEC’s, and the satisfaction by P of its local IC’s. After that, having a precise definition of the intended solution instances for a peer P, the *peer consistent answers* (PCAs) from peer P to query Q are defined as those answers that can be retrieved from *every* possible solution for P.

The definition of solution for P may suggest that P may physically change other peers’ data, but this is not the case. The notion of solution is used as an auxiliary notion to characterize the correct answers from P’s point of view. Ideally, P should be able to obtain its peer consistent answers just by querying the already available local instances. This resembles the approach to *consistent query answering* (CQA) in databases, where consistent answers to a query posed to a database, which is possibly inconsistent wrt to a given set of IC’s, are defined as those answers that can be retrieved from every minimally repaired version of the original instance. Methods have been developed for computing these answers without having to physically repair the given instance [1, 8, 9].

Our work goes in the direction of semantic approaches to peer-to-peer data exchange [26, 28, 27, 13, 20, 7, 14, 22]. In [7] trust relationships were introduced for the first time in this scenario and the notions of solution instance for a peer and of peer consistent answer to a query were introduced. The case of a peer and its immediate neighbors was considered and investigated. However, the situation where a peer is related by logical transitivity via DEC’s to other, non neighboring peers was not modelled. Actually, in [7] it was only indicated that giving a semantics to solution instances for a peer that consider the transitive relationships to other peers could be done by using logic programs with stable model semantics. We fully develop this idea here. First, we provide a general model-theoretic definition of solutions for a peer, including the transitive case, and next, we specify the solutions as the models of disjunctive logic programs with stable model semantics [23], aka. *answer set programs* [24, 3].

Logic programs can capture the different ways the system stabilizes after satisfying the DEC’s, the trust relationships, and the local IC’s. Disjunctive programs allow for the specification of alternative virtual updates on data sources under certain conditions. We propose appropriate logic programs and then we establish that there is a one-to-one correspondence between the set of solution instances for a peer and the set of stable models of the program. Logic programs provide an expressive language for representing and specifying the alternative solutions for a peer, and become executable specifications for computing peer consistent answers. This approach has been exploited in CQA, where database repairs are specified as stable models of a program [2, 25, 4, 16].

## 2 A Semantics for PDESs

We consider peers that have mutually disjoint relational schemas; but all of them share a possibly infinite database domain  $\mathcal{U}$ . Peers are denoted by A, B, P, Q, ...

**Definition 1.** A *peer data exchange system*  $\mathfrak{P}$  consists of: (a) A finite set  $\mathcal{P}$  of peers, with each peer P owning a relational database schema  $\mathcal{R}(P)$ , and

a database instance  $D(\mathcal{P})$  conforming to schema  $\mathcal{R}(\mathcal{P})$ . The schemas determine FO languages, e.g.  $\mathcal{L}(\mathcal{P})$ ,  $\mathcal{L}(\mathcal{P}, \mathcal{Q})$ . (b) For each peer  $\mathcal{P}$ , collections  $\Sigma(\mathcal{P}, \mathcal{Q})$  of sentences of  $\mathcal{L}(\mathcal{P}, \mathcal{Q})$ , which contain the DEC's between  $\mathcal{P}$  and a peer  $\mathcal{Q}$ . Here,  $\Sigma(\mathcal{P}) := \bigcup_{\mathcal{Q}} \Sigma(\mathcal{P}, \mathcal{Q})$  and  $\Sigma := \bigcup_{\mathcal{P} \in \mathcal{P}} \Sigma(\mathcal{P})$ . (c) For each peer  $\mathcal{P}$ , a set of  $\mathcal{L}(\mathcal{P})$ -sentences  $IC(\mathcal{P})$  that are ICs on  $\mathcal{R}(\mathcal{P})$ . Here,  $IC = \bigcup_{\mathcal{P} \in \mathcal{P}} IC(\mathcal{P})$ . (d) A relation  $trust \subseteq \mathcal{P} \times \{less, same\} \times \mathcal{P}$ , with exactly one triple of the form  $\langle \mathcal{P}, \cdot, \mathcal{Q} \rangle$  for each non empty  $\Sigma(\mathcal{P}, \mathcal{Q})$ .  $\square$

The intended semantics of  $(\mathbf{A}, less, \mathbf{B}) \in trust$  is that peer  $\mathbf{A}$  trusts itself less than  $\mathbf{B}$ ; while  $(\mathbf{A}, same, \mathbf{B}) \in trust$  indicates that  $\mathbf{A}$  trusts itself the same as  $\mathbf{B}$ .

**Definition 2.** (a) A *universal data exchange constraint* (UDEC) between peers  $\mathcal{P1}, \mathcal{P2}$  is a first-order (FO) formula of form:

$$\forall \bar{x} \left( \bigwedge_{i=1}^n R_i(\bar{x}_i) \longrightarrow \left( \bigvee_{j=1}^m Q_j(\bar{y}_j) \vee \varphi \right) \right); \quad (1)$$

where the  $R_i, Q_j$  are relations in  $\mathcal{R}(\mathcal{P1}) \cup \mathcal{R}(\mathcal{P2})$ ,  $\varphi$  is a formula containing built-in atoms<sup>1</sup> only, and  $\bar{x}_i, \bar{y}_j \subseteq \bar{x}$ . (b) A *referential data exchange constraint* (RDEC) between peers  $\mathcal{P1}, \mathcal{P2}$  is an  $\mathcal{L}(\mathcal{P1}, \mathcal{P2})$ -sentence of the form:

$$\forall \bar{x} (R(\bar{x}) \longrightarrow \exists \bar{y} Q(\bar{x}', \bar{y})); \quad (2)$$

where  $R, Q \in \mathcal{R}(\mathcal{P1}) \cup \mathcal{R}(\mathcal{P2})$ , and  $\bar{x}' \subseteq \bar{x}$ .  $\square$

Notice that the sets of DEC's  $\Sigma(\mathcal{P1}, \mathcal{P2})$  and  $\Sigma(\mathcal{P2}, \mathcal{P1})$  can be different. When exchanging or repairing data, the existential quantifier in RDEC's will be interpreted as a null value. By having one database atom in the consequent, we avoid existential joins that, when filled with nulls, do not have a clear semantics (cf. [11] for a discussion and a FO semantics of nulls in SQL databases).<sup>2</sup>

*Example 1.* Consider a PDES  $\mathfrak{P}$  with four peers and  $\mathcal{R}(\mathcal{P1}) = \{R^1(\cdot, \cdot)\}$ ,  $\mathcal{R}(\mathcal{P2}) = \{R^2(\cdot, \cdot), S^2(\cdot, \cdot)\}$ ,  $\mathcal{R}(\mathcal{P3}) = \{R^3(\cdot, \cdot)\}$ ,  $\mathcal{R}(\mathcal{P4}) = \{R^4(\cdot, \cdot, \cdot)\}$ , and DEC's:

$$\begin{aligned} \Sigma(\mathcal{P1}, \mathcal{P2}) &= \{\forall xy (R^2(x, y) \rightarrow R^1(x, y))\}, \\ \Sigma(\mathcal{P2}, \mathcal{P3}) &= \{\forall xy (R^2(x, y) \wedge R^3(x, y) \rightarrow \mathbf{false})\}, \\ \Sigma(\mathcal{P4}, \mathcal{P2}) &= \{\forall xyz (R^2(x, y) \wedge S^2(y, z) \rightarrow R^4(x, y, z))\}, \\ \Sigma(\mathcal{P4}, \mathcal{P3}) &= \{\forall xy (R^3(x, y) \rightarrow \exists z R^4(x, y, z))\}. \end{aligned}$$

Here **false** in the second DEC is a built-in atom that is false in every instance, so the DEC specifies that relations  $R^2$  and  $R^3$  are disjoint. The DEC's in  $\Sigma(\mathcal{P1}, \mathcal{P2})$ ,  $\Sigma(\mathcal{P2}, \mathcal{P3})$  and  $\Sigma(\mathcal{P4}, \mathcal{P2})$  are UDEC's and the one in  $\Sigma(\mathcal{P4}, \mathcal{P3})$  is a RDEC. Finally, we could have  $trust = \{(\mathcal{P1}, less, \mathcal{P2}), (\mathcal{P2}, same, \mathcal{P3}), (\mathcal{P4}, less, \mathcal{P2}), (\mathcal{P4}, less, \mathcal{P3})\}$ . Peer  $\mathcal{P1}$  trusts  $\mathcal{P2}$  more than itself, and it will import all the data from  $R^2$  into its table  $R^1$ . On the other hand, peer  $\mathcal{P2}$  trusts peer  $\mathcal{P3}$  as much as itself, and its DEC states that it is not possible to have the same tuple in both  $R^3$  and  $R^2$ .  $\square$

Local IC's  $IC(\mathcal{P})$  are also of the form in Definition 2, but with all database predicates in  $\mathcal{R}(\mathcal{P})$ . So, we can identify  $IC(\mathcal{P})$  with  $\Sigma(\mathcal{P}, \mathcal{P})$ . All the common IC's found in database practice can be accommodated into these syntactic classes.

<sup>1</sup> For example,  $x \neq 3$ ,  $y = z$  and  $z > 3$ .

<sup>2</sup> Our framework can be easily adapted to cases where, instead of *null*, one uses arbitrary elements of the database domain [12] or labelled nulls [29].

In particular, using the atom **false**, denial constraints are of the form (1). In Example 1, we could have  $IC(P2) = \{\forall x \forall y (R^2(x, y) \wedge S^2(x, y) \rightarrow \mathbf{false})\}$ .


Each peer  $P$  is responsible for maintaining its material instance consistent with respect to its ICs  $IC(P)$ , independently from other peers; and we may assume that this is the case. However, our semantics works the same if we allow peers with locally inconsistent data. According to our semantics presented below, unsatisfied local ICs are considered when solution instances are specified, which is done by using a repair semantics and techniques developed for CQA [9, 6].

When a peer  $P$  is posed a query, it may have to submit queries to other peers according to its DEC, using data in other peers' relations appearing in them. Data brought from other peers will possibly cause virtual updates on  $P$ 's data, which may create virtual violations of  $P$ 's local ICs, which has to be considered.

From the perspective of a peer  $P$ , its own database may be inconsistent with respect to the data owned by another peer  $Q$  it trusts more or the same, and the DEC in  $\Sigma(P, Q)$ . When  $P$  queries its own database, the answers from  $P$  should be consistent with  $\Sigma(P, Q)$  and its own ICs  $IC(P)$ . In principle,  $P$ , which is not allowed to change other peers' data, could try to physically repair its database in order to satisfy  $\Sigma(P) \cup IC(P)$ . This is not a realistic approach. Rather,  $P$  should solve its semantic conflicts at query time. This leads to a set of virtual instances, the minimal repairs of  $P$ 's local database, where  $P$ 's DEC and ICs are satisfied, while respecting  $P$ 's trust relationships to other peers. The answers returned by  $P$  to the user are those that are true of all these instances.

The solution instances of a peer will be determined not only by its relationships with its neighbors, but also by the neighbors of its neighbors, etc.

**Definition 3.** (a) The *accessibility graph*  $\mathcal{G}(\mathfrak{P})$  of a PDES  $\mathfrak{P}$  contains a vertex for each peer  $P \in \mathcal{P}$  and a directed edge from  $P_i$  to  $P_j$  if  $\Sigma(P_i, P_j)$  is non empty. An edge from  $P_i$  to  $P_j$  is labelled with “ $<$ ” when  $(P_i, \text{less}, P_j) \in \text{trust}$ , and with “ $=$ ” when  $(P_i, \text{same}, P_j) \in \text{trust}$ .<sup>3</sup> (b) Peer  $P'$  is *accessible* from  $P$  if there is a path in  $\mathcal{G}(\mathfrak{P})$  from  $P$  to  $P'$  or if  $P' = P$ . Peer  $P'$  is a *neighbor* of  $P$  if there is an edge from  $P$  to  $P'$  in  $\mathcal{G}(\mathfrak{P})$ , or if  $P' = P$ . With  $\mathcal{AC}(P)$  and  $\mathcal{N}(P)$  we denote the sets of peers that are accessible from  $P$  and the neighbors of  $P$ , respectively. For  $P \in \mathcal{P}$ ,  $\mathcal{G}(P)$  is the restriction of  $\mathcal{G}(\mathfrak{P})$  to  $\mathcal{AC}(P)$ .  $\square$



**Fig. 1.** Graphs for Example 2

*Example 2.* (Example 1 continued) In this system,  $\mathcal{AC}(P1) = \{P1, P2, P3\}$ ,  $\mathcal{AC}(P2) = \{P2, P3\}$ ,  $\mathcal{AC}(P3) = \{P3\}$ ,  $\mathcal{AC}(P4) = \{P2, P3, P4\}$ ,  $\mathcal{N}(P1) = \{P1, P2\}$ ,  $\mathcal{N}(P2) = \{P2, P3\}$ ,  $\mathcal{N}(P3) = \{P3\}$ , and  $\mathcal{N}(P4) = \{P2, P3, P4\}$ .  $\square$

<sup>3</sup> In case a peer  $P$  trusts itself more than another peer, the information of the latter is irrelevant to  $P$ .

The data distributed across different peers has to be appropriately gathered to build solution instances for a peer, and different semantics may emerge as candidates, depending on the granularity of the data sent between peers. We develop one of them,<sup>4</sup> according to which, the data that a peer  $P$  receives from a neighbor  $Q$  to build its own solutions is the *intersection of the solutions* for  $Q$ . After  $P$  collects this data, only  $P$ 's DECs and ICs are considered. This is a recursive definition since the solutions for the neighbors have to be determined, under the same semantics. Base cases of the recursion are peers with no relevant DECs. In consequence, this semantics requires an acyclic accessibility graph.

In [26] problematic cases involving cyclic dependencies through DECs are identified, which implicitly involve a cyclic accessibility graph. For example, we may have a PDES  $\mathfrak{P}$  with  $\mathcal{P} = \{P1, P2, P3\}$ , with relations  $R^1(\cdot), R^2(\cdot), R^3(\cdot)$ , resp., and DECs  $\Sigma(P1) = \{\forall x(R^2(x) \rightarrow R^1(x))\}$ ,  $\Sigma(P2) = \{\forall x(R^3(x) \rightarrow R^2(x))\}$ ,  $\Sigma(P3) = \{\forall x(R^1(x) \rightarrow R^3(x))\}$ , each of them satisfied only by importing data into the peer who owns the DEC. The implicit trust relation  $\{(P1, less, P2), (P2, less, P3), (P3, less, P1)\}$  makes  $\mathcal{AC}(\mathfrak{P})$  cyclic. In [26] it is assumed that no cycles of this kind appear. In the following, we will also assume that  $\mathcal{G}(\mathfrak{P})$  is acyclic, and then for each particular peer  $P$ ,  $\mathcal{G}(P)$  is acyclic.

Null values will be used to satisfy referential DECs and local referential ICs; and the repair semantics based on introduction of null values, presented and developed in [11, 10] for single relational databases and RICs, can be adapted here, but now taking into account the trust relationships. Data sources at the peers' sites may contain null values that, as those used to satisfy referential constraints, will have a semantics that corresponds to the way nulls are handled by DBMSs that follow the SQL standard. In particular, there is only one constant, *null*, that is used as the null value.<sup>5</sup> We use a semantics for IC satisfaction in the presence of nulls that generalizes the one implemented in DBMSs, and coincides with the first-order notion of formula satisfaction in databases without nulls [10].

A formal development of the notion of constraint satisfaction, denoted  $D \models \psi$ , with  $D$  possibly containing *null* and  $\psi$  a constraint, can be found in [11, 10] (see Appendix B for a review). What matters most for the rest of this paper is that the satisfaction of a constraint  $\psi$  depends upon the presence of *null* in attributes of database relations which are also *relevant attributes* for  $\psi$ .

Solutions for a peer should stay close to its original physical instance while satisfying the DECs and local ICs. We do not want to import or give up more data than strictly required to satisfy the constraints. To formalize this idea, we first need to compare tuples that may contain *null*. A constant  $c$  provides *less or equal information* than a constant  $d$ , denoted  $c \sqsubseteq d$ , iff  $c$  is *null* or  $c = d$  [30]. A tuple  $\bar{t}_1 = (c_1, \dots, c_n)$  provides less or equal information than  $\bar{t}_2 = (d_1, \dots, d_n)$ , denoted  $\bar{t}_1 \sqsubseteq \bar{t}_2$ , iff  $c_i \sqsubseteq d_i$  for every  $i = 1, \dots, n$ . Finally,  $\bar{t}_1 \sqsubset \bar{t}_2$  means  $\bar{t}_1 \sqsubseteq \bar{t}_2$

<sup>4</sup> In [10] also two other alternative semantics are fully developed and compared, in particular establishing some conditions under which they coincide or differ. The other semantics assume that more detailed information, such as mappings and trust relationships, can be sent between peers.

<sup>5</sup> This null is obviously different from the multiple labelled null values that are considered in data exchange for satisfying existential quantifiers (cf. [29] for a survey).

and  $\bar{t}_1 \neq \bar{t}_2$ . In the following, a database instance is identified with a finite set of ground database atoms; and  $\Delta(\cdot, \cdot)$  denotes the symmetric difference of sets.

**Definition 4.** Let  $D, D', D''$  be database instances for the same schema. It holds that  $D' \leq_D D''$  iff for every  $P(\bar{a}) \in \Delta(D, D')$ , there exists  $P(\bar{a}')$ , such that: (a)  $P(\bar{a}') \in \Delta(D, D'')$ ; (b)  $\bar{a} \sqsubseteq \bar{a}'$ ; and (c) if  $\bar{a} \sqsubset \bar{a}'$ , then  $P(\bar{a}') \notin \Delta(D, D')$ . Finally,  $D'' <_D D'$  means  $D'' \leq_D D'$  but not  $D' \leq_D D''$ .  $\square$

If  $D' \leq_D D''$ , we say that  $D'$  is closer to  $D$  than  $D''$ . Condition (c) in Def. 4 ensures that a database that adds to  $D$  a tuple with *null* is closer to  $D$  than other that adds other constant. This condition will later ensure that the satisfaction of RDECs is enforced by using *null*.

For an instance  $D$  of a schema  $\mathcal{S}$ , and  $\mathcal{S}'$  a subschema of  $\mathcal{S}$ ,  $D|\mathcal{S}'$  denotes the restriction of  $D$  to  $\mathcal{S}'$ . Thus, if  $R$  is a predicate in  $\mathcal{S}$  and  $D$  is an instance for  $\mathcal{S}$ ,  $D|\{R\}$  denotes the extension of  $R$  in  $D$ . If  $\mathcal{R}(\mathcal{P}) \subseteq \mathcal{S}$ ,  $D|\mathcal{P}$  is the restriction of  $D$  to  $\mathcal{R}(\mathcal{P})$ . A neighborhood solution for  $\mathcal{P}$  and a database for its whole neighborhood is a closest database that satisfies  $\mathcal{P}$ 's DECs, ICs, and trust relationships.

**Definition 5.** Given a peer  $\mathcal{P}$  in a PDES  $\mathfrak{P}$  and instances  $D, D'$  on schema  $\bigcup_{\mathcal{Q} \in \mathcal{N}(\mathcal{P})} \mathcal{R}(\mathcal{Q})$ ,  $D'$  is a *neighborhood solution for  $\mathcal{P}$  and  $D$*  if: (a)  $D' \models \bigcup_{\mathcal{Q} \in \mathcal{N}(\mathcal{P})} \Sigma(\mathcal{P}, \mathcal{Q}) \cup IC(\mathcal{P})$ . (b)  $D'|\{R\} = D|\{R\}$  for every predicate  $R \in \mathcal{R}(\mathcal{Q})$  with  $(\mathcal{P}, \text{less}, \mathcal{Q}) \in \text{trust}$ . (c) There is no instance  $D''$  that satisfies (a) and (b), and such that  $D'' <_D D'$ .  $\square$

We do not require in (a)  $IC(\mathcal{Q})$  to be satisfied, because  $\mathcal{Q}$  will move data to  $\mathcal{P}$ 's site, where inconsistencies will be solved locally, according to Definition 6, where  $S(\mathcal{P})$  denotes the set of solutions for peer  $\mathcal{P}$ .

**Definition 6.** Given a peer  $\mathcal{P}$  in a PDES  $\mathfrak{P}$  with local instance  $D(\mathcal{P})$ , an instance  $D$  over  $\mathcal{R}(\mathcal{P})$  is a *solution instance* for  $\mathcal{P}$  if: (a)  $D = D(\mathcal{P})$  and  $\Sigma(\mathcal{P}) = \emptyset$ ; or (b)  $\Sigma(\mathcal{P}) \neq \emptyset$ ,  $D = \overline{D}|\mathcal{P}$  where  $\overline{D}$  is a neighborhood solution for  $\mathcal{P}$  and the database instance  $D(\mathcal{P}) \cup \bigcup_{\mathcal{Q} \in (\mathcal{N}(\mathcal{P}) \setminus \{\mathcal{P}\})} \bigcap_{I \in S(\mathcal{Q})} I$  over schema  $\bigcup_{\mathcal{Q} \in \mathcal{N}(\mathcal{P})} \mathcal{R}(\mathcal{Q})$ .  $\square$

Intuitively, before constructing  $\mathcal{P}$ 's solutions,  $\mathcal{P}$  has its local instance  $D(\mathcal{P})$  and each of its neighbors has as local instance the intersection of its own solutions. This produces a combined database. After that, the solutions for  $\mathcal{P}$  are obtained by restricting to  $\mathcal{P}$  the neighborhood solutions for the combined instance. The neighborhood solution captures the minimal virtual updates that are necessary to satisfy the DECS and local ICs. As there may be several neighborhood solutions, several solutions for a peer are possible.

*Example 3.* (Example 1 and 2 continued) Consider the following instances of peers  $\mathcal{P}1, \mathcal{P}2$  and  $\mathcal{P}3$ :  $D(\mathcal{P}1) = \{R^1(a, 2)\}$ ,  $D(\mathcal{P}2) = \{R^2(c, 4), R^2(d, 5)\}$ ,  $D(\mathcal{P}3) = \{R^3(c, 4)\}$  and  $D(\mathcal{P}4) = \{R^4(d, 5, 1)\}$ . If we want the solutions for  $\mathcal{P}1$ , the solutions for  $\mathcal{P}2$  are needed, who will need in turn the solutions for  $\mathcal{P}1$ . Since  $\mathcal{P}3$  has no DECs with other peers, its only neighborhood solution is its local instance  $D(\mathcal{P}3)$ . This data is sent back to  $\mathcal{P}2$ , who needs to repair  $\{R^2(c, 4), R^2(d, 5), R^3(c, 4)\}$  now wrt  $\Sigma(\mathcal{P}2, \mathcal{P}3)$ . As  $\mathcal{P}2$  trusts  $\mathcal{P}3$  the same as itself, it can modify its own data or the data it got from  $\mathcal{P}3$ . There are two neighborhood solutions for  $\mathcal{P}2$ :  $\{R^2(c, 4), R^2(d, 5)\}$  and  $\{R^2(d, 5), R^3(c, 4)\}$ , that lead to two solutions for  $\mathcal{P}2$ :  $\{R^2(c, 4), R^2(d, 5)\}$  and  $\{R^2(d, 5)\}$ . Peer  $\mathcal{P}2$  will send to  $\mathcal{P}1$  the intersection of its solutions:  $\{R^2(d, 5)\}$ . Now,  $\mathcal{P}1$  has to repair  $\{R^1(a, 2), R^2(d, 5)\}$  wrt

$\Sigma(\text{P1}, \text{P2}) = \{\forall xy (R^2(x, y) \rightarrow R^1(x, y))\}$ . Since P1 trusts its own data less than the data of P2, it will solve inconsistencies by modifying its own data. There is only one neighborhood solution,  $\{R^1(a, 2), R^2(d, 5), R^1(d, 5)\}$ , and the solution for P1 is  $\{R^1(a, 2), R^1(d, 5)\}$ .

To compute the solutions for P4, the solutions of P2 and P3 are computed as shown before. Neighborhood solutions for P4 are obtained by repairing  $\{R^4(d, 5, 1), R^2(d, 5), R^3(c, 4)\}$  wrt  $\Sigma(\text{P4}, \text{P2})$ , and  $\Sigma(\text{P4}, \text{P3})$ . The DECs in  $\Sigma(\text{P4}, \text{P2})$  are already satisfied, but not the ones in  $\Sigma(\text{P4}, \text{P3})$ . Since P4 trusts the data in P3 more, a repair is obtained by adding a tuple with *null* into P4. The unique neighborhood solution for P4 is  $\{R^4(d, 5, 1), R^2(d, 5), R^3(c, 4), R^4(c, 4, \text{null})\}$ . Consequently,  $S(\text{P4}) = \{\{R^4(d, 5, 1), R^4(c, 4, \text{null})\}\}$ .  $\square$

The peer consistent answers are the semantically correct answers to a query returned by a peer who consistently considers the data of- and trust relationships with its neighbors.

**Definition 7.** Given a FO query  $\mathcal{Q}(\bar{x}) \in \mathcal{L}(\text{P})$  posed to P, a ground tuple  $\bar{t}$  is a *peer consistent answer* (PCA) to  $\mathcal{Q}$  from P iff  $D \models \mathcal{Q}(\bar{t})$  for every solution instance  $D$  for P.  $\square$

*Example 4.* (Example 3 continued) If P2 is posed the query  $\mathcal{Q}: R^2(x, y)$ , from its first solution instance we get  $\{(c, 4), (d, 5)\}$ , and from the second,  $\{(d, 5)\}$ . Therefore, the only PCA from P2 is  $\{(d, 5)\}$ .  $\square$

Even in the absence of cycles in  $\mathcal{G}(\mathfrak{P})$ , there may be no solutions for a peer (cf. Example 10 in Appendix A). Furthermore, still with acyclic  $\mathcal{G}(\mathfrak{P})$ , the *decision problem of peer consistent answering*, i.e. deciding if a tuple is a PCA to a query, may be undecidable if consistency wrt RDECs is achieved using arbitrary values in the domain.<sup>6</sup> However, using *null* instead avoids this problem, making the problem decidable. Actually, by reduction from CQA to PCA and known results on the *data complexity* of CQA [11], we obtain

**Theorem 1.** The problem of peer consistent answering is  $\Pi_2^P$ -complete.  $\square$

### 3 Answer Set Programs and the Solutions for a Peer

In order to define the solutions for a peer P, we have to consider P's relevant peers, which are those in  $\mathcal{AC}(\text{P})$ . The presence of cycles, through trust relationships or constraints (DECs or ICs), have an impact on the semantics. The former cycles appear in a cyclic  $\mathcal{G}(\text{P})$ . The latter appear when the DECs and local ICs of peers in  $\mathcal{AC}(\text{P})$  put together present cycles through the implications that involve an RDEC or a local referential IC (a RIC). Sets of local ICs of this kind are called *RIC-cyclic* in [11]. For example,  $IC = \{\forall x(S(x) \rightarrow Q(x)), \forall x(Q(x) \rightarrow S(x)), \forall x(Q(x) \rightarrow \exists y T(x, y))\}$  is not RIC-cyclic, whereas  $IC' = IC \cup \{\forall xy (T(x, y) \rightarrow Q(y))\}$  is, because there is a cycle involving the RIC  $\forall x(Q(x) \rightarrow \exists y T(x, y))$ . RIC-cyclicity at the level of local ICs may lead to more solutions than intended when capturing the repair semantics by means of logic programs [11].

<sup>6</sup> The undecidability result for CQA in [12] can be reconstructed in our framework, because even with  $\mathcal{G}(\mathfrak{P})$  acyclic, DECs can have ref-cycles (cf. Example 5).

In order to deal with the new issues arising in PDESs, we will assume that, for each peer  $P$ ,  $IC(P)$  is RIC-acyclic (cf. Section 4 for a discussion). Cycles through DECs and ICs will be crucial for a logic programming-based specification of solutions for a peer. We will say that the PDES  $\mathfrak{P}$  is *ref-acyclic* when in  $\Sigma \cup IC$  there are no cycles that involve an RDEC or a RIC.

As Example 5 below shows, even assuming the acyclicity of  $\mathcal{G}(\mathfrak{P})$ , and RIC-acyclicity at the level of local ICs (or no local ICs at all), we may have ref-cycles in the set of all DECs. This is due to the generality of DECs, where we can have relations of any of the two peers on both sides of the implication.

*Example 5.* Peers  $P1, P2$  have relations  $R^1, R^2$ , resp.  $\Sigma(P1) = \{\forall xz(R^1(x, z) \rightarrow \exists yR^2(x, y)), \forall xz(R^2(x, z) \rightarrow \exists yR^1(x, y))\}$ ,  $\Sigma(P2) = \emptyset$ ; and  $(P1, less, P2) \in trust$ . Here,  $\mathcal{AC}(\mathfrak{P})$  is acyclic, but  $\Sigma(P1) \cup \Sigma(P2)$  has a ref-cycle.  $\square$

Now we will show how to specify solutions for a peer, given instances for the other peers, as the stable models of disjunctive logic programs. These programs use annotation constants to indicate the atoms that may become true or false (virtually inserted or deleted) in order to satisfy the DECs and local ICs. For each database predicate  $P$  we generate a new copy  $P$  with an extra argument to accommodate the annotation. In  $P(\bar{a}, \mathbf{t}_a)$ , annotation  $\mathbf{t}_a$  means that the atom is advised to be made true; and  $\mathbf{f}_a$ , that the atom should be made false. For each DEC and local IC  $\psi$ , a rule captures through its disjunctive head the alternative virtual updates that can be performed to satisfy  $\psi$  (cf. rules 2. and 3. in Definition 9 for DECs, and 4. and 5. for local ICs).

Annotation  $\mathbf{t}^*$  indicates that the atom is true or becomes true in the program. It is introduced in order to execute a sequence of virtual updates that is needed due to interacting DECs and ICs. Finally, atoms annotated with  $\mathbf{t}^{**}$  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).

The relevant attributes of a constraint are those where the occurrence of *null* is relevant for its satisfaction [11], and then, they receive a special treatment in the logic programs. For the DEC in  $\Sigma(P4, P3)$  in Example 1, the two attributes of  $R^3$  and the first two of  $R^4$  are relevant, but not the third attribute of  $R^4$ .

**Definition 8.** For a constraint  $\psi \in \mathcal{L}(\mathcal{R})$  and a variable or a domain constant  $t$ ,  $pos^R(\psi, t)$  is the set of positions in predicate  $R \in \mathcal{R}$  where  $t$  appears in  $\psi$ . The set of *relevant variables* for  $\psi$  is  $\mathcal{V}(\psi) = \{x \mid x \text{ is a repeated variable in } \psi\}$ . The set of *relevant attributes* for  $\psi$  is  $\mathcal{A}(\psi) = \{R[i] \mid x \in \mathcal{V}(\psi) \text{ and } i \in pos^R(\psi, x)\} \cup \{R[i] \mid c \text{ is a constant in } \psi \text{ and } i \in pos^R(\psi, c)\}$ , where  $R[i]$  denotes the attribute in position  $i$  in  $R$ .  $\square$

**Definition 9.** Consider a PDES  $\mathfrak{P}$ , a peer  $P \in \mathcal{P}$  with  $\mathcal{N}(P) = \{P, P1, \dots, Pn\}$ , and  $\mathcal{I} = \{I_1, \dots, I_n\}$ , where  $I_j$  is a database instance over the schema of  $Pj$ . The *solution program*  $\Pi(\mathfrak{P}, P, \mathcal{I})$  for  $P$  contains:

1.  $dom(a)$ , for every  $a \in (\mathcal{U} \setminus \{null\})$ .  $R(\bar{a})$ , for each atom  $R(\bar{a}) \in D(P)$ .  $R(\bar{a})$ , for each  $R(\bar{a}) \in I$  with  $I \in \mathcal{I}$ .
2. For every UDEC  $\psi \in \Sigma(P, Pj)$  of the form (1) with  $Pj \in \mathcal{N}(P)$  and  $(P, \{same \text{ or } less\}, Pj) \in trust$ , the rule:

$$\bigvee_{R \in R_P} R(\bar{x}_i, \mathbf{f}_a) \vee \bigvee_{Q \in Q_P} Q(\bar{y}_j, \mathbf{t}_a) \leftarrow \bigwedge_{i=1}^n R_i(\bar{x}_i, \mathbf{t}^*), \bigwedge_{j=1}^m Q_j(\bar{y}_j, \mathbf{f}^*), \bigwedge_{x_l \in \mathcal{A}(\psi)} x_l \neq null, \bar{\varphi},$$



where  $\mathcal{A}(\psi)$  is the set of relevant attributes of  $\psi$ ,  $\bar{\varphi}$  is a conjunction of built-ins that is equivalent to the negation of  $\varphi$ ; and, given  $\mathcal{R} = \{R_i \mid i = 1, \dots, n, R_i \text{ appears in } (1)\}$ ,  $R_{\mathbf{P}}$  is defined by  $R_{\mathbf{P}} = \mathcal{R} \cap \mathcal{R}(\mathbf{P})$  if  $(\mathbf{P}, \text{less}, \mathbf{Pj}) \in \text{trust}$ ; and  $R_{\mathbf{P}} = \mathcal{R}$  if  $(\mathbf{P}, \text{same}, \mathbf{Pj}) \in \text{trust}$ .  $Q_{\mathbf{P}}$  is defined analogously in terms of the  $Q_j$  predicates in (1).

3. For every RDEC  $\psi \in \Sigma(\mathbf{P}, \mathbf{Pj})$  of the form (2) such that  $\mathbf{Pj} \in \mathcal{N}(\mathbf{P})$  and  $(\mathbf{P}, \{\text{same or less}\}, \mathbf{Pj}) \in \text{trust}$ :
    - (a) If  $(\mathbf{P}, \text{same}, \mathbf{Pj}) \in \text{trust}$ , the rule:
$$R(\bar{x}, \mathbf{f}_{\mathbf{a}}) \vee Q(\bar{x}', \overline{\text{null}}, \mathbf{t}_{\mathbf{a}}) \leftarrow R(\bar{x}, \mathbf{t}^*), \text{ not } \text{aux}_{\psi}(\bar{x}'), \bar{x}' \neq \text{null}.$$
    - (b) If  $(\mathbf{P}, \text{less}, \mathbf{Pj}) \in \text{trust}$  and  $R \in \mathcal{R}(\mathbf{P})$ , the rule:
$$R(\bar{x}, \mathbf{f}_{\mathbf{a}}) \leftarrow R(\bar{x}, \mathbf{t}^*), \text{ not } \text{aux}_{\psi}(\bar{x}'), \bar{x}' \neq \text{null}.$$
    - (c) If  $(\mathbf{P}, \text{less}, \mathbf{Pj}) \in \text{trust}$  and  $Q \in \mathcal{R}(\mathbf{P})$ , the rule:
$$Q(\bar{x}', \overline{\text{null}}, \mathbf{t}_{\mathbf{a}}) \leftarrow R(\bar{x}, \mathbf{t}^*), \text{ not } \text{aux}_{\psi}(\bar{x}'), \bar{x}' \neq \text{null}.$$
- Plus the auxiliary rules:
- $$\text{aux}_{\psi}(\bar{x}') \leftarrow Q(\bar{x}', \overline{\text{null}}), \text{ not } Q(\bar{x}', \overline{\text{null}}, \mathbf{f}_{\mathbf{a}}), \bar{x}' \neq \text{null}.$$
- For every  $y_i \in \bar{y}$ :
- $$\text{aux}_{\psi}(\bar{x}') \leftarrow Q(\bar{x}', \bar{y}, \mathbf{t}^*), \text{ not } Q(\bar{x}', \bar{y}, \mathbf{f}_{\mathbf{a}}), \bar{x}' \neq \text{null}, y_i \neq \text{null}.$$
4. For every UIC  $\psi \in \text{IC}(\mathbf{P})$  of the form (1), the rule:
$$\bigvee_{i=1}^n P_i(\bar{x}_i, \mathbf{f}_{\mathbf{a}}) \vee \bigvee_{j=1}^m Q_j(\bar{y}_j, \mathbf{t}_{\mathbf{a}}) \leftarrow \bigwedge_{i=1}^n P_i(\bar{x}_i, \mathbf{t}^*), \bigwedge_{j=1}^m Q_j(\bar{y}_j, \mathbf{f}^*), \bigwedge_{x_l \in \mathcal{A}(\psi)} x_l \neq \text{null}, \bar{\varphi}.$$
  5. For every RIC  $\psi \in \text{IC}(\mathbf{P})$  of the form (2), the rules:
$$P(\bar{x}, \mathbf{f}_{\mathbf{a}}) \vee Q(\bar{x}', \overline{\text{null}}, \mathbf{t}_{\mathbf{a}}) \leftarrow P(\bar{x}, \mathbf{t}^*), \text{ not } \text{aux}_{\psi}(\bar{x}'), \bar{x}' \neq \text{null}.$$

$$\text{aux}_{\psi}(\bar{x}') \leftarrow Q(\bar{x}', \overline{\text{null}}), \text{ not } Q(\bar{x}', \overline{\text{null}}, \mathbf{f}_{\mathbf{a}}), \bar{x}' \neq \text{null}.$$

For every  $y_i \in \bar{y}$ :

$$\text{aux}_{\psi}(\bar{x}') \leftarrow Q(\bar{x}', \bar{y}, \mathbf{t}^*), \text{ not } Q(\bar{x}', \bar{y}, \mathbf{f}_{\mathbf{a}}), \bar{x}' \neq \text{null}, y_i \neq \text{null}.$$
  6. For each predicate  $R \in \mathcal{R}(\mathcal{N}(\mathbf{P}))$ , the annotation rules:
$$R(\bar{x}, \mathbf{f}^*) \leftarrow \text{dom}(\bar{x}), \text{ not } R(\bar{x}). \quad R(\bar{x}, \mathbf{f}^*) \leftarrow R(\bar{x}, \mathbf{f}_{\mathbf{a}}).$$

$$R(\bar{x}, \mathbf{t}^*) \leftarrow R(\bar{x}). \quad R(\bar{x}, \mathbf{t}^*) \leftarrow R(\bar{x}, \mathbf{t}_{\mathbf{a}}).$$
  7. For each predicate  $R \in \mathcal{R}(\mathcal{N}(\mathbf{P}))$ , the program constraint:
$$\leftarrow R(\bar{x}, \mathbf{t}_{\mathbf{a}}), R(\bar{x}, \mathbf{f}_{\mathbf{a}}).$$
  8. For each predicate  $R \in \mathcal{R}(\mathbf{P})$ , the interpretation rule:
$$R(\bar{x}, \mathbf{t}^{**}) \leftarrow R(\bar{x}, \mathbf{t}^*), \text{ not } R(\bar{x}, \mathbf{f}_{\mathbf{a}}). \quad \square$$

In bodies of rules associated to DEC's or IC's  $\psi$ , the conditions of the form  $x \neq \text{null}$ , with  $x$  a variable appearing in a relevant attribute of  $\psi$  are used to capture the semantics of null values as used in SQL (cf. [11] for details). An atom of the form  $P(\bar{x}, \overline{\text{null}}, \dots)$  in the program represents an atom with possibly several occurrences of *null*, not necessarily in its last arguments, e.g.  $P(x, \text{null}, y, \text{null}, \dots)$ . For  $\bar{x} = x_1, \dots, x_n$ ,  $\bar{x} \neq \text{null}$  abbreviates  $x_1 \neq \text{null}, \dots, x_n \neq \text{null}$ . The program constraints in 7. discard models where an atom is both inserted and deleted.

The facts of the program are those in instance  $D(\mathbf{P})$  of  $\mathbf{P}$  and those in instances  $I(\mathbf{Pi})$  for  $\mathbf{P}$ 's neighbors  $\mathbf{Pi}$ . The instances  $I(\mathbf{Pi})$  used in the program may not coincide with the physical instances  $D(\mathbf{Pi})$ . Actually, as shown below, if each  $I(\mathbf{Pi})$  is the intersection of the solutions of  $\mathbf{Pi}$ , then the stable models of the program are in one-to-one correspondence with the solutions of peer  $\mathbf{P}$ .

Since virtual updates are executed on the peers' instances, their local IC's have to be kept satisfied. That is the role of rules 4. (for universal IC's) and 5.

(for referential ICs) above. We adopt the stable model semantics for the solution programs [23], i.e. their intended models are their stable models.

*Example 6.* (Example 1 continued) Consider  $D(\text{P1}) = \{R^1(a, 2)\}$ , and the instance  $I_2 = \{R^2(d, 5)\}$  for P2, the neighbor of P1. The solution program  $\Pi(\mathfrak{P}, \text{P1}, \{I_2\})$  contains:  $dom(a). \quad dom(c). \quad \dots \quad R^1(a, 2). \quad R^2(d, 5).$   
 $R^1(x, y, \mathbf{t}_a) \leftarrow R^2(x, y, \mathbf{t}^*), R^1(x, y, \mathbf{f}^*), x \neq null, y \neq null.$   
 $R^1(x, y, \mathbf{t}^*) \leftarrow R^1(x, y, \mathbf{t}_a). \quad R^1(x, y, \mathbf{t}^*) \leftarrow R^1(x, y).$   
 $R^1(x, y, \mathbf{f}^*) \leftarrow R^1(x, y, \mathbf{f}_a). \quad R^1(x, y, \mathbf{f}^*) \leftarrow dom(x), dom(y), not R^1(x, y).$   
 $\leftarrow R^1(x, y, \mathbf{t}_a), R^1(x, y, \mathbf{f}_a). \quad R^1(x, y, \mathbf{t}^{**}) \leftarrow R^1(x, y, \mathbf{t}^*), not R^1(x, y, \mathbf{f}_a).$   
 With similar rules to the 2nd-6th for  $R^2$ . The first rule makes sure that an  $R^2$ -tuple that is not in  $R^1$ , is also virtually inserted into  $R^1$ . In this case, since P1 trusts P2 more than itself, virtual changes affect only peer P1.  $\square$

*Example 7.* Consider a PDES  $\mathfrak{P}$  with  $\mathcal{R}(\text{P1}) = \{R^1(\cdot, \cdot)\}$ ,  $D(\text{P1}) = \{R^1(s, t), R^1(a, null)\}$ ,  $\mathcal{R}(\text{P2}) = \{R^2(\cdot, \cdot)\}$ ,  $D(\text{P2}) = \{R^2(c, d), R^2(a, e)\}$ ,  $\Sigma(\text{P1}, \text{P2}) = \{\forall xy (R^2(x, y) \rightarrow \exists z R^1(x, z))\}$ ,  $IC(\text{P1}) = \{\forall xyz (R^1(x, y) \wedge R^1(x, z) \rightarrow y = z)\}$ ,  $\Sigma(\text{P2}, \text{P1}) = IC(\text{P2}) = \emptyset$ , and  $trust = \{\text{P1}, less, \text{P2}\}$ . The program  $\Pi(\mathfrak{P}, \text{P1}, \{I_2\})$  for P1 needs an instance for P2 that may be different from  $D(\text{P2})$ , but in this case we choose  $I_2 = D(\text{P2})$ , obtaining the following program (omitting rules 6., 7.):  $dom(a). \quad dom(b). \quad \dots \quad R^1(a, null). \quad R^1(s, t). \quad R^2(c, d). \quad R^2(a, e).$

$$\begin{aligned} R^1(x, null, \mathbf{t}_a) &\leftarrow R^2(x, \mathbf{t}^*), not aux(x), x \neq null. \\ aux(x) &\leftarrow R^1(x, null), not R^1(x, null, \mathbf{f}_a). \\ aux(x) &\leftarrow R^1(x, y, \mathbf{t}^*), not R^1(x, y, \mathbf{f}_a), x \neq null, y \neq null. \\ R^1(x, y, \mathbf{f}_a) \vee R^1(x, z, \mathbf{f}_a) &\leftarrow R^1(x, y, \mathbf{t}^*), R^1(x, z, \mathbf{t}^*), x \neq null, y \neq z. \\ R^1(x, y, \mathbf{t}^{**}) &\leftarrow R^1(x, y, \mathbf{t}^*), not R^1(x, y, \mathbf{f}_a). \end{aligned}$$

The first rule has the role of satisfying the RDEC by introducing a *null* into  $R^1$ . The fourth rule takes care of the local functional dependency.  $\square$

The atoms annotated with  $\mathbf{t}^{**}$  in a stable model of P's program have predicates of P only. They define a database instance for P. In Example 7, only  $R^1$ -atoms become annotated with  $\mathbf{t}^{**}$ . The program has only one stable model, with associated instance  $\{R^1(a, null), R^1(s, t), R^1(c, null)\}$ .

**Definition 10.** The database instance for peer P associated to a stable model  $\mathcal{M}$  of program  $\Pi(\mathfrak{P}, \text{P}, \mathcal{I})$  is  $D_{\mathcal{M}} = \{R(\bar{a}) \mid R(\bar{a}, \mathbf{t}^{**}) \in \mathcal{M}\}$ .  $\square$

**Theorem 2.** Given a PDES  $\mathfrak{P}$ ,  $\text{P} \in \mathcal{P}$ ,  $\mathcal{N}(\text{P}) = \{\text{P}, \text{P1}, \dots, \text{Pn}\}$ ,  $n \geq 0$ ,  $D(\text{P})$  an instance for P, and  $\mathcal{I}^* = \{I_1, \dots, I_n\}$  instances for P1,  $\dots$ , Pn, resp. If  $\Sigma \cup IC$  is ref-acyclic and each of the  $I_i$  is the intersection of the solution instances for peer Pi, then the instances of the form  $D_{\mathcal{M}}$ , where  $\mathcal{M}$  is a stable model of  $\Pi(\mathfrak{P}, \text{P}, \mathcal{I}^*)$ , are the all and only solution instances for P.  $\square$

In Example 7, given that  $D(\text{P2})$  is already the only solution instance for P2 (P2 has neither DEC's nor local ICs) and  $\Sigma(\text{P1}, \text{P2})$  is ref-acyclic, the only solution instance for P is  $D = \{R^1(s, t), R^1(a, null), R^1(c, null)\}$ . However, if there are ref-cycles, the stable models may correspond to a strict superset of the solutions for a peer (c.f. Example 12 in Appendix A). In this case, post-processing that deletes models corresponding to non-minimal "solutions" is necessary.

Under the assumption that we have already computed the (intersection of the) solution instances for the neighbors of  $P$ , the program for  $P$  allows us to compute its solution instances. This generates a recursive process that can be applied because  $\mathcal{G}(P)$  is acyclic. The terminal peers  $P_t$ , i.e. those with no outgoing edges in  $\mathcal{G}(P)$ , will become the base cases for the recursion. If a peer  $P'$  has  $P_t$  as one of its neighbors, the instance  $I_t$  to be used for  $P_t$  in the program for  $P'$  is simply  $D(P_t)$ ,  $P_t$ 's original local instance.

Theorem 2 still holds if peer  $P$ , instead of collecting the intersection  $I_i$  of the solutions of a neighbor  $P_i$ , uses the intersection of the solutions for  $P_i$  restricted to the subschema of  $P_i$  that contains  $P_i$ 's relations that appear in  $\Sigma(P, P_i)$ , which are those  $P$  needs to run its program.

With a solution program for  $P$ , PCAs for a query  $Q$  posed to  $P$  can be obtained by running under the stable model semantics a query program  $\Pi(Q)$  that represents the query in combination with  $\Pi(\mathfrak{P}, P, \mathcal{I}^*)$ .

*Example 8.* (Example 6 continued) In order to obtain the PCAs to the query  $Q_1 : R^1(x, y)$ , asking for the tuples in  $R^1$ , the rule  $ans_1(x, y) \leftarrow R^1(x, y, \mathbf{t}^{**})$  has to be added to  $\Pi(\mathfrak{P}, P_1, \{I_2\})$  (assuming that  $I_2$  is the intersection of the solution instances for  $P_2$ ). The ground  $ans_1$ -atoms in the intersection of all stable models correspond to the PCAs. For the query  $Q_2 : \exists y R^1(x, y)$ , the projection of  $R^1$  on its first attribute, the query rule is  $ans_2(x) \leftarrow R^1(x, y, \mathbf{t}^{**})$ .  $\square$

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  $P'$  queries of the form  $Q : R(\bar{x})$ , where  $R$  is a relation of  $P'$  that appears in  $\Sigma(P, P')$ .  $P$  expects to receive from  $P'$  the set of PCAs to  $Q$ , because they corresponds to the extension of  $R$  in the intersection of solutions for  $P'$ . In order to return to  $P$  the PCAs to its queries, the neighboring peers have to run their own programs (except for the facts, each peer has a fixed solution program that can be used with any query). As before, they will need PCAs from their own neighbors; etc. This recursion will eventually reach peers that have no DEC's (and its local IC's will be satisfied), who will offer answers from their original instances to queries by other peers. Now, propagation of PCAs goes backwards until reaching  $P$ , and  $P$  gets the facts to run its program and obtain the PCAs to the original query.

*Example 9.* (Example 6 continued) Consider local instances  $D(P_1) = \{R^1(a, 2)\}$ ,  $D(P_2) = \{R^2(c, 4), R^2(d, 5)\}$ , and  $D(P_3) = \{R^3(c, 4)\}$ . A user poses the query  $Q_0 : R^1(x, y)$  to  $P_1$ , expecting its PCAs. To run its program,  $P_1$  needs the intersection of the solutions of peer  $P_2$ . So,  $P_1$  sends to  $P_2$  the queries  $Q_1^1 : R^2(x, y)$  and  $Q_2^1 : S^2(x, y)$  (actually,  $P_1$  does not need the latter,  $S^2$  is not relevant to  $P_1$ ). In order to peer-consistently-answer these queries,  $P_2$  needs from  $P_3$  the PCAs to  $Q_3^2 : R^3(x, y)$ . Since  $P_3$  has no neighboring peers, it returns to  $P_2$  the entire extension in its local database of relation  $R^3$ :  $I_3 = D(P_3) = \{R^3(c, 4)\}$  is given to  $P_2$ . Now,  $P_2$  can run its solution program  $\Pi(\mathfrak{P}, P_2, \{I_3\})$ , containing:  $dom(c)$ .  $dom(d)$ .  $\dots$   $R^2(c, 4)$ .  $R^2(d, 5)$ .  $R^3(c, 4)$ .

$$\begin{aligned} R^2(x, y, \mathbf{f}_a) \vee R^3(x, y, \mathbf{f}_a) &\leftarrow R^2(x, y, \mathbf{t}^*), R^3(x, y, \mathbf{t}^*), x \neq null, y \neq null. \\ R^2(x, y, \mathbf{t}^{**}) &\leftarrow R^2(x, y, \mathbf{t}^*), \text{ not } R^2(x, y, \mathbf{f}_a). \\ S^2(x, y, \mathbf{t}^{**}) &\leftarrow S^2(x, y, \mathbf{t}^*), \text{ not } S^2(x, y, \mathbf{f}_a). \end{aligned}$$

We obtain two solutions:  $\{R^2(d, 5)\}$  and  $\{R^2(c, 4), R^2(d, 5)\}$ . So, the intersection of P2's solutions is  $I_2 = \{R^2(d, 5)\}$ . Finally, the program  $\Pi(\mathfrak{P}, P1, \{I_2\})$  given in Example 6 is run. It has only one solution, namely  $\{R^1(a, 2), R^1(d, 5)\}$ . Therefore, the peer consistent answers to  $\mathcal{Q}_0$  are  $(a, 2)$  and  $(d, 5)$ .  $\square$

## 4 Discussion

The domain predicate, *dom*, in the solution programs can always be instantiated in a finite active domain; actually *dom* can be eliminated by adding rules [10].

In the most common PDESs, let us call them the *unrestricted import case*, peers P only import data from other peers they trust more than themselves; and using DEC's  $\Sigma(P)$  of the form (1) or (2) that have only one database predicate in the consequent which belongs to  $\mathcal{R}(P)$ , and all predicates in the antecedent belonging to another peer's schema. In this case, the relevant part of the intersection of the solutions of each neighbor can be obtained as the PCAs to a single conjunctive query, namely the one in the antecedent of the DEC.

**Proposition 1.** For the unrestricted import case of PDES, a peer P with an empty set  $IC(P)$  of local ICs always has a solution instance.  $\square$

This result still holds under rather weak conditions, e.g. if the ICs in  $IC(P)$ : (a) have a consequent that contains at least one database predicate (not a built-in); or (b) if only built-ins appear in the consequent, e.g. **false**, there is a predicate in the antecedent of the IC that does not appear in any of P's DEC's.

If a PDES has no local ICs, then it is easy to see that the solution program is head-cycle free [17], and we obtain

**Proposition 2.** In the unrestricted import case, the solution program for a peer with an empty set of local ICs is equivalent to a non-disjunctive program.  $\square$

The hypothesis on local ICs in this result can be much weakened by assuming that local ICs are of the form identified in [5, 11], that lead to head-cycle free repair programs. Since cautious reasoning from normal logic programs is *coNP*-complete [17], peer consistent answering is in *coNP* in data complexity.

We have assumed that  $\mathcal{G}(\mathfrak{P})$  is acyclic. However, the peers, not being aware of being in a cyclic  $\mathcal{G}(\mathfrak{P})$ , could attempt to do data exchange as described above. In order not to detect an infinite loop, for each query a unique identifier can be created and kept in all the queries that have origin in it.

The assumption of acyclicity of the accessibility graph is quite cautious in the sense that it excludes cases where a reasonable semantics could be given and the logic programs would work correctly, because the cycles in  $\mathcal{G}(\mathfrak{P})$  are not relevant (cf. Example 11 in Appendix A).

We have also assumed that the sets  $IC(P)$  of local ICs of peers P are each ref-acyclic. Even under this assumption, and also with  $\Sigma(P)$  ref-acyclic,  $IC(P) \cup \Sigma(P)$  can have ref-cycles. For example, with  $IC(P1) = \{\forall x(R^1(x) \rightarrow S^1(x))\}$ ,  $\Sigma(P1, P2) = \{\forall x(S^1(x) \rightarrow \exists yR^2(x, y)), \forall x(R^2(x, y) \rightarrow R^1(x))\}$ . There are also cases with an acyclic  $\mathcal{G}(\mathfrak{P})$ , but with ref-cycles in the DEC's, where the logic programming counterpart of the semantics is correct due to the role of the trust relationships (cf. Example 13 in Appendix A).

It becomes clear that it is possible to find more relaxed conditions, both on the accessibility graph and ref-cycles, under which a sensible semantics for

solutions and semantically corresponding logic programs can be given. Also, with general cyclic accessibility graphs, *super peers* [31] could be used to detect cycles and prune certain DEC's, making the graph acyclic if necessary; and then our semantics could be applied.

In [13, 15, 14], the semantics of a PDES is given in terms of epistemic logic. The mappings (our DEC's) are of the form  $cq_i \rightarrow cq_j$ , with  $cq_i$  and  $cq_j$  conjunctive queries over  $P_i$  and  $P_j$ 's schemas, resp. Those DEC's keep the schemas separate. It is implicitly assumed that peers trust themselves less than other peers. The semantics can be applied in the presence of cycles in the accessibility graph.

The treatment of local IC's differs from ours in two ways : (a) A peer that is inconsistent wrt its local IC's is not considered for data exchange, while in our case such a peer may apply a repair semantics, as in CQA. (b) Atoms are imported into a peer by interaction with other peers only if this does not produce a local IC violation. In our case, under the same trust conditions, the data is accepted and the peer applies again a local repair semantics.

In order to answer a query [14], a peer traverses the network eventually collecting at its site all DEC's, IC's 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 [10].

The semantics in [20, 21] coincides with the epistemic semantics in [15]. They 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.

## 5 Conclusions

We have introduced a framework for peer data exchange with trust relationships. Each peer solves its data and semantic conflicts at query time, when querying its own and other peers' data.

Logic programs can be used to specify solutions for a peer and to obtain peer consistent answers. Techniques to partially compute the solution instances can be useful, since we are not interested in them per se, but in the PCAs. Techniques used in CQA, such as magic sets for stable model semantics [19] and identification of predicates that are relevant to queries and constraints, could also be used in this setting, to restrict the number of rules and the amount of data that are needed to run the program [16, 18].

The problem of query evaluation from disjunctive programs is  $\Pi_2^P$ -complete [17], which matches the complexity of PCA. In spite of this, it is possible to identify syntactic classes of PDES's for which peer consistent query answering has a lower complexity, and specifically tailored mechanisms to solve this problem could be developed, as for CQA (cf. [9] for a survey).

The concepts and results presented in this paper smoothly extend the semantics for *local solutions* for a peer as introduced in [7] to the transitive case.

Basically, those local solutions correspond to the neighborhood solutions we introduced above. No general solution programs were presented in [7].

Semantics for PDESs have been introduced and analyzed in [26, 20, 28, 27, 21, 13, 15, 14, 22], but without considering trust relationships. In them, if there is a DEC from P to Q, it is implicitly assumed that P trusts itself less than Q. Also, all the research so far, has concentrated on the unrestricted import case. In our setting, a DEC may also restrict the data that can belong to a peer.

**Acknowledgements:** Research supported by NSERC and a CITO/IBM-CAS Student Internship. L. Bertossi is Faculty Fellow of IBM CAS (Toronto Lab.). Part of this work was done when L. Bertossi was visiting the Database Group at Edinburgh University. Their hospitality is much appreciated.

## References

- [1] Arenas, M., Bertossi, L. and Chomicki, J. Consistent Query Answers in Inconsistent Databases. In *Proc. ACM Symposium on Principles of Database Systems (PODS'99)*. ACM Press, 1999, pp. 68-79.
- [2] Arenas, M., Bertossi, L., Chomicki, J. Answer Sets for Consistent Query Answers. *Theory and Practice of Logic Programming*, 2003, 3(4&5):393-424.
- [3] Baral, C. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press, 2003.
- [4] Barcelo, P. and Bertossi, L. Logic Programs for Querying Inconsistent Databases. In *Proc. Practical Aspects of Declarative Languages (PADL'03)*. Springer LNCS 2562, 2003, pp. 208-222.
- [5] Barceló, P., Bertossi, L. and Bravo, L. Characterizing and Computing Semantically Correct Answers from Databases with Annotated Logic and Answer Sets. In *Semantics in Databases*. Springer LNCS 2582, 2003, pp. 7-33.
- [6] Bertossi, L. and Bravo, L. Consistent Query Answers in Virtual Data Integration Systems. In *Inconsistency Tolerance*. Springer LNCS 3300, 2004, pp. 42-83.
- [7] Bertossi, L. and Bravo, L. Query Answering in Peer-to-Peer Data Exchange Systems. In *Proc. EDBT Workshop on Peer-to-Peer Computing and Databases (P2P&DB'04)*. Springer LNCS 3268, 2004, pp. 476-485.
- [8] Bertossi, L. and Chomicki, J. Query Answering in Inconsistent Databases. In *Logics for Emerging Applications of Databases*. Springer, 2003, pp. 43-83.
- [9] Bertossi, L. Consistent Query Answering in Databases. *ACM Sigmod Record*, June 2006, 35(2):68-76.
- [10] Bravo, L. Handling Inconsistency in Databases and Data Integration Systems. PhD. Thesis, Carleton University, Department of Computer Science, 2007. <http://homepages.inf.ed.ac.uk/lbravo/Publications.htm>
- [11] Bravo, L., Bertossi, L. Semantically Correct Query Answers in the Presence of Null Values. In *Proc. EDBT WS on Inconsistency and Incompleteness in Databases (IIDB'06)*. Springer LNCS 4254, 2006, pp. 336-357.
- [12] Cali, A., Lembo, D. and Rosati, R. On the Decidability and Complexity of Query Answering over Inconsistent and Incomplete Databases. In *Proc. ACM Symposium on Principles of Database Systems (PODS'03)*. ACM Press, 2003, pp. 260-271.
- [13] Calvanese, D., De Giacomo, G., Lenzerini, M. and Rosati, R. Logical Foundations of Peer-To-Peer Data Integration. In *Proc. ACM Symposium on Principles of Database Systems (PODS'04)*. ACM Press, 2004, pp. 241-251.
- [14] Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M. and Rosati, R. Inconsistency Tolerance in P2P Data Integration: An Epistemic Logic Approach. In *Proc. International Symposium on Database Programming Languages (DBPL'05)*. Springer LNCS 3774, 2005, pp. 90-105.

- [15] Calvanese, D., Damaggio, E., De Giacomo, G., Lenzerini, M. and Rosati, R. Semantic Data Integration in P2P Systems. In *Proc. VLDB International Workshop on Databases, Information Systems and Peer-to-Peer Computing (DBISP2P'03)*. Springer LNCS 2944, 2004, pp. 77-90.
- [16] Caniupan, M. and Bertossi, L. Optimizing Repair Programs for Consistent Query Answering. In *Proc. International Conference of the Chilean Computer Science Society (SCCC'05)*. IEEE Computer Society Press, 2005, pp. 3-12.
- [17] Dantsin, E., Eiter, T., Gottlob, G. and Voronkov, A. Complexity and Expressive Power of Logic Programming. *ACM Computing Surveys*, 2001, 33(3): 374-425.
- [18] Eiter, T., Fink, M., Greco, G. and Lembo, D. Efficient Evaluation of Logic Programs for Querying Data Integration Systems. In *Proc. 19th International Conference on Logic Programming (ICLP 03)*. Springer LNCS 2916, 2003, pp. 163-177.
- [19] Faber W., Greco G., Leone N. Magic Sets and their Application to Data Integration". *J. Comp. and Sys. Sciences*, 2007, 73(4):584-609.
- [20] Franconi, E., Kuper, G., Lopatenko, A. and Serafini, L. A Robust Logical and Computational Characterisation of Peer-to-Peer Database Systems. In *Proc. VLDB Workshop on Databases, Information Systems and P2P Computing (DBISP2P'03)*. Springer LNCS 2944, 2004, pp. 64-76.
- [21] Franconi, E., Kuper, G., Lopatenko, A. and Zaihrayeu, I. A Distributed Algorithm for Robust Data Sharing and Updates in P2P Database Networks. In *Proc. EDBT Workshop on Peer-to-peer Computing and Databases (P2P&DB'04)*, Springer LNCS 3268, 2004, pp. 446-455.
- [22] Fuxman, A., Kolaitis, Ph., Miller, R. and Tan, W. Peer Data Exchange. *ACM Trans. Database Systems*, 2006, 31(4): 1454-1498.
- [23] Gelfond, M. and Lifschitz, V. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 1991, 9:365-385.
- [24] Gelfond, M. and Leone, N. Logic Programming and Knowledge Representation: The A-Prolog Perspective. *Artificial Intelligence*, 2002, 138(1-2):3-38.
- [25] Greco, G., Greco, S. and Zumpano, E. A Logical Framework for Querying and Repairing Inconsistent Databases. *IEEE Transactions on Knowledge and Data Engineering*, 2003, 15(6):1389-1408.
- [26] Halevy, A., Ives, Z., Suciu, D. and Tatarinov, I. Schema Mediation in Peer Data Management Systems. In *Proc. International Conference on Data Engineering (ICDE'03)*. IEEE Computer Society, 2003, pp. 505-518.
- [27] Halevy, A., Ives, Z., Madhavan, J., Mork, P., Suciu, D. and Tatarinov, I. The Piazza Peer Data Management System. *IEEE Transactions on Knowledge and Data Engineering*, 2004, 16(7):787-798.
- [28] Kementsietsidis, A., Arenas, M., and Miller, R. Mapping Data in Peer-to-Peer Systems: Semantics and Algorithmic Issues. In *Proc. ACM International Conference on Management of Data (SIGMOD '03)*. ACM Press, 2003, pp. 325-336.
- [29] Kolaitis, Ph. Schema Mappings, Data Exchange, and Metadata Management. In *Proc. of ACM Symposium on Principles of Database Systems (PODS '05)*. ACM Press, 2005, pp. 61-75.
- [30] Levene, M. and Loizou, G. Null Inclusion Dependencies in Relational Databases. *Information and Computation*, 1997, 136(2):67-108.
- [31] Yang, B. and Garcia-Molina, H. Designing a Super-Peer Network. In *Proc. International Conference on Data Engineering (ICDE'03)*, IEEE Computer Society, 2003, p. 49.





$\emptyset$  respectively. Only the first instance corresponds to a solution instance. The second model is the result of cycles through weak negation ( *not* ). The cycle creates the self justification of facts as follows: (i) If we choose  $R^2(a, c, \mathbf{f}_a)$  to be true, then by the second and third rule,  $aux_1(a)$  is false. (ii) Then, the first rule can be satisfied, by making  $R^1(a, b, \mathbf{f}_a)$  true. (iii) By rules five and six,  $aux_2(a)$  is false. (iv) This justifies making  $R^2(a, b, \mathbf{f}_a)$  true, thus, closing the cycle. Notice, that in the whole justification the changes were not determined by inconsistencies.  $\square$

*Example 13 (Ref-cycles, complete and sound program).* (example 12 continued) If we replace  $(P1, same, P2) \in trust$  by  $(P1, less, P2) \in trust$ , rules 1. to 3. of the solution program for peer P1 are:  $dom(a). dom(b). \dots R^1(a, b). R^2(a, c).$

$$\begin{aligned}
R^1(x, y, \mathbf{f}_a) &\leftarrow R^1(x, y, \mathbf{t}^*), \text{ not } aux_1(x), x \neq null. \\
aux_1(x) &\leftarrow R^2(x, null), \text{ not } R^2(x, null, \mathbf{f}_a). \\
aux_1(x) &\leftarrow R^2(x, y, \mathbf{t}^*), \text{ not } R^2(x, y, \mathbf{f}_a), x \neq null, y \neq null. \\
R^1(x, null, \mathbf{t}_a) &\leftarrow R^2(x, y, \mathbf{t}^*), \text{ not } aux_2(x), x \neq null. \\
aux_2(x) &\leftarrow R^1(x, null), \text{ not } R^1(x, null, \mathbf{f}_a). \\
aux_2(x) &\leftarrow R^1(x, y, \mathbf{t}^*), \text{ not } R^1(x, y, \mathbf{f}_a), x \neq null, y \neq null.
\end{aligned}$$

Since P1 trusts more peer P2 than itself, it will modify only its own data. This program computes exactly the solutions for peer P1, i.e.  $\{R^1(a, b)\}$ , even though the DEC's exhibit ref-cycles.  $\square$

## B Satisfaction of ICs in Databases with NULL

In [11] and [10, chapter 4], a precise, uniform, and logic-based definition of IC satisfaction in SQL databases is given. Loosely speaking, in a commercial DBMS like IBM DB2, a constraint is satisfied if any of the relevant attributes in it has *null* or the constraint is satisfied in the traditional way (i.e. as first-order satisfaction with *null* treated as any other constant).

*Example 14.* Consider a foreign-key constraint  $P[A, B] \subseteq R[A, B]$  and  $D$ :

$P$	$A$	$B$	$C$
	$a$	$5$	$d$
	$b$	<i>null</i>	$a$

$R$	$A$	$B$
	$a$	$5$
	$a$	$2$

DBMSs implement the so-called simple semantics of the SQL standard for satisfaction of ICs, according to which the database  $D$  above satisfies the foreign-key constraint. This is because, for every tuple  $\bar{t}$  in  $P$ , if  $\bar{t}[A]$  and  $\bar{t}[B]$  are different from *null*, there is a tuple  $\bar{t}'$  in  $R$  such that  $\bar{t}[A, B] = \bar{t}'[A, B]$ . In this case, the attributes that are relevant for checking the satisfaction of the IC are  $A$  and  $B$  from both  $P$  and  $R$ .

Now, the version in first-order logic of the IC is a universal integrity constraint:  $\psi : \forall xyz (P(x, y, z) \rightarrow R(x, y))$ . Since  $x$  and  $y$  appear twice in  $\psi$ , by Definition 8,  $\mathcal{A}(\psi) = \{P[1], R[1], P[2], R[2]\}$ . We can see that the value for  $C$  is not relevant to check the satisfaction of the constraint, which makes sense, because we only want to make sure that the values in the first two attributes in  $P$  also appear in  $R$ .

Now, if we try to insert tuple  $(c, d, null)$  into  $P$ , the DBMS will reject the tuple since none of the attributes that are relevant for checking the constraint are *null*, and there is no tuple  $(c, d)$  in  $R$ .  $\square$

For the formal definition of IC satisfaction, we need some concepts. Given a set of attributes  $\mathcal{A}$  and a predicate  $P \in \mathcal{R}$ , we denote by  $P^{\mathcal{A}}$  the predicate  $P$  restricted to (or projected onto) the attributes in  $\mathcal{A}$ .  $D^{\mathcal{A}}$  denotes the database  $D$  with all its database atoms projected onto the attributes in  $\mathcal{A}$ , i.e.,  $D^{\mathcal{A}} = \{P^{\mathcal{A}}(\Pi_{\mathcal{A}}(\bar{t})) \mid P(\bar{t}) \in D\}$ , where  $\Pi_{\mathcal{A}}(\bar{t})$  is the projection on  $\mathcal{A}$  of tuple  $\bar{t}$ .

**Definition 11.** A constraint  $\psi$  of the form:

$$\forall \bar{x} \left( \bigwedge_{i=1}^n P_i(\bar{x}_i) \longrightarrow \exists \bar{z} \left( \bigvee_{j=1}^m Q_j(\bar{y}_j, \bar{z}_j) \vee \varphi \right) \right), \quad (3)$$

is satisfied in the database instance  $D$ , denoted  $D \models_N \psi$ , iff  $D^{\mathcal{A}(\psi)} \models \psi^N$ , where  $\psi^N$  is

$$\forall \bar{x} \left( \bigwedge_{i=1}^m P_i^{\mathcal{A}(\psi)}(\bar{x}_i) \rightarrow \left( \bigvee_{v_j \in \mathcal{V}(\psi)} v_j = null \vee \bigvee_{j=1}^n Q_j^{\mathcal{A}(\psi)}(\bar{y}_j) \vee \varphi \right) \right), \quad (4)$$

and  $\bar{x} = \cup_{i=1}^m \bar{x}_i$ . Here,  $D^{\mathcal{A}(\psi)} \models \psi^N$  refers to classical first-order satisfaction, where *null* is treated as any other constant in the domain.  $\square$

When we use  $D \models \psi$  in the main body of this paper, we are really mean  $D \models_N \psi$ , as just defined. Notice that UICs, RICs, UDECs and RDECs are special cases of Formula (3). We can see from Definition 11 that there are basically two cases for constraint satisfaction: (a) If *null* is in any of the relevant attributes in the antecedent, then the constraint is satisfied. (b) If *null* does not appear in the relevant attributes, then the second disjunct in the consequent of formula (4) has to be checked, i.e., the consequent of the IC restricted to the relevant attributes. This can be done as usual, treating *null* as any other constant.

*Example 15.* (example 14 continued) In order to check if  $D \models_N \forall xyz (P(x, y, z) \rightarrow R(x, y))$ , we need to determine if  $D^{\mathcal{A}(\psi)} \models \forall xy (P^{\mathcal{A}(\psi)}(x, y) \rightarrow (x = null \vee y = null \vee R^{\mathcal{A}(\psi)}(x, y)))$ , with:

$D^{\mathcal{A}(\psi)} :$	<table style="border-collapse: collapse; display: inline-table;"> <tr><th style="border: none;"><math>P^{\mathcal{A}_1}</math></th><th style="border: none;"><math>A</math></th><th style="border: none;"><math>B</math></th></tr> <tr><td style="border: none;"></td><td style="border: none; text-align: center;"><math>a</math></td><td style="border: none; text-align: center;"><math>5</math></td></tr> <tr><td style="border: none;"></td><td style="border: none; text-align: center;"><math>b</math></td><td style="border: none; text-align: center;"><i>null</i></td></tr> </table>	$P^{\mathcal{A}_1}$	$A$	$B$		$a$	$5$		$b$	<i>null</i>	<table style="border-collapse: collapse; display: inline-table;"> <tr><th style="border: none;"><math>R^{\mathcal{A}_1}</math></th><th style="border: none;"><math>D</math></th><th style="border: none;"><math>E</math></th></tr> <tr><td style="border: none;"></td><td style="border: none; text-align: center;"><math>a</math></td><td style="border: none; text-align: center;"><math>d</math></td></tr> </table>	$R^{\mathcal{A}_1}$	$D$	$E$		$a$	$d$
$P^{\mathcal{A}_1}$	$A$	$B$															
	$a$	$5$															
	$b$	<i>null</i>															
$R^{\mathcal{A}_1}$	$D$	$E$															
	$a$	$d$															

For  $x = a$  and  $y = 5$ ,  $D^{\mathcal{A}_1} \models P^{\mathcal{A}(\psi)}(a, 5)$ , but none of them is *null*, therefore we need to check if  $D^{\mathcal{A}(\psi)} \models R^{\mathcal{A}(\psi)}(a, 5)$ . This is true, therefore the constraint is satisfied for  $x = a$  and  $y = 5$ . Now, for  $x = b$  and  $y = null$ ,  $D^{\mathcal{A}(\psi)} \models P^{\mathcal{A}(\psi)}(b, null)$ , and since  $y = null$ , the constraint is satisfied.

If we add tuple  $P(c, d, null)$  to  $D$ , it would become inconsistent with respect to constraint (a), because  $D^{\mathcal{A}(\psi)} \not\models (P^{\mathcal{A}(\psi)}(c, d) \rightarrow (c = null \vee d = null \vee R^{\mathcal{A}(\psi)}(c, d)))$ .  $\square$

The semantics of IC satisfaction in Definition 11 coincides, extends, and puts on a solid logical ground the notion of IC satisfaction as implemented in DBMSs that follow the SQL standard. More details can be found in [10, chapter 4].