Answer Sets for Consistent Query Answering in Inconsistent Databases

Marcelo Arenas*
Pontificia Universidad Catolica de Chile, Departamento de Ciencia de Computacion, Santiago, Chile. marenas@ing.puc.cl

Leopoldo Bertossi
Carleton University, School of Computer Science, Ottawa, Canada. bertossi@scs.carleton.ca

Jan Chomicki
State University of New York at Buffalo, Dept. of Computer Science and Engineering, Buffalo, NY, USA. chomicki@cse.buffalo.edu

Abstract

A relational database is *inconsistent* if it does not satisfy a given set of integrity constraints. Nevertheless, it is likely that most of the data in it is consistent with the constraints. In this paper we apply logic programming based on answer sets to the problem of retrieving consistent information from a possibly inconsistent database. Since consistent information persists from the original database to every of its minimal repairs, the approach is based on a specification of database repairs using *disjunctive logic programs with exceptions*, whose answer set semantics can be represented and computed by systems that implement stable model semantics. These programs allow us to declare persistence by defaults and repairing changes by exceptions. We concentrate mainly on logic programs for binary integrity constraints, among which we find most of the integrity constraints found in practice.

1 Introduction

Integrity constraints (IC) capture an important normative aspect of every database application, whose aim is to guarantee the consistency of its data. However, it is very difficult, if not impossible, to always have a consistent database instance. Databases may become inconsistent with respect to a given set of integrity constraints. This may happen due, among others, to the following factors: (1) Certain ICs cannot be expressed/maintained by existing DBMSs. (2) Transient inconsistencies caused by the inherent non-atomicity of database transactions. (3) Delayed updates of a datawarehouse. (4) Integration of heterogeneous databases, in particular with duplicated information. (5) Inconsistency with respect to *soft* integrity constraints,

* Current address: University of Toronto, Department of Computer Science, Toronto, Canada. marenas@cs.toronto.edu

where transactions in violation of their conditions are not prevented from executing. (6) Legacy data on which one wants to impose semantic constraints. (7) The consistency of the database will be restored by executing further transactions. (8) User constraints than cannot be checked or maintained (9) No permission to restore consistency. (10) Inconsistent information can be useful. (11) Restoring consistency can be a complex and non deterministic process.

An inconsistent database may be the only source of data, and we may still want or need to use it. In these cases, one faces the important problem of characterizing and retrieving the *consistent* information from the database. Furthermore, the database can still provide us with correct answers to certain queries, making the problem of determining what kinds of queries and query answers are consistent with the integrity constraints a worthwhile effort. These problems have been studied in the context of relational databases and recent publications deal with some of the issues that arise from trying to define and retrieve consistent information from an inconsistent relational database.

More specifically, in order to approach the problem of *consistent query answering* (CQA), in (Arenas *et al.*, 1999), a model theoretic definition of consistent answer to a query posed to an inconsistent database was introduced. The notion is based on the minimal repairs of the inconsistent database: an answer $\bar{t}$ to a first-order query $Q(\bar{x})$ is *consistent* if it is an answer to the same query in every minimal repair of the database. A repair must be minimal in the sense that the set of inserted or deleted tuples (to restore inconsistency) is minimal under set inclusion.

A possible computational mechanisms for retrieving consistent answers is as follows: Given a first-order query Q and an inconsistent database instance r, instead of explicitly computing all the repairs of r and querying all of them, a new query $T(Q)$ is computed and posed to r, the only available database. The answers to the new query are expected to be the consistent answers to Q. Iterative query transformation operators were introduced and analyzed with respect to soundness, completeness and termination in (Arenas *et al.*, 1999; Celle and Bertossi, 2000).

Nevertheless, the query rewriting approach has some limitations. The operators introduced in (Arenas *et al.*, 1999; Celle and Bertossi, 2000) work for some particular classes of queries and constraints, e.g. they cannot be applied to disjunctive or existential queries.

Furthermore, what we have so far is a *semantic*, model based definition of consistent query answer, based on database repairs, plus a computational mechanism. Nevertheless, we do not have a specification of the database repairs in a logical language of the class of all the repairs of a given database instance relative to a fixed set of ICs. Such a description is natural and useful if we want to express the explicit or implicit properties that are shared by all database repairs, e.g. consistent answers, without constructing those repairs explicitly. Actually, as shown in (Arenas *et al.*, 2001), it is easy to find situations where there is an exponential number of database repairs wrt the original database instance. In consequence, the specification would be a compact way of representing the whole class of database repairs.

From such a specification, say *Spec*, we could: (1) Reason from *Spec*. (2) Consistently answer queries by asking: *Spec* $\models Q(\bar{x})$?. (3) Obtain the intended models of *Spec* that should correspond to the database repairs. (4) Derive algorithms for consistent query answering. (5) Analyze complexity issues related to consistent query answering. In this paper, we are motivated mainly by the possibility of retrieving consistent answers to general first-order queries, extending the possibilities we developed in (Arenas *et al.*, 1999). We are also interested in the possibility of obtaining the models of the specification, i.e. the database repairs. Having the repairs explicitly available, allows us to analyze different ways to restore the consistency of the database. That is, a mechanism for computing database repairs could be used for conflict resolution.

Notice that consistent answers are non-monotonic in the sense that adding information to the original database, may cause loosing previous consistent answers. In consequence, a non-monotonic semantics for *Spec* (or their consequences) should be expected.

In this direction, in (Arenas *et al.*, 2000b) a specification in annotated predicate calculus of the database repairs was presented. That specification was used to derive algorithms for consistent query answering and to obtain some complexity results. As expected, the database repairs correspond to certain minimal models of the specifications. This approach is based on a non-classical logic.

In the same spirit, in (Arenas *et al.*, 2000a) an alternative specification of the database repairs was presented. It is based on extended disjunctive logic programs with an answer sets semantics. The the database repairs correspond to the intended models of the program. This paper extends the results presented in (Arenas *et al.*, 2000a), addressing several new issues.

Our main contributions are: (1) The introduction and application of extended disjunctive logic programs with exceptions to the specification of database repairs and consistent query answering. (2) A detailed analysis of the correspondence between answer sets and database repairs for binary integrity constraints. (3) The use of the specification of database repairs to retrieve consistent answers to general first order queries. (4) An analysis of the applicability of the disjunctive well-founded semantics to consistent query answering. (5) Application of the *DLV* system (Eiter *et al.*, 1998) to obtain database repairs and consistent answers. (6) The use of weak constraints to capture database repairs based on minimal *number* of changes. (7) Extensions to of the methodology to more general universal constraints and to referential integrity constraints.

This paper is structured as follows. In section 2 we introduce the notions of database repair and consistent answer to a query, and the query language. Sections 3 introduces extended disjunctive logic programs with exceptions. In section 4 we show a complete, but informal example that describes the main ideas behind our approach to consistent query answering and specification of database repairs by means of logic programs. In section, 5 we present the repair programs in their general form for binary integrity constraints. In section 6 we analyze the well-founded interpretation as an approximation to the set of consistent answers, and we identify cases where it provides the exact solution. In section 7, we discuss how to evaluate the queries. In section 8 we show some examples using the *DLV* system to obtain database repairs and consistent answers. In section 9 we shown how database repairs based on minimal number of changes can be specified by introducing weak constraints in our repair programs. In section 10 we show by means of examples how to extend the methodology to deal with more general universal constraints and referential integrity constraints. In section 11 we discuss related work, mention some open issues, and draw conclusions.

2 Consistent Query Answers

A relational database instance r can be represented by a finite set of finite relations whose names are part of a database schema. A database schema can be represented by a typed language $\mathcal{L}$ of first-order predicate logic, that contains a finite set of sorted predicates and a fixed infinite set of constants D. The language contains a predicate for each database relation and the constants in D correspond to the elements in the database domain, that we will also denote by D. In consequence, every database instance has an infinite domain D.

The active domain of a database instance r is the set of those elements of D that explicitly appear in r. The active domain is always finite and we denote it by $Act(r)$. We may also have a set of built-in (or evaluable) predicates, like equality, order relations, arithmetical relations, etc. In this case, we have the language $\mathcal{L}$ possibly extended with these predicates. In all database instances each of these predicates has a fixed and possibly infinite extension. Of course, since we defined database instances as finite sets of ground atoms, we are not considering these built-in atoms as members of database instances.

In addition to the database schema and instances, we may also have a set of integrity constraints IC expressed in a language $\mathcal{L}$. These are first-order formulas which the database instances are expected to satisfy.

If a database instance r satisfies IC, what is denoted by $r \models IC$, we say that it is consistent (wrt IC), otherwise we say it is inconsistent[1]. In any case, we will assume from now on that IC is consistent as a set of first-order sentences.

The original motivation in (Arenas *et al.*, 1999) was to consistently answer first-order queries. We will call them *basic queries* and are defined by the grammar

$$B \ ::= \ Atom \mid B \wedge B \mid \neg B \mid \exists x \ B.$$

One way of explicitly asking at the object level about the consistent answers to a first-order query consists in introducing a new logical operator $\mathcal{K}$, in such a way that $\mathcal{K}\varphi(\bar{x})$, where $\varphi(\bar{x})$ is a basic query, asks for the values of $\bar{x}$ that are consistent answers to $\varphi(\bar{x})$ (or whether φ is consistently true, i.e. true in all repairs, when φ is a sentence). The $\mathcal{K}$-*queries* are similarly defined:

$$A \ ::= \ \mathcal{K}B \mid A \wedge A \mid \neg A \mid \exists x. \ A.$$

In this paper we will concentrate mostly on answering *basic* $\mathcal{K}$-*queries* of the form

[1] It is natural to see a relational database instance also as a structure for interpretation of language $\mathcal{L}$, in consequence, the notation $r \models IC$ makes sense.

$\mathcal{K}B$, where B is a basic query, but in section 7 we sketch how to handle general $\mathcal{K}$-queries.

Definition 1
(a) (Arenas *et al.*, 1999) Given a database instance r (seen as a set of ground atomic formulas) and a set of integrity constraints, IC, a *repair* of r wrt IC is a database instance r', over the same schema and domain, that satisfies IC and such that $r\Delta r'$, the symmetric difference of r and r', is minimal under set inclusion.
(b) (Arenas *et al.*, 1999) A tuple $\bar{t}$ is a *consistent answer* to a first-order query $Q(\bar{x})$, or equivalently, an answer to the query $\mathcal{K}Q(\bar{x})$, in a database instance r iff $\bar{t}$ is an answer to query $Q(\bar{x})$ in every repair r' of r wrt IC. In symbols:

$$r \models \mathcal{K}Q[\bar{t}] \quad \Longleftrightarrow \quad r' \models Q[\bar{t}] \quad \text{for every repair} \quad r' \quad \text{of} \quad r.$$

(c) If Q is a general $\mathcal{K}$-query, then $r \models Q$ is defined as usual, (b) being the base case. □

Example 1
Assume there is the following database instance *Salary*

Salary	*Name*	*Amount*
	V.Smith	5000
	V.Smith	8000
	P.Jones	3000
	M.Stone	7000

and *FD* is the functional dependency $Name \rightarrow Amount$, meaning that *Name* functionally determines *Amount*, that is violated by the table *Salary*, actually by the tuples with the value *V.Smith* in attribute *Name*.

The possible repairs are

$Salary_1$	*Name*	*Amount*
	V.Smith	5000
	P.Jones	3000
	M.Stone	7000

$Salary_2$	*Name*	*Amount*
	V.Smith	8000
	P.Jones	3000
	M.Stone	7000

Here, $\bar{t}_3 = (P.Jones, 3000)$ is a consistent answer to the query $Salary(\bar{x})$, i.e. $r \models \mathcal{K}\ Salary(x,y)[(P.Jones, 3000)]$, but $r \not\models \mathcal{K}\ Salary(x,y)[(V.Smith, 8000)]$. It also holds $r \models \mathcal{K}\ (Salary(V.Smith, 5000) \vee Salary(V.Smith, 8000))$, and $r \models \mathcal{K}\ \exists X(Salary(V.Smith, X) \wedge X > 4000)$. □

Notice that the definition of consistent query answer depends on our definition of repair. In section 9.1 we will consider an alternative definition based on minimal *number* of changes instead of minimal *set* of changes.

Computing consistent answer through generation of all possible repairs is not a natural and feasible alternative (Arenas *et al.*, 2001). Instead, an approach based on querying the available, although inconsistent, database is much more natural. In (Arenas *et al.*, 1999) a query rewriting iterative operator T was introduced, that applied to a given query Q produces a new query $T(Q)$ whose (ordinary) answers in an instance r are the consistent answers to Q in r.

Example 2
(example 1 continued) The functional dependency FD can be expressed by means of the formula $\forall xyz\ (\neg Salary(x,y)\ \vee\ \neg Salary(x,z)\ \vee\ y=z)$. Given the query $Q(x,y):\ Salary(x,y)$, the consistent answers are: $(P.Jones, 3000)$, $(M.Stone,\ 7000)$, but not $(V.Smith,\ 5000)$, $(V.Smith,\ 8000)$.

The consistent answers can be obtained by means of the transformed query $T(Q(x,y)) := Salary(x,y)\ \wedge\ \forall z\ (\neg Salary(x,z)\ \vee\ y=z)$ posed to the given instance. □

This rewriting approach is not complete for disjunctive or existential queries, like $\exists Y\ Salary(V.Smith, Y)$. We would like to be able to obtain consistent answers to basic $\mathcal{K}$-queries at least.

3 Logic Programs with Exceptions

Logic programs with exceptions (LPEs) were first introduced in (Kowalski and Sadri, 1991). They are built with definite extended clauses, that is, with clauses where the (non-disjunctive) head and the body are literals (with classical negation) and weak negation (or negation as failure) may appear in the bodies (Gelfond and Lifschitz, 1991). Among those clauses, in a LPE there are positive *default* rules, that is clauses with positive heads, whose conclusions can be overridden by conclusions derived from *exception* rules, which are clauses with negative heads.

The idea is that exceptions have priority over defaults. To capture this intuition, a new semantics is introduced, *E-answer sets.*

Example 3
As an example of LPE, we present here a program Π that cleans a database instance $r(X,Y)$ from tuples participating in the violation of the $FD\ \ X \rightarrow Y$. We start by introducing a new predicate $r'(X,Y)$ that will store the tuples in the clean version of the database.

1. Default rule: $r'(X,Y) \longleftarrow r(X,Y)$.
   It says that every tuple (X,Y) passes from r to r'.
2. Negative exception rule: $\neg r'(X,Y) \longleftarrow r(X,Y), r(X,Z), \mathit{not}\ Y=Z$.
   It says that tuples (X,Y) where X is associated to different values are not accepted in the clean table.
3. Facts: the contents of r plus $X = X \longleftarrow$.

Intuitively, rule 2. should have a priority over rule 1. □

The semantics of the program should give an account of the priorities; they should be reflected in the intended models of the program.

The semantics is constructed as follows. First, instantiate the program Π in the database domain, making it ground. Now, let S be a set of ground literals $S = \{L, \dots\}$. In example 3, S could be something like $\{r(a,b), \neg r(a,b), r'(a,b), \neg r'(a,b), a = b, \neg a = b, ...\}$. This S is a candidate to be a model, a guess to be verified, and accepted if properly justified.

Next, generate a new set of ground rules ${}^S\Pi$ according to the following steps:

(a) Delete every rule in Π containing *not* L in the body, with $L \in S$.
(b) Delete from the remaining clauses every condition *not* L in the body, when $L \notin S$.
(c) Delete every rule having a positive conclusion L with $\neg\, L \in S$.

The result is a ground extended logic program without *not*. Now, we say that S is an *e-answer set* of the original program if S is the smallest set of ground literals, such that

- For any clause $L_0 \longleftarrow L_1, \cdots, L_m$ in ${}^S\Pi$, if $L_1, \cdots, L_m \in S$, then $L_0 \in S$.
- If S contains two complementary literals, then S is the set of all literals.

The e-answer sets are the intended models of the original program. In example 3, the only e-answer set is essentially the cleaned instance, what is reflected in the existence of only one e-answer set, where the extension of r' is that of r, but without its conflicting tuples. Notice that this instance is not a repair, but the intersection of all repairs.

Above, (a) - (b) are as in the *answer sets semantics* for extended logic programs (Gelfond and Lifschitz, 1991), but now (c) gives an account of exceptions.

In order to specify database repairs, we need to extend the LPEs as presented in (Kowalski and Sadri, 1991) in order to accommodate also negative defaults. i.e. defaults with negative conclusions that can be overridden by positive exceptions, and extended disjunctive exceptions, i.e. rules of the form

$$L_1 \vee \cdots \vee L_k \;\longleftarrow\; L_{k+1}, \dots, L_r, not\ L_{r+1}, \dots, not\ L_k,$$

where the L_is are literals. In our application scenario we will need disjunctive exceptions rules, but not disjunctive defaults.

The e-answer semantics is extended as follows. The ground program is pruned according to a new version of the rule (c) we had in the previous section:

(c') Delete every (positive) default having a positive conclusion L, with $\neg L \in S$; and every (negative) default having a negative conclusion $\neg L$, with $L \in S$.

Applying (a), (b) and (c') to the ground program, we are left with a ground disjunctive logic program without *not*. If the candidate set of literals S belongs to $\alpha({}^S\Pi)$, the set of minimal models of program ${}^S\Pi$, then we say that S is an *e-answer set*.

As described in (Kowalski and Sadri, 1991) for non-disjunctive programs with negative exceptions, there is a one to one correspondence between the e-answer sets of a disjunctive logic program with exceptions and the answer sets of an extended disjunctive logic program (Gelfond and Lifschitz, 1991). It is easy to show a general program transformation that establishes this correspondence. We will give this transformation in section 5.1, remark 1, but only for the programs we will use. As shown in (Gelfond and Lifschitz, 1991), the extended disjunctive program obtained after the transformation can be transformed in its turn into a disjunctive normal program (without classical negation) with a stable model semantics.

4 Logic Programs for CQA: An Example

Now, we will use logic programs with exceptions for answering basic $\mathcal{K}$-queries.

Given a set of ICs and an inconsistent database instance r, the first step consists of writing a logic program $\Pi(r)$ having as e-answer sets the repairs of the original database instance. For this purpose, we will use the *disjunctive logic program with exceptions* (DLPEs) introduced in section 3. With those *repair programs* we can specify the class of all the repairs of a given inconsistent database instance. Next, if a first-order query is posed with the intention of retrieving all and only its consistent answers, then a new logic program, a *query program*, that expresses the query, is run together with the repair program. In this way we can pose and consistently answer queries we can not handle with the query rewriting approach presented in (Arenas *et al.*, 1999).

4.1 The repair program

Program $\Pi(r)$ captures the fact that when a database instance r is repaired most of the data persists, except for some tuples. In consequence, default rules are introduced: everything persists from the instance r to the repairs. It is also necessary to introduce exception rules: everything persists, as stated by the defaults, unless the ICs are violated and have to be satisfied.

We illustrate repair programs by means of an example.

Example 4
Consider an inclusion dependency $IC:\ \forall xy\ (P(x,y) \rightarrow Q(x,y))$, stating that every tuple in table P has to be also a tuple in table Q, and the inconsistent database instance $r = \{P(a,b), Q(b,c)\}$. The repairs of the database are specified by a DLPE $\Pi(r)$ obtained as follows: we introduce new predicates P', Q' corresponding to the repaired versions of the original tables, plus the following program clauses:

1. *Facts*: $P(a,b),\ Q(b,c)$.

2. *Triggering exception*: $\neg P'(X,Y) \vee Q'(X,Y) \ \leftarrow\ P(X,Y),\ \mathit{not}\ Q(X,Y)$.

   This rule gives an account of the first two possible steps leading to a repair of the DB: in order to "locally" repair the (in this case, single) IC, either eliminate (X,Y) from P or insert (X,Y) into Q. The semantics of these DLPEs gives the disjunction an exclusive interpretation. We use weak negations in the body of the last rule in order to give an account of the closed world assumption.

3. *Stabilizing exceptions*: $Q'(X,Y) \leftarrow P'(X,Y); \quad \neg P'(X,Y) \leftarrow \neg Q'(X,Y)$.

   This rule states that eventually the IC has to be satisfied in the repairs, this kind of exceptions are necessary if there are interacting ICs and local repairs alone are not sufficient. The contrapositive is introduced for technical reasons.

4. *Persistence defaults*: $P'(X,Y) \leftarrow P(X,Y); \quad Q'(X,Y) \leftarrow Q(X,Y)$;

   $$\neg P'(X,Y) \leftarrow \textit{not } P(X,Y); \quad \neg Q'(X,Y) \leftarrow \textit{not } Q(X,Y).$$

   This means that, by default, everything from r is put into r' and nothing else.

Rules 2. and 3. have priority over rule 4.

It is possible to verify that the e-answer sets of the program are the expected database repairs: $\{\underline{\neg P'(a,b)}, \neg Q'(a,b), Q'(b,c), P(a,b), Q(b,c)\}$; $\{P'(a,b), \underline{Q'(a,b)}, Q'(b,c), P(a,b), Q(b,c)\}$. The underlined literals represent the deletion of $P(a,b)$ in one repair and the insertion of $Q(a,b)$, in the other, resp. Notice that $P(a,b)$ and $Q(b,c)$ do not change, because there is no rule able to do that. □

4.2 The query program

In order to obtain the consistent answers to a first-order query $\varphi(\bar{x})$, this query is translated into a stratified logic program $\Pi(\varphi)$ with new query goal $\mathit{Query}(\bar{x})$ using a standard methodology (Lloyd, 1987). The predicates appearing in the query program $\Pi(\varphi)$ will be the repaired, primed versions of the original database tables, more precisely the set of consistent answers to $\varphi(\bar{x})$ will be the set $\{\bar{t} \mid \mathit{Query}(\bar{t}) \in M$, for every e-answer set M of $\Pi(r) \cup \Pi(\varphi)\}$.

Example 5
(example 4 continued) The query $\varphi_1(x): P(x,a) \vee Q(a,x)$, asking for consistent values of x in the database instance, can be transformed in the following query program $\Pi(\varphi_1)$:

$$\begin{array}{rcl} \mathit{Query}(X) & \longleftarrow & P'(X,a) \\ \mathit{Query}(X) & \longleftarrow & Q'(a,X) \end{array}$$

In order to obtain consistent answers it is necessary to evaluate the query goal $\mathit{Query}(\bar{x})$ wrt the combined program $\Pi(\varphi_1) \cup \Pi(r)$. Each of the e-answer sets of the combined program will contain a a set of ground *Query*-atoms. Those *Query*-atoms (rather their tuple arguments) that are present simultaneously in all the e-answer sets will be the consistent answers to the original query.

As another example of query, consider $\varphi_2(y): \exists x Q(x,y)$. In order to obtain the consistent answers, we keep $\Pi(r)$, but we run it in combination with the new query program $\Pi(\varphi_2)$:

$$\mathit{Query}(Y) \longleftarrow Q'(X,Y).$$ □

Queries like the ones in the previous example can not be handled by the rewriting methodology presented in (Arenas *et al.*, 1999).

5 DLPEs for Binary Integrity Constraints

In this section we will introduce the DLPEs for specifying database repairs, and will give a careful analysis of those programs for consistent query answering wrt binary integrity constraints (BICs).

We represent integrity constraints in the *standard format* (Arenas *et al.*, 1999)

$$\bigvee_{i=1}^{n} p_i(\bar{x}_i) \vee \bigvee_{i=1}^{m} \neg q_i(\bar{y}_i) \vee \varphi, \tag{1}$$

where the p_i, q_i are atomic database formulas. where φ is a first-order formula containing only built-in predicates[2] only and whose variables are among the $\bar{x}_i, \bar{y}_i$s; and there is an implicit universal quantification in front.

Binary integrity constraints are in this standard format, but they have the restricted syntactic form

$$\forall \bar{x}\ (L_1 \vee L_2 \vee \varphi), \tag{2}$$

where L_1, L_2 are database literals associated to database tables, i.e. atomic or negations of atomic formulas whose predicates are part of the database schema. That is, in (1) the conditions $0 \leq n, m$, $1 \leq n+m \leq 2$ hold.

We represent BICs in this form, as a particular case of the general standard format, because it is easy to generalize the program we will give next for BICs to the general case. Nevertheless, in this paper we will concentrate mainly on BICs.

BICs with one database literal plus possibly a formula containing built-ins are called *unary ICs*. Several interesting classes of ICs (Abiteboul *et al.*, 1995) can be represented by BICs: (a) range constraints, e.g. $P(x,y) \rightarrow x > 5$; (b) non existential inclusion constraints (example 4), (b) functional dependencies (examples 1, 2), etc. Nevertheless, for referential ICs, like in $P(x,y) \rightarrow \exists z Q(x,z)$, we need existential quantifiers or Skolem functions (Fitting, 1996). We briefly consider existential inclusion dependencies in section 10.

5.1 Finite domain databases

In this section we will first analyze the case of finite domain databases. That is, in this section, we will momentarily depart from our initial assumption that databases have an infinite domain D (see section 1). The reason is that in the general case, we will be interested in *domain independent* BICs, for which only the active domain is relevant (and finite).

5.1.1 The change program

In the following, in order to analyze the behavior of DLPEs for BICs, we will separate the default rules of the programs from the other rules that represent exceptions.

[2] Built-in predicates have a fixed extension in every database, in particular, in every repair; so they are not subject to repairs. More details can be found in (Arenas *et al.*, 1999).

We will concentrate first on the program without the defaults, that we will denote by $\Pi_\Delta(r)$. This is the part of the program responsible for the changes.

Splitting the program in this way makes its analysis easier. In addition, we will see that keeping $\Pi_\Delta(r)$, but using different form of defaults, we can capture different kinds of repairs. In this section, we will use defaults (def. 6) that lead to our notion of repair based on minimal set of changes (def. 1). In section 9.1, we will use other defaults that lead to repairs based on minimal number of changes.

Definition 2

Given a set of BICs IC and an instance r, the change DLPE, $\Pi_\Delta(r)$, contains the following rules:

1. Facts:
   (a) For every atomic database formula $p(\bar{a})$ such that $r \models p(\bar{a})$, the fact $p(\bar{a})$.
   (b) For every a in D, the fact $dom(a)$.
2. For every IC of the form (1), the triggering rule
$$\bigvee_{i=1}^{n} p_i'(\bar{X}_i) \vee \bigvee_{i=1}^{m} \neg q_i'(\bar{Y}_i) \longleftarrow$$
$$dom(\bar{X}_1, \ldots, \bar{X}_n), \ \bigwedge_{i=1}^{n} not\ p_i(\bar{X}_i), \ \bigwedge_{i=1}^{m} q_i(\bar{Y}_i), \quad not\ \varphi.$$
3. For every $1 \leq j \leq n$, the stabilizing rule
$$\bigvee_{i=1}^{j-1} p_i'(\bar{X}_i) \vee \bigvee_{i=j+1}^{n} p_i'(\bar{X}_i) \vee \bigvee_{i=1}^{m} \neg q_i'(\bar{Y}_i) \longleftarrow$$
$$dom(\bar{X}_1, \ldots, \bar{X}_n, \bar{Y}_1, \ldots, \bar{Y}_m), \ \neg p_j'(\bar{X}_j), \quad not\ \varphi.$$
   For every $1 \leq j \leq m$, the stabilizing rule
$$\bigvee_{i=1}^{n} p_i'(\bar{X}_i) \vee \bigvee_{i=1}^{j-1} \neg q_i'(\bar{Y}_i) \vee \bigvee_{i=j+1}^{m} \neg q_i'(\bar{Y}_i) \longleftarrow$$
$$dom(\bar{X}_1, \ldots, \bar{X}_n, \bar{Y}_1, \ldots, \bar{Y}_m), \ q_j'(\bar{Y}_j), \quad not\ \varphi.$$

In these rules, $dom(\bar{X}_1, \ldots, \bar{X}_n)$ is an abbreviation for the conjunction of cases of membership to dom of all the components in the $\bar{X}_i$s. Of course, depending on the syntax, it may be necessary to unfold the formula φ appearing in the bodies in additional program rules, but φ will usually be a conjunction of literals. □

Notice that the sets of rules 2. and 3. in example 4 have the form of these general triggering and stabilizing rules, respectively.

It is always the case that for BICs, the stabilizing rules in $\Pi_\Delta(r)$ do not contain disjunctions in the heads. This can also be seen in example 4. Only triggering rules are properly disjunctive.

Definition 3

A *model* of a DLPE, Π, is a set of ground literals, S, that does not contain complementary literals and satisfies Π in the usual logical sense, but with weak negation interpreted as not being an element of S. □

Definition 4

Given a model S of $\Pi_\Delta(r)$, we define the database instance corresponding to S by

$$I(S) = \{p(\bar{a}) \mid p'(\bar{a}) \in S\} \ \cup \ \{p(\bar{a}) \mid p(\bar{a}) \in S \text{ and } \neg p'(\bar{a}) \notin S)\}.$$ □

Notice that, for a given model S of the change program, $I(S)$ merges in one new instance of the schema all the positive primed tuples with all the old, non primed tuples that persisted, i.e. that their negative primed version do not belong to the model. Since we do not have persistence defaults in $\Pi_\Delta(r)$, persistence is captured and imposed through $I(S)$, by keeping in it all the atoms from the original database that were not discarded via the primed predicates.

Proposition 1

Given a database instance r and a set of BICs IC, if S is a model of $\Pi_\Delta(r)$, then $I(S)$ satisfies IC. □

Definition 5

Given database instances r and r' over the same schema and domain, we define

$$\begin{aligned} S(r, r') \quad = \quad & \{p(\bar{a}) \mid r \models p(\bar{a})\} \ \cup \ \{p'(\bar{a}) \mid r' \models p(\bar{a})\} \ \cup \ \{\neg p'(\bar{a}) \mid r' \not\models p(\bar{a})\} \\ & \cup \ \{dom(a) \mid a \in D\}. \end{aligned}$$ □

$S(r, r')$ collects the maximal set of literals that can be obtained from two database instances. It contains everything from both r and r'. The atoms corresponding to the second argument are primed. Negative literals corresponding to the first argument, intended to be the original databases instance, are not considered, because we will apply weak negation to them.

Proposition 2

Given a database instance r and a set of BICs IC, if r' satisfies IC, then $S(r, r')$ is a model of $\Pi_\Delta(r)$. □

In the following we will be considering subsets of $S(r, r')$. The previous result tells us that its subsets can be potential models of the program. $S(r, r')$ can be a large model, in the sense that the difference between r and r' may not be minimal.

Proposition 3

For BICs, the change program $\Pi_\Delta(r)$ has an answer set. □

5.1.2 The repair program

Program $\Pi_\Delta(r)$ gives an account of the changes in the original instance that are needed to produce the repairs, but the actual repairs contain data that persists from the original instance. This can be captured by adding persistence defaults.

Definition 6
The repair program $\Pi(r)$ consists of the rules in program $\Pi_\Delta(r)$ (def. 2) plus the following two rules for each predicate p in the original database:

4. Persistence *defaults*

$$\begin{array}{rcl} p'(\bar{X}) & \longleftarrow & p(\bar{X}) \\ \neg p'(\bar{X}) & \longleftarrow & dom(\bar{X}),\ p(\bar{X}). \end{array}$$ □

Remark 1
As shown in (Kowalski and Sadri, 1991), the DLPE 1.– 4., which has an e-answer semantics, can be transformed into a disjunctive extended logic program with answer set semantics, by transforming the persistence defaults into

4'. Persistence *rules*

$$\begin{array}{rcl} p'(\bar{X}) & \longleftarrow & p(\bar{X}),\ \ not\ \neg p'(\bar{X}) \\ \neg p'(\bar{X}) & \longleftarrow & dom(\bar{X}),\ \ not\ p(\bar{X}),\ \ not\ p'(\bar{X}). \end{array}$$

As shown in (Gelfond and Lifschitz, 1991), the resulting extended disjunctive normal program can be further transformed into a disjunctive normal program with a stable model semantics, with a one to one correspondence between answer sets and stable models. For this reason, we will interchangeably use the terms answer sets and stable models. □

Proposition 4
Given a database instance r over a finite domain, and a set of BICs IC, if S_M is an answer set[3] of $\Pi_\Delta(r)$, then

$$S = S_M \cup \{p'(\bar{a}) \mid p(\bar{a}) \in S_M \text{ and } \neg p'(\bar{a}) \notin S_M\} \cup \\ \{\neg p'(\bar{a}) \mid p(\bar{a}) \notin S_M \text{ and } p'(\bar{a}) \notin S_M\} \quad (3)$$

is an answer set of $\Pi(r)$. □

In order to establish the correspondence between the answer sets of the repair program $\Pi(r)$ and the repairs of r, we need the following lemma. It says that whenever we build an answer set S with literals taken from $S(r, r')$, and r' satisfies the ICs and is already as close as possible to r, then in S we can get essentially r' only. The condition that S is contained in $S(r, r')$ makes sure that the literals in S are taken from the right, maximal set of literals.

[3] In the programs we are considering so far, namely $\Pi_\Delta(r)$ and $\Pi(r)$ with persistence rules instead of persistence defaults, we do not find any defaults. In consequence, we can talk about answer sets as in (Gelfond and Lifschitz, 1991) instead of e-answer sets (Kowalski and Sadri, 1991).

Lemma 1
Let r and r' be database instances over the same schema and domain, and IC, a set of BICs. Assume that $r' \models IC$ and the symmetric difference $\Delta(r, r')$ is a minimal element under set inclusion in the set $\{\Delta(r, r^*) \mid r^* \models IC\}$. Then, for every answer set S of $\Pi_\Delta(r)$ contained in $S(r, r')$, it holds $r' = I(S)$. □

Theorem 1
If $\Pi(r)$ is the program $\Pi_\Delta(r)$ plus rules 4'., for a finite domain database instance r and a set of BICs IC, then it holds:

1. For every repair r' of r wrt IC, there exists an answer set S of $\Pi(r)$ such that $r' = \{p(\bar{a}) \mid p'(\bar{a}) \in S\}$.
2. For every answer set S of $\Pi(r)$, there exists a repair r' of r wrt IC such that $r' = \{p(\bar{a}) \mid p'(\bar{a}) \in S\}$. □

In the case of finite domain databases, the domain can be and has been declared. In this situation, we can handle any set of binary ICs, without caring about their safeness or domain independence.

Example 6
Let us take $D = \{a, b, c\}$, $r = \{p(a)\}$ and $IC = \{\forall x p(x)\}$. In this case, the program $\Pi(r)$ is

$$
\begin{array}{rcl}
p'(X) & \longleftarrow & p(X), not\ \neg p'(X) \\
\neg p'(X) & \longleftarrow & dom(X), not\ p(X), not\ p'(X) \\
p'(X) & \longleftarrow & dom(X), not\ p(X) \\
p'(X) & \longleftarrow & dom(X) \\
dom(a) & \longleftarrow & \\
dom(b) & \longleftarrow & \\
dom(c) & \longleftarrow & \\
p(a) & \longleftarrow &
\end{array}
$$

The only answer set is $\{dom(a), dom(b), dom(c), p(a), p'(a), p'(b), p'(c)\}$, that corresponds to the only repair $r' = \{p(a), p(b), p(c)\}$. □

In this example, the IC demands that every element in domain D belongs to table p; and this is possible to satisfy because the domain is finite. Nevertheless, if the domain D were infinite this would not be possible, because relational tables contain finitely many tuples. So, this kind of ICs cannot be handled in infinite domains.

5.2 Infinite domain databases

Now we consider only ICs that are *domain independent* (Ullman, 1988). For these ICs only the active domain matters. In particular, checking their satisfaction in an instance r can be done considering the elements of $Act(r)$ only. The IC in example 6 is not domain independent.

For domain independent BICs all previous lemmas and theorems still hold if we have infinite database domains D. To obtain them, all we need to do is to use a predicate $act_r(x)$, standing for the active domain $Act(r)$ of instance r, instead of predicate $dom(x)$. This is because, for domain independent BICs, the database domain can be considered to be the finite domain $Act(r)$. Furthermore, in this case we can omit the *dom* facts in 1. of $\Pi_\Delta(r)$ (definition 2). In consequence, we have the following theorem, first stated in (Arenas *et al.*, 2000a).

Theorem 2

For a set of domain independent binary integrity constraints and a database instance r, there is a one to one correspondence between the answers sets of the repair program $\Pi(r)$ and the repairs of r. $\square$

As a consequence of this general result, we have that our DLPEs $\Pi(r)$ correctly specify the repairs of relational databases that violate usual integrity constraints like range constraints, key constraints, functional dependencies, and non-existential inclusion dependencies.

6 Well-Founded Consistent Answers

Computing the stable model semantics for disjunctive programs is Π_2^P-complete in the size of the ground program[4]. In some cases, computing consistent answers can be done more efficiently.

The intersection of all answer sets of a extended disjunctive logic program contains the well-founded interpretation for such programs (Leone *et al.*, 1997), that can be computed in polynomial time in the size of the ground program. This interpretation may be partial and not necessarily a model of the program. Actually, it is a total interpretation if and only if it is the only answer set.

The well-founded interpretation, $W_{\Pi(r)} =< W^+, W^- >$, of program $\Pi(r)$, where W^+, W^-, W^u are the sets of true positive, negative, unknown literals, resp., is the given by the fixpoint $\mathcal{W}^{\omega}_{\Pi(r)}(\emptyset)$ of operator $\mathcal{W}_{\Pi(r)}$, that maps interpretations to interpretations (Leone *et al.*, 1997). More precisely, assuming that we have the ground instantiation of the repair program $\Pi(r)$, $\mathcal{W}_{\Pi(r)}(I)$ is defined on interpretations I that are sets of ground literals (without pairs of complementary literals) by: $\mathcal{W}_{\Pi(r)}(I) := T_{\Pi(r)}(I) \cup \neg.GUS_{\Pi(r)}(I)$.

Intuitively, $T_{\Pi(r)}$ is the immediate consequence operator that declares a literal true whenever there is ground rule containing it in the head, the body is true in I and the other literals in the (disjunctive) head are false in I. $\neg.GUS_{\Pi(r)}(I)$ denotes the set of complements of the literals in $GUS_{\Pi(r)}(I)$, being the latter the largest set of unfounded literals, those that definitely cannot be derived from the program and the set I of assumptions; in consequence their complements are declared true. The

[4] See (Dantsin *et al.*, 200?) for a review of complexity results in logic programming.

well-founded interpretation, $W_{\Pi(r)}$, is the least fixpoint $\mathcal{W}^{\omega}_{\Pi(r)}(\emptyset)$ of $\mathcal{W}_{\Pi(r)}$. More details can be found in (Leone *et al.*, 1997).

The intersection of all answer sets of $\Pi(r)$ is

$$Core(\Pi(r)) := \bigcap\{S \mid S \text{ is an answer set of } \Pi(r)\}.$$

Interpretation W_{Π}, being a subset of $Core(\Pi(r))$, can be used as an approximation from below to the core that can be computed more efficiently than all database repairs, or their intersection in the general case. Nevertheless, it is possible to identify classes of ICs for which $W_{\Pi}(r)$ coincides with $Core(\Pi(r))$. In these cases, the core is no longer approximated by $W_{\Pi(r)}$, but computed exactly.

In order to prove these results, we will assume, as in section 5.1, that we have a finite database domain D. We know that the results obtained under this hypothesis still hold for infinite domain databases and domain independent integrity constraints. In consequence, program $\Pi(r)$ contains domain predicates. The domain facts belong to every answer set and are obtained after the first iteration of the well-founded operator.

Proposition 5

For a database instance r, and a set of ICs containing functional dependencies and unary ICs only[5], the $Core(\Pi(r))$ of program $\Pi(r)$ coincides with $W_{\Pi(r)}$, the well-founded interpretation of program $\Pi(r)$. □

As corollary of this proposition and results presented in (Leone *et al.*, 1997) about the computational complexity of the disjunctive well-founded interpretation, we obtain that, for FDs and unary constraints, $Core(\Pi(r))$ can be computed in polynomial time in the size of the ground instantiation of $\Pi(r)$, a result first established in (Arenas *et al.*, 2001) for FDs. In particular, we can consistently answer non-existential conjunctive queries in polynomial time, because we can use $Core(\Pi(r))$ only. Furthermore, in (Arenas *et al.*, 2001), for the case of functional dependencies, some conditions on queries are identified under which one can take advantage of computations on the core to answer aggregate queries more efficiently.

For programs of the kind we may have for BICs, it is not always the case that the core coincides with the well-founded interpretation.

Example 7

Consider the BICs $IC = \{q \vee r,\ s \vee \neg q,\ s \vee \neg r\}$ and the empty database instance. The program $\Pi(r)$ wrt IC is

[5] Remember that they are BICs with at most one database literal in the standard format (2), plus built-ins. They include range constraints, e.g. $stock(x) \rightarrow 100 \leq x$, stating that products in stock may not go below 100 units.

Triggering rules:

$$\begin{array}{rcl} q' \vee r' & \longleftarrow & not\ q, not\ r \\ s' \vee \neg q' & \longleftarrow & not\ s, q \\ s' \vee \neg r' & \longleftarrow & not\ s, r \end{array}$$

Stabilizing rules:

$$\begin{array}{rcl} q' & \longleftarrow & \neg r' \\ r' & \longleftarrow & \neg q' \\ s' & \longleftarrow & q' \\ \neg q' & \longleftarrow & \neg s \\ s' & \longleftarrow & r' \\ \neg r' & \longleftarrow & \neg s' \end{array}$$

Persistence rules:

$$\begin{array}{rcl} q' & \longleftarrow & q, not\ \neg q' \\ s' & \longleftarrow & s, not\ \neg s' \\ r' & \longleftarrow & r, not\ \neg r' \\ \neg q' & \longleftarrow & not\ q, not\ q' \\ \neg s' & \longleftarrow & not\ s, not\ s' \\ \neg r' & \longleftarrow & not\ r, not\ r'. \end{array}$$

The answer sets are: $\{q', s', \neg r'\}$ and $\{\neg q', s', r'\}$. Then $Core(\Pi(r)) = \{s'\}$, but $W_{\Pi(r)} = \emptyset$. □

The results obtained so far in this section apply to the repair program $\Pi(r)$. Nevertheless, when we add an arbitrary query program $\Pi(Q)$ to $\Pi(r)$, obtaining program Π, then it is possible that $Core(\Pi)$ properly extends the well-founded interpretation of Π, even for FDs.

Example 8
Consider $r = \{P(a,b), P(a,c)\}$, the FD: $P(x,y) \vee P(x,z) \vee y = z$, and the query $Q(x):\ \exists y\ P(x,y)$. Program Π is:

$$\begin{array}{rcl} dom(a) & \leftarrow & \\ dom(b) & \leftarrow & \\ dom(c) & \leftarrow & \\ P(a,b) & \leftarrow & \\ P(a,c) & \leftarrow & \\ Query(X) & \leftarrow & P'(X,Y) \\ P'(X,Y) & \leftarrow & P(X,Y), not\ \neg P'(X,Y) \\ \neg P'(X,Y) & \leftarrow & dom(X), dom(Y), not\ P(X,Y), not\ P'(X,Y) \end{array}$$

$$\begin{aligned}
\neg P'(X,Y) \vee \neg P'(X,Z) &\leftarrow P(X,Y), P(X,Z), Y \neq Z \\
\neg P'(X,Y) &\leftarrow dom(X), dom(Y), dom(Z), P'(X,Z), Y \neq Z \\
\neg P'(X,Z) &\leftarrow dom(X), dom(Y), dom(Z), P'(X,Y), Y \neq Z.
\end{aligned}$$

The answer sets of Π are $S_1 = \{Query(a), P'(a,b), P(a,b), P(a,c), \dots\}$ and $S_2 = \{Query(a), P'(a,c), P(a,b), P(a,c), \dots\}$. The well-founded interpretation is $W_\Pi = < W^+, W^-, W^u >$, with $W^+ = \{P(a,b), P(a,c), dom(a), \dots\}$, $W^- = \{\neg P'(a,a), \dots\}$, and implicitly, the set of undetermined literals $W^u = \{P'(a,b), P'(a,c), Query(a)\}$. In particular, $Query(a) \in Core(\Pi)$, but $Query(a) \notin W^+$. □

We know, by complexity results presented in (Arenas *et al.*, 2001) for functional dependencies that, unless $P = NP$, consistent answers to first-order queries cannot be computed in polynomial time. In consequence, we cannot expect to compute the $Core(\Pi)$ of the program that includes the query program by means of the well-founded interpretation of Π only.

7 Evaluating $\mathcal{K}$-queries

The results in section 5 provide the underpinning of a general method of evaluating $\mathcal{K}$-queries. Assume r is a database instance and the set of integrity constraints IC is given. We show how to evaluate queries of the form $\beta \equiv \mathcal{K}\alpha$ where α is a basic query. First, from α we obtain a stratified logic program $\Pi(\alpha)$ (this is a standard construction (Lloyd, 1987; Abiteboul *et al.*, 1995)) in terms of the new, primed predicates. One of the predicate symbols, $Query_\alpha$, of $\Pi(\alpha)$ is designated as the query predicate. This is illustrated in section 4.2. Second, determine all the answers sets $S_1, \dots, S_k$ of the logic program $\Pi = \Pi(\alpha) \cup \Pi(r)$. Third, compute the intersection $r_\beta = \bigcap_{1 \leq i \leq k} S_i / Query_\alpha$, where $S_i / Query_\alpha$ is the extension of $Query_\alpha$ in S_i. The set of tuples r_β is the set of answers to β in r.

Notice that the set U consisting of all the ground primed database literals, $(\neg)p'(a)$, and all the ground non primed database literals, $(\neg)p(a)$, form a *partition* for the program Π, because whenever a literal in U appears in a head, all the literals in the body also appear in U (Lifschitz and Turner, 1994). The set U partitions the program precisely into the two expected parts, $\Pi(r)$ and $\Pi(Q)$, because the literals in U do not appear in heads of rules in $\Pi(Q)$ (for $\Pi(Q)$, literals in U are like extensional literals). From (Lifschitz and Turner, 1994), we know that every answer set of Π can be represented as the union of an answer set of $\Pi(r)$ and an answer set of $\Pi(Q)$, where each answer set for $\Pi(r)$ acts as an extensional database for the computation of the answer sets of $\Pi(Q)$. Program $\Pi(Q)$ is stratified, in consequence, for each answer set for $\Pi(r)$, there will only one answer set for $\Pi(Q)$.

To obtain query answers to general $\mathcal{K}$-queries the above method needs to be combined with some method of evaluating first-order queries. For example, safe-range first-order queries (Abiteboul *et al.*, 1995) can be translated to relational algebra. The same approach can be used for $\mathcal{K}$ queries with the subqueries of the form $\mathcal{K}\alpha$

replaced by new relation symbols. Then when the resulting relational algebra query is evaluated and the need arises to materialize one of the new relations, the above method can be used to accomplish that goal.

8 Computational Examples

As shown in section 5.1, our disjunctive programs with exceptions can be transformed (Kowalski and Sadri, 1991) into extended disjunctive logic programs with an answer set semantics (Gelfond and Lifschitz, 1991). Once this transformation has been performed, obtaining program $\Pi(r)$, it is possible to use any implementation of extended disjunctive logic programs with answer set semantics. In this section, we give some examples that show the application of the *DLV* system (Eiter *et al.*, 1998) to the computation of database repairs and consistent query answers.

In (Leone *et al.*, 1997) it is shown how to compute the answer sets of a program starting from the well-founded interpretation, that can be efficiently computed and is contained in the intersection of the answer sets. This is what *DLV* basically does, but instead of starting from the well-founded interpretation, it starts from the also efficiently computable set of *deterministic consequences* of the program, that is still contained in the intersection of all answer sets, and in its turn, contains the well-founded interpretation (Leone, 2000). Actually, *DLV* can be explicitly asked to return the set of deterministic consequences of the program[6], and it can be also used as an approximation from below to the intersection of all answer sets.

8.1 Computing database repairs with DLV

Example 9
Consider a database schema $Emp(Name, SSN)$. Each person should have just one SSN and different persons should have different SSNs. That is, the following functional dependencies are expected to hold: $Name \rightarrow SSN$, $SSN \rightarrow Name$. The following is an inconsistent instance:

Emp	*Name*	*SSN*
	Irwin Koper	677-223-112
	Irwin Koper	952-223-564
	Michael Baneman	334-454-991

In order to consistently query this database, we can generate the following *DLV* program, where the prime predicate Emp', containing the repaired extension, we had before is now denoted by *emp_p*

```
% domains of the database
dom_name("Irwin Koper").         dom_name("Michael Baneman").
dom_number("677-223-112").       dom_number("952-223-564").
dom_number("334-454-991").
```

[6] By means of its option `- det`.

```
% initial database
emp("Irwin Koper", "677-223-112").
emp("Irwin Koper", "952-223-564").
emp("Michael Baneman", "334-454-991").

% default rules
emp_p(X,Y)  :- emp(X,Y), not -emp_p(X,Y).
-emp_p(X,Y) :- dom_name(X), dom_number(Y), not emp(X,Y), not emp_p(X,Y).

% triggering rules
-emp_p(X,Y) v -emp_p(X,Z) :- emp(X,Y), emp(X,Z), Y!=Z.
-emp_p(Y,X) v -emp_p(Z,X) :- emp(Y,X), emp(Z,X), Y!=Z.

% stabilizing rules.
-emp_p(X,Y) :- emp_p(X,Z), dom_number(Y), Y!=Z.
-emp_p(Y,X) :- emp_p(Z,X), dom_name(Y), Y!=Z.
```

DLV running on this program delivers two answer sets, corresponding to the two following repairs:

Emp	*Name*	*SSN*
	Irwin Koper	952-223-564
	Michael Baneman	334-454-991

Emp	*Name*	*SSN*
	Irwin Koper	677-223-112
	Michael Baneman	334-454-991

In order to pose the query $Emp(X, Y)$?, asking for the consistent tuples in table *Employee*, it is necessary to add a new rule to the program:

```
query(X,Y) :- emp_p(X,Y).
```

Now, the two answer sets of the program will contain `query` literals, namely

```
{ ..., query("Irwin Koper","952-223-564"),
                          query("Michael Baneman","334-454-991")}

{ ..., query("Irwin Koper","677-223-112"),
                          query("Michael Baneman","334-454-991")}
```

In order to obtain the consistent answers to the query, it is sufficient to choose all the ground `query` atoms that are in the intersection of all answer sets of the program extended by the query rule. In this case, we obtain as only consistent answer the tuple: `X="Michael Baneman"`, `Y="334-454-991"`. Here we had a non-quantified conjunctive query. In other cases, the `query` predicate will be defined by a more complex program $\Pi(Q)$. □

9 An Alternative Semantics

As discussed in (Arenas *et al.*, 1999), our notion of database repair coincides with that of revision model obtained with the "possible model approach" introduced by Winslett in (Winslett, 1988) (see also (Chou and Winslett, 1994)) in the context of belief update, when the database instance (a model) is updated by the ICs, generating a new set of models, in this case, the database repairs. In consequence, we have shown in section 5.1 that our repair program $\Pi(r)$ (cf. theorem 1) has as its answer sets the Winslett's revision models of r wrt IC.

9.1 Cardinality based repairs and weak constraints

Winslett's revision models are based on minimal *set* of changes. In (Dalal, 1988), Dalal presents an alternative notion of revision model based on minimal *number* of changes.

Definition 7
Given a database instance r, an instance r' is a *Dalal repair* of r wrt to IC iff $r' \models IC$ and $|\Delta(r, r')|$ is a minimal element of $\{|\Delta(r, r^*)| \mid r^* \models IC\}$. □

We could give a definition of *Dalal consistent answer* exactly in the terms of definition 1, but replacing "repair" by "Dalal repair".

It is possible to specify Dalal repairs using the same repair programs we had in section 5, but with the persistence defaults replaced by *weak constraints* (Buccafurri *et al.*, 2000). The latter will not be a sort of weak version of the original, database ICs. Rather they will be new constraints imposed on the answer sets of the repair program, actually on $\Pi_\Delta(r)$, the part of the repair program of section 5 that is responsible for the changes.

As described in (Buccafurri *et al.*, 2000), weak constraints are written in the form $\Leftarrow L_1, \ldots, L_n$, where the L_i's are literals containing strong or weak negation. They are added to an extended disjunctive program. Their semantics is such, that, when violated in a model of the program, they do not necessarily "kill" the model. The models of the program that minimize the *number* of violated ground instantiations of the given weak constraints are kept.

In order to capture the Dalal repairs we need a very simple form of weak constraint. The program $\Pi^D(r)$ that specifies the Dalal repairs of a database instance r wrt a set of BICs consists of program $\Pi_\Delta(r)$ of section 5.1 (rules 1. – 3.) plus

4". For every database predicate p, the weak constraints

$$\begin{aligned} &\Leftarrow \quad p'(\bar{x}),\ not\ p(\bar{X}), \\ &\Leftarrow \quad \neg p'(\bar{x}),\ p(\bar{X}). \end{aligned} \tag{4}$$

These constraints say that the contents of the original database and of each repair are expected to coincide. Since they are weak constraints, they allow violations, but only a minimum *n*umber of tuples that belong to the repair and not to the original instance, or the other way around, will be accepted.

The results for the change program $\Pi_\Delta(r)$ still hold here. The program obtained by the combination of the change program with the weak constraints 4". in (4) will have answers sets that correspond to repairs that are minimal under set inclusion and under number of changes, i.e. Dalal repairs only. So, for BICs and finite domain databases we have

Theorem 3
Given a (finite domain) database instance r and a set of BICs IC:

1. For every Dalal repair r' of r wrt IC, there exists an answer set S of $\Pi^D(r)$ such that $I(S) = r'$.
2. For every answer set S of $\Pi^D(r)$, there exists a Dalal repair r' of r wrt IC such that $I(S) = r'$. □

As with Winslett's repairs, the theorem still holds for infinite domain databases when the BICs are domain independent.

Example 10
Let $D = \{a\}$, $r = \{p(a)\}$ and $IC = \{\neg p(x) \vee q(x), \neg q(x) \vee r(x)\}$. In this case, $\Pi^D(r)$ is

$$\begin{array}{rcl}
dom(a) & \longleftarrow & \\
p(a) & \longleftarrow & \\
\neg p'(X) \vee q'(X) & \longleftarrow & p(X),\ not\ q(X) \\
q'(X) & \longleftarrow & p'(X) \\
\neg p'(X) & \longleftarrow & \neg q'(X) \\
\neg q'(X) \vee r'(X) & \longleftarrow & q(X),\ not\ r(X) \\
r'(X) & \longleftarrow & q'(X) \\
\neg q'(X) & \longleftarrow & \neg r'(X) \\
 & \Leftarrow & p'(X),\ not\ p(X) \\
 & \Leftarrow & \neg p'(X), p(X) \\
 & \Leftarrow & q'(X),\ not\ q(X) \\
 & \Leftarrow & \neg q'(X), q(X) \\
 & \Leftarrow & r'(X),\ not\ r(X) \\
 & \Leftarrow & \neg r'(X), r(X).
\end{array}$$

Weak constraints are implemented in *DLV*[7]. Running *DLV* on this program we obtain the answer set $\{dom(a), p(a), \neg p'(a)\}$, corresponding to the empty database repair, but not the other Winslett's repair $\{p(a), q(a), r(a)\}$, whose set of changes wrt r has two elements, whereas the first repair differs from r by one change only. □

[7] They are and are specified by :∼ *Conj*., where *Conj* is a conjunction of (possibly negated) literals. See *DLV*'s user manual in http://www.dbai.tuwien.ac.at/proj/dlv/man.

Notice that, by construction of $\Pi^D(r)$, in most cases the primed predicates in the answer sets of the program will not contain all the information (we replaced the persistence defaults 4. by the weaks constraints 4".). In consequence, we will have to interpret the result, and the answers to queries will be obtained by using negation as failure, as we shown in the following table:

original query	query in the program
$p(\bar{x})$	query(X̄) ← p′(X̄). query(X̄) ← p(X̄), not ¬p′(X̄).
$\neg p(\bar{x})$	query(X̄) ← ¬p′(X̄). query(X̄) ← *dom*(X̄), not p(X̄), not p′(X̄).

10 Extensions

In this section we will show how the specifications of database repairs given for binary integrity constraints can be extended to referential integrity constraints and to universal constrains with more than two database literals in the standard format. We will only consider the case of minimal repairs under set inclusion.

10.1 Referential integrity constraints

By appropriate representation of existential quantifiers as program rules it is possible to apply the methodology for universal constraints presented in the previous sections to handle referential integrity constraints (RICs).

Consider the *RIC*: $\forall \bar{x}\ (P(\bar{x}) \rightarrow \exists \bar{y}\ R(\bar{x}, \bar{y}))$, and the inconsistent database instance $r = \{P(\bar{a}), P(\bar{b}), R(\bar{b}, \bar{a})\}$. For things to work properly, it is necessary to assume that there is an underlying database domain D. The repair program has the persistence default rules

$$P'(\bar{X}) \leftarrow P(\bar{X}); \quad \neg P'(\bar{X}) \leftarrow \mathit{not}\ P(\bar{X});$$

$$R'(\bar{X}, \bar{Y}) \leftarrow R(\bar{X}, \bar{Y}); \quad \neg R'(\bar{X}, \bar{Y}) \leftarrow \mathit{not}\ R(\bar{X}, \bar{Y}).$$

In addition, it has the stabilizing exceptions

$$\neg P'(\bar{X}) \leftarrow \mathit{not}\ \mathit{aux}'(\bar{X}), \neg R'(\bar{X}, \mathit{null}), \tag{5}$$

$$R'(\bar{X}, \mathit{null}) \leftarrow P'(\bar{X}), \mathit{not}\ \mathit{aux}'(\bar{X}); \tag{6}$$

with

$$\mathit{aux}'(\bar{X}) \leftarrow R'(\bar{X}, \bar{Y});$$

and the triggering exception

$$\neg P'(\bar{X}) \vee R'(\bar{X}, \mathit{null}) \leftarrow P(\bar{X}), \mathit{not}\ \mathit{aux}(\bar{X}), \tag{7}$$

with $\mathit{aux}(\bar{X}) \leftarrow R(\bar{X}, \bar{Y})$.

The variables in this program range over D, that is, they do not take the value *null*. This is the reason for the last literal in clause (5). The last literal in clause (6) is necessary to insert a null value only when it is needed; this clause relies on the fact that variables range over D only. Instantiating variables on D only[8], the only two answer sets are the expected ones, namely delete $P(\bar{a})$ or insert $R(\bar{a}, null)$.

It would be natural to include here the functional dependency $\bar{X} \rightarrow \bar{Y}$ on R, expressing that $\bar{X}$ is a primary key in R and a foreign key in P. This can be done without problems, actually the two constraints would not interact, that is, reparing one of them will not cause violations of the other one.

Finally, it should be clear how to modify the specification above if the only admissible changes are elimination of tuples, but not introduction of null values. For example, the triggering exception (7) would have to be changed into $\neg P'(\bar{X}) \leftarrow P(\bar{X}), not\ aux(\bar{X})$.

10.2 Referential ICs and strong constraints

It has been possible to use *DLV* to impose some preferences on the repairs via an appropriate representation of constraints, obtaining, for example for RICs, introduction of null values, a cascade policy, ...

Example 11
(example 9 cont'd) Consider the same schema and *FD*s as before, but now extended with a unary table *Person*(*Name*). Now, we have the following instance

Emp	*Name*	*SSN*
	Irwin Koper	677-223-112
	Irwin Koper	952-223-564
	Michael Baneman	952-223-564

Person	*Name*
	Irwin Koper
	Michael Baneman

The *DLV* repair program, without considering any change on table *Person*, is as in example 9, but with:

```
dom_number("677-223-112").          dom_number("952-223-564").
% initial database
emp("Irwin Koper", "677-223-112"). emp("Irwin Koper",
"952-223-564"). emp("Michael Baneman", "952-223-564").
```

If *DLV* is run with this program as input, we obtain the answer sets:

```
{ ..., emp_p("Irwin Koper","677-223-112"),
      -emp_p("Irwin Koper","952-223-564"),
       emp_p("Michael Baneman","952-223-564"),
      -emp_p("Michael Baneman","677-223-112")}
```

[8] A simple way to enforce this at the object level is to introduce the predicate D in the clauses, to force variables to take values in D only, excluding the null value.

```
{ ...,-emp_p("Irwin Koper","677-223-112"),
       emp_p("Irwin Koper","952-223-564"),
      -emp_p("Michael Baneman","952-223-564"),
      -emp_p("Michael Baneman","677-223-112")}
```

corresponding to the database repairs:

Emp	*Name*	*SSN*
	Irwin Koper	677-223-112
	Michael Baneman	952-223-564

Emp	*Name*	*SSN*
	Irwin Koper	952-223-564

In the second repair Michael Baneman does not have a SSN, in consequence it is not a consistent answer that he has a SSN. Actually, it is possible to ask with *DLV* for those persons who have a SSN by computing the answer sets of the program extended by the query rule: `query(X):- emp_p(X,Y).` Two answer sets are obtained:

```
{ ..., query("Irwin Koper"), query("Michael Baneman")}, { ...,
query("Irwin Koper")}
```

From them, we can say that only Irwin Koper has a SSN (in all repairs). If we want every person to have a SSN, then we may enforce the RIC, $\forall\, x(Person(x) \rightarrow \exists\, yEmp(x,y))$, stating that every person must have a SSN.

In section 10.1, we repaired the database introducing the RIC as a part of the program, producing either the introduction of null values or cascading deletions. We may not want any of these options (we do not want null values in the key *SSN*) or we do not want to delete any employees (in this case, M. Baneman from *Person*). An alternative is to use *DLV*'s possibility of specifying *strong constraints*, that have the effect of pruning those answer sets that do not satisfy them. This can be done in *DLV* by introducing the denial `:- dom_name(X), not has_ssn(X).`, with `has_ssn(X) :- emp_p(X,Y).`. The answer sets of the original program that do not satisfy the ICs are filtered out.

Now, only one repair is obtained:

```
{ ..., emp_p("Irwin Koper","677-223-112"),
      -emp_p("Irwin Koper","952-223-564"),
       emp_p("Michael Baneman","952-223-564"),
      -emp_p("Michael Baneman","677-223-112"),
       has_ssn("Irwin Koper"),  has_ssn("Michael Baneman"),
       query("Irwin Koper"),  query("Michael Baneman")}
```

In it, every person has a SSN (according to the `has_ssa` predicate). As expected, the answers to the original query are `X="Irwin Koper"` and `X="Michael Baneman"`.

10.3 Extensions to other universal ICs

In order to handle universal ICs with a larger number of database literals, we can use the same program $\Pi(r)$, but with $\Pi_\Delta(r)$, as introduced in definition 2 for BICs, now without the condition $n + m \leq 2$ in (1). Now m and n can be arbitrary. This will produce properly disjunctive stabilizing rules, by passing in turns just one of the disjuncts to the body in it complementary form. Furthermore, we need to extend $\Pi_\Delta(r)$ with more disjunctive stabilizing rules. In essence, as indicated in (Arenas *et al.*, 2000a), we need to consider all possible subsets of the database literals appearing in the standard format (1) and put them in a disjunction in the head, passing to the body the remaining literals and the formula φ. We will show this by means of some examples.

Example 12
Consider the DB instance $r = \{P(a), Q(a), R(a)\}$ and the following set of ternary integrity constraints $IC = \{\neg P(x) \vee \neg Q(x) \vee R(x), \neg P(x) \vee \neg Q(x) \vee \neg R(x), \neg P(x) \vee Q(x) \vee \neg R(x), P(x) \vee \neg Q(x) \vee \neg R(x), \neg P(x) \vee Q(x) \vee R(x), P(x) \vee \neg Q(x) \vee R(x), P(x) \vee Q(x) \vee \neg R(x)\}$. The repair program contains the usual persistence defaults for P, Q, R, plus triggering exceptions, e.g. for the first IC in IC:

$$\neg P'(x) \vee \neg Q'(x) \vee R'(x) \quad \leftarrow \quad P(x), Q(x), not\ R(x),$$

and the stabilizing rules directly obtained from definition 2, e.g. for the first IC

$$\begin{aligned} \neg P'(x) \vee \neg Q'(x) &\leftarrow \neg R'(x), \\ \neg P'(x) \vee R'(x) &\leftarrow Q'(x), \\ \neg Q'(x) \vee R'(x) &\leftarrow P'(x). \end{aligned} \tag{8}$$

We add to the program the other possible combinations, e.g. for the first IC:

$$\begin{aligned} \neg P'(x) &\leftarrow Q'(x), \neg R'(x), \\ R'(x) &\leftarrow P'(x), Q'(x), \\ \neg Q'(x) &\leftarrow P'(x), \neg R'(x). \end{aligned} \tag{9}$$

In this case we obtain as answer set the only repair, namely the empty instance, represented by $\{P(a), Q(a), R(a), \neg P'(a), \neg Q'(a), \neg R'(a)\}$. It is easy to verify that without using the disjunctive stabilizing rules (8), but with the rules (9) as the only stabilizing rules, the empty repair cannot be obtained. □

Example 13
Let $D = \{a, b, c\}$, $r = \{p(a,b), p(b,c)\}$ and IC the transitivity constraint $\neg p(x,y) \vee \neg p(y,z) \vee p(x,z)$. Here, the only repairs are $r_1 = \{p(a,b)\}$, $r_2 = \{p(b,c)\}$ and $r_3 = \{p(a,b), p(b,c), p(a,c)\}$. The repair program $\Pi(r)$ contains the facts $dom(a), dom(b), dom(c), p(a,b), p(b,c)$ plus the rules

$$\begin{aligned} p'(X,Y) &\leftarrow p(X,Y), not\ \neg p'(X,Y) \\ \neg p'(X,Y) &\leftarrow dom(X), dom(Y), not\ p(X,Y), not\ p'(X,Y) \\ \neg p'(X,Y) \vee \neg p'(Y,Z) \vee p'(X,Z) &\leftarrow p(X,Y), p(Y,Z), not\ p(X,Z) \end{aligned}$$

$$\begin{array}{rcl}\neg p'(Y,Z) \vee p'(X,Z) & \leftarrow & dom(Z), p'(X,Y) \\ \neg p'(X,Y) \vee p'(X,Z) & \leftarrow & dom(X), p'(Y,Z) \\ \neg p'(X,Y) \vee \neg p'(Y,Z) & \leftarrow & dom(Y), \neg p'(X,Z) \\ \neg p'(Y,Z) & \leftarrow & dom(Z), p'(X,Y), \neg p'(X,Z) \\ p'(X,Z) & \leftarrow & p'(X,Y), p'(Y,Z) \\ \neg p(X,Y) & \leftarrow & dom(X), p'(Y,Z), \neg p'(X,Z). \end{array} \tag{10}$$

With this program we will get as answer sets the three repairs, but also six other models corresponding to non-minimal repairs, e.g. $\{dom(a),\ dom(b),\ dom(c),\ p(a,b),\ p(b,c),\ p'(a,b),\ p'(b,c),\ p'(a,c),\ p'(b,b),\ p'(c,b),\ p'(c,c),\ p'(a,a),\ p'(b,a),\ p'(c,a)\}$, meaning that we are repairing the database by inserting the tuple (a,c) into p, but also the tuples $(b,b),(c,b)$, etc. The reason is that a rule like (10) for example is inserting tuples (x,z), for essentially any z, independently of the condition in the body. In order to solve this problem, one can add extra conditions in the body of (10) that relate the possible values for z with the values in the original database. In consequence, that rule should be replaced by $\neg p'(Y,Z) \vee p'(X,Z) \leftarrow p(Y,Z), not\ p(X,Z), p'(X,Y)$. The domain predicate, introduced in (10) to make the rule safe, is no longer needed. Similar changes have to be performed in the other stabilizing rules. □

11 Conclusions

There are several interesting open issues related to computational implementation of the methodology we have presented.

The existing implementations of stable models semantics are based on grounding the rules, what, in database applications, may lead to huge ground programs. In addition, those implementations are geared to computing stable models, possibly not all of them, and answering ground queries. At the same time, in database applications the possibility of posing and answering open queries (with variables) is much more natural. In addition, consistent query answering requires, at least implicitly, having all stable models.

It would be useful to implement a consistent query answering system based on the interaction of our repairs logic programs with relational DBMS. For this purpose, some functionalities and front-ends included in *DLV*'s architecture (Eiter *et al.*, 2000) could be used.

Another interesting issue has to do with the possibility of having the consistent query answering mechanism guided by the query, so that irrelevant computations are avoided.

There, are other open problems that could be considered: (a) Analyzing conditions under which simpler programs can be obtained. (b) A more detailed treatment of referential ICs (and other existential ICs). (c) Identification of other classes of ICs for which the well-founded interpretation and the intersection of all database repairs coincide. (d) Preferences for certain kinds of repair actions. In principle they could be captured by choosing the right disjuncts in the triggering rules.

With respect to related work, the closest approach to ours is presented in (Greco *et al.*, 2001) (see also (Greco and Zumpano, 2000; Greco and Zumpano, 2001)). There disjunctive programs are used to specify the sets of changes under set inclusion that lead to database repairs in the sense of (Arenas *et al.*, 1999). They present a compact schema for generating repair programs for universal integrity constraints. The application of such a schema leads to programs that involve essentially all possible disjunctions of database literals in the heads. They concentrate mainly on producing the set of changes, rather than the repaired databases explicitly. In particular, they do no have persistence rules in the program. In consequence, the program cannot be used directly to obtain consistent answers. They also introduce "repair constraints" to specify preferences for certain kinds of repairs.

Another approach to database repairs based on logic programming semantics consists of the *revision programs* (Marek and Truszczynski, 1998). The rules in those programs explicitly declare how to enforce the satisfaction of an integrity constraint, rather than explicitly stating the ICs, e.g.

$$in(a) \leftarrow in(a_1), \ldots , in(a_k), out(b_1), \ldots , out(b_m)$$

has the intended procedural meaning of inserting the database atom a whenever $a_1, \ldots , a_k$ are in the database, but not $b_1, \ldots , b_m$. They also give a declarative, stable model semantics to revision programs. Preferences for certain kinds of repair actions can be captured by declaring the corresponding rules in program and omitting rules that could lead to other forms of repairs. Revision programs could be used, as the programs in (Greco *et al.*, 2001), to obtain consistent answers, but not directly, because they give an account of the changes only.

Acknowledgments: Work supported by FONDECYT Grant 1000593, NSF Grant INT-9901877/CONICYT Grant 1998-02-083, NSF Grant IIS-0119186, Carleton University Start-Up Grant 9364-01. We are grateful to Francisco Orchard for informative paper presentations, discussions, and experiments with DLV. We are grateful to Nicola Leone for kindly answering all our questions about *DLV*.

References

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

Arenas, M.; Bertossi, L.; and Chomicki, J. Consistent Query Answers in Inconsistent Databases. In *Proc. ACM Symposium on Principles of Database Systems (ACM PODS'99, Philadelphia)*, ACM Press, 1999, pp. 68–79.

Arenas, M.; Bertossi, L.; and Chomicki, J. Specifying and Querying Database Repairs using Logic Programs with Exceptions. In *Flexible Query Answering Systems. Recent Developments*, H.L. Larsen, J. Kacprzyk, S. Zadrozny, H. Christiansen (eds.), Springer, 2000, pp. 27–41.

Arenas, M.; Bertossi, L.; and Chomicki, J. Scalar Aggregation in FD-Inconsistent Databases. In *Database Theory - ICDT 2001* (Proc. International Conference on

Database Theory, ICDT'2001), Springer, Lecture Notes in Computer Science 1973, 2001, pp. 39 – 53.

Arenas, M.; Bertossi, L.; and Kifer, M. Applications of Annotated Predicate Calculus to Querying Inconsistent Databases. In *'Computational Logic - CL 2000'. Stream: 6th International Conference on Rules and Objects in Databases (DOOD'2000)*, Springer, Lecture Notes in Artificial Intelligence 1861, 2000, pp. 926–941.

Buccafurri, F.; Leone, N.; Rullo, P. Enhancing Disjunctive Datalog by Constraints. *IEEE Transactions on Knowledge and Data Engineering*, 2000, 12(5): 845-860.

Celle, A.; and Bertossi, L. Querying Inconsistent Databases: Algorithms and Implementation. In *'Computational Logic - CL 2000'. Stream: 6th International Conference on Rules and Objects in Databases (DOOD'2000)*, pp. 942–956. Springer Lecture Notes in Artificial Intelligence 1861.

Chou, T.; and Winslett, M. A Model-Based Belief Revision System. *Journal of Automated Reasoning*, 1994, 12:157–208.

Dalal, M. Investigations into a Theory of Knowledge Base Revision: preliminary report. In *Proc. Seventh National Conference on Artificial Intelligence (AAAI'88)*, 1988, pp. 475–479.

Dantsin, E., Eiter, T., Gottlob, G. and Voronkov, A. Complexity and Expressive Power of Logic Programming. To appear in ACM Computing Surveys.

Eiter, T.; Leone, N.; Mateis, C.; Pfeifer, G.; and Scarcello, F. The Knowledge Representation System DLV: Progress Report, Comparisons, and Benchmarks. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning, KR98, Trento, Italy, June 1998.* Morgan Kaufman.

Eiter, T.; Faber, W.; Leone, N.; Pfeifer, G. Declarative Problem-Solving in DLV. In *Logic-Based Artificial Intelligence*, J. Minker (ed.), Kluwer, 2000, pp. 79–103.

Fitting, M. *First Order Logic and Automated Theorem Proving.* Texts and Monographs in Computer Science. Springer Verlag, 2nd ed., 1996.

Gelfond, M.; and Lifschitz, V. The Stable Model Semantics for Logic Programming. In *Logic Programming, Proceedings of the Fifth International Conference and Symposium*, R. A. Kowalski and K. A. Bowen (eds.), MIT Press 1988, pp. 1070–1080.

Gelfond, M.; and Lifschitz, V. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 1991, 9:365–385.

Greco, S.; and Zumpano, E. Querying Inconsistent Databases. In Proc. 7th International Conference on Logic for Programming and Automated Reasoning (LPAR'2000), Springer LNCS 1955, 2000, pp. 308-325.

Greco, S.; and Zumpano, E. Computing Repairs for Inconsistent Databases. Proc. The Third International Symposium on Cooperative Database Systems for Advanced Applications (CODAS01), Beijing, April 23-24, 2001.

Greco, G.; Greco, S.; and Zumpano, E. A Logic Programming Approach to the Integration, Repairing and Querying of Inconsistent Databases. In Proc. 17th International Conference on Logic Programming, ICLP'01, Ph. Codognet (ed.), LNCS 2237, Springer, 2001, pp. 348–364.

Kowalski, R.; and Sadri, F. Logic Programs with Exceptions. *New Generation Computing*, 1991, 9:387–400.

Leone, N.; Rullo, P.; and Scarcello, F. Disjunctive Stable Models: Unfounded Sets, Fixpoint Semantics, and Computation. *Information and Computation*, 1997, 135(2):69-112.

Leone, N. Personal communication, 2000.

Lifschitz, V. and Turner, H. Splitting a Logic Program. In Proceedings of the Eleventh

International Conference on Logic Programming, Pascal van Hentenryck (ed.), MIT Press, 1994, pp. 23-37.

Lloyd, J.W. *Foundations of Logic Programming.* Springer Verlag, 1987.

Marek, V. W.; and Truszczynski, M. Revision Programming. *Theoretical Computer Science*, 1998, 190(2):241–277.

Ullman, J. *Principles of Database and Knowledge-Base Systems, Vol. I.* Computer Science Press, 1988.

Winslett, M. Reasoning about Action using a Possible Model Approach. In *Proc. Seventh National Conference on Artificial Intelligence (AAAI'88)*, 1988, pp. 89–93.

Appendix: Proofs

Proof of Proposition 1

Consider an arbitrary element in IC:

$$\bigvee_{i=1}^{n} p_i(\bar{x}_i) \vee \bigvee_{i=1}^{m} \neg q_i(\bar{y}_i) \vee \varphi.$$

We have to prove that $I(S)$ satisfies any instantiation of this formula, that is

$$I(S) \models \bigvee_{i=1}^{n} p_i(\bar{a}_i) \vee \bigvee_{i=1}^{m} \neg q_i(\bar{b}_i) \vee \varphi \tag{11}$$

We need to consider two cases.

(I) If r does not satisfy this ground constraint, then S satisfies the body of the ground triggering rule:

$$\bigvee_{i=1}^{n} p'_i(\bar{a}_i) \vee \bigvee_{i=1}^{m} \neg q'_i(\bar{b}_i) \leftarrow dom(\bar{a}_1, \ldots, \bar{a}_n), \ \bigwedge_{i=1}^{n} not\, p_i(\bar{a}_i), \ \bigwedge_{i=1}^{m} q_i(\bar{b}_i), \ \ not\, \varphi.$$

Thus, there exists $p'_i(\bar{a}_i) \in S$ or $\neg q'_j(\bar{b}_j) \in S$. But, if $p'_i(\bar{a}_i) \in S$, then $I(S) \models p_i(\bar{a}_i)$, and therefore (11) holds. If $\neg q'_j(\bar{b}_j) \in S$, then $q'_j(\bar{b}_j) \notin S$. Thus $I(S) \models \neg q_j(\bar{b}_j)$, and therefore (11) holds.

(II) If r satisfies the ground constraint, then r could satisfy φ, and in this case $I(S) \models \varphi$. Otherwise, r satisfies some $p_j(\bar{a}_j)$ or some $\neg q_j(\bar{b}_j)$.
Assume that (11) is not true. In this case, $I(S) \not\models \varphi$, and therefore r must satisfy some $p_j(\bar{a}_j)$ or some $\neg q_j(\bar{b}_j)$. If r satisfies $p_j(\bar{a}_j)$, then $p_j(\bar{a}_j) \in S$. But, $I(S) \not\models p_j(\bar{a}_j)$, since (11) does not holds, and therefore $\neg p'_j(\bar{a}_j) \in S$, by definition of $I(S)$. But, in this case S satisfies the body of the ground stabilizing rule:

$$\bigvee_{i=1}^{j-1} p'_i(\bar{a}_i) \vee \bigvee_{i=j+1}^{n} p'_i(\bar{a}_i) \vee \bigvee_{i=1}^{m} \neg q'_i(\bar{b}_i) \leftarrow$$
$$dom(\bar{a}_1, \ldots, \bar{b}_m), \ \neg p'_j(\bar{a}_j), \ \ not\ \varphi.$$

Thus, by using an argument analogous to the argument given in (I), we conclude that (11) holds, a contradiction.
The case in which r satisfies $\neg q_j(\bar{b}_j)$ can be handled in a similar way. $\square$

Proof of Proposition 2

In order to prove that $S(r,r')$ satisfies $\Pi_\Delta(r)$, we need to consider only the four different kinds of ground stabilizing rules (the satisfaction of the other rules follows from the fact that r' satisfies IC).

If $S(r,r')$ satisfies the body of the rule $q'(\bar{b}) \leftarrow p'(\bar{a}),\ \ \mathit{not}\ \varphi$, then r' must satisfy $p(\bar{a})$ and $\neg\varphi$. But $r' \models q(\bar{b}) \vee \neg p(\bar{a}) \vee \varphi$, because $\forall\bar{x}\forall\bar{y}(q(\bar{x}) \vee \neg p(\bar{y}) \vee \varphi) \in IC$, and therefore, $r' \models q(\bar{b})$. Thus, $q'(\bar{b}) \in S(r,r')$.

In the same way, it is possible to prove that $S(r,r')$ satisfies all the rules of the form:

$$\begin{array}{rcl} \neg q'(\bar{b}) & \leftarrow & p'(\bar{a}),\ \ \mathit{not}\ \varphi, \\ q'(\bar{b}) & \leftarrow & \neg p'(\bar{a}),\ \ \mathit{not}\ \varphi, \\ \neg q'(\bar{b}) & \leftarrow & \neg p'(\bar{a}),\ \ \mathit{not}\ \varphi. \end{array}$$

□

Proof of Proposition 3

From the previous proposition, we know that the change program has models. Now, that program can be partitioned into two programs, the first one containing the stabilizing rules and modified versions of the triggering rules, where the literals of the form *not* p in the bodies are replaced by $p^\star$. The other one contains the domain and database facts plus the new rules $p^\star(\bar{X}) \leftarrow \mathit{not}\ p(\bar{X})$. By a result in (Lifschitz and Turner, 1994), the partitioned program has as answer sets the unions of the answer sets of the second program and the answer sets of the first one, where the atoms $p^\star$ are treated as extensional database predicates. The second program is stratified and has one answer set. The first one does not contain weak negation, it is a positive program in that sense, and its minimal models coincide with its answer sets. □

Proof of proposition 4

Notice that the two sets added to S_M on the right-hand side are expected to give an account of the persistence rules that are not included in $\Pi_\Delta(r)$.

Let S'_M be the set added to S_M:

$$\{p'(\bar{a}) \mid p(\bar{a}) \in S_M \text{ and } \neg p'(\bar{a}) \notin S_M\} \cup \{\neg p'(\bar{a}) \mid p(\bar{a}) \notin S_M \text{ and } p'(\bar{a}) \notin S_M\}.$$

It is easy to verify that ${}^S\Pi(r) = {}^{S_M}\Pi_\Delta(r)$. Then, since S_M is an answer set of $\Pi_\Delta(r)$, in order to prove that S is an answer set of $\Pi(r)$, it suffices to prove (I) and (II) below.

(I) $S'_M \subseteq \cap\alpha({}^S\Pi(r))$.
Let $l(\bar{a})$ be an element of S'_M. If $l(\bar{a}) = p'(\bar{a})$, then $p(\bar{a}) \in S_M$ and $\neg p'(\bar{a}) \notin S_M$, and, therefore, $p(\bar{a})$ and $p'(\bar{a}) \leftarrow p(\bar{a})$ are rules in ${}^S\Pi(r)$. Thus, $p'(\bar{a})$ is in $\cap\alpha({}^S\Pi(r))$. If $l(\bar{a}) = \neg p'(\bar{a})$, then $p(\bar{a}) \notin S_M$ and $p'(\bar{a}) \notin S_M$, and, therefore, $\neg p'(\bar{a}) \leftarrow \mathit{dom}(\bar{a})$ is a reduced ground persistence rule in ${}^S\Pi(r)$. Thus, $\neg p'(\bar{a})$ is in $\cap\alpha({}^S\Pi(r))$.

(II) From S'_M is not possible to deduce an element that is not included in S.
Assume that $q'(\bar{Y}) \leftarrow \mathit{dom}(\bar{Y}),\ p'(\bar{X}),\ \ \mathit{not}\ \varphi$ is a rule in $\Pi_\Delta(r)$, and

$q'(\bar{b}) \leftarrow dom(\bar{b}),\ p'(\bar{a})$ is a rule in ${}^{S}\Pi(r)$. If $p'(\bar{a}) \in S'_M$, we need to show that $q'(\bar{b}) \in S$. By contradiction, suppose that $q'(\bar{b}) \notin S$. Then $q'(\bar{b}) \notin S_M$ and $q'(\bar{b}) \notin S'_M$. Therefore, $q(\bar{b}) \notin S_M$ or $\neg q'(\bar{b}) \in S_M$. If $q(\bar{b})$ is not in S_M, then S_M satisfies the body of the rule

$$q'(\bar{b}) \vee \neg p'(\bar{a}) \leftarrow dom(\bar{b}),\ p(\bar{a}),\ \ not\ q(\bar{b}),\ \ not\ \varphi,$$

because $p'(\bar{a}) \in S'_M$. In consequence, $p(\bar{a}) \in S_M$. But, this implies that $q'(\bar{b}) \in S_M$, a contradiction, or $\neg p'(\bar{a}) \in S_M$, also a contradiction (since $p'(\bar{a}) \in S'_M$). Otherwise, if $\neg q'(\bar{b}) \in S_M$, then by using the rule $\neg p'(\bar{a}) \leftarrow dom(\bar{a}),\ \neg q'(\bar{b})$, we can conclude that $\neg p'(\bar{a})$ is in S_M, a contradiction.
Analogously, it is possible to prove the same property for any of the rules:

$$\begin{array}{rcl} \neg q'(\bar{Y}) & \leftarrow & dom(\bar{Y}),\ p'(\bar{X}),\ \ not\ \varphi, \\ q'(\bar{Y}) & \leftarrow & dom(\bar{Y}),\ \neg p'(\bar{X}),\ \ not\ \varphi, \\ \neg q'(\bar{Y}) & \leftarrow & dom(\bar{Y}),\ \neg p'(\bar{X}),\ \ not\ \varphi. \end{array}$$ □

Proof of Lemma 1
Let S be a answer set of $\Pi_\Delta(r)$ such that S is a subset of $S(r,r')$. First, we will prove that $\Delta(r, I(S)) \subseteq \Delta(r,r')$.

If $p(\bar{a}) \in \Delta(r, I(S))$, then one of the following cases holds.

(I) $r \models p(\bar{a})$ and $I(S) \not\models p(\bar{a})$. In this case, $p(\bar{a}) \in S$ and $p'(\bar{a}) \notin S$. Thus, by definition of $I(S)$ we conclude that $\neg p'(\bar{a}) \in S$. Therefore, $\neg p'(\bar{a}) \in S(r,r')$. But this implies that $r' \not\models p(\bar{a})$. Thus, $p(\bar{a}) \in \Delta(r,r')$.

(II) $r \not\models p(\bar{a})$ and $I(S) \models p(\bar{a})$. In this case, $p(\bar{a}) \notin S$ (S is a minimal model and $p(a)$ does not need to be in S if it was not in r). Thus, by definition of $I(S)$ we conclude that $p'(\bar{a}) \in S$. Therefore, $p'(\bar{a}) \in S(r,r')$. But this implies that $r' \models p(\bar{a})$. Thus, $p(\bar{a}) \in \Delta(r,r')$.

Thus, $\Delta(r, I(S)) \subseteq \Delta(r,r')$. But, by proposition 1, $I(S)$ satisfies IC, and therefore, $\Delta(r, I(S))$ must be equal to $\Delta(r,r')$, since $\Delta(r,r')$ is minimal under set inclusion in $\{\Delta(r,r^*) \mid r^* \models IC\}$. Then, we conclude that $I(S) = r'$. □

Proof of Theorem 1
We will prove the first part of this theorem. The second one can be proved analogously.

Given a repair r' of r, by lemma 1, $r' = I(S_M)$, where S_M is an answer set of $\Pi_\Delta(r)$, with $S_M \subseteq S(r,r')$. Define S from S_M as in (3). We obtain that S is an answer set of $\Pi(r)$. By construction of S, $I(S) = I(S_M)$. Furthermore, $I(S) = \{p(a) \mid p'(a) \in S\}$. □

Proof of Proposition 5
Since it is always the case that $W_{\Pi(r)} \subseteqq Core(\Pi(r))$ (Leone *et al.*, 1997), we only need to show that $Core(\Pi(r)) \subseteqq W_{\Pi(r)}$. In consequence, it is necessary to check that whenever a literal $(\neg)p'(a)$ belongs to $Core(\Pi(r))$, where a is tuple of elements in

the domain D and p is a database predicate, $(\neg)p'(a)$ can be fetched into $\mathcal{W}^n_{\Pi(r)}(\emptyset)$ for some finite integer n.

Each literal L in the original database r, in its primed version, will become L or its complement $\overline{L}$ in the answer sets[9]. We will do the proof by cases, considering for a literal $L:\ (\neg)p'(a)$ contained in $Core(\Pi(r))$ all the possible transitions from the original instance to the core: (a) negative to positive. (b) positive to positive. (c) negative to negative. (d) positive to negative. We will prove only the first two cases, the other two are similar.

For each case, again several cases have to be verified according to the different ground program rules that could have made $p'(a)$ get into $Core(\Pi(r))$.

(I) Assume $p'(a) \in Core(\Pi(r))$. To prove: $p'(a) \in W_{\Pi(r)}$. Two cases

1. $p(a) \notin r$. Since FDs can only produce deletions $p'(a)$ has to be true because an unary constraint was false for $p(a)$: $(p(a) \vee \varphi(a)) \in IC_D$ is false, where IC_D is the instantiation of the ICs in the domain D. Then, $\varphi(a)$ is false. In the ground program we find the rule $p'(a) \longleftarrow dom(a), \neg\varphi(a)$. The second subgoal becomes true of $\emptyset$. Since $dom(a) \in \mathcal{W}^1_{\Pi(r)}(\emptyset)$. Then, $p'(a) \in \mathcal{W}^2_{\Pi(r)}(\emptyset)$.
2. $p(a) \in r$. Intuitively, $p(a)$ persisted. This means there is no ground IC of the form $p'(a) \vee \varphi(a)$ that is false, nor of the form $\neg p'(a) \vee \psi(a)$, with $\psi(a)$ false. Otherwise, $p'(a)$ could not be in the core. The second case implies that we can never obtain $\neg p'(a)$ by means of a rule of the form $\neg p'(a) \longleftarrow dom(a), \neg\psi(a)$. Some cases need to be examined:
   (a) There is $(p(a) \vee \varphi(a))$ with $\varphi(a) \in IC_D$ false. Then, with the rule $p'(a) \longleftarrow dom(a), \neg\varphi(a)$, with $\varphi(a)$. In this case, as in case 1., $p'(a) \in \mathcal{W}^2_{\Pi(r)}(\emptyset)$.
   (b) If there is no ground constraint as in the previous item, either because there is no $(p(a) \vee \varphi(a))$ in IC_D or the $\varphi(a)$s are true, then there is no applicable rule of the form $p'(a) \longleftarrow dom(a), \neg\varphi(a)$ in the ground program. Since rules associated to FDs delete tuples only, we could obtain $p'(a)$ due to a default rule $p'(a) \longleftarrow dom(a), p(a), not\ \neg p'(a)$ only, via the unfoundedness of $\neg p'(a)$, or directly via the unfoundedness of $\neg p'(a)$ in the ground program. If the $\mathcal{W}_{\Pi(r)}$ operator declares $\neg p'(a)$ unfounded, then $p'(a)$ will belong to $W_{\Pi(r)}$. So, we have to concentrate on the unfoundeness of $\neg p'(a)$.
      i We know that we can never get $\neg p'(a)$ from rules of the form $\neg p'(a) \longleftarrow dom(a), \neg\psi(a)$.
      ii $\neg p'(a)$ cannot be obtained via the default rule, because $\neg p'(a) \longleftarrow dom(a), not\ p(a), not\ p'(a)$ has the second subgoal false.
      iii $\neg p'(a)$ cannot be obtained via a possible unfoundedness of $p'(a)$, because $p'(a)$ belongs to answer sets.
      iv We are left with rules associated to FDs. Assume that $(\neg p(a) \vee \neg p(b) \vee$

[9] Actually only positive literals appear in r, but we are invoking the CWA. All the literals in the original instance will belong to $Core(\Pi(r))$.

$c = d) \in IC_D$. There are only two alternatives: $c = d$, in which case, the associated triggering rule cannot be applied; or $c \neq d$ and $\neg p(b)$ is true and there is no $(p(b) \vee \chi(b)) \in IC_D$ with $\chi(b)$ false (otherwise, $\neg p'(a)$ would have to be true).
The rule $\neg p'(a) \vee \neg p'(b) \longleftarrow dom(a), dom(b), p(a), p(b), c \neq d$ cannot be applied, because $p(b)$ is false.
We have to analyze the stabilizing rule $\neg p'(a) \longleftarrow p'(b), c \neq d$. If $c = d$, the rule does not apply. Otherwise, we have $p(b) \notin r$, and then $\neg p(b) \in \mathcal{W}^{(}_{\Pi(r)}\emptyset)$. $p'(b)$ cannot be obtained from the default $p'(b) \longleftarrow dom(b), p(b), not\ \neg p'(b)$, because $p(b)$ is false. Neither can it be obtained from a rule $p'(b) \longleftarrow dom(b), \neg\chi(b)$, because $\chi(b)$ would have to be true.
In consequence, $p'(b)$ is unfounded, i.e. $\neg p'(b) \in \mathcal{W}^2_{\Pi(r)}(\emptyset)$, then $\neg p'(a)$ turns out to be unfounded: $p'(a) \in \mathcal{W}^3_{\Pi(r)}(\emptyset)$.

(II) $\neg p'(a) \in Core(\Pi(r))$. The other (similar) cases are:

1. $p(a) \notin r$ and $\neg p'(a) \in Core(\Pi(r))$.
2. $p(a) \in r$ and $\neg p'(a) \in C$.

It is possible to show that always $Core(\Pi(r)) \subseteq \mathcal{W}^3_{\Pi(r)}(\emptyset)$. □