A Appendix

A.1 Proofs of Section 4

Lemma 1. Given a database D and a set IC consisting of UICs, RICs and non-conflicting NNCs, for every repair $D' \in Rep(D, IC)$, $adom(D') \subseteq adom(D) \cup \{null\}$. $\square$

Proof: By contradiction, let us assume there is a repair $D' \in Rep(D, IC)$ with an atom $R(\bar{a}, c)$ such that $c \notin (adom(D) \cup \{null\})$. Since $c \notin adom(D)$, $R(\bar{a}, c) \notin D$ and, therefore $R(\bar{a}, c) \in \Delta(D, D')$. $R(\bar{a}, c)$ could have been added to restore consistency wrt a RIC or UIC (it cannot be for a NNC, because only tuple deletion is used in this case). We have two cases: (a) Constant c is in a existentially quantified attribute. For database $D'' = (D' \cup \{R(\bar{a}, null)\}) \smallsetminus \{R(\bar{a}, c)\}$, we would have $D'' \models IC$ and $D'' <_D D'$, and D' would not be a repair. We have a contradiction. (b) Constant c is in a universally quantified attribute of ICs. Since constraints have form 1, it implies that the atom(s) that created the inconsistency had c in at least one of its attributes. If this atom was part of D, then $c \in adom(D)$, and this is a contradiction. If this atom was added to solve an inconsistency we go back to the beginning of this argument. $\square$

Proof of Proposition 1: First let us prove that $Rep(D, IC)$ is non-empty. Since the set of ICs is finitely and logically consistent, there exists an instance D' compatible with Σ, the schema of D, such that $D' \models IC$. If there is no other instance D'' over Σ, such that $D'' \models IC$ and $D'' \leq_D D'$, then D' is a repair. On the other hand, if there is D'' such that $D'' \models IC$ and $D'' <_D D'$, we can repeat the same argument; etc. Since the instances involved have all finite extensions for the database relations, it is not difficult to show that the partial order $\leq_D$ is well-founded, so there won't be an infinite decreasing chain.

Now let us prove that every $D' \in Rep(D, IC)$ is finite and the number of repairs is also finite. From Lemma 1 we have that the active domain of the repairs is a subset of $adom(D) \cup \{null\}$ (which is finite). Since the number of predicates is also finite, the set of databases that can be candidates to repairs are obtained from all the possible instantiations and combinations of predicates and atoms and therefore the number of repairs is finite and each of them is finite too. $\square$

A.2 Proofs of Section 5

In the proofs of this section we will refer to the rules of definition 8 only by their numbers.

Definition 11. A model $\mathcal{M}$ of a disjunctive program P is stable iff $\mathcal{M}$ is a minimal model of $P^{\mathcal{M}}$, where $P^{\mathcal{M}}$ is defined as $\{A_1, \ldots, A_k \leftarrow B_1, \ldots, B_n \mid A_1, \ldots, A_k \leftarrow B_1, \ldots, B_n,$ $not\ C_1, \ldots,\ not\ C_m$ is a ground instance of a clause in P and, for every $1 \leq i \leq m$, $\mathcal{M} \not\models C_i\}$ $\square$

Lemma 2. If $\mathcal{M}$ is a stable model of $\Pi(D, IC)$, i.e. a minimal model of $(\Pi(D, IC))^{\mathcal{M}}$, then exactly one of the following cases holds:

1. $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t^\star})$ and $P\_(\bar{a}, \mathbf{t^{\star\star}})$ belong to $\mathcal{M}$, and no other $P\_(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$.
2. $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t^\star})$ and $P\_(\bar{a}, \mathbf{f_a})$ belong to $\mathcal{M}$, and no other $P\_(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$.
3. $P(\bar{a}) \notin \mathcal{M}$ and $P\_(\bar{a}, \mathbf{t_a})$, $P\_(\bar{a}, \mathbf{t^\star})$ and $P\_(\bar{a}, \mathbf{t^{\star\star}})$ belong to $\mathcal{M}$, and no other $P\_(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$.
4. $P(\bar{a}) \notin \mathcal{M}$ and no $P\_(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$.

Proof: For an atom $P(\bar{a})$ we have two possibilities:

– $P(\bar{a}) \in \mathcal{M}$. Then, from rule 5, $P\_(\bar{a}, \mathbf{t^\star}) \in \mathcal{M}$. Two cases are possible now: $P\_(\bar{a}, \mathbf{f_a}) \notin \mathcal{M}$ or $P\_(\bar{a}, \mathbf{f_a})b \in \mathcal{M}$. For the first case, since $\mathcal{M}$ is minimal, $P\_(\bar{a}, \mathbf{t_a}) \notin \mathcal{M})$ and $P\_(\bar{a}, \mathbf{t^{\star\star}}) \in \mathcal{M}$. For the second case, because of rule 7, $P\_(\bar{a}, \mathbf{t_a}) \notin \mathcal{M}$. These cases cover the first two items in the lemma.

– $P(\bar{a}) \notin \mathcal{M}$. Two cases are possible now: $P\_(\bar{a}, \mathbf{t_a}) \in \mathcal{M}$ or $P\_(\bar{a}, \mathbf{t_a}) \notin \mathcal{M}$. For the first one we also have $P\_(\bar{a}, \mathbf{t}^{\star\star})$, $P\_(\bar{a}, \mathbf{t}^{\star}) \in \mathcal{M}$ because of rules 5 and 6 and $P\_(\bar{a}, \mathbf{f_a}) \notin \mathcal{M}$ because of rule 7. For the second one, $P\_(\bar{a}, \mathbf{t}^{\star}) \notin \mathcal{M}$ (since $\mathcal{M}$ is minimal), $P\_(\bar{a}, \mathbf{f_a}) \notin \mathcal{M}$ (because $P\_(\bar{a}, \mathbf{t}^{\star}) \notin \mathcal{M}$ and $\mathcal{M}$ is minimal). These cases cover the last two items in the lemma.

□

From two database instances we can define a structure.

Definition 12. For two database instances D_1 and D_2 over the same schema and domain and a set of ICs IC, $\mathcal{M}^{\star}_{IC}(D_1, D_2)$ is the Herbrand structure $\langle D, I_{\mathcal{P}}, I_{\mathcal{B}} \rangle$, where $\mathcal{U}$ is the domain of the database[6] and I_P, I_B are the interpretations for the database predicates (extended with annotation arguments) and the built-ins, respectively. I_P is inductively defined as follows:

1. If $P(\bar{a}) \in D_1$ and $P(\bar{a}) \in D_2$, then $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^{\star})$ and $P\_(\bar{a}, \mathbf{t}^{\star\star}) \in I_P$.
2. If $P(\bar{a}) \in D_1$ and $P(\bar{a}) \notin D_2$, then $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^{\star})$ and $P\_(\bar{a}, \mathbf{f_a}) \in I_P$.
3. If $P(\bar{a}) \notin D_1$ and $P(\bar{a}) \notin D_2$, then $P(\bar{a}, v) \notin I_{\mathcal{P}}$ for all annotated constants v.
4. If $P(\bar{a}) \notin D_1$ and $P(\bar{a}) \in D_2$, then $P\_(\bar{a}, \mathbf{t_a})$, $P\_(\bar{a}, \mathbf{t}^{\star})$ and $P\_(\bar{a}, \mathbf{t}^{\star\star}) \in I_{\mathcal{P}}$.
5. For every RIC $ic \in IC$ of the form $\forall \bar{x} P(\bar{x}) \rightarrow \exists \bar{y} Q(\bar{x}', \bar{y})$:
   – If $P\_(\bar{a}, \mathbf{t}^{\star\star}) \in I_{\mathcal{P}}$, for $\bar{a} \neq null$ and there exists a $\bar{b} \neq null$ such that $Q\_(\bar{a}', \bar{b}, \mathbf{t}^{\star\star}) \in I_{\mathcal{P}}$, then $aux_1(\bar{a}') \in I_{\mathcal{P}}$.
   – If $Q\_(\bar{a}', \bar{b}, \mathbf{t}^{\star\star}) \in I_{\mathcal{P}}$ with $\bar{b}$ not necessarily different from $null$, $aux_2(\bar{a}') \in I_{\mathcal{P}}$.

The interpretation $I_{\mathcal{B}}$ is defined as expected: if q is a built-in, then $Q(\bar{a}) \in I_{\mathcal{B}}$ iff $Q(\bar{a})$ is true in classical logic, and $Q(\bar{a}) \notin I_{\mathcal{B}}$ iff $Q(\bar{a})$ is false. □

Notice that the database associated to $\mathcal{M}^{\star}_{IC}(D_1, D_2)$ corresponds exactly to D_2, i.e. $D_{\mathcal{M}^{\star}_{IC}(D_1,D_2)} = D_2$.

Lemma 3. If $D' \models IC$, then there is a model $\mathcal{M}$ of the program $(\Pi(D, IC))^{\mathcal{M}}$ such that $D_{\mathcal{M}} = D'$. Furthermore, $\mathcal{M}^{\star}_{IC}(D, D')$ is this model.

Proof: As $\mathcal{M}^{\star}_{IC}(D, D') = D'$, we only need to show $\mathcal{M}^{\star}_{IC}(D, D')$ satisfies all the rules of $(\Pi(D, IC))^{\mathcal{M}}$. It is clear, that by construction, rules 1, 5 and 6 are satisfied by $\mathcal{M}^{\star}_{IC}(D, D')$. For every UIC in IC there is a set of rules of form 2. If the body of the rule is satisfied we have that the atoms $P_{i}\_(\bar{a}_i, \mathbf{t}^{\star}) \in \mathcal{M}^{\star}_{IC}(D, D')$ and $Q_{i}\_(\bar{b}_i, \mathbf{f_a}) \in \mathcal{M}^{\star}_{IC}(D, D')$ or $Q_i(\bar{b}_i) \notin \mathcal{M}^{\star}_{IC}(D, D')$. Also, since the constraint is satisfied, we know that at least one of the $P_i(\bar{a}_i)$ is not in D' or one of the $Q_i(\bar{b}_i)$ is in D'. By construction of $\mathcal{M}^{\star}_{IC}(D, D')$ this implies that at least one $P_{i}\_(\bar{a}_i, \mathbf{f_a})$ or $Q_{i}\_(\bar{b}_i, \mathbf{t_a})$ is in $\mathcal{M}^{\star}_{IC}(D, D')$. Therefore, the head of the rule is also satisfied and the whole rule is satisfied. For every RIC in IC there is a set of rules of form 3. By construction of $\mathcal{M}^{\star}_{IC}(D, D')$ the rules that define $aux_1(\bar{x})$ and $aux_2(\bar{x})$ are satisfied. If the body of the first rule in 3 is true in $\mathcal{M}^{\star}_{IC}(D, D')$, it means that the integrity constraint is not satisfied in the original database or at some point of the repair process. Since the constraint is satisfied in D' the satisfaction had to be restored by adding $Q(\bar{x}, \overline{null})$ or by deleting $P(\bar{x})$. This implies that $Q\_(\bar{x}, \overline{null}, \mathbf{t_a}) \in \mathcal{M}^{\star}_{IC}(D, D')$ or $P\_(\bar{x}, \mathbf{f_a}) \in \mathcal{M}^{\star}_{IC}(D, D')$ and therefore that the first (or second) rule is satisfied. For every NNC in IC there is a rule of form 4. If the body of the rule is true, e.g $P\_(\bar{a}, null, \mathbf{t}^{\star}) \in \mathcal{M}^{\star}_{IC}(D, D')$, the constraint is not satisfied at some point in the repair process. Since it is satisfied in D', $P(\bar{a}, null) \notin D'$. Then, by construction of $\mathcal{M}^{\star}_{IC}(D, D')$, $P\_(\bar{a}, null, \mathbf{f_a}) \in \mathcal{M}^{\star}_{IC}(D, D')$ and the head of the rule is also satisfied. □

The next lemma shows that if $\mathcal{M}$ is a minimal model of the program $(\Pi(D, IC))^{\mathcal{M}}$, then $D_{\mathcal{M}}$ satisfies the constraints.

Lemma 4. If $\mathcal{M}$ is a stable model of the program $\Pi(D, IC)$ then $D_{\mathcal{M}} \models IC$.

[6] Strictly speaking, the domain $\mathcal{U}$ now also contains the annotations values.

Proof: We want to show $D_{\mathcal{M}} \models ic$, for every constraint $ic \in IC$. There are three cases to consider:

- If ic is a UIC. Since $\mathcal{M}$ is a model of $(\Pi(D, IC))^{\mathcal{M}}$, we have that $\mathcal{M}$ satisfies rules 2 of $\Pi(D, IC)$. Then, at least one of the following cases holds:
  - $\mathcal{M} \models P_{i\_}(\bar{a}, \mathbf{f_a})$. Then, $\mathcal{M} \not\models P_{i\_}(\bar{a}, \mathbf{t}^{\star\star})$ and $P(\bar{a}) \notin D_{\mathcal{M}}$ (by lemma 2). Hence, $D_{\mathcal{M}} \models \neg P_i(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}} \models \bigwedge_{i=1}^{n} P_i(\bar{x_i}) \rightarrow \bigvee_{j=1}^{m} Q_j(\bar{y_j}) \vee \varphi$ holds.
  - $\mathcal{M} \models Q_{j\_}(\bar{a}, \mathbf{t_a})$. It is symmetrical to the previous one.
  - It is not true that $\mathcal{M} \models \bar{\varphi}$. Then $\mathcal{M} \models \varphi$. Hence, φ is true, and $D_{\mathcal{M}} \models \bigwedge_{i=1}^{n} P_i(\bar{x_i}) \rightarrow \bigvee_{j=1}^{m} Q_j(\bar{y_j}) \vee \varphi$ holds.
  - $\mathcal{M} \not\models P_{i\_}(\bar{a}, \mathbf{t}^{\star})$. Given the model is minimal, just the last item in Lemma 2 holds. This means $\mathcal{M} \not\models P_{i\_}(\bar{a}, \mathbf{t}^{\star\star})$, $P_i(\bar{a}) \notin D_{\mathcal{M}}$ and $D_{\mathcal{M}} \models \neg P_i(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}} \models \bigwedge_{i=1}^{n} P_i(\bar{x_i}) \rightarrow \bigvee_{j=1}^{m} Q_j(\bar{y_j}) \vee \varphi$ holds.
  - $\mathcal{M} \not\models Q_{j\_}(\bar{a}, \mathbf{f_a})$ or $\mathcal{M} \models Q_j(\bar{a})$. Given the model is minimal, just the first item in lemma 2 holds. Then, $\mathcal{M} \models Q_{j\_}(\bar{a}, \mathbf{t}^{\star\star})$, $Q_j(\bar{a}) \in D_{\mathcal{M}}$ and $D_{\mathcal{M}} \models Q_j(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}} \models \bigwedge_{i=1}^{n} P_i(\bar{x_i}) \rightarrow \bigvee_{j=1}^{m} Q_j(\bar{y_j}) \vee \varphi$ holds. □
- If ic is a RIC. Since $\mathcal{M}$ is a model of $(\Pi(D, IC))^{\mathcal{M}}$, we have that $\mathcal{M}$ satisfies rules 3 of $\Pi(D, IC)$. Then, at least one of the following cases holds:
  - $\mathcal{M} \models P_{\_}(\bar{a}, \mathbf{f_a})$. Then, $\mathcal{M} \not\models P_{i\_}(\bar{a}, \mathbf{t}^{\star\star})$ and $P(\bar{a}) \notin D_{\mathcal{M}}$ (by lemma 2). Hence, $D_{\mathcal{M}} \models \neg P_i(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}} \models (P(\bar{x}) \rightarrow Q(\bar{x}', \bar{y}))$ holds.
  - $\mathcal{M} \models Q_{\_}(\bar{a}', \overline{null}, \mathbf{t_a})$. It is symmetrical to the previous one.
  - $\mathcal{M} \not\models P_{\_}(\bar{a}, \mathbf{t}^{\star})$. Given the model is minimal, just the last item in Lemma 2 holds. This means $\mathcal{M} \not\models P_{\_}(\bar{a}, \mathbf{t}^{\star\star})$, $P(\bar{a}) \notin D_{\mathcal{M}}$ and $D_{\mathcal{M}} \models \neg P(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}} \models (P(\bar{x}) \rightarrow Q(\bar{x}', \bar{y}))$ holds.
  - $\mathcal{M} \models aux_1(\bar{a}')$. This means that $P_{\_}(\bar{a}, \mathbf{t}^{\star\star}) \in \mathcal{M}$ and that there exists $\bar{b} \neq null$ such that $Q_{\_}(\bar{a}', \bar{b}, \mathbf{t}^{\star\star}) \in \mathcal{M}$ and therefore that $P(\bar{a}) \in D_{\mathcal{M}}$ and $Q(\bar{a}', \bar{b}) \in D_{\mathcal{M}}$. Then, the constraint is satisfied.
  - $\mathcal{M} \models aux_2(\bar{a}')$. This means that there exists $\bar{b}$ (not necessarily different from $null$) such that $Q_{\_}(\bar{a}', \bar{b}, \mathbf{t}^{\star\star}) \in \mathcal{M}$ and therefore that $Q(\bar{a}', \bar{b}) \in D_{\mathcal{M}}$. Then, the constraint is satisfied
- If ic is a NNC. Since $\mathcal{M}$ is a model of $(\Pi(D, IC))^{\mathcal{M}}$, we have that $\mathcal{M}$ satisfies rules 4 of $\Pi(D, IC)$. Then, at least one of the following cases holds:
  - $\mathcal{M} \models P_{\_}(\bar{a}, \mathbf{f_a})$. Then, $\mathcal{M} \not\models P_{\_}(\bar{a}, \mathbf{t}^{\star\star})$ and $P(\bar{a}) \notin D_{\mathcal{M}}$ (by lemma 2). Hence, $D_{\mathcal{M}} \models \neg P(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}}$ satisfies the constraint.
  - $\mathcal{M} \not\models P_{\_}(\bar{a}, \mathbf{t}^{\star})$. Given the model is minimal, just the last item in Lemma 2 holds. This means $\mathcal{M} \not\models P_{\_}(\bar{a}, \mathbf{t}^{\star\star})$, $P(\bar{a}) \notin D_{\mathcal{M}}$ and $D_{\mathcal{M}} \models \neg P(\bar{a})$. Since the analysis was done for an arbitrary value $\bar{a}$, $D_{\mathcal{M}}$ satisfies the constraint.
  - $\mathcal{M} \models (a_i \neq null)$. Hence $D_{\mathcal{M}}$ satisfies the constraint.

□

Lemma 5. Consider two database instances D and D' over the same schema and domain. If $\mathcal{M}$ is a minimal model of $(\Pi(D, IC))^{\mathcal{M}^\star(D,D')}$, such that $\mathcal{M} \subsetneq \mathcal{M}^\star(D, D')$, then there exists model $\mathcal{M}'$ such that $\mathcal{M}'$ is a minimal model of $(\Pi(D, IC))^{\mathcal{M}'}$ and $D_{\mathcal{M}'} <_D D'$.

Proof: Since $\mathcal{M}$ is a minimal model of $(\Pi(D, IC))^{\mathcal{M}^\star_{IC}(D,D')}$, we have that $P(\bar{a}) \in \mathcal{M}$ iff $P(\bar{a}) \in D$. By how we defined $\mathcal{M}^\star_{IC}(D, D')$ and given $\mathcal{M} \subsetneq \mathcal{M}^\star(D, D')$, the only two ways that both models can differ is that, for some $P(\bar{a}) \in D$, $\{P_{\_}(\bar{a}, \mathbf{f_a})\} \subseteq \mathcal{M}^\star(D, D')$ and none of these atoms belong to $\mathcal{M}$, or for some $P(\bar{a}) \notin D$, $\{P_{\_}(\bar{a}, \mathbf{t_a}), P_{\_}(\bar{a}, \mathbf{t}^{\star}), P_{\_}(\bar{a}, \mathbf{t}^{\star\star})\} \subseteq \mathcal{M}^\star_{IC}(D, D')$ and none of these atoms belong to $\mathcal{M}$. Now, if we use the interpretation rules over $\mathcal{M}$, we will construct a model $\mathcal{M}'$ that is a minimal model of $(\Pi(D, IC))^{\mathcal{M}'}$. From $\mathcal{M}$ the model $\mathcal{M}'$ is constructed as follows:

- If $P(\bar{a}) \in \mathcal{M}$ and $P\_(\bar{a}, \mathbf{f_a}) \notin \mathcal{M}$, then $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{t}^{\star\star}) \in \mathcal{M}'$.
- If $P(\bar{a}) \in \mathcal{M}$ and $P\_(\bar{a}, \mathbf{f_a}) \in \mathcal{M}$, then $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{f_a}) \in \mathcal{M}'$.
- If $P(\bar{a}) \notin \mathcal{M}$ and $P\_(\bar{a}, \mathbf{t_a}) \notin \mathcal{M}$, then nothing is added to $\mathcal{M}'$.
- If $P(\bar{a}) \notin \mathcal{M}$ and $P\_(\bar{a}, \mathbf{t_a}) \in \mathcal{M}$, then $P\_(\bar{a}, \mathbf{t_a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{t}^{\star\star}) \in \mathcal{M}'$.

It is clear that $\mathcal{M}'$ is a coherent and minimal model of $(\Pi(D, IC))^{\mathcal{M}'}$. It just rests to prove that $D_{\mathcal{M}'} <_D D'$. First, we will prove $D_{\mathcal{M}'} \leq_D D'$. Let us suppose $P(\bar{a}) \in \Delta(D, D_{\mathcal{M}'})$. Then, either $P(\bar{a}) \in D$ and $P(\bar{a}) \notin D_{\mathcal{M}'}$ or $P(\bar{a}) \notin D$ and $P(\bar{a}) \in D_{\mathcal{M}'}$. In the first case, $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{f_a})$ are in $\mathcal{M}'$. These atoms are also in $\mathcal{M}$ and, by our assumption, they are also in $\mathcal{M}^\star_{IC}(D, D')$. Hence, $P(\bar{a}) \in \Delta(D, D')$. In the second case, $P\_(\bar{a}, \mathbf{t_a})$ and $P\_(\bar{a}, \mathbf{t}^\star)$ are in $\mathcal{M}'$. These atoms are also in $\mathcal{M}$ and, by our assumption, these are also in $\mathcal{M}^\star_{IC}(D, D')$. Hence, $P(\bar{a}) \in \Delta(D, D')$.

We will now prove $D_{\mathcal{M}'} <_D D'$. We know that for some fact $P(\bar{a})$ there is an annotated version of it which is in $\mathcal{M}^\star_{IC}(D, D')$ and which is not in $\mathcal{M}$. One possible case is $P\_(\bar{a}, \mathbf{f_a})$ is in $\mathcal{M}^\star_{IC}(D, D')$ and not in $\mathcal{M}$. Then, $P(\bar{a}) \in \Delta(D, D')$, but $P(\bar{a}) \notin \Delta(D, D_{\mathcal{M}'})$. The other possible case is that $P\_(\bar{a}, \mathbf{t_a})$ and $P\_(\bar{a}, \mathbf{t}^\star)$ are in $\mathcal{M}^\star_{IC}(D, D')$ and not in $\mathcal{M}$. Then, $P(\bar{a}) \in \Delta(D, D')$, but $P(\bar{a}) \notin \Delta(D, D_{\mathcal{M}'})$ and therefore $D_{\mathcal{M}'} <_D D'$ □

Proposition 2. If D' is a repair of D with respect to IC, then there is a stable model $\mathcal{M}$ of the program $(\Pi(D, IC))^{\mathcal{M}}$ such that $D_{\mathcal{M}} = D'$. Furthermore, the model $\mathcal{M}$ corresponds to $\mathcal{M}^\star_{IC}(D, D')$.

Proof: By Lemma 3 we have $\mathcal{M}^\star_{IC}(D, D')$ is a model of the program $\Pi(D, IC)^{\mathcal{M}^\star_{IC}(D,D')}$. We just have to show it is minimal. Let us suppose by contradiction that there exists a model $\mathcal{M}$ of $(\Pi(D, IC))^{\mathcal{M}^\star_{IC}(D,D')}$ such that it is the case that $\mathcal{M} \subsetneq \mathcal{M}^\star_{IC}(D, D')$. Since $\mathcal{M} \subsetneq \mathcal{M}^\star_{IC}(D, D')$, the model $\mathcal{M}$ contains the atom $P(\bar{a})$ iff $P(\bar{a}) \in D$. Then, we can assume without loss of generality that $\mathcal{M}$ is minimal (if it is not minimal, we can always generate from it a minimal model $\mathcal{M}'$, such that $\mathcal{M}' \subsetneq \mathcal{M}$, by deleting its non-supported atoms).

By Lemma 5, there exists model $\mathcal{M}'$ such that $D_{\mathcal{M}'} <_D D'$) and $\mathcal{M}'$ is a minimal model of $(\Pi(D, IC))^{\mathcal{M}'}$. By Lemma 4, $D_{\mathcal{M}'} \models IC$. This contradicts our fact that D' is a repair. □
□

Proposition 3. If $\mathcal{M}$ is a stable model of $\Pi(D, IC)$ then $D_{\mathcal{M}}$ is a repair of D with respect to IC.

Proof: From Proposition 4, we have $D_{\mathcal{M}} \models IC$. We only need to prove that it is $\leq_D$*-minimal*. Let us suppose there is a database instance D', such that D is a repair of D wrt IC and $D' \leq_D D_{\mathcal{M}}$. From Proposition 2 we have that $\mathcal{M}^\star_{IC}(D, D')$ is a stable model of $\Pi(D, IC)$ and that $D_{\mathcal{M}^\star_{IC}(D,D')} = D'$.

For $D' \leq_D D_{\mathcal{M}}$ to hold, there should be an atom $P(\bar{a})$, with $\bar{a} \in (\mathcal{U} \cup \{null\})$, in $\Delta(D, D_{\mathcal{M}})$ and not in $\Delta(D, D')$ or an atom $P(\bar{a}', \bar{b}) \in \Delta(D, D_{\mathcal{M}})$, with $\bar{a}', \bar{b} \in (\mathcal{U} \smallsetminus \{null\})$, and an atom $P(\bar{a}', \overline{null}) \in \Delta(D, D')$.

1. $P(\bar{a}) \in \Delta(D, D_{\mathcal{M}})$ and $P(\bar{a}) \notin \Delta(D, D')$. Since $P(\bar{a}) \in \Delta(D, D_{\mathcal{M}})$ we have that $P(a, t_a)$ or $P(a, f_a)$ belongs to $\mathcal{M}$. From Lemma 2 we have that there are two options:
   - $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{f_a})$ belong to $\mathcal{M}$, and no other $P(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$. $P(\bar{a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{t}^{\star\star})$ belong to $\mathcal{M}^\star$, and no other $P(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}^\star$.
   - $P\_(\bar{a}, \mathbf{t_a})$, $P\_(\bar{a}, \mathbf{t}^\star)$ and $P\_(\bar{a}, \mathbf{t}^{\star\star})$ belong to $\mathcal{M}$, and no other $P(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}$. No $P(\bar{a}, v)$, for v an annotation value, belongs to $\mathcal{M}^\star$.

   If an atom belongs to a model, e.g. $P\_(\bar{a}, \mathbf{f_a})$, and there is another model in which it is not present, then there must be in this last model an atom annotated with t_a or f_a in order to satisfy the rule that was satisfied in the other model by $P\_(\bar{a}, \mathbf{f_a})$. This implies that $\mathcal{M}^\star$ has an atom annotated with t_a or f_a that does not belong to $\mathcal{M}$. This implies that there is an atom that belongs to $\Delta(D, D')$ and that does not belong to $\Delta(D, D_{\mathcal{M}})$. We have reached a contradiction because $\Delta(D, D')$ is a subset of $\Delta(D, D_{\mathcal{M}})$.

2. $P(\bar{a}',\bar{b}) \in \Delta(D, D_{\mathcal{M}})$ and $P(\bar{a}',\overline{null}) \in \Delta(D, D')$. If $\mathcal{M} \not\models P(\bar{a}',\bar{b})$, we have that $\mathcal{M} \models P\_(\bar{a}',\bar{b},\mathbf{t_a})$, that $\mathcal{M} \not\models P(\bar{a}',\overline{null})$ and that $\mathcal{M} \not\models P\_(\bar{a}',\overline{null},\mathbf{t_a})$. As we want $P(\bar{a}',\overline{null}) \in \Delta(D, D')$ we have that $\mathcal{M}' \models P\_(\bar{a}',\overline{null},\mathbf{t_a})$ which implies that there is a rule representing a RIC in $\mathcal{M}'$ such that $P\_(\bar{a}',\overline{null},\mathbf{t_a})$ is the only atom true in the head, so in $\mathcal{M}$ there must be another atom satisfied by it such that it is in the head of that rule. As $P(\bar{a}',\bar{b})$ can not be that atom (because of the structure of the RICs) then the requirement can not be fulfilled.

Therefore, it is not possible to have $D' <_D D_{\mathcal{M}}$. This implies that $D_{\mathcal{M}}$ is a repair of D. $\Box$

Proof of Theorem 1: From propositions 2 and 3. $\Box$

Proof of Theorem 2: From Theorem 1 we know that the repair program $\Pi(D,\, IC)$ calculates exactly the repairs of database D wrt IC. In order to answer the queries using the repair program, we use the query program. By coupling the query program with $\Pi(D, IC)$ we store in predicate Ans the answers to the query, therefore in order to retrieve the CQA we need to use skeptical semantics of stable models semantics. In [18] it was proven that checking if an atom belongs to every stable models is in $\Sigma_2^p - hard$ for a finite propositional disjunctive logic program. In our case, the repair program together with the query program can be instantiated over $\mathcal{U}$ and since the number of rules in $\Pi(D, IC)$ is finite and the number of elements in $\mathcal{U}$ too, the obtained program is a finite propositional disjunctive logic program and skeptical semantics is decidable $\Box$

A.3 Proofs of Section 6

In order to prove Theorem 3, we need to introduce the following lemma.

Lemma 6. For a set IC of UICs, RICs and NNCs, if there is a cycle in its dependency graph, then there exists at least one bilateral predicate. Furthermore, all the atoms in the cycle correspond to bilateral predicates $\Box$

Proof: First let us analyze which are the relationships between atoms in the dependency graph depending on the type of the constraints. First, note that the database atoms, aux_1, aux_2 and atoms with constant $\mathbf{t}^{\star\star}$, will never be involved in a cycle, because they are exclusively in the head of rules (maybe negated in the body) or exclusively in the body. The only predicates that can be in a cycle are the ones with constants $\mathbf{t}^{\star}$, $\mathbf{t_a}$ and $\mathbf{f_a}$. We will concentrate in this atoms in what follows. The possible edges in the dependency graph between two different atoms in a UIC of form 2 are:

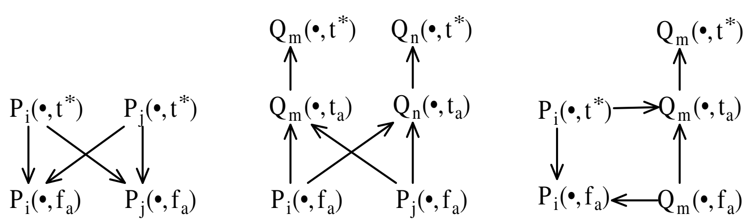

(a) (b) (c)

For a RIC of form 3 the edges are as in case (c) for UICs. For a NNC, since there is unique database atom in it, the relationship is a simplified version of (a) with only predicate P_i. It is clear from the figures that the only way we can have a cycle is by having a a predicate in the consequent of a constraint (as a Q_m) and as a antecedent (as a P_i). It is also easy to see that the predicates of all the atoms in the cycle will be bilateral $\Box$

Proof of of Theorem 3: First let us assume by contradiction that ic has no bilateral predicates but it is not HCF. This implies there is a cycle involving a pair of atoms in the head of a rule of

$\Pi(D, IC)$. But, from Lemma 6 we know that if there is a cycle there is a bilateral predicate. We have reached a contradiction.

Now, let us assume by contradiction ic has exactly one occurrence of a bilateral predicate (without repetitions) but it is not HCF. This implies there is a cycle involving a pair of atoms in the head of a rule of $\Pi(D, IC)$. From Lemma 6 we know then that both atoms should be bilateral predicates. We have reached a contradiction. $\square$