

Corrigendum to “Counting Database Repairs that Satisfy Conjunctive Queries with Self-Joins”

Jef Wijsen

University of Mons, Belgium

Abstract

The helping Lemma 7 in [Maslowski and Wijsen, ICDT, 2014] is false. The lemma is used in (and only in) the proof of Theorem 3 of that same paper. In this corrigendum, we provide a new proof for the latter theorem.

1 The Flaw

The helping Lemma 7 in [MW14] is false. A counterexample is given next.

Example 1. For $\mathbf{S} = \{R, S\}$ and $q = \{R(\underline{x}, y), S(\underline{y})\}$, we have $\text{enc}_{\mathbf{S}}(q) = \{N(\underline{R}, x, y), N(\underline{S}, y, 0)\}$. From [MW14, Lemma 8], it follows that $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}(q))$ is $\#\mathbf{P}$ -hard. From [MW13, Theorem 4], it follows that $\#\text{CERTAINTY}(q)$ is in \mathbf{FP} . Consequently, assuming $\#\mathbf{P} \neq \mathbf{FP}$, there exists no polynomial-time many-one reduction from $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}(q))$ to $\#\text{CERTAINTY}(q)$. Lemma 7 in [MW14] is thus false. \square

The first part in the proof of Lemma 7 in [MW14] is correct; it shows a polynomial-time many-one reduction from $\#\text{CERTAINTY}(q)$ to $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}(q))$. However, the second part in that proof is flawed when it claims “We can compute in polynomial time the (unique) database \mathbf{db}'_0 with schema \mathbf{S} such that $\text{enc}_{\mathbf{S}}(\mathbf{db}'_0) = \mathbf{db}_0$.” The flaw is that the database \mathbf{db}'_0 does not generally exist, as shown next. Let $\mathbf{S} = \{R, S\}$ and $q = \{R(\underline{x}, y), S(\underline{y})\}$, as in Example 1. Then, $\text{enc}_{\mathbf{S}}(q) = \{N(\underline{R}, x, y), N(\underline{S}, y, 0)\}$. A legal input to $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}(q))$ is $\mathbf{db}_0 = \{N(\underline{R}, b, c), N(\underline{S}, c, 0), N(\underline{S}, c, 1)\}$. However, there exists no database \mathbf{db}'_0 such that $\text{enc}_{\mathbf{S}}(\mathbf{db}'_0) = \mathbf{db}_0$. Indeed, for every database \mathbf{db}'_0 with schema \mathbf{S} , if $N(\underline{S}, c, s) \in \text{enc}_{\mathbf{S}}(\mathbf{db}'_0)$, then $s = 0$.

2 The Solution

The following treatment is relative to a database schema \mathbf{S} . Let k, m be non-negative integers such that every relation name in \mathbf{S} has at most k primary-key positions, and at most m non-primary-key positions. We define a new function $\text{enc}_{\mathbf{S}}^*(q)$ which encodes Boolean conjunctive queries q into unirelational Boolean conjunctive queries. For $\text{enc}_{\mathbf{S}}^*(q)$, we use a fresh relation name N with $k + 1$ primary-key positions, and m non-primary-key positions. For every atom $R(\underline{x}, \underline{y})$ in q , the query $\text{enc}_{\mathbf{S}}^*(q)$ will contain some atom $N(\underline{R}, \underline{x}, \underline{0}, \underline{y}, \underline{z})$, where $\underline{0}$ is a sequence of padding zeros, and \underline{z} is a sequence of padding fresh variables, all distinct and not occurring elsewhere. This encoding is different from [MW14, Definition 3] where a sequence of padding zeros was used instead of \underline{z} .

Example 2. We illustrate the difference between the old encoding $\text{enc}_{\mathbf{S}}(\cdot)$ of [MW14, Definition 3] and the newly proposed encoding $\text{enc}_{\mathbf{S}}^*(\cdot)$. For $q_0 = \{R(\underline{x}, y), S(\underline{y})\}$, we have

$$\begin{aligned}\text{enc}_{\mathbf{S}}(q_0) &= \{N(\underline{R}, x, y), N(\underline{S}, y, 0)\}, \\ \text{enc}_{\mathbf{S}}^*(q_0) &= \{N(\underline{R}, x, y), N(\underline{S}, y, z)\}.\end{aligned}$$

We recall from [MW14, p. 156] that the *complex part* of a Boolean conjunctive query contains every atom $F \in q$ such that some non-primary-key position in F contains either a variable with two or more occurrences in q or a constant. Note that $N(\underline{S}, y, 0)$ belongs to the complex part of $\text{enc}_{\mathbf{S}}(q_0)$, while $N(\underline{S}, y, z)$ is not in the complex part of $\text{enc}_{\mathbf{S}}^*(q_0)$. \square

Definition 1. We define **skBCQ** as the class of Boolean conjunctive queries in which all relation names are simple-key. We say that a query $q \in \text{skBCQ}$ is *minimal* if both

- q contains no two distinct atoms $R_1(\underline{x}_1, \vec{y}_1), R_2(\underline{x}_2, \vec{y}_2)$ such that $R_1 = R_2$ and $x_1 = x_2$; and
- there exists no substitution θ over $\text{vars}(q)$ such that $\theta(q) \subsetneq q$.

We define **cxBCQ** as the class of *unirelational* Boolean conjunctive queries q whose relation name has signature $[n, 2]$ (for some $n \geq 2$) such that for every $F \in q$, the first position of F is a constant.

Definition 2. The *intersection graph* of a Boolean conjunctive query is an undirected graph whose vertices are the atoms of q . There is an undirected edge between any two atoms that have a variable in common.

Lemma 1. Assume $\#\mathbf{P} \neq \mathbf{FP}$. For every minimal query q in **skBCQ**, if $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q))$ is $\#\mathbf{P}$ -hard, then so is $\#\text{CERTAINTY}(q)$.

Proof. Let q be a minimal query in **skBCQ** such that $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q))$ is $\#\mathbf{P}$ -hard. Note that q does not need to be unirelational or self-join-free. The query $\text{enc}_{\mathbf{S}}^*(q)$, which is unirelational, is a legal input to the function **IsEasy** of [MW14, p. 163].[†] Since $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q))$ is $\#\mathbf{P}$ -hard, the function **IsEasy** will return **false** on input $\text{enc}_{\mathbf{S}}^*(q)$. This function will repeat, as long as possible, the following step: pick some atom $N(\underline{R}, c, \vec{y})$ and some variable $y \in \text{vars}(\vec{y})$, with R some relation name (treated as a constant) and c some constant, and replace all occurrences of y with an arbitrary constant. Let \bar{q} be the query that results from these steps. Clearly, for every atom $N(\underline{R}, s, \vec{t})$ in \bar{q} , either s is a constant or \vec{t} is variable-free. Since **IsEasy** returns **false** on input \bar{q} , it follows that \bar{q} does not satisfy the premise of [MW14, Lemma 5]. Therefore, it must be the case that \bar{q} contains two distinct atoms $N(\underline{R}, x, \vec{u})$ and $N(\underline{S}, y, \vec{w})$ that are connected in the intersection graph of \bar{q} such that

- R and S are relation names (serving as constants), not necessarily distinct;
- x and y are distinct variables; and
- neither \vec{u} nor \vec{w} is exclusively composed of variables occurring only once in the query. That is, $N(\underline{R}, x, \vec{u})$ and $N(\underline{S}, y, \vec{w})$ belong to the complex part of \bar{q} .

[†]For uniformity of notation, we will assume that the unirelational query uses relation name N .

For every relation name R that appears in q , we assume fresh relation names R_1, R_2, R_3, \dots with the same signature as R . Using these relation names, we can construct a self-join-free Boolean conjunctive query q' such that $|q'| = |q|$ and for every atom $R(\underline{x}, \vec{y})$ in q , the query q' contains some atom $R_i(\underline{x}, \vec{y})$. For example, if $q = \{R(\underline{x}, y), R(\underline{y}, z), S(\underline{z}, x)\}$, then we can let $q' = \{R_1(\underline{x}, y), R_2(\underline{y}, z), S_1(\underline{z}, x)\}$. It can now be shown that the function `IsSafe` in [MW14, p. 158] will return **false** on input q' , and thus $\sharp\text{CERTAINTY}(q')$ is $\sharp\mathbf{P}$ -hard. Indeed, whenever `IsEasy` picked $N(\underline{R}, c, \vec{y})$ and some variable $y \in \text{vars}(\vec{y}) \cap \text{vars}(q)$, the function `IsSafe` can execute SE3 on the corresponding R_i -atom of q' . This eventually leads to a query whose complex part contains two atoms $R_i(\underline{x}, \vec{u}')$ and $S_j(\underline{y}, \vec{w}')$, $x \neq y$, that are connected in the intersection graph, at which point `IsSafe` will return **false**. In this reasoning, one needs that non-primary-key positions are padded with fresh variables occurring only once, as can be seen from Example 2.

In the remainder of this proof, we show the existence of a polynomial-time many-one reduction from $\sharp\text{CERTAINTY}(q')$ to $\sharp\text{CERTAINTY}(q)$. We incidentally note that the remaining reasoning, which generalizes the proof of [MW14, Lemma 2], does not require that relation names are simple-key.

Let f be a mapping from facts to facts such that for every atom $R_i(x_1, \dots, x_n) \in q'$, for every R_i -fact $A := R_i(a_1, \dots, a_n)$, $f(A) := R(\langle a_1, x_1 \rangle, \dots, \langle a_n, x_n \rangle)$. Notice that f maps R_i -facts to R -facts. Here, every couple $\langle a_i, x_i \rangle$ denotes a constant such that $\langle a_i, x_i \rangle = \langle a_j, x_j \rangle$ if and only if both $a_i = a_j$ and $x_i = x_j$. Moreover, if c is a constant, then $\langle c, c \rangle := c$. Since no two distinct atoms of q agree on both their relation name and primary key, it will be the case that for all facts A and B , $A \sim B$ if and only if $f(A) \sim f(B)$, where \sim denotes “is key-equal-to.”

We extend the function f in the natural way to databases \mathbf{db} that use only relation names from q' : $f(\mathbf{db}) := \{f(A) \mid A \in \mathbf{db}\}$. Clearly, $f(\mathbf{db})$ can be computed in polynomial time in the size of \mathbf{db} . Let \mathbf{db} be a set of facts with relation names in q' . It can be easily seen that $|\text{rset}(\mathbf{db})| = |\text{rset}(f(\mathbf{db}))|$ and $\text{rset}(f(\mathbf{db})) = \{f(\mathbf{r}) \mid \mathbf{r} \in \text{rset}(\mathbf{db})\}$. Let \mathbf{r} be an arbitrary repair of \mathbf{db} . It suffices to show that

$$\mathbf{r} \models q' \iff f(\mathbf{r}) \models q.$$

For the implication \implies , assume that $\mathbf{r} \models q'$. We can assume a valuation θ over $\text{vars}(q')$ such that $\theta(q') \subseteq \mathbf{r}$. Let μ be the valuation such that for every variable $x \in \text{vars}(q')$, $\mu(x) = \langle \theta(x), x \rangle$. By our construction of q' and f , it will be the case that $\mu(q) \subseteq f(\mathbf{r})$, thus $f(\mathbf{r}) \models q$.

For the implication \impliedby , assume that $f(\mathbf{r}) \models q$. We can assume a valuation θ over $\text{vars}(q)$ such that $\theta(q) \subseteq f(\mathbf{r})$. Notice that if c is a constant in q , then it must be the case that $\theta(c) = \langle c, c \rangle := c$. We define θ_L as the substitution that maps every variable x in $\text{vars}(q)$ to the first coordinate of $\theta(x)$; and θ_R maps every x to the second coordinate of $\theta(x)$. It is convenient to think of L and R as references to the Left and the Right coordinates, respectively. Thus, by definition, $\theta(x) = \langle \theta_L(x), \theta_R(x) \rangle$.

By inspecting the right-hand coordinates of couples $\langle a_i, x_i \rangle$ in $f(\mathbf{r})$, it can be easily seen that $\theta(q) \subseteq f(\mathbf{r})$ implies $\theta_R(q) \subseteq q$. Since the query q is minimal, it follows that $\theta_R(q) = q$, i.e., θ_R is an automorphism. Since the inverse of an automorphism is an automorphism, θ_R^{-1} is an automorphism as well. Note that θ_R will be the identity on constants that appear in q . We now define $\mu := \theta_L \circ \theta_R^{-1}$ (i.e., μ is the composed function θ_L after the inverse of θ_R), and show that $\mu(q') \subseteq \mathbf{r}$, which implies the desired result that $\mathbf{r} \models q'$. To this extent, let $R_i(x_1, \dots, x_n)$ be an arbitrary atom of q' . It suffices to show $R_i(\mu(x_1), \dots, \mu(x_n)) \in \mathbf{r}$, which can be proved as follows. From $R_i(x_1, \dots, x_n) \in q'$, it follows $R(x_1, \dots, x_n) \in q$. Thus, since θ_R^{-1} is an automorphism,

$$R(\theta_R^{-1}(x_1), \dots, \theta_R^{-1}(x_n)) \in q.$$

Since $\theta(q) \subseteq f(\mathbf{r})$,

$$R(\theta(\theta_R^{-1}(x_1)), \dots, \theta(\theta_R^{-1}(x_n))) \in f(\mathbf{r}).$$

Since, for every symbol s , $\theta(s) = \langle \theta_L(s), \theta_R(s) \rangle$ and $\theta_R(\theta_R^{-1}(s)) = s$, we obtain

$$R(\langle \theta_L(\theta_R^{-1}(x_1)), x_1 \rangle, \dots, \langle \theta_L(\theta_R^{-1}(x_n)), x_n \rangle) \in f(\mathbf{r}).$$

That is, by our definition of μ ,

$$R(\langle \mu(x_1), x_1 \rangle, \dots, \langle \mu(x_n), x_n \rangle) \in f(\mathbf{r}).$$

From this, it is correct to conclude that $R_i(\mu(x_1), \dots, \mu(x_n)) \in \mathbf{r}$. This concludes the proof. \square

Lemma 2. *For every Boolean conjunctive query q , there exists a polynomial-time many-one reduction from $\sharp\text{CERTAINTY}(q)$ to $\sharp\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q))$.*

Proof. Let q be a Boolean conjunctive query. Let R be a relation name that occurs in q . Let $\{R(\vec{x}_i, \vec{y}_i)\}_{i=1}^m$ be the set of R -atoms of q . Then, $\text{enc}_{\mathbf{S}}^*(q)$ will contain, for every $i \in \{1, \dots, m\}$, some atom $N(\underline{R}, \vec{x}_i, \vec{0}, \vec{y}_i, \vec{z}_i)$, where \vec{z}_i is a (possibly empty) sequence of distinct fresh variables not occurring elsewhere. For every R -fact $A := R(\vec{a}, \vec{b})$, we define $f(A) := N(\underline{R}, \vec{a}, \vec{0}, \vec{b}, \vec{0})$. Note here that $f(A)$ depends on the signatures of R and N , but not on the R -atoms of q . The mapping f is defined similarly for all relation names that appear in q . It can be easily seen that for all facts A and B whose relation names appear in q , $A \sim B$ if and only if $f(A) \sim f(B)$.

If \mathbf{db} is an instance of $\sharp\text{CERTAINTY}(q)$, we can assume without loss of generality that every relation name in \mathbf{db} also appears in q . We extend the function f to such instances \mathbf{db} of $\sharp\text{CERTAINTY}(q)$: $f(\mathbf{db}) := \{f(A) \mid A \in \mathbf{db}\}$. Obviously, $f(\mathbf{db})$ can be computed in polynomial time in the size of \mathbf{db} . It is also obvious that $|\text{rset}(\mathbf{db})| = |\text{rset}(f(\mathbf{db}))|$ and $\text{rset}(f(\mathbf{db})) = \{f(\mathbf{r}) \mid \mathbf{r} \in \text{rset}(\mathbf{db})\}$. It suffices to show that for every repair \mathbf{r} of \mathbf{db} ,

$$\mathbf{r} \models q \iff f(\mathbf{r}) \models \text{enc}_{\mathbf{S}}^*(q).$$

For the implication \implies , assume $\mathbf{r} \models q$. We can assume a valuation θ over $\text{vars}(q)$ such that $\theta(q) \subseteq \mathbf{r}$. Let θ' be the valuation that extends θ from $\text{vars}(q)$ to $\text{vars}(\text{enc}_{\mathbf{S}}^*(q))$ such that $\theta'(z) = 0$ for every variable z that appears in $\text{enc}_{\mathbf{S}}^*(q)$ but not in q . By the construction of f , it will be the case that $\theta'(\text{enc}_{\mathbf{S}}^*(q)) \subseteq f(\mathbf{r})$. Indeed, if $\text{enc}_{\mathbf{S}}^*(q)$ contains $N(\underline{R}, \vec{x}_i, \vec{0}, \vec{y}_i, \vec{z}_i)$, then \mathbf{r} will contain $R(\theta(\vec{x}_i), \theta(\vec{y}_i))$, hence $f(\mathbf{r})$ will contain $N(\underline{R}, \theta(\vec{x}_i), \vec{0}, \theta(\vec{y}_i), \vec{0})$ and $\theta'(\vec{z}_i) = \vec{0}$.

For the implication \impliedby , assume $f(\mathbf{r}) \models \text{enc}_{\mathbf{S}}^*(q)$. We can assume a valuation θ over $\text{vars}(\text{enc}_{\mathbf{S}}^*(q))$ such that $\theta(\text{enc}_{\mathbf{S}}^*(q)) \subseteq f(\mathbf{r})$. It is straightforward to see that $\theta(q) \subseteq \mathbf{r}$. \square

We now give the new proof for Theorem 3 in [MW14].

Theorem 1 ([MW14, Theorem 3]). *The set $\{\sharp\text{CERTAINTY}(q) \mid q \in \text{skBCQ}\}$ exhibits an effective FP- $\sharp\text{P}$ -dichotomy.*

New proof. Let $q \in \text{skBCQ}$. It can be decided whether q can be satisfied by a consistent database. If q cannot be satisfied by a consistent database, then for every database \mathbf{db} , the number of repairs of \mathbf{db} satisfying q is 0. An example is $q = \{R(\underline{x}, 0), R(\underline{x}, 1)\}$. Assume next that q can be satisfied by a consistent database. Then, we can compute a minimal query q_m such that for every database,

the number of repairs satisfying q_m is equal to the number of repairs satisfying q . That is, the problems $\#\text{CERTAINTY}(q_m)$ and $\#\text{CERTAINTY}(q)$ are identical.

Then, $\text{enc}_{\mathbf{S}}^*(q_m)$ belongs to $\text{c}\times\text{BCQ}$. By [MW14, Lemma8], the set $\{\#\text{CERTAINTY}(q) \mid q \in \text{c}\times\text{BCQ}\}$ exhibits an effective \mathbf{FP} - $\#\mathbf{P}$ -hard dichotomy. If the problem $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q_m))$ is in \mathbf{FP} , then $\#\text{CERTAINTY}(q)$ is in \mathbf{FP} by Lemma 2; and if $\#\text{CERTAINTY}(\text{enc}_{\mathbf{S}}^*(q_m))$ is $\#\mathbf{P}$ -hard, then $\#\text{CERTAINTY}(q)$ is $\#\mathbf{P}$ -hard by Lemma 1. Consequently, $\#\text{CERTAINTY}(q)$ is in \mathbf{FP} or $\#\mathbf{P}$ -hard, and it is decidable which of the two cases applies. \square

References

- [MW13] Dany Maslowski and Jef Wijsen. A dichotomy in the complexity of counting database repairs. *J. Comput. Syst. Sci.*, 79(6):958–983, 2013.
- [MW14] Dany Maslowski and Jef Wijsen. Counting database repairs that satisfy conjunctive queries with self-joins. In Nicole Schweikardt, Vassilis Christophides, and Vincent Leroy, editors, *Proc. 17th International Conference on Database Theory (ICDT), Athens, Greece, March 24-28, 2014.*, pages 155–164. OpenProceedings.org, 2014.