[28] W. W. McCune and L. J. Henschen. Maintaining state constraints in relational databases: A proof theoretic basis. *Journal of the ACM* 36 (1989), 46–68.

[29] J.-M. Nicolas. Logic for improving integrity checking in relational data bases. *Acta Informatica*, 18:227–253, 1982.

[30] J. Nurmonen. On winning strategies with unary quantifiers. *J. Logic and Computation*, 6 (1996), 755–778.

[31] X. Qian. An effective method for integrity constraint simplification. In *Fourth International Conference on Data Engineering*, 1988.

[32] X. Qian. An axiom system for database transactions. *Information Processing Letters* 36 (1990), 183–189.

[33] X. Qian. The expressive power of the bounded-iteration construct. *Acta Informatica* 28(7):631–656, October 1991.

[34] J. G. Rosenstein. *Linear Orderings*. Academic Press, New York, 1982.

[35] T. Sheard and D. Stemple. Automatic verification of database transaction safety. *ACM Trans. Database Systems* 14 (1989), 322–368.

[36] T. Schwentick. On winning Ehrenfeucht games and monadic NP. *Annals of Pure and Applied Logic* 79 (1996), 61–92.

[37] D. Stemple, S. Mazumdar and T. Sheard. On the modes and meaning of feedback to transaction designers. In *Proceedings of ACM-SIGMOD 1987 International Conference on Management of Data*, pages 374–386, 1987.

[11] E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivations of programs. *Comm. ACM* 18 (1975), 453–457.

[12] E. W. Dijkstra. *A Discipline of Programming.* Prentice-Hall, 1976

[13] K. Etessami. Counting quantifiers, successor relations, and logarithmic space. *Journal of Computer and System Sciences* 54 (1997), 400–411.

[14] R. Fagin. Finite model theory – a personal perspective. *Theoretical Computer Science* 116 (1993), 3–31.

[15] R. Fagin. Easier ways to win logical games. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, American Mathematical Society, vol. 31, 1997, pages 1–32.

[16] R. Fagin. Comparing the power of games on graphs. *Mathematical Logic Quarterly* 43 (1997), 431–455.

[17] R. Fagin, L. Stockmeyer, and M. Vardi. On monadic NP vs. monadic co-NP. *Information and Computation* 120 (1995), 78–92.

[18] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium '81*, pages 105–135. North-Holland, 1982.

[19] S. Grumbach and C. Tollu. On the expressive power of counting. *Theoretical Computer Science*, 149 (1995), 67–99.

[20] Y. Gurevich. Toward logic tailored for computational complexity. In *Proceedings of Computation and Proof Theory*, Springer Lecture Notes in Mathematics, vol. 1104, 1984, pages 175–216.

[21] L. J. Henschen, W. W. McCune, and S. A. Naqvi. Compiling constraint-checking programs from first-order formulas. In H. Gallaire, J. Minker, and J. Nicolas, editors, *Advances in Database Theory*, pages 145–170. Plenum Press, 1984.

[22] A. Hsu and T. Imielinski. Integrity checking for multiple updates. In *Proceedings of ACM-SIGMOD 1985 International Conference on Management of Data*, pages 152–168, 1985.

[23] N. Immerman. Languages that capture complexity classes. *SIAM J. Comput.* 16 (1987), 760–778.

[24] N. Immerman and E. Lander. Describing graphs: a first-order approach to graph canonization. In A. Selman, ed., *Complexity Theory Retrospective*, Springer, Berling, 1990, pages 59–81.

[25] N. Immerman, S. Patnaik and D. Stemple. The expressiveness of a family of finite set languages. *Theoretical Computer Science*, 155(1) (1996), 111–140.

[26] D. Karabeg and V. Vianu. Simplification rules and complete axiomatization for relational update transactions. *ACM Trans. Database Systems* 16 (1991), 439–475.

[27] L. Libkin and L. Wong. Query languages for bags and aggregate functions. *Journal of Computer and System Sciences* 55 (1997), 241–272.

try to test its safety.

As mentioned in the introduction, we are interested in transforming a verifiable transaction T into a safe transaction $\underline{\mathsf{if}}\ \mathsf{wpc}(T,\alpha)\ \underline{\mathsf{then}}\ T\ \underline{\mathsf{else}}\ \mathit{abort}$ which will maintain α as an invariant. As pointed out in [31, 21, 22, 28, 29], assuming that α is always true, it may be possible to find a Δ, which is much simpler than $\mathsf{wpc}(T,\alpha)$, such that $\alpha \rightarrow (\Delta \leftrightarrow \mathsf{wpc}(T,\alpha))$. Using this we can transform T to $\underline{\mathsf{if}}\ \Delta\ \underline{\mathsf{then}}\ T\ \underline{\mathsf{else}}\ \mathit{abort}$ which is more efficient. We are interested in studying classes of transactions for which such a simplification can be efficiently carried out.

Acknowledgements: We would like to thank Neil Immerman and Scott Weinstein for their helpful comments during earlier stages of this work. We are grateful to the reviewers for a number of valuable suggestions; we especially thank the reviewer who found an error in an earlier version of the proof of Theorem 3. Thanks to Ron Fagin for clarifying the use of the Ajtai-Fagin games and pointing out [16], and to Moshe Vardi for bringing the results of [30] to our attention.

References

[1] S. Abiteboul, R. Hull and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1994.

[2] S. Abiteboul and V. Vianu. A transaction-based approach to relational database specification. *Journal of the ACM* 36 (1989), 758–789. Extended abstract in *Proceedings of the 4th Symposium on Principles of Database Systems*, 1985, pages 193–204.

[3] S. Abiteboul and V. Vianu. Procedural languages for database queries and updates. *Journal of Computer and System Sciences* 41 (1990), 181–229.

[4] S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences* 43 (1991), 62–124.

[5] M. Ajtai and R. Fagin. Connectivity is harder for directed than for undirected graphs. *Journal of Symbolic Logic* 55 (1990), 113–150.

[6] R.J.R. Back. Proving total correctness of nondeterministic programs in infinitary logic. *Acta Informatica* 15 (1981), 233–249.

[7] M. Benedikt, G. Dong, L. Libkin and L. Wong. Relational expressive power of constraint query languages. *Journal of the ACM*, 45 (1998), to appear. Extended abstract in *Proceedings of the 15th Symposium on Principles of Database Systems*, 1996, pages 5–16.

[8] M. Benedikt, T. Griffin and L. Libkin. Verifiable properties of database transactions. In *Proceedings of the 15th Symposium on Principles of Database Systems*, 1996, pages 117–127.

[9] R. Berghammer, B. Elbl and U. Schmerl. Formalizing Dijkstra's predicate transformer wp in weak second-order logic. *Theoretical Computer Science* 146 (1995), 185–197.

[10] V. Breazu-Tannen and R. Subrahmanyam. Logical and computational aspects of programming with sets/bags/lists. In *LNCS 510: Proc. of 18th ICALP, Madrid, Spain, July 1991*, pages 60–75. Springer Verlag, 1991.

where each $Q_k z_k$ is a quantifier binding variable z_k, and each B_i is a conjunction of literals A_{ij} such that any A_{ij} with a variable x free is of the form $x = \tau(\vec{z})$ or $x \neq \tau(\vec{z})$, where x is not free in $\tau(\vec{z})$. Let Γ be all terms that appear in some B_i.

Claim 2: For every graph G and any (x, y) in $\mathcal{U} \times \mathcal{U}$, if $G \models \beta(x, y)$ then x and y are in $\Gamma(G)$.

Proof of Claim 2. If this is not the case, fix G, x_0 and y_0, with (WLOG) y_0 not in $\Gamma(G)$. Let y_1 be any other element of $\mathcal{U} - \Gamma(G)$. Then $G \models \beta(x_0, y_1)$, since for each $\tau \in \Gamma$ and for any $z_1, \ldots z_k$ nodes of G, $G \models y_0 = \tau(z_1, \ldots, z_k)$ if and only if $G \models y_1 = \tau(z_1, \ldots, z_k)$. But this gives infinitely many edges satisfying $\beta(x, y)$, contradicting the properties of β. This proves Claim 2.

Claim 2 shows that Γ and $\phi_E = \beta$ satisfy the conclusion of Lemma 9. Now Lemma 9 along with Claim 1 complete the proof of the theorem. □

Corollary 5 *$\mathcal{PR}(\mathrm{FO}_c(\Omega))$ is the maximal robustly verifiable language over $\mathrm{FO}_c(\Omega)$.* □

6 Concluding remarks

In this paper we have looked at the problem of verifying transaction safety *before* transactions are executed. The main results can be summarized as follows. If integrity constraints are specified in first-order logic, then for first-order transaction languages it is possible to compute both weakest preconditions and prerelations. However, if a transaction language has a mechanism for doing recursion, then such ability is generally lost. There are still some transactions that are not first-order definable but have weakest preconditions. However, if we require that the ability to calculate weakest preconditions be independent of extensions to the language, then we can use only languages that admit prerelations.

The last statement, which is a reformulation of corollary 5 can be interpreted as follows. If we are interested in designing a "nice" transaction language that is verifiable over FO_c, and the verification algorithm can be extended to encompass additions to the signature, then we cannot hope that the language will be more expressive than the first-order transaction language defined in [32, 33].

In this paper we only discussed whether it is possible to compute preconditions or prerelations, and what are the implications of our ability to compute them. The algorithmic aspects of computing preconditions were left unexplored. In particular, we would like to address in our future work the problem of *efficiently* computing preconditions. We believe it is of practical importance to identify fragments of transaction languages such as those in [3, 4, 32] for which computing preconditions can be done efficiently (say, in polynomial time).

Computing preconditions may depend on a form in which transactions are specified. For simple relational transactions it is possible to reason about their equivalence [26]. Combining this reasoning together with the theorem proving approach of [35] is an interesting direction to pursue. That is, first we may try to find a better analyzable transaction which is equivalent to the original one, and then

To finish the proof, it therefore will suffice to prove that

$$\langle T(G), p_1 \cup (\mathit{dom}(p) \bigotimes \{(x,y)\})\rangle \models \mathit{Hasmax} \Leftrightarrow \langle T(G), p_1 \cup ((\mathit{dom}(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \mathit{Hasmax}$$

But since (x,y) is not an edge of $T(G)$, any pair (s,t) is p-maximal in either of the above models if and only if it is p-maximal in $\langle T(G), p_1\rangle$. Hence both the above statements are equivalent to $\langle T(G), p_1\rangle \models \mathit{Hasmax}$. So $\langle G, x, y\rangle$ fails the second disjunct in *EDEF* as well as the first, and the proof of claim 1 is complete. □

To complete the proof, we need the following lemma.

Lemma 9 *EDEF is definable by a formula $\phi_E(x,y)$ along with a set of terms Γ. That is, there is a formula $\phi_E(x)$ and a set of $\mathrm{FO}_c(\Omega)$-terms Γ such that*

$$\langle G, x, y\rangle \in \mathit{EDEF} \;\Leftrightarrow\; x, y \in \Gamma(G) \;\; \textit{and} \;\; G \models \phi_E(x,y)$$

Proof of lemma. $T \in \mathcal{WPC}(\mathrm{FO}_c(\Omega))$ allows us to write out the first disjunct of *EDEF* in L as

$$\bigvee_{(a,b)\in \mathit{dom}(p)} \mathsf{wpc}(T, \exists x \exists y.(x = a \wedge y = b)) \wedge (x = a \wedge y = b).$$

The second disjunct in *EDEF* is composed of two conjuncts. The first conjunct is just the negation of $\bigvee_{(a,b)\in \mathit{dom}(p)}(x = a \wedge y = b)$. To handle the second conjunct we first obtain a L-formula $\gamma(x,y)$ such that

$$\langle G, p_1 \cup (\mathit{dom}(p) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow G \models \gamma(x,y)$$

by taking ϕ and replacing every atomic formula of the form $\underline{\mathrm{P}}(\tau_1, \tau_2, \tau_3, \tau_4)$ by $(\bigvee_{(a,b,c,d)\in p_1}(\tau_1 = a) \wedge (\tau_2 = b) \wedge (\tau_3 = c) \wedge (\tau_4 = d)) \vee (\bigvee_{(e,f)\in \mathit{dom}(p)}(\tau_1 = e) \wedge (\tau_2 = f) \wedge (\tau_3 = x) \wedge (\tau_4 = y))$.

One can analogously get a L-formula $\kappa(x,y)$ such that

$$\langle G, p_1 \cup ((\mathit{dom}(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow\; G \models \kappa(x,y)$$

Putting these two together we can express the second conjunct in L, which shows that *EDEF* is L-definable by some formula $\beta(x,y)$.

Notice that the above construction yields a $\beta(x,y)$ with the following properties. First, every atomic subformula containing one of the variables x or y is of the form $x = \tau(\vec{z})$ or $y = \tau(\vec{z})$ for some term τ,where $\vec{z}$ contains no variables free in β. Second, $\beta(x,y)$ has only finitely many (x,y) satisfying it for every graph G. Now we will use these properties of β to construct the required set Γ.

Assume *EDEF* is nonempty. We can write $\beta(x,y)$ as

$$Q_1 z_1 \ldots Q_n z_n. \bigvee_{i<m} B_i$$

In the above, $A \bigotimes B = \{(x,y,w,z) \in \mathcal{U} \times \mathcal{U} \times \mathcal{U} \times \mathcal{U} \mid (x,y) \in A \wedge (w,z) \in B\}$, $p_1 = p^{-1}(1)$ and DB stands for the set of finite graphs whose nodes are from $\mathcal{U}$.

We will show that $EDEF$ is definable by a formula $\phi_E(x,y)$ and set of terms Γ and that $EDEF$ is exactly the prerelation for T of of G. That is, we prove the following.

Claim 1: $\langle G, x, y\rangle \in EDEF \Leftrightarrow (x,y) \in T(G)$

Proof of Claim 1. First suppose $(x,y) \in T(G)$. We will show that $\langle G, x, y\rangle \in EDEF$.

If $(x,y) \in dom(p)$, then by the first disjunct in $EDEF$, $\langle G, x, y\rangle \in EDEF$, so assume $(x,y) \notin dom(p)$. Since $(x,y) \notin dom(p)$, the characteristic function of $p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})$ gives an extension of p. Since p is ϕ-good, this means that

$$\langle G, p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow\; \langle T(G), p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models Hasmax$$

and

$$\langle G, p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow\; \langle T(G), p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models Hasmax$$

Note that since $(x,y) \in T(G)$, we have for each $(s,t) \in dom(p)$,

$$\langle T(G), p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \neg P\text{-}maximal(s,t)$$

and

$$\langle T(G), p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models \neg P\text{-}maximal(s,t)$$

Also note that since $(x,y) \notin dom(p)$, and $dom(p) \neq \emptyset$, (x,y) is P-maximal in the model $\langle T(G), p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle$. Furthermore, (x,y) will not be P-maximal in the model $\langle T(G), p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle$, since here we have $\underline{P}(x,y,x,y)$ holding.

Using the above we can verify that

$$\langle T(G), p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models Hasmax$$

and

$$\langle T(G), p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \neg Hasmax$$

since (x,y) is the only P-maximal element in the first model, and there are no P-maximal elements in the second model. This shows that the second disjunct in $EDEF$ holds. Hence, $\langle G, x, y\rangle \in EDEF$.

Now suppose $(x,y) \notin T(G)$. We will show $\langle G, x, y\rangle \notin EDEF$. Clearly $\langle G, x, y\rangle$ does not satisfy the first disjunct in $EDEF$. If $(x,y) \in dom(p)$, then it fails the second disjunct as well, and we are done. So suppose $(x,y) \notin dom(p)$. Once again we have

$$\langle G, p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow\; \langle T(G), p_1 \cup (dom(p) \bigotimes \{(x,y)\})\rangle \models Hasmax$$

and

$$\langle G, p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models \phi \;\Leftrightarrow\; \langle T(G), p_1 \cup ((dom(p) \cup \{(x,y)\}) \bigotimes \{(x,y)\})\rangle \models Hasmax$$

Let p' be a finite partial function such that $p' \succ p$ and

$$(2) \qquad \begin{array}{lll} (\tau_1(\vec{s_1}), \ldots, \tau_4(\vec{s_4})) \in P & \rightarrow & p'(\tau_1(\vec{s_1}), \ldots, \tau_4(\vec{s_4})) = 1 \\ (\tau_1(\vec{s_1}), \ldots, \tau_4(\vec{s_4})) \notin P & \rightarrow & p'(\tau_1(\vec{s_1}), \ldots, \tau_4(\vec{s_4})) = 0 \end{array}$$

Here each $\tau_i(\vec{x})$, for $i \in \{1,2,3,4\}$, ranges over terms contained in an atomic subformulae of ϕ or atomic subformulae of *Hasmax*, and $\vec{s_i}$ range over the vectors from $\mathcal{U}$ such that the length of $\vec{s_i}$ equals the arity of τ_i, and each element of $\vec{s_i}$ is a vertex of G or a vertex of $T(G)$.

Then p' will be ϕ-terrible. Condition (2) guarantees that for every P' with $\chi_{P'} \succ p'$, we have

$$\langle T(G), P' \rangle \models \mathit{Hasmax} \;\Leftrightarrow\; \langle T(G), P \rangle \models \mathit{Hasmax}$$

and also

$$\langle T(G), P' \rangle \models \phi \;\Leftrightarrow\; \langle T(G), P \rangle \models \phi$$

This finishes the proof of the proposition. □

To finish Case 1, we inductively construct a function p defining a characteristic function which is ϕ-terrible for every ϕ. That is, let $\langle \phi_i \rangle$ be an enumeration of all the sentences in L^+, and let p_0 be empty, and p_{n+1} be any finite $p' \succ p_n$ such that p' is ϕ_n-terrible. The construction of this sequence can be carried out effectively by Proposition 6. We can also ensure that $\bigcup_n p_n$ defines a total function by throwing the nth element of $\mathcal{U}$ into the domain of p_n if it is not there already.

Let P be the unique set such that $\chi_P \succ p_n$ for all n. Clearly, P is recursive, since the construction of p_n is recursive. If we unwind the definition of what it means to be terrible, we see that for P as above, there can be no sentence of L^+ that holds in a graph G exactly when $\langle T(G), P \rangle \models \mathit{Hasmax}$. Hence T does not admit weakest preconditions over L^+, and this contradicts the fact that T is robustly verifiable, completing the proof in Case 1.

Getting a prerelation if the diagonalization is stalled. This corresponds to the negation of Case 1. That is:

Case 2: There is a ϕ and a p that is ϕ-good.

Fix such a p and ϕ. Without loss of generality, we can assume that p is not empty. Let $\mathit{dom}(p)$ be the set of pairs (s,t) such that there exists $\vec{t}$ in the domain of p that contains both s and t.

Our goal will be to show that there is a formula $\phi_E(x,y)$ and a set of terms Γ which defines the edges of $T(G)$. Consider the set:

$$\begin{array}{lll} \mathit{EDEF} & \stackrel{\text{def}}{=} & \{\langle G, x, y\rangle \in DB \times \mathcal{U} \times \mathcal{U} \mid \\ & & (\ (x,y) \in \mathit{dom}(p) \wedge\ \ (x,y) \in T(G)\ ) \vee \\ & & (\neg(\ (x,y) \in \mathit{dom}(p)\ ) \wedge \\ & & \neg(\langle G, p_1 \cup ((\mathit{dom}(p) \cup \{(x,y)\})\ ) \bigotimes \{(x,y)\})\rangle \models \phi \leftrightarrow \langle G, p_1 \cup (\mathit{dom}(p) \bigotimes \{(x,y)\})\rangle \models \phi)\ )\} \end{array}$$

Let $P\text{-}maximal(s,t)$ be the L^+ formula $(\exists x \exists y.E(x,y) \wedge \underline{\mathrm{P}}(x,y,s,t)) \wedge (\forall x \forall y.\,(E(x,y) \rightarrow \neg \underline{\mathrm{P}}(s,t,x,y)))$. Let

$$Hasmax \;\equiv\; \exists s \exists t.E(s,t) \wedge P\text{-}maximal(s,t)$$

That is, *Hasmax* is the L^+-sentence asserting that there is a P-maximal pair.

If p and q are partial functions, then $q \succ p$ means that q extends p as a function. For any set A, let χ_A be its characteristic function. We use this notation in the following definitions.

Suppose p is a finite partial function from $\mathcal{U}^4$ into $\{0,1\}$. If ϕ is an L^+ sentence, say that:

[1] p is ϕ-good if for all finite P with $\chi_P \succ p$, and for all graphs G,

$$\langle G, P\rangle \models \phi \;\Leftrightarrow\; \langle T(G), P\rangle \models Hasmax$$

[2] say p is ϕ-poor if p is not ϕ-good.

[3] p is ϕ-terrible if there is a graph G such that for all finite P with $\chi_P \succ p$,

$$(*) \qquad \neg(\langle G, P\rangle \models \phi \Leftrightarrow \langle T(G), P\rangle \models Hasmax)$$

Intuitively, a finite partial function p used in the definitions above is to be interpreted as an approximation to a predicate P. The idea of the construction that will follow is that we want our final P not to satisfy $(\langle G, P\rangle \models \phi \Leftrightarrow \langle T(G), P\rangle \models Hasmax)$ for any L^+ sentence ϕ. Then *Hasmax* will not have a definable weakest precondition in L^+, contradicting the assumption that T is robustly verifiable. We construct our solution by inductively constructing a sequence of approximations p_n that guarantee $(*)$ for each ϕ in turn. We can show that we can continue this construction successfully provided that the p we have constructed thus far is not ϕ-good. We show that if we are "stuck" with a ϕ-good p then T has a $\langle \Gamma, \phi_E(x,y)\rangle$ as required.

The proof of the theorem now falls into two cases.

Deriving a contradiction if the diagonalization succeeds. The diagonalization succeeding corresponds to the following being true:

Case 1: Every finite partial function p is ϕ-poor for each ϕ.

We will now show that this leads to a contradiction of robust verifiability.

Proposition 6 *For every L^+ sentence ϕ and every ϕ-poor p, there exists a partial function p' with a finite domain such that $p' \succ p$ and p' is ϕ-terrible. Furthermore, p' can be found effectively given p.*

Proof. Since p is poor for ϕ, there is finite P with $\chi_P \succ p$, and a graph G such that $(*)$ holds. Since we can test property $(*)$, we can find such G and P effectively by listing out all pairs $\langle G, P\rangle$ until we find one such that $(*)$ holds.

The induction steps for boolean connectives go through routinely:

$$\begin{array}{ll} & T(G) \models \gamma_1 \wedge \gamma_2 \\ \Leftrightarrow & T(G) \models \gamma_1 \text{ and } T(G) \models \gamma_2 \\ \Leftrightarrow & G \models \mathrm{WPC}[\gamma_1] \text{ and } G \models \mathrm{WPC}[\gamma_2] \quad \text{by induction hypothesis} \\ \Leftrightarrow & G \models \mathrm{WPC}[\gamma_1] \wedge \mathrm{WPC}(\gamma_2) \\ \Leftrightarrow & G \models \mathrm{WPC}[\gamma_1 \wedge \gamma_2]0 \end{array}$$

For the quantifier case $\gamma \equiv \exists x.\phi(x, \vec{z})$, let $\phi'(x, \vec{z}) = \mathrm{WPC}(\phi)$. For each G and for each $\vec{t} \subseteq \Gamma(G)$, we get

$$\begin{array}{ll} & T(G) \models \exists x.\phi(x, \vec{t}) \\ \Leftrightarrow & \exists v \in \Gamma(G).T(G) \models \phi(v, \vec{t}) \\ \Leftrightarrow & \exists v \in \Gamma(G).G \models \phi'(v, \vec{t}) \text{ by induction} \\ \Leftrightarrow & \exists v \in \Gamma(G), \exists v' \in \Gamma(G) \text{ such that } \ G \models \phi'(v, \vec{t}) \text{ and } \ ((v, v') \in T(G) \text{ or } (v', v) \in T(G)) \\ \Leftrightarrow & G \models \exists \vec{v}.(\bigvee_{\tau \in \Gamma} \phi'(\tau(\vec{v}), \vec{t}) \wedge \exists \vec{v'}.\bigvee_{\rho \in \Gamma}(\phi_E(\tau(\vec{v}), \rho(\vec{v'})) \vee \phi_E(\rho(\vec{v'}), \tau(\vec{v})))) \\ \Leftrightarrow & G \models \mathrm{WPC}[\gamma](\vec{t}) \end{array}$$

which completes the proof of correctness of the WPC algorithm.

For the other direction, assume that T is a robustly verifiable transaction. We will prove that T admits a formula $\phi_E(x, y)$ and a set of terms Γ defining the prerelation as in the definition of $\mathcal{PR}(\mathrm{FO}_c(\Omega))$. That is, we will show that there are Γ and $\phi_E(x, y)$ such that $T(G)$ has edge set $\{(x, y) \mid x, y \in \Gamma(G), G \models \phi_E(x, y)\}$.

Let $\underline{\mathrm{P}}$ be an additional 4-place predicate symbol. Let L^+ be the set of all first-order formulae in the language with symbols for all constants in the universe, all elements of Ω, and the symbol $\underline{\mathrm{P}}$. Let L be $\mathrm{FO}_c(\Omega)$.

Our proof will proceed by attempting to inductively construct a recursive interpretation P of $\underline{\mathrm{P}}$ such that T does not admit definable weakest preconditions over the extension of L by P. Our attempt to construct such a 'bad' P will be a straightforward diagonalization: we will pick a certain sentence *Hasmax* below, and for each possible sentence ϕ we will pick an initial segment of P designed to prevent ϕ from being the precondition of *Hasmax*. Since we know that T is in $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$, this diagonalization cannot succeed; that is, at some point we have constructed a finite initial segment of P that cannot be extended further to prevent a particular sentence ϕ from being the precondition of *Hasmax*. The heart of the argument shows that we can use this initial segment and the sentence ϕ to get a first-order definition of T, thus showing that T is in $\mathcal{PR}(\mathrm{FO}_c(\Omega))$.

The proof is now organized as follows: first, we define the framework for our inductive diagonalization argument. We then show that if this diagonalization were to succeed, we would have a contradiction of robust verifiability for T. Finally, we show that from a partial function at which the diagonalization fails we can construct a prerelation for T.

First, we introduce some new notation. For $\phi(t)$ an L^+ formula, P a subset of $\mathcal{U} \times \mathcal{U} \times \mathcal{U} \times \mathcal{U}$, G a graph (that is, a graph on a finite subset of $\mathcal{U}$), and $t \in \mathcal{U}$, we write $\langle G, P \rangle \models \phi(t)$ if $G \models \phi(t)$, when $\underline{\mathrm{P}}$ is interpreted as P.

language that captures this class. The existence of such a language was proved in proposition 3.

Theorem 8 *For every Ω, the transaction language $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ is robustly verifiable over $\mathrm{FO}_c(\Omega)$, and every transaction robustly verifiable over $\mathrm{FO}_c(\Omega)$ is equivalent to a transaction in $\mathcal{PR}(\mathrm{FO}_c(\Omega))$.*

Proof. As in the previous results, without loss of generality we will assume that the signature consists of a single binary relation. We first show that transactions in $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ are robustly verifiable. Let T be given by a set of terms Γ and formula $\phi_E(x,y)$ giving the prerelation of $E(x,y)$ with respect to T. Without loss of generality, we can assume all terms in Γ have the same arity n. We give an algorithm WPC transforming every formula $\gamma(\vec{z})$ in $\mathrm{FO}_c(\Omega')$ for any extension $\Omega' \supseteq \Omega$ into a $\mathrm{FO}_c(\Omega')$ formula $\mathrm{WPC}[\gamma](\vec{z})$ such that $G \models \mathrm{WPC}[\gamma](\vec{z})$ iff $T(G) \models \gamma(\vec{z})$.

The algorithm is recursive in the logical complexity of γ. If γ is quantifier-free, WPC proceeds by replacing every occurrence of $E(\tau_1,\tau_2)$ where each τ_i is a term, by the formula

$$\begin{array}{ll} & (\exists \vec{y}_1 . \bigvee_{\tau\in\Gamma} \tau(\vec{y}_1) = \tau_1(\vec{z})) \\ \wedge & (\exists \vec{y}_2 . \bigvee_{\tau\in\Gamma} \tau(\vec{y}_2) = \tau_2(\vec{z})) \\ \wedge & \phi_E(\tau_1(\vec{z}), \tau_2(\vec{z})) \end{array}$$

Then define

$$\begin{array}{rcl} \mathrm{WPC}[\neg\gamma] & = & \neg\mathrm{WPC}[\gamma] \\ \mathrm{WPC}[\gamma_1 \vee \gamma_2] & = & \mathrm{WPC}[\gamma_1] \vee \mathrm{WPC}[\gamma_2], \\ \mathrm{WPC}[\gamma_1 \wedge \gamma_2] & = & \mathrm{WPC}[\gamma_1] \wedge \mathrm{WPC}[\gamma_2], \end{array}$$

If $\gamma \equiv \exists x.\phi(x,\vec{z})$, define $\mathrm{WPC}[\gamma]$ to be

$$\begin{array}{l} \exists \vec{y_1} .(\bigvee_{\tau\in\Gamma} \phi'(\tau(\vec{y_1}),\vec{z}) \ \wedge \\ \quad \exists \vec{y_2} .(\ \bigvee_{\rho\in\Gamma} \phi_E(\rho(\vec{y_2}),\tau(\vec{y_1})) \vee \phi_E(\rho(\vec{y_2}),\tau(\vec{y_1})))), \end{array}$$

where $\phi'(x,\vec{z}) = \mathrm{WPC}[\phi]$. First, we must show that the algorithm WPC is correct. We will show that for each T in $\mathcal{PR}(\mathrm{FO}_c(\Omega))$, given by Γ, ϕ_E, and for each formula $\gamma(\vec{z})$ in $\mathrm{FO}_c(\Omega')$, and for each graph $G = \langle X, E\rangle$, and a vector $\vec{t}$ of elements from $\Gamma(G)$, the following holds:

$$G \models \mathrm{WPC}[\gamma](\vec{t}) \ \Leftrightarrow \ T(G) \models \gamma(\vec{t})$$

Taking γ to be a sentence, we get as a corollary to this lemma that WPC witnesses the effective verifiability of $\mathcal{PR}(\mathrm{FO}_c(\Omega))$. The proof proceeds by induction on the complexity of γ.

If $\gamma(\vec{t})$ is atomic, then $\gamma(\vec{t})$ is either of the form $E(t_1,t_2)$, with each t_i an element of $\Gamma(G)$ or contains no occurrences of E. In the first case we get the required equivalence immediately from the definition of $\mathcal{PR}$:

$$T(G) \models E(t_1,t_2) \ \Leftrightarrow \ t_1, t_2 \in \Gamma(G) \text{ and } G \models \phi_E(t_1,t_2) \ \Leftrightarrow G \ \models \mathrm{WPC}[E(t_1,t_2)]$$

In the second case, $\mathrm{WPC}[\gamma] = \gamma$, so the equivalence is trivial.

any constants), $G \models \beta(a,b)$ iff $\pi(G) \models \beta(\pi(a),\pi(b))$. Since $\pi(X) \cap C = \emptyset$, $\pi(G) \models \beta(\pi(a),\pi(b))$ iff $(\pi(a),\pi(b)) \in T(\pi(G))$ which, by genericity, is equivalent to $(a,b) \in T(G)$. Thus, $G \models \beta(a,b)$ iff $(a,b) \in T(G)$, proving that T admits prerelations. □

Since the formula β constructed in the proof above does not mention any constant, we can also conclude that there are no generic transactions in $\mathcal{WPC}(\mathrm{FO}_c) - \mathcal{PR}(\mathrm{FO})$.

5 Robustly verifiable transactions

Summarizing the results of the previous sections, we can say that many non-first-order transactions do not have weakest preconditions. Nevertheless, it is possible to find some non-first-order transactions that do have them. However, there is a gap between FO and FO_c. While we can find "nice" transactions that are not first-order definable and have weakest preconditions over FO, no such transactions exist when FO is replaced by FO_c. In particular, the Datalog$^\neg$-definable transaction that provided the separation for $\mathcal{WPC}(\mathrm{FO})$ and $\mathcal{PR}(\mathrm{FO})$ no longer has preconditions if the language is FO_c. While this follows from proposition 4, we give a direct proof that illustrates the problems caused by the presence of constants.

Proposition 5 *The transaction $T \in \mathcal{WPC}(\mathrm{FO}) - \mathcal{PR}(\mathrm{FO})$ from the proof of theorem 7 is not in $\mathcal{WPC}(\mathrm{FO}_c)$.*

Proof. Let c be a constant symbol for some element of the universe. Assume that $T \in \mathcal{WPC}(\mathrm{FO}_c)$. Let α be a formula in the language that consists of E and c defined as $(\exists x \exists y.E(x,y) \wedge (x \neq y)) \wedge (\forall x.\neg E(x,c) \wedge \neg E(c,x))$; this formula says that the graph has at least one edge which is not a loop, and c is not one of its nodes. Let $\beta = \mathsf{wpc}(T,\alpha)$ and $\gamma \equiv \beta \wedge (\exists x.E(x,c) \vee E(c,x))$. Then $G \models \gamma$ iff $G \models \Psi_{\mathrm{c\&c}}$ (since $T(G)$ has an edge which is not a loop), c is one of its nodes but c is not a node in $T(G)$. That is, G is a C&C-graph but not a chain. In other words, for the graphs in which one node is the interpretation of c, testing for chains would be FO_c-definable. However, this is impossible, and hence $T \notin \mathcal{WPC}(\mathrm{FO}_c)$. □

Proposition 5 shows that a transaction verifiable over a language is not necessarily verifiable over even a simple extension of the language. This motivates the following definition. We call a transaction robustly verifiable if it is verifiable in every extension of the language. Below, we show that for FO_c as a specification language (in fact, even for $\mathrm{FO}_c(\Omega)$), a precise characterization of robust transactions is possible – these are exactly the $\mathrm{FO}_c(\Omega)$ definable ones.

Formally, a transaction T is *robustly verifiable* over $\mathrm{FO}_c(\Omega)$ if for every extension of Ω to Ω' with a recursive collection of recursive predicates and functions, it is the case that $T \in \mathcal{WPC}(\mathrm{FO}_c(\Omega'))$. A transaction language $\mathcal{TL}$ is robustly verifiable if all its transactions are.

The following result, together with proposition 3, gives the proof of Theorem E. In what follows, we use $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ to mean both the class of transactions that admit prerelations and the transaction

Sequential languages that define the same classes of transactions can be found in the literature, see for example [32, 33].

From proposition 3 and theorem 5 we obtain

Corollary 4 *For any Ω, there are transactions in $\mathcal{WPC}(\mathrm{FO}_c(\Omega)) - \mathcal{PR}(\mathrm{FO}_c(\Omega))$.* □

However, in contrast to pure first-order logic, there are no generic separating transactions if constants are present in the language.

Proposition 4 *There are no generic transactions in $\mathcal{WPC}(\mathrm{FO}_c) - \mathcal{PR}(\mathrm{FO}_c)$.*

Proof. As before, we assume without loss of generality that we are dealing with the language of graphs. Let T be a generic transaction in $\mathcal{WPC}(\mathrm{FO}_c)$. Our goal is to show that T admits prerelations. First notice that for any graph G with the set of nodes X, the set of nodes of $T(G)$, Y, is a subset of X. Indeed, assuming the existence of $a \in Y - X$, pick any $b \in \mathcal{U} - X$, $b \neq a$, and consider a permutation π that permutes a and b and leaves everything else fixed. Since $G = \pi(G)$, by genericity we obtain that b is also a node of $T(G)$. Since $\mathcal{U}$ is infinite, this means that Y is infinite, which is impossible. Thus, $Y \subseteq X$. Consequently, defining prerelations for T, we can take Γ to contain a single variable. In other words, it is enough to find a formula $\beta(x, y)$ such that $G \models \beta(a, b) \Leftrightarrow (a, b) \in T(G)$ for every $a, b \in X$.

Let c and d be arbitrary distinct elements of the universe, and let Ψ be a FO_c-sentence expressing $\mathsf{wpc}(T, E(c, d))$. Let Φ be a sentence expressing $\mathsf{wpc}(T, E(c, c))$. Now we consider formulae $\psi(x, y)$ and $\phi(x)$ in two and one free variables respectively that are obtained from Ψ and Φ by replacing each occurrence of c by x and each occurrence of d by y. Define

$$\gamma(x, y) \equiv (x = y \wedge \phi(x)) \vee (\neg(x = y) \wedge \psi(x, y)).$$

Let C be the set of all constants that occur in γ. Note that $c, d \notin C$.

We now claim that for any graph G whose set of nodes X is disjoint from C, the formula γ defines prerelations. That is, $G \models \gamma(a, b)$ iff $(a, b) \in T(G)$ for every $a, b \in X$. Consider the case $a \neq b$; the case of $a = b$ is similar. Let π be a permutation of the universe that permutes a and c, b and d, and leaves everything else fixed. Then by genericity $(a, b) \in T(G)$ iff $(c, d) \in T(\pi(G))$. Since $\Psi = \psi(c, d)$ is a weakest precondition, $(c, d) \in T(\pi(G))$ iff $\pi(G) \models \psi(c, d)$, which is equivalent to $G \models \psi(a, b)$, since π leaves everything outside of $\{a, b, c, d\}$ fixed and ψ does not use any of these constants. This proves our claim.

Now consider a new formula $\beta(x, y)$ obtained from $\gamma(x, y)$ as follows: any atomic subformula involving a constant k from C (which is of the form $z = k$ or $E(k, z)$ or $E(z, k)$ where z is either x, or y, or a bound variable) is replaced by *false*. Note that β is now a FO-formula. Then for any graph G with the node set X such that $X \cap C = \emptyset$, we have $G \models \gamma(a, b)$ iff $G \models \beta(a, b)$ for any $a, b \in X$.

Let G now be an arbitrary graph on X. Consider $a, b \in X$ and any permutation π of $\mathcal{U}$ such that $\pi(X) \cap C = \emptyset$ (it exists because $\mathcal{U}$ is infinite). Since β is a FO-formula (in particular, does not mention

Let $N_1 \cup N_2$ be the partition of $\{0, 1, \ldots, n-1\}$ such that for all $i \in N_1$: $L_i \models \alpha'$ and for all $j \in N_2$: $L_j \models \neg\alpha'$. Let p_i^0 be the sentence saying that the chain part of the input has precisely i elements; it can be defined as $p_i \wedge \neg p_{i+1}$. Then β is found as follows:

$$\beta \equiv (\bigvee_{i \in N_1} p_i^0 \wedge \bigwedge_{j \in N_2} \neg p_j^0) \vee p_n$$

if for all $m \geq n$: $L_m \models \alpha'$, and

$$\beta \equiv (\bigvee_{i \in N_1} p_i^0 \wedge \bigwedge_{j \in N_2} \neg p_j^0) \wedge \neg p_n$$

if for all $m \geq n$: $L_m \models \neg\alpha'$.

Finally, we observe that the weakest precondition β given above is computable from α. Indeed, it is possible to verify if $\psi^{(r)}(x)$ holds in one-node diagonal graphs, and for case 3 it is known [20] that n can be taken to be 2^k, where k is the quantifier rank of α'. Hence, all that is needed for constructing the precondition in case 3 is verifying if α' holds for L_j, where $j = 0, \ldots, 2^k + 1$. The theorem is proved. □

The next result shows that the preconditions constructed by the algorithm in the proof of theorem 7 are rather complex.

Corollary 3 *Let T be the transaction from theorem 7. Then, for any $n > 1$, there exists a sentence α of quantifier rank n, such that if m is the quantifier rank of $\mathsf{wpc}(T, \alpha)$ computed by the algorithm of theorem 7, then $m \geq 2^n$.*

Proof. We may take α to be already a Gaifman sentence. Then, according to the algorithm of theorem 7, one of the Boolean components of $\mathsf{wpc}(T, \alpha)$ is p_{2^n} which means that the quantifier rank of $\mathsf{wpc}(T, \alpha)$ is at least 2^n. □

We do not know if a better algorithm is possible. In fact, it remains open whether polynomial time weakest precondition algorithms exist for any transaction in $\mathcal{WPC}(\mathrm{FO})$ that is not first-order definable.

The separation result continues to hold for more powerful logics, although it is impossible to find a separating transaction that will have such a nice description as above. First we observe the following.

Proposition 3 *For any Ω, $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ can be captured by a transaction language.*

Proof. Since all functions and predicates in Ω are recursive, the set $\Gamma(D)$ is computable. The transactions in the language capturing $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ are given by tuples $(\Gamma, \alpha_1, \ldots, \alpha_k)$ where Γ is a finite set of terms, and each α_i has n_i free variables. The generated transaction is defined to be such that the tuple $(\Gamma, \alpha_1, \ldots, \alpha_k)$ constitutes a prerelation. This fully determines the transaction, and the definition of prerelation implies that such a language captures $\mathcal{PR}(\mathrm{FO}_c(\Omega))$. □

$d(x,y) \geq i$ denote that the distance between x and y is bigger than i. For every constant i, $d(x,y) > i$ can be expressed in first order logic. It is known that every first order sentence is equivalent to a Boolean combination of Gaifman sentences of the following form [18]:

$$(1) \qquad \alpha \equiv \exists x_1 \ldots \exists x_s.(\psi^{(r)}(x_1) \wedge \ldots \wedge \psi^{(r)}(x_s) \wedge \bigwedge_{i,j=1,\ldots,s,i\neq j} d(x_i,x_j) > 2r)$$

Hence, it is enough to show that with respect to T, every α of form (1) has a weakest precondition. Without loss of generality, assume $s \geq 1$ (otherwise α is constant and the proof is trivial).

Now observe that $T(G)$ has the following property. It is either a diagonal graph (having loops on all its nodes and nothing else), in which case $N_r(x) = \{x\}$ for any x, or it is a finite nonreflexive linear order (transitive closure of a chain) in which case $N_1(x)$ contains all nodes for any x.

It is enough to find a sentence β such that $G \models \beta$ iff $T(G) \models \alpha$ provided that G is a graph that satisfies $\Psi_{\text{c\&c}}$. Since all neighborhoods in diagonal graphs consist of one point, we have two cases. If $\psi^{(r)}(x)$ holds in the neighborhood consisting of x only, where the corresponding graph has a loop on x, then α on the diagonal graph reduces to the statement that there exist at least s distinct nodes. Hence, in this case a precondition for α is $(\Psi_{\text{c\&c}} \wedge \beta) \vee (\neg\Psi_{\text{c\&c}} \wedge \mu_s)$ where μ_s is a first order sentence saying that there are at least s distinct nodes. If $\psi^{(r)}(x)$ fails in the one-point neighborhoods, then $\Psi_{\text{c\&c}} \wedge \beta$ is a precondition.

Now, assuming that G satisfies $\Psi_{\text{c\&c}}$, we construct a precondition for α by considering the following cases.

Case 1: $s > 1$ and $r \geq 1$. Then α cannot be satisfied on a finite linear order because this would imply the existence of two nodes such that the distance between them is at least 2. Hence, for case 1 we take β to be *false*.

Case 2: $s \geq 1$, $r = 0$. Here we consider two subcases. First, suppose that $\psi^{(r)}(x)$ fails in one-point neighborhoods where the only edge is the loop on that point. Then α becomes *false* and hence β is taken to be *false*. Second, assume that $\psi^{(r)}(x)$ is satisfied in such one-point neighborhoods. Then α becomes equivalent to $\exists x_1 \ldots \exists x_s.(\bigwedge_{i,j=1,\ldots,s,i\neq j} d(x_i,x_j) > 0)$. Thus, it will hold in $T(G)$ iff the chain part of G has at least s points. The sentence p_s verifying this property is given below; then β is taken to be p_s.

$$p_s \equiv \exists y_1 \ldots \exists y_s.((\forall z.\neg E(z,y_1)) \wedge E(y_1,y_2) \wedge \ldots E(y_{s-1},y_s))$$

Case 3: $s = 1$, $r \geq 1$. Then α becomes $\exists x.\psi^{(r)}(x)$. Since we only consider α on graphs of the form L_n, a linear order on n nodes, an r-neighborhood of any $x \in L_n$ is the whole of L_n, which means that there exists a (non-local) sentence α' that is equivalent to α on all L_n (it is obtained from α by removing range restriction from local quantifiers). Now we need the following claim, proved in [20, 34] by using Ehrenfeucht-Fraissé games.

Claim. For every first order sentence α', it is possible to find $n \in \mathbb{N}$ such that either for all $m \geq n$: $L_m \models \alpha'$, or for all $m \geq n$: $L_m \models \neg\alpha'$.

Using this claim, we can find the precondition β as follows. Look at α' and find the n as stated above.

4 Preconditions vs. prerelations

The goal of this section is to examine the relationship between $\mathcal{WPC}(\mathrm{FO})$ and $\mathcal{PR}(\mathrm{FO})$. In particular, we find a transaction that separates them. However, we are not interested in just any transaction. Usually transformations expressed by database languages are required to satisfy certain properties. In particular, one is most often interested in transactions and queries computable in polynomial time. In addition, queries and transformations are sometimes required to be *generic*, that is, invariant under any permutation of the universe.

There is a large body of research dealing with languages capable of expressing polynomial-time generic database queries or transformations. Languages in which such transformations can be expressed include datalog with negation [1], languages based on structural recursion [10], loops [27], nondeterministic *reduce* operators [25], *while* queries [1] etc.

Hence, our first goal is to produce a generic polynomial time computable transaction that separates $\mathcal{WPC}(\mathrm{FO})$ from $\mathcal{PR}(\mathrm{FO})$. The transaction exhibited below can be expressed in all languages mentioned in the previous paragraph. The result below provides the proof of Theorem D from the introduction. After proving theorem 7 we proceed to show that $\mathcal{PR}(\mathrm{FO}_c(\Omega))$ can be captured by a transaction language, which gives us a separation result for $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$ and $\mathcal{PR}(\mathrm{FO}_c(\Omega))$. We close the section by showing that $\mathcal{WPC}(\mathrm{FO}_c)$ and $\mathcal{PR}(\mathrm{FO}_c)$ coincide on generic transactions.

Theorem 7 *There exists a polynomial-time computable generic transaction T such that*

$$T \in \mathcal{WPC}(\mathrm{FO}) - \mathcal{PR}(\mathrm{FO}).$$

Proof. Given a C&C-graph G, let $chain(G)$ be the connected component of G which is a chain. We consider the following transaction T on graphs $G = \langle X, E\rangle$:

$$T(G) = \begin{cases} tc(chain(G)) & \text{if } G \models \Psi_{\mathrm{c\&c}} \\ \langle X, \{(x,x) \mid x \in X\}\rangle & \text{if } G \models \neg\Psi_{\mathrm{c\&c}} \end{cases}$$

It is not hard to see that T is generic and PTIME-computable. Now assume that T is in $\mathcal{PR}(\mathrm{FO})$. Since Ω is empty, there exists a first-order formula $\beta(x,y)$ in free variables x and y such that (x,y) belongs to the output of $T(G)$ if $G \models \beta(x,y)$ when x and y are interpreted as elements of X. In particular, $T(G)$ can be calculated as the result of a first order query. Notice that if G is a chain, then $T(G)$ is its transitive closure. However, according to the bounded degree property of first order logic [27], there is no first order definable query on graphs which computes transitive closure if its input happens to be a chain. Hence, $T \notin \mathcal{PR}(\mathrm{FO})$.

The proof of $T \in \mathcal{WPC}(\mathrm{FO})$ is more involved. Recall the following definitions and results from [18]. In a graph, an r-neighborhood of a node x, denoted by $N_r(x)$, is the set of nodes of that graph that can be reached from x by a non-oriented path of length $\leq r$. By $\psi^{(r)}(x)$ we denote a formula in which x is the only free variable and all quantifiers are of form $\forall y \in N_k(x)$ or $\exists y \in N_k(x)$ where $k \leq r$. Let

Let I be the range of the enumeration $\langle C_m \rangle_{n \in \omega}$. For any graph G, let $[G]$ denote the (unique) graph C_n that is isomorphic to G. Given any mapping T from I to I, T expands to a mapping $[T]$ from graphs to graphs by letting $[T](G) = T([G])$.

Note that every finite collection of isomorphism classes can be expressed by a sentence of FO. Let $\langle \phi_i \rangle$, as before, enumerate the FO sentences. We can now define a transaction T from I to I exactly as we defined T before, but using C_i in place of G_i.

Now we claim that $[T]$ is in $\mathcal{WPC}(\mathrm{FO})$, but is unequal to any T_n. The proof that $[T]$ is recursive and diagonalizes the T_n's is exactly as before, using the fact that for each n there is an $=_n$ class that contains infinitely many nonisomorphic graphs.

To see that that $[T]$ is in $\mathcal{WPC}(\mathrm{FO})$, show that for each positive integer n and for each $i > P(n)$, $T(C_i) =_n C_i$: this follows exactly as before. A weakest precondition algorithm for $[T]$ is constructed as follows. Given a sentence ϕ, find an n such that $\phi = \phi_n$. Then find an m such that $T(C_i) =_n C_i$ for each $i > m$. Determine by testing which elements of $\{C_i : i \leq m\}$ have the property that $T(C_i) \models \phi$, and generate a sentence χ of FO that defines the union of the finite set of isomorphism classes of this set. Let ψ be a sentence of FO that describes the set of isomorphism classes of all C_i with $i \leq m$. Then exactly the same proof as above shows that $\chi \vee (\neg\psi \wedge \phi)$ is a weakest precondition for ϕ. This completes the proof of Theorem 5. □

As the last result of this section, we show that one cannot find a condition on degrees of nodes in graphs that describes the transactions with weakest preconditions. We are motivated by the following property of first-order queries, established in [27]. For a graph G, its *degree count* $dc(G)$ is the number of different in- and out-degrees of nodes of G. Then, for any first-order query q, $dc(q(G))$ is bounded by a number that depends only on q and the maximal possible in- or out-degree in G. For example, if G is an arbitrary binary tree, then an upper bound on $dc(q(G))$ is fully determined by q.

One may ask whether the class $\mathcal{WPC}(\mathrm{FO})$, which includes all first-order definable transactions, satisfies this property. We can show that this is not the case. In fact, a degree-based characterization is impossible for $\mathcal{WPC}(\mathrm{FO})$. Given any function $f : \mathbb{N} \to \mathbb{N}$ such that $f(n) > 0$, let $\mathcal{Q}_f$ be the family of graph queries q such that $dc(q(G)) \leq f(dc(G))$ for every G.

Corollary 2 *For any f, both $\mathcal{Q}_f - \mathcal{WPC}(\mathrm{FO})$ and $\mathcal{WPC}(\mathrm{FO}) - \mathcal{Q}_f$ are nonempty.*

Proof. Let q return the diagonal if its input graph is connected, and the complete graph on the input's nodes if it is not. Since connectivity is not first-order, $q \notin \mathcal{WPC}(\mathrm{FO})$. On the other hand, $q \in \mathcal{Q}_{\lambda x.1}$. Conversely, $T \in \mathcal{WPC}(\mathrm{FO})$ that will be described later in the proof of theorem 7 is such that when its input is a chain, it constructs its transitive closure. Hence, T is not in $\mathcal{Q}_f$ for any f. □

function, it is in $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$, and it is different from each T_i in $\mathcal{TL}$.

To prove the first claim, notice that $=_n$ is an equivalence relation with only finitely many equivalence classes. We now show that $H(m,n)$ is a total function. That is, for every m and n there are $\langle i,j\rangle$ such that $m < i < j$ and $G_j =_n G_i$ and $G_j \neq G_i$. If the above were not the case, then fix a counterexample m and n. Then all G_is with $i > m$ are pairwise nonequivalent with respect to $=_n$. But this gives us infinitely many $=_n$ classes. From this, we see that $H(m,n)$ is recursive, since we can check each pair $\langle i,j\rangle$ in turn to see if $i \leq j$ and $G_j =_n G_i$ and $G_j \neq G_i$ until we find a pair for which this is true.

From this we see easily, that P and Q are total recursive. Next we note that since both components of $H(m,n)$ are above m, the functions $P(m)$ and $Q(m)$ are (strictly) monotonically increasing. From this it follows that we can test recursively whether an i is in the range of P or not, that $P^{-1}(n)$ is well-defined for every n in the range, and that the range of P is infinite.

Note also that $G_{P(n+1)} =_n G_{Q(n+1)}$ and $P(n+1) \neq Q(n+1)$, using the definition of $H(m,n)$.

Since $P(n)$ and $Q(n)$ are distinct for every n, we get that for each m and i and each n it is true that one of $G_{P(n)}, G_{Q(n)}$ must be unequal to $T_m(G_i)$. Since P is monotonic, we can calculate the unique m such that $P(m) = i$ whenever i happens to be in the range of P. Together the last two sentences show that T is total and recursive.

We now turn to showing the second claim: that T admits a weakest precondition algorithm. We do this by verifying that for each positive integer n and for each $i > P(n)$, $T(G_i) =_n G_i$, and then citing Lemma 6 above.

Consider $T(G_i)$ for $i > P(n)$. If i is not in the range of P or $i = 1$, then $T(G_i) = G_i$ and so clearly $T(G_i) =_n G_i$. If i is in the range of P, then let $j = P^{-1}(i)$. Then $T(G_i)$ is one of either $G_{Q(j)}$ or $G_{P(j)}$, so in either case, $T(G_i) =_j G_{P(j)}$, since, as commented above $G_{Q(j)} =_j G_{P(j)}$. But then $T(G_i) =_j G_i$, since $P(j) = i$. Since P is monotonic and $i > P(n)$, we have $P^{-1}(i) > n$. This says that $j > n$, and hence $T(G_i) =_n G_i$, which completes the proof of the second hypothesis of Lemma 6, and hence the proof that T is in $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$

Finally we show that for each $m > 0$, $T \neq T_m$. Let $i = P(m)$ (so $m = P^{-1}(i)$). Clearly, i is in the range of P (and the strict monotonicity of P guarantees $i > 1$), so by definition, $T(G_i)$ has the property that $T(G_i) \neq T_{P^{-1}(i)}(G_i)$. But this means $T(G_i) \neq T_m(G_i)$, and this completes the proof that no transaction language captures $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$.

To prove that no transaction language captures $\mathcal{WPC}(\mathrm{FO})$ we need a slight modification of the proof above: we need to also ensure that T is generic in order for T to have a chance to be definable in pure first-order logic, while the above construction did not do anything to ensure genericity.

We again make use of an enumeration $\langle C_n\rangle$ of graphs such that no two graphs in the enumeration are isomorphic, and every graph is isomorphic to one of the C_n's. That is, the enumeration contains representatives for every isomorphism class of graphs. We can get such a recursive enumeration by enumerating the first graph G_1, then enumerating graphs until we come upon one nonisomorphic to any previously enumerated graph, etc.

enumerate the set of all graphs. All sentences of $\mathrm{FO}_c(\Omega)$ will be enumerated as $\phi_0, \phi_1, \ldots$, and we define the equivalence relation $=_n$ on graphs by $G =_n G'$ iff $G \models \phi_i \Leftrightarrow G' \models \phi_i$ for all $i \leq n$.

The idea of the proof is a diagonal argument. We will be building a transaction T, and for each m we will find a graph G such $T(G) \neq T_m(G)$. In order to ensure that the transaction T we build is in $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$, we will ensure that for each positive integer n, there is an integer $P(n)$ such that for all $i > P(n)$, $T(G_i) =_n G_i$. That is, for each n and G, eventually T does not change the $=_n$ class. To see that this last suffices, we note the following.

Lemma 6 *Suppose T is a computable transaction and there is a recursive function P on the integers such that for each positive integer n and for all $i > P(n)$, $T(G_i) =_n G_i$. Then T is in $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$*

Proof of the lemma: We construct a weakest precondition algorithm for T as follows. Given a sentence ϕ, find an n such that $\phi = \phi_n$. Then apply the hypothesis to find an $m\,(= P(n))$ such that $T(G_i) =_n G_i$ for each $i > m$. Determine by testing which elements of $\{G_i : i \leq m\}$ have the property that $T(G_i) \models \phi$, and generate a sentence χ of $\mathrm{FO}_c(\Omega)$ that defines this finite set. Let ψ be a sentence that describes the set of all G_i with $i \leq m$. Now output the sentence $\Phi = \chi \vee (\neg\psi \wedge \phi)$.

We claim that Φ is a weakest precondition for ϕ. If we have G_i with $i \leq m$, then G_i satisfies Φ if and only if it satisfies χ, and by the definition of χ this is true if and only if $T(G_i) \models \phi$. If we have G_i with $i \geq m$, then G_i satisfies Φ if and only if it satisfies ϕ. Since $T(G_i) =_n G_i$ for $i \geq m$ and $\phi = \phi_n$, we get that for each $i \geq m$, $T(G_i) \models \phi \Leftrightarrow G_i \models \phi$. Combining the previous two sentences we get that for $i \geq m$, $T(G_i) \models \phi \Leftrightarrow G_i \models \Phi$, which completes the proof that Φ is a weakest precondition for ϕ. Lemma is proved. □

The construction, then, is intuitively as follows: start enumerating graphs until arriving at a graph G_j whose $=_1$ class has many elements in it. Let T be the identity on $G_1 \ldots G_{j-1}$, and then let $T(G_j)$ be an element $=_1$ equivalent to G_j but not equal to $T_1(G_j)$. Then enumerate all graphs above G_j until arriving at a G_k whose $=_2$ class has many elements in it. Let T be the identity on $G_{j+1} \ldots G_{k-1}$, and then let $T(G_k)$ be an element $=_2$ equivalent to G_k but not equal to $T_2(G_k)$. Continue this process so as to diagonalize each T_i, while preserving the $=_n$ class for progressively larger n.

We now start the formal construction of T. We define a function $H(m, n)$ by letting $H(m, n)$ be lexicographically least pair $\langle i, j\rangle$ such that $m < i < j$ and $G_j =_n G_i$ and $G_j \neq G_i$.

Then let $P(n)$ and $Q(n)$ be defined inductively by $P(0) = Q(0) = 1$, and by letting $P(n+1)$ be the first component of $H(P(n), n)$, and $Q(n+1)$ be the second component of $H(P(n), n)$.

Finally, define the following transaction T. If $i = 1$ or i is not in the range of P, then $T(G_i) = G_i$. If $i > 1$ is in the range of P, consider the graph $G' = T_{P^{-1}(i)}(G_i)$. Since $G_{P(k)} \neq G_{Q(k)}$ for every $k > 0$, we know that for $j = Q(P^{-1}(i))$ it is the case that $G_i \neq G_j$. Now we define $T(G_i)$ in this case to be the one of G_i, G_j that is not equal to G', and if both are unequal to G', then we define $T(G_i)$ to be $G_{\min(i,j)}$.

The theorem now follows from the following three claims: the transaction T is a total recursive

$G_{n,n}$, we obtain that there exist two internal nodes a, b in the same branch, such that a and b have isomorphic d-neighborhoods, b is at a distance j from a, and for every node on a path from a to b, the isomorphism type of its d-neighborhood occurs at least $j + m$ times among the internal nodes of this branch.

Then in step 3, the duplicator selects the graph G' obtained from $G_{n,n}$ by collapsing b to a, that is, by removing all the nodes starting from the successor of a up to b. Note that $G' \in \mathit{Tree} - \mathcal{G}$. Then G' is colored by the duplicator, and the coloring in inherited from G_1. We call the resulting colored graph G_2.

It remains to show that $G_1 \approx_{d,m} G_2$, since this will imply that the duplicator can win the k-round Ehrenfeucht-Fraïssé game on G_1 and G_2 by Claim 1. First notice that there is no d-neighborhood that is present in one graph but is absent in the other. Furthermore, the only neighborhoods that have different number of realizers in G_1 and G_2 are those of the nodes removed in order to construct G_2 from G_1. For each of those neighborhood types, the number of nodes that realized them and that were removed, does not exceed j, and at the same time we know that each of those neighborhood types was realized at least $j + m$ times in G_1. Thus, each of of those neighborhood types is realized at least m times in both G_1 and G_2, which proves $G_1 \approx_{d,m} G_2$.

This concludes proving that no $q \in \mathcal{SG}_{tree}$ is verifiable over monadic Σ^1_1, and completes the proof of the theorem. □

Since the sentences not having weakest preconditions in the proof of theorem 3 are all in FO, we obtain

Corollary 1 *Let $\mathcal{L}$ be either $\mathrm{FO}_{\mathrm{count}}$, or $\mathrm{FO}_c(\Omega)$ for an arbitrary Ω, or monadic Σ^1_1. Then neither transitive closure, nor deterministic transitive closure, nor any query from $\mathcal{SG}_{tree}$ is in $\mathcal{WPC}(\mathcal{L}, \mathrm{FO})$.* □

3.3 The structure of transactions with weakest preconditions

Here we prove the result showing that the class of verifiable transactions cannot be captured by a transaction language. (By capturing we mean that a language expresses exactly the transactions from a given set.)

Theorem 5 *There is no transaction language that captures $\mathcal{WPC}(\mathrm{FO})$. Furthermore, for an arbitrary recursive signature Ω, there is no transaction language that captures $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$.*

Proof. We first prove that no transaction language captures $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$. We assume without loss of generality that the schema consists of a binary relational symbol $E(\cdot, \cdot)$. The proof below will automatically apply to any nonempty relational schema. Given a transaction language $\mathcal{TL}$, we show that it cannot capture $\mathcal{WPC}(\mathrm{FO}_c(\Omega))$. We will assume that the transactions in $\mathcal{TL}$ are enumerated as $\langle T_i \rangle_{i>0}$. This is possible because the syntax of any language is a recursive set. We let $\langle G_i \rangle_{i>0}$

realize this type, or both have at least m nodes that realize this type. Note that in a graph colored with c colors, a neighborhood of a point is structure in the vocabulary $\langle E, U_1, \ldots, U_c, a\rangle$, where E is the binary edge relation, U_is are unary relations that are interpreted as sets of nodes colored with the ith color, and a is a constant interpreted as a point around which the neighborhood is taken.

Claim 1 Let k be a positive integer. Then there exist positive integers d and m such that, whenever G_1 and G_2 are two colored trees of outdegree at most 2, and $G_1 \approx_{d,m} G_2$, then the duplicator has a winning strategy in the k-round Ehrehfeucht-Fraïssé game on G_1 and G_2.

Claim 1 is just a special case of Theorem 4.3 of [17], when the structures are trees, and the maximum degree is 2.

Before we prove that the duplicator has a winning strategy in the (c,k) Ajtai-Fagin game, we need the following combinatorial lemma.

Lemma 4 *For every positive integers p and l, there exists a positive integer $N[p,l]$ such that, for any $N > N[p,l]$ and any partition of the set $\{1,\ldots,N\}$ into l sets, there exist two numbers i_1, i_2 that belong to the same class, such that for any $i_1 \le i \le i_2$, it is the case that i belongs to a partition class that has at least $p + i_2 - i_1$ elements.*

Proof of the lemma. If $l = 1$, the statement is trivial, so we assume $l > 1$. Let $f = \max\{p,l\}$. We claim that $N[p,l]$ can be taken to be $4f^4+f(f+1)+1$. There exists a number $N_0 \in [4f^4+1, 4f^4+f(f+1)+1]$ such that N_0 is divisible by l and $l+1$. Let $k = \frac{N_0}{l}$ and $s = \frac{N_0}{l+1}$. Consider partitioning of $\{1,\ldots,N_0\}$ into l classes. Then at least one class, say X, contains k elements, $x_1,\ldots,x_k$. Let $d_i = x_{i+1} - x_i$, $i = 1,\ldots,k-1$. Then at most s of d_is are greater than or equal to $l+1$. Hence, at least $m_0 = \frac{N_0}{l(l+1)} - 1$ of d_is are less than $l+1$. Let $I = \{i \mid x_{i+1} - x_i \le l\}$; we know that $|I| \ge m_0$.

If at least one $d_i = 1$, then we are done, since $k \ge p+1$. So assume that $d_i > 1$ for all $i \in I$. Assume that the conclusion of the lemma is false. Then, for every $i \in I$, there exists y_i such that $x_i < y_i < x_{i+1}$ and y_i belongs to a partition class that has fewer than $p + d_i$ elements. In particular, such a class has fewer than $p + l$ elements. Since there are l classes, we obtain that the number of such y_is is at most $l(p+l)$, which implies that the cardinality of I is at most $l(p+1)$, that is, $m_0 \le l(p+l)$. This implies $N_0 \le l^2(l+1)(p+l)+1$, which is easily seen to contradict the assumption that $N_0 \ge 4f^4+1$. This contradiction proves the lemma.

Using Lemma 4, we conclude the proof as follows. Given c and k, let d and m be given by Claim 1. Let $l = (2d+1)^c$, and let $N = N[m,l]$. Then in step 1 the duplicator selects $G_{n,n}$ where $n > N + 2(d+1)$. Let us call a node *internal* if it is at the distance at least $d+1$ from the root and the leaves. A d-neighborhood of such a node is a $2d+1$-element chain. The choice of n guarantees that both branches of $G_{n,n}$ have at least $N+1$ internal nodes.

Now the spoiler colors the graph with $G_{n,n}$ with c colors. Let G_1 be this colored graph. Note that there are at most $l = (2d+1)^c$ of isomorphism types of d-neighborhoods of internal nodes in G_1. Thus, coloring corresponds to partitioning the internal nodes into l classes, given by the isomorphism types of their d-neighborhoods. Applying Lemma 4 to $N+1$ internal nodes in one of the branches of

of the form $G_{n,m}$ with the set of nodes A. Let $f : A \to X$ be a monotone injective map from A to X, and let $f(G)$ be the image of G under this map. Then, since γ_2 is a $FO(\prec)$ sentence, we get $G \models \gamma_2$ iff $f(G) \models \gamma_2$ iff $|n - m| = 1$. Thus, the assumption that $q \in \mathcal{WPC}(\mathrm{FO}_c(\Omega))$ implies that in first-order logic with built-in order relation, we can test if $G_{n,m}$ satisfies $|n - m| = 1$.

Now we are going to show that such a test is impossible. Assume that the first-order language contains only the symbol $\prec$ to be interpreted as a linear order. For each x, define

$$\begin{array}{rcl} E_x(u,v) & \equiv & [(v \prec u) \wedge ((u \prec x) \vee (u = x)) \wedge (\forall z.\neg(v \prec z) \vee \neg(z \prec u))] \\ & \vee & [(u \prec v) \wedge ((x \prec u) \vee (u = x)) \wedge (\forall z.\neg(u \prec z) \vee \neg(z \prec v))] \end{array}$$

That is, E_x defines a relation that is "successor backwards" for elements under x and the successor relation (associated with the ordering) for elements above x. Such a relation is isomorphic to $G_{n,m}$ where n is the number of elements under x and m is the number of elements above x. Let $\gamma_2'(x)$ be a first-order formula in the language of $\prec$ obtained from γ_2 by replacing each occurrence of $E(u,v)$ with $E_x(u,v)$. Then a finite linear ordering $(L, <)$ satisfies $\exists x.\gamma_2'(x)$ iff there a "middle" element such that the number of elements above it is n and the number of elements below it is $n-1$ (or vice versa). This happens iff the size of the finite universe is even. However, since a $\mathrm{FO}(\prec)$ sentence of quantifier rank k cannot distinguish two linear orderings of size $> 2^k$ [34], we have a contradiction that shows that q is not verifiable over $\mathrm{FO}_c(\Omega)$.

Finally, we must show that q is not verifiable over monadic Σ_1^1. Assume such a q is verifiable. Let $\mathcal{G} = \{G_{n,n} \mid n \geq 1\}$. Then there exists a monadic Σ_1^1 sentence β such that, if G is a tree, then $G \models \beta$ iff $G \in \mathcal{G}$. To see this, note that the first-order sentence α_0 saying that a graph has one root of outdegree 2, two leaves on indegree 1, and that every other node has both in- and outdegree 1, defines a graph one of whose connected components is $G_{n,m}$ for some n, m, and all other connected components are cycles. Let α_1 be the first-order sentence $\exists!x.E(x,x) \wedge (\forall y.E(x,y) \vee E(y,x) \to x = y)$; that is, α_1 states that there exists a unique isolated point. Then, if G is a tree, then $G \models \alpha_0 \wedge \mathsf{wpc}(q, \alpha_1)$ iff $G \in \mathcal{G}$, since for any tree of the form $G_{n,m}$, where $n \neq m$, there exist at least two isolated points in $q(G_{n,m})$.

To complete the proof, we must show that there is no monadic Σ_1^1 sentence β such that, if G is a tree, then $G \models \beta$ iff $G \in \mathcal{G}$. Let *Tree* be the class of all trees. According to [16, Theorem 5.5], it suffices to show that for any positive integers c and k, the duplicator can win the (c,k) Ajtai-Fagin game for $\mathcal{G}$ and $\mathit{Tree} - \mathcal{G}$. The game is played as follows.

Step 1. The duplicator selects a graph $G \in \mathcal{G}$.
Step 2. The spoiler colors the nodes of G with c colors.
Step 3. The duplicator selects a graph $G' \in \mathit{Tree} - \mathcal{G}$ and colors its nodes with c colors.
Step 4. The spoiler and the duplicator play k rounds of the Ehrehfeucht-Fraïssé game on colored G and G'.

The winner is determined as the winner in Step 4. For more details on games, see [15, 16]. To determine a winner in Step 4, we shall use the criterion below, that follows immediately from Theorem 4.3 of [17]. We shall use the notation $G_1 \approx_{d,m} G_2$ if G_1 and G_2 are two colored graphs, and for every isomorphism type of a d-neighborhood of a node, either both graphs have the same number $n \leq m$ of nodes that

Consequently, no first-order sentence distinguishes the families $\{G_{n,n} \mid n > 1\}$ and $\{G_{n-1,n+1} \mid n > 1\}$. This finishes the proof of Claim 3 and the theorem. □

3.2 Weakest preconditions for more powerful logics

Here we extend the result of the previous subsection for more powerful logics defined in section 2. By doing so, we complete the proof of Theorem B from the introduction.

Theorem 3 *Let $\mathcal{L}$ be either $\mathrm{FO}_{\mathrm{count}}$, or $\mathrm{FO}_c(\Omega)$ for an arbitrary Ω, or monadic Σ^1_1. Then neither transitive closure, nor deterministic transitive closure, nor any query in $\mathcal{SG}_{tree}$ is verifiable over $\mathcal{L}$.*

Proof. We start by showing that the transitive closure and deterministic transitive closure are not verifiable over $\mathcal{L}$. The proof for $\mathrm{FO}_{\mathrm{count}}$ and $\mathrm{FO}_c(\Omega)$ is essentially the same as the proof of theorem 2, if we can show that connectivity and testing for chain are not definable in those logics. For $\mathrm{FO}_{\mathrm{count}}$ this follows immediately from [27] (or [13]), and for $\mathrm{FO}_c(\Omega)$ this follows from the result of [7] that generic Boolean queries definable in $\mathrm{FO}_c(\Omega)$ are definable in relational calculus with the order relation, since directed connectivity is not definable in the latter [36]. That testing for chain is not definable in first-order logic with built-in order relation follows from [13], which shows that this problem is first-order complete for deterministic logspace.

In the case of monadic Σ^1_1, the proof for *tc* is the same as before, since connectivity is not expressible in monadic Σ^1_1 [5]. For *dtc*, we have to show that testing whether a graph is a chain is not monadic Σ^1_1-definable. It was shown in [17] that for any monadic Σ^1_1 sentence ψ there exists a number n such that ψ can not distinguish C^1_n, one cycle of length $2n$, from C^2_n, two cycles of length n. However, C^1_n can be distinguished from C^2_n for any $n > 2$ if testing for chain is available. Indeed, consider the sentence $\gamma \equiv \forall x \forall y.E(x,y) \rightarrow \mathit{chain}[E'_{xy}/E]$. Here *chain* is the sentence that tests whether a graph is a chain, and everywhere in this sentence we replace the basic predicate $E(z,v)$ with $E'_{xy}(z,v) \equiv E(z,v) \wedge (\neg(x = z) \vee \neg(y = v))$. Then γ holds iff for every edge (x,y), the graph resulting from removing this edge is a chain. Such a sentence γ clearly distinguishes C^1_n from C^2_n. This finishes the proof for the case of (deterministic) transitive closure.

Now let $q \in \mathcal{SG}_{tree}$; we show that q is not verifiable over $\mathcal{L}$. When $\mathcal{L} = \mathrm{FO}_{\mathrm{count}}$, we can apply the proof of claim 3 in theorem 2, since, by the result of [30] (see also [15]), for each k it is possible to find a number r such that any two structures that realize the same number of all r-neighborhoods cannot be distinguished by a $\mathrm{FO}_{\mathrm{count}}$ sentence of quantifier rank k.

For $\mathrm{FO}_c(\Omega)$, first define an order relation $\prec$ on $\mathcal{U}$ that is isomorphic to ω (that is, the usual order on $\mathbb{N}$). Now we show that q is not verifiable over $\mathrm{FO}_c(\Omega \cup \{\prec\})$. Assume it is verifiable. Consider α_2 from the proof of theorem 2 (asserting that there are two isolated points) and let $\beta_2 = \mathsf{wpc}(q, \alpha_2)$. Note that $G_{n,m} \models \beta_2$ iff $n - m = 1$ or $m - n = 1$.

According to [7, Proposition 3], there exists an infinite set $X \subseteq \mathcal{U}$ and a $FO(\prec)$ sentence γ_2 such that for all graphs G whose sets of nodes are in X, we have $G \models \gamma_2$ iff $G \models \beta_2$. In particular, for any graph of the form $G_{n,m}$ with nodes from X, $G_{n,m} \models \gamma_2$ iff $|\, n - m \,| = 1$. Now consider an arbitrary graph G

Claim 2: $dtc \notin \mathcal{WPC}(\mathrm{FO})$.

Proof of Claim 2. Assume $dtc \in \mathcal{WPC}(\mathrm{FO})$. Let $\alpha \equiv \forall x \forall y.(x \neq y \rightarrow E(x,y) \vee E(y,x))$, and let $\beta = \mathsf{wpc}(dtc, \alpha)$. Define $\gamma \equiv \Psi_{\mathrm{c\&c}} \wedge \beta$. Then $G \models \gamma$ iff $G \models \Psi_{\mathrm{c\&c}}$ and in $dtc(G)$ there are no two unconnected nodes. For any graph G satisfying $\Psi_{\mathrm{c\&c}}$ we have $tc(G) = dtc(G)$, so $G \models \gamma$ iff G is a chain. However, it is known (see [13, 27]) that testing for chain is not FO-definable. Thus, $dtc \notin \mathcal{WPC}(\mathrm{FO})$.

Claim 3: If $q \in \mathcal{SG}_{tree}$, then $q \notin \mathcal{WPC}(\mathrm{FO})$.

Proof of Claim 3. We define a family of graphs, $\{G_{n,m} \mid n, m \geq 1\}$ that will be used in this proof as well as in the proof of theorem 3. A graph $G_{n,m}$ is shown on the picture below. It is a tree, with two branches, each being a chain. The subtree rooted at one of the root's children is an n-node chain, and the other is an m-node chain.

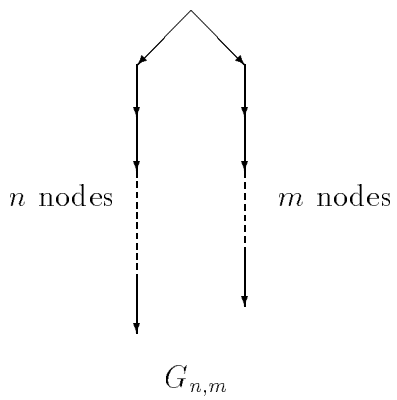

Now assume that $q \in \mathcal{WPC}(\mathrm{FO})$. If G is a tree, then $q(G)$ consists of a number of connected components, each of them being a complete graph. In particular, if for some node x there is no other node in the same generation as x, then x is an isolated node in $q(G)$. Let α_i, where $i \geq 1$, be a first-order sentence saying that there exist exactly i isolated nodes (*i.e.* those with a loop and no other incoming or outgoing edge). Let $\beta_i = \mathsf{wpc}(q, \alpha_i)$. Then it is easy to see that $G_{n,m} \models \beta_i$ iff $\mid n - m \mid = i - 1$. Thus, for any $n \geq 1$, $G_{n-1,n+1} \models \beta_3$ and $G_{n,n} \models \neg\beta_3$.

To prove claim 3, we now have to show that for any first-order sentence γ there exists a number N such that for all $n > N$, $G_{n-1,n+1} \models \gamma$ iff $G_{n,n} \models \gamma$. To prove this statement, we use Hanf's technique modified for the finite case by [17]. Given a graph G, an r-neighborhood of a node a, $N_r(a)$, is the subgraph of G on the set of all nodes reachable from a by unordered paths of length at most r. An r-type of a is the isomorphism type of $N_r(a)$, with a as a distinguished node. It can be seen from the structure of the graphs $G_{n,m}$ that for every r, and every $n > 2r + 1$, the graphs $G_{n,n}$ and $G_{n-1,n+1}$ have the same number of neighborhoods of every given r-type.

Now if γ of quantifier rank k is given, take r to be 3^k. Then for any $n > 2r + 1$, $G_{n,n}$ and $G_{n-1,n+1}$ have the same number of r-neighborhoods that realize each r-type. That is, using the terminology from [17], these two graphs are r-equivalent. According to [17], 3^k-equivalent structures cannot be distinguished by sentences of quantifier rank k; hence $G_{n,n}$ and $G_{n-1,n+1}$ cannot be distinguished by γ.

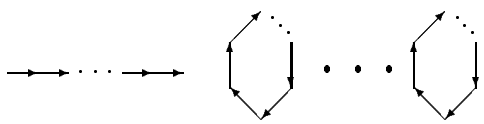

Lemma 1 *The class of C&C-graphs is first-order definable. Neither chains nor cycles are first-order definable.*

Proof. A graph is a C&C-graph iff it has exactly one root (node with in-degree zero), exactly one endpoint (node with out-degree zero), the root has out-degree 1, the endpoint has in-degree one, and all other nodes have both in- and out-degrees 1. That is:

$$\begin{array}{lll} & \forall x \forall y \forall z. E(x,y) \wedge E(x,z) \rightarrow z = y & \text{outdegrees are at most 1} \\ \wedge & \forall x \forall y \forall z. E(y,x) \wedge E(z,x) \rightarrow z = y & \text{indegrees are at most 1} \\ \wedge & \exists! x \forall y. \neg E(y,x) & \text{unique root} \\ \wedge & \exists! x \forall y. \neg E(x,y) & \text{unique endpoint} \end{array}$$

defines the C&C-graphs.

That chains and cycles are not first-order definable can be proved straightforwardly using standard techniques (cf. [18, 27]). □

For the rest of the paper, $\Psi_{\text{c\&c}}$ denotes the first-order sentence in the language of $E(\cdot,\cdot)$ above that defines C&C-graphs.

3.1 Languages not verifiable over FO

In this subsection we prove part of Theorem B from the introduction. That is,

Theorem 2 *Assume that $\mathcal{TL}$ can express one of the following:*

- *transitive closure;*
- *deterministic transitive closure;*
- *any query from $\mathcal{SG}_{tree}$.*

Then $\mathcal{TL}$ is not verifiable over FO.

Proof follows from the following three claims.

Claim 1: $tc \notin \mathcal{WPC}(\text{FO})$.

Proof of Claim 1. Assume $tc \in \mathcal{WPC}(\text{FO})$ and let $\alpha \equiv \forall x \forall y. E(x,y)$. Then a graph satisfies $\mathsf{wpc}(tc, \alpha)$ iff it is connected, which means connectivity is FO-definable. This contradiction proves $tc \notin \mathcal{WPC}(\text{FO})$.

whether a given tuple will be in the database afterwards. We need the set Γ of terms to be able to define transactions that extend the active domain of a database (for example, insertion of a new tuple may extend the domain). The set $\Gamma(D)$ is a superset of the active domain of $T(D)$, and the formulae pre_i^T are used to determine when a tuple from $\Gamma(D)$ belongs to the ith relation of $T(D)$.

Prerelations can be viewed as providing the ability to verify local properties of databases as well as global properties. Prerelations have been used extensively in [32]. If $\mathcal{L}$ is pure first-order logic (in particular, Ω is empty and the only terms are variables), then the notion of having prerelations reduces to the familiar notion of a first-order definable transaction.

For all the languages we shall consider in this paper it is the case that $\mathcal{PR}(\mathcal{L}) \subseteq \mathcal{WPC}(\mathcal{L})$. Indeed, to construct $\mathsf{wpc}(T, \alpha)$ for any $T \in \mathcal{PR}(\mathcal{L})$ we can use Γ and pre_i^T to simulate the new database state after the transaction, and then substitute all symbols for R_i in α by the formulae defining the new state. It is easy to see that the power of first-order logic is sufficient for doing this. Later in the paper we examine this inclusion and show that for most languages of interest it is strict.

3 Weakest preconditions

The goal of this section is to study weakest preconditions for first-order logic and its extensions. We first prove that a language that can express (deterministic) transitive closure or the same-generation query cannot have weakest preconditions over FO. We then extend this result to more powerful specification languages, and finally prove that no transaction language captures the class of transactions having weakest preconditions over first-order logic.

We need a few technical definitions. Given a graph $G = \langle V, E \rangle$, its *deterministic transitive closure* (cf. [23]), $dtc(G)$, is the graph $\langle V, E' \rangle$ where $(x, y) \in E'$ iff $(x, y) \in E$ or there is a sequence of nodes $x = x_1, x_2, \ldots, x_n = y$ such that each $(x_i, x_{i+1}) \in E$, and furthermore, each x_i has out-degree 1, $i = 1, \ldots, n-1$. The transitive closure of a graph is denoted by $tc(G)$.

Given a graph $G = \langle V, E \rangle$, the *same-generation* query returns the graph $sg(G) = \langle V, E' \rangle$ where $(x, y) \in E'$ iff there is a node $v \in V$ and two walks in G from v to x and from v to y that have the same length. Usually the same-generation query is asked for trees. So we define the class of queries

$$\mathcal{SG}_{tree} \;=\; \{q \mid \text{ if } \; G \; \text{ is a tree, then } q(G) = sg(G)\}$$

In particular, $sg \in \mathcal{SG}_{tree}$.

A *chain* is a graph of the form $\{(x_1, x_2), \ldots, (x_{n-1}, x_n)\}$ where all x_is are distinct. A *simple cycle* is a graph of the form $\{(x_1, x_2), \ldots, (x_{n-1}, x_n), (x_n, x_1)\}$ where all x_is are distinct. A *chain-and-cycle graph*, or a *C&C-graph* consists of $n \geq 1$ connected components, of which exactly one is a chain, and others (perhaps zero) are simple cycles. An example of a C&C-graph is shown below.

2) We say that a transaction language $\mathcal{TL}$ is $\mathcal{L}_1$-verifiable over $\mathcal{L}_2$ *if every $\mathcal{TL}$ transaction T is in $\mathcal{WPC}(\mathcal{L}_1, \mathcal{L}_2)$. If $\mathcal{L} = \mathcal{L}_1 = \mathcal{L}_2$, we speak of a language verifiable over $\mathcal{L}$.*

The proposition below shows that while $\mathcal{WPC}(\cdot,\cdot)$ is monotone in the first argument and antimonotone in the second, $\mathcal{WPC}(\cdot)$ need not be (anti)monotone in its one argument. We use the notation $\mathcal{L} \preccurlyeq \mathcal{L}'$ to mean that $\mathcal{L}$ is a sublanguage of $\mathcal{L}'$.

Proposition 2 *(a) Let $\mathcal{L}_1' \preccurlyeq \mathcal{L}_1$ and $\mathcal{L}_2 \preccurlyeq \mathcal{L}_2'$. Then $\mathcal{WPC}(\mathcal{L}_1', \mathcal{L}_2') \subseteq \mathcal{WPC}(\mathcal{L}_1, \mathcal{L}_2)$.*
(b) It is possible to find languages $\mathcal{L} \preccurlyeq \mathcal{L}'$ and a transaction T such that $T \in \mathcal{WPC}(\mathcal{L}) - \mathcal{WPC}(\mathcal{L}')$.
(c) It is possible to find languages $\mathcal{L} \preccurlyeq \mathcal{L}'$ and a transaction T such that $T \in \mathcal{WPC}(\mathcal{L}') - \mathcal{WPC}(\mathcal{L})$.

Proof. (a) follows immediately from the definitions. For (b) and (c), take T to be the transitive closure of a graph. For (b), take $\mathcal{L}'$ to be first-order logic augmented with constants for each element $u \in \mathcal{U}$. Take $\mathcal{L}$ to be the family of Boolean combinations of formulae $\psi_u \equiv \exists x.(E(x,u) \vee E(u,x))$ meaning that u is a node with either incoming or outcoming edge in the graph given by the predicate E. Then $D \models \psi_u$ iff $T(D) \models \psi_u$, and hence $T \in \mathcal{WPC}(\mathcal{L})$. On the other hand, $T \notin \mathcal{WPC}(\mathcal{L}')$ (see theorem 3). For (c), we take $\mathcal{L}$ to be first-order logic, and $\mathcal{L}'$ to be first-order logic with fixpoint. Since T is definable in the latter, $T \in \mathcal{WPC}(\mathcal{L}')$, but $T \notin \mathcal{WPC}(\mathcal{L})$, see theorem 2. □

Similar notions have been studied in a more general programming language context. For example, weakest preconditions over first-order logic for a simple while loop language can be expressed in infinitary logic [6] or in weak second-order logic [9]. However, proposition 2 shows that it is not always possible to extend results of this kind to less expressive specification languages. In particular, according to Theorem A, weakest preconditions for transitive closure (which is definable with while loops) are not first-order expressible.

We will also use *prerelations* which allow for testing of local conditions. We are now ready to formalize the definition we presented informally in the introduction. Let Ω be a signature. By $\textsc{Term}(\Omega)$ we denote the set of all Ω-terms which are built up from variables by using the symbols from Ω. (We treat constants as functions of arity zero.) Given a database D, its (active) *domain* $\mathrm{dom}(D)$ is the set of all elements of $\mathcal{U}$ that occur in the database. For a set $\Gamma \subseteq \textsc{Term}(\Omega)$ and a database D we define $\Gamma(D)$ to be the set of all $x \in \mathcal{U}$ such that $x = \tau(y_1, \ldots, y_n)$ for some n-ary term $\tau \in \Gamma$ and some $y_1, \ldots, y_n \in \mathrm{dom}(D)$.

Definition *Let T be a transaction. We say that T admits* prerelations *over $\mathcal{L}$ if there exists a finite collection of terms Γ and a collection of $\mathcal{L}$-formulae $\mathsf{pre}_1^T, \ldots, \mathsf{pre}_k^T$, pre_i^T having n_i free variables, such that for any R_i in the schema and any $d_1, \ldots, d_{n_i} \in \mathcal{U}$, it is the case the for every database D:*

$$\begin{array}{ll} & D \models \mathsf{pre}_i^T(d_1, \ldots, d_{n_i}) \text{ and } d_1, \ldots, d_{n_i} \in \Gamma(D) \\ \Leftrightarrow & T(D) \models R_i(d_1, \ldots, d_{n_i}) \end{array}$$

The collection of all transactions that admit prerelations over $\mathcal{L}$ is denoted by $\mathcal{PR}(\mathcal{L})$.

The intuition behind the definition above is this. Since $T(D) \models R_i(d_1, \ldots, d_{n_i})$ says that the tuple $(d_1, \ldots, d_{n_i})$ belongs to R_i after T is applied, using prerelations we can test, *before* T is committed,

database, and returns either another database encoding or "error", provided $s \in S$. Examples include relational algebra and calculus, SQL, first-order transaction language of [32], a variety of languages from [3, 4] and so on. We shall always use the same symbol for both syntactic and semantic objects. For example, for the language given by $\langle S, M \rangle$, $T \in S$ and a database D, we will write $T(D)$ to denote the result of M on T and the encoding of D.

We are interested in maintaining database integrity. To formalize this problem, let $\mathcal{L}$ be a specification language and $\mathcal{TL}$ a transaction language. We define the following problem:

PRESERVE($\mathcal{TL}, \mathcal{L}$): Given $T \in \mathcal{TL}$ and α in $\mathcal{L}$, is it the case that $D \models \alpha$ implies $T(D) \models \alpha$ for every database D?

According to Fact A, this problem is undecidable even for transactions specified by the simplest form of SQL statements and for first-order constraints. This can be derived from a result of [2]. Since the transaction model of [2] differs from the one that we consider, we present a simple proof below for the sake of completeness.

Proposition 1 *Let $\mathcal{L}$ be* FO, *and let $\mathcal{TL}$ contain the transactions given by the select-project-join expressions of the relational algebra. Then* PRESERVE($\mathcal{TL}, \mathcal{L}$) *is undecidable.*

Proof. Assume that PRESERVE is decidable, and consider two transactions on graphs: T_1 takes a graph on a finite set of nodes V and produces its diagonal ($\{(x, x) \mid x \in V\}$) and T_2 produces the complete graph without loops ($\{(x, y) \mid x, y \in V, x \neq y\}$). $T_1(E)$ can be implemented as $\Pi_{1,3}(\sigma_{1=3}(E \times E))$ and T_2 as $\Pi_{1,3}(\sigma_{1\neq 3}(E \times E))$ (we assume that V is the union of the first and the second projections of E).

It suffices to restrict our attention only to nonempty graphs. Observe that if a first-order sentence ϕ is not satisfied in any $T(D)$, then PRESERVE(T, ϕ) iff $\neg\phi$ is valid. Now let $\psi \equiv \exists x.E(x, x)$, and let β be an arbitrary first-order sentence in the language of $E(\cdot, \cdot)$. Then β is valid in every finite graph if and only if both $\neg\psi \vee \beta$ and $\psi \vee \beta$ are valid in every finite graph. Using the observation above, we derive that $\beta \vee \psi$ is valid iff PRESERVE($T_1, \neg\beta \wedge \neg\psi$), because $\neg\beta \wedge \neg\psi$ fails in all nonempty outputs of T_1. Similarly, $\beta \vee \neg\psi$ is valid iff PRESERVE($T_2, \neg\beta \wedge \psi$). Thus, if PRESERVE is decidable, then it is decidable whether an arbitrary sentence β is valid in all finite graphs. Since undecidability of the latter is known (cf. [14]), we have a contradiction. □

As we explained in the introduction, a way around this is to introduce weakest preconditions. Below is the main definition.

Definition *1) Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two specification languages. We say that a transaction T in a transaction language $\mathcal{TL}$ has* $\mathcal{L}_1$-weakest preconditions over $\mathcal{L}_2$ *if for every $\mathcal{L}_2$-sentence α there exists an $\mathcal{L}_1$ sentence $\mathsf{wpc}(T, \alpha)$, computable from T and α, such that for every database D on which T is defined,*

$$D \models \mathsf{wpc}(T, \alpha) \quad \textit{if and only if} \quad T(D) \models \alpha.$$

The class of transactions having $\mathcal{L}_1$-weakest preconditions over $\mathcal{L}_2$ is denoted by $\mathcal{WPC}(\mathcal{L}_1, \mathcal{L}_2)$. We also write $\mathcal{WPC}(\mathcal{L})$ for $\mathcal{WPC}(\mathcal{L}, \mathcal{L})$.

2 Formal setting

We fix the domain to be a countably infinite set $\mathcal{U}$. A *relational schema* is a nonempty set $SC = \langle R_1, \ldots, R_k \rangle$ of predicates, with a finite arity $n_i > 0$ associated with each R_i. A *database* over SC is an interpretation of each R_i as a finite subset of $\mathcal{U}^{n_i}$. Most often we will use the schema that consists of a single binary predicate E. Then databases are interpreted as finite graphs whose nodes are elements of $\mathcal{U}$.

We shall use two notions – *specification language*, in which constraints are specified, and *transaction language*. A specification language $\mathcal{L}$ is a recursive subset of a set of strings with an associate subset $sent(\mathcal{L})$ of *sentences*, which is also recursive. As stated in the introduction, we wish to focus on the specification languages that are relevant to the integrity constraints. In particular, we will study

- FO – (pure) first-order logic (that is, first-order logic over the relational schema SC).
- FO_c — first-order logic over the signature that consists of SC supplemented with constant symbols for all elements of $\mathcal{U}$.
- $\mathrm{FO}_c(\Omega)$ — first-order logic over the signature that consists of SC supplemented with constant symbols as above and a recursive collection Ω of recursive functions and predicates over $\mathcal{U}$.

We shall also consider the following two logics, which are more powerful than the first-order logic, and play an important role in finite model theory and descriptive complexity.

- $\mathrm{FO}_{\mathrm{count}}$ — first-order logic with counting. For precise definition, see [13, 19, 24]. This is a two-sorted logic, with the second sort of natural numbers whose universe is $\{1, \ldots, n\}$, where n is the size of the universe of the first sort. First-order logic is extended with counting quantifiers, $\exists i x.\psi(x)$ meaning that here are at least i elements of the first-sort universe satisfying ψ. This quantifier binds x but not i. The order, the constants for 1 and the maximal element, and the **bit** predicate are available on natural numbers, where $\mathbf{bit}(i, j)$ is true iff the jth bit in the binary representation of i is one. For example, $\exists i \exists i x.\psi(x) \wedge \mathbf{bit}(i, 1) \wedge (\forall j(\exists j x.\psi(x)) \rightarrow j \leq i)$ says that there are odd number of elements satisfying ψ. Another example of non-first-order property definable in $\mathrm{FO}_{\mathrm{count}}$ is equal cardinality.
- Monadic Σ^1_1, whose formulae are of the form $\exists A_1 \ldots \exists A_k.\Psi$ where A_i's are monadic (unary) predicates and Ψ is a first-order formula over $SC \cup \langle A_1, \ldots, A_k \rangle$.

For all specification languages it will be the case that we shall have a binary relation $\models$ for validity between databases and sentences of the languages. For all the languages above, validity is defined in the usual way.

A transaction language consists of 1) Syntax: a recursive subset of the set S of strings over a finite alphabet, and 2) Semantics: a total recursive function M that takes a string s, and an encoding of a

$$D \models \mathsf{pre}_R^T(d_1, \ldots, d_n) \quad \text{and} \quad d_1, \ldots, d_n \in \Gamma(D)$$
$$\Leftrightarrow \quad T(D) \models R(d_1, \ldots, d_n)$$

Here $\Gamma(D)$ is a finite set that extends the active domain of the database. The precise definition will be given in the next section.

The condition $T(D) \models R(d_1, \ldots, d_n)$ means that t belongs to the relation R after the transaction T is applied to the database D. Thus, prerelations allow us to test membership before transactions are committed.

In papers like [32] prerelations are used in place of weakest preconditions. The question of the relationship of these two notions, however, has not been addressed. We shall see that for most specification languages of interest, it is the case that any transaction that admits prerelations also admits weakest preconditions. The converse, however, is not true even in very simple settings.

Theorem D *There exists a transaction T in $\mathcal{WPC}(\mathrm{FO})$ that does not admit prerelations. Moreover, T can be chosen to be* Datalog$^\neg$*-definable.*

However, this result is not robust: if FO is extended by adding a constant for each element of the domain, then the transaction from Theorem D no longer provides a separation example. Motivated by this, we say that a transaction language $\mathcal{TL}$ is *robustly verifiable* over a first-order signature Ω if, for any extension of Ω to a signature Ω' via a set of recursive predicates and functions, all transactions in $\mathcal{TL}$ remain verifiable over the corresponding extension of first-order logic. We then prove a positive result that characterizes maximal robustly verifiable languages when integrity constraints can refer to the individual elements of the domain. Let FO_c be first-order logic over the relational schema supplemented with symbols for all constants.

Theorem E *There exists a transaction language that captures the class of transactions admitting prerelations over FO_c. Moreover, this language is the maximal robustly verifiable language over FO_c. In particular, every transaction robustly verifiable over FO_c is equivalent to a transaction that admits prerelations.*

This has implications for transaction language design. If we are interested in a "nice" transaction language that is verifiable in an extensible way over all first-order languages, then we cannot hope that the language will be more expressive than the first-order transaction language defined in [32, 33].

Organization We rigorously define our formal setting in the next section. In section 3 we study the limitations of weakest preconditions. Section 4 studies the relationship between weakest preconditions and prerelations. Robustly verifiable transaction languages are the subject of section 5. Concluding remarks are given in section 6.

The extended abstract of this paper appeared in [8].

The main goal of this paper is to study the properties of weakest preconditions for a variety of transaction and specification languages. In particular, we concentrate on specification languages that are relevant to integrity constraints, such as first-order logic over the database schema. Below we give a brief overview of the main results of the paper.

Assume that we have two languages $\mathcal{L}_1$ and $\mathcal{L}_2$ in which constraints may be specified. Also assume that we have a transaction language $\mathcal{TL}$. We say that $\mathcal{TL}$ is *$\mathcal{L}_1$-verifiable over $\mathcal{L}_2$* iff for every $T \in \mathcal{TL}$ and every $\alpha \in \mathcal{L}_2$, it is the case that a weakest precondition for α with respect to T is in $\mathcal{L}_1$. If $\mathcal{L} = \mathcal{L}_1 = \mathcal{L}_2$, we speak of *$\mathcal{L}$-verifiable* transaction languages. Similar concepts have been studied in the more general programming language context, see [6, 9].

Verifiable transaction languages are exactly those for which the weakest-precondition approach to integrity maintenance is possible. An example of a language verifiable over first-order logic is the first-order transaction language of [32]. However, no results exist for transaction languages that possess more power than first-order logic (for example, for languages that allow some forms of recursion). The first result of this paper is negative: it shows that adding recursion destroys first-order verifiability.

Theorem B *Let $\mathcal{TL}$ be a transaction language that expresses at least one of the following:*

- *transitive closure;*
- *deterministic transitive closure;*
- *same-generation (even restricted to trees).*

Then $\mathcal{TL}$ is not verifiable over any of the following languages: first-order logic, first-order logic with counting, first-order logic supplemented with a recursive collection of recursive functions and predicates, monadic Σ^1_1.

We next fix our specification language to be first-order logic, and look at the set $\mathcal{WPC}(\mathrm{FO})$ of transactions (that is, total maps from databases to databases) that have weakest preconditions for first-order sentences. We show that this collection cannot be captured by any transaction language.

Theorem C *There is no transaction language that expresses exactly all the transactions from $\mathcal{WPC}(\mathrm{FO})$. This remains true if* FO *is replaced by first-order logic over a signature containing a recursive collection of interpreted recursive functions and predicates.*

The ability to statically check global conditions — that is, conditions that must be satisfied by the database as a whole — is, as noted above, extremely helpful. One might also demand a stronger "local" version of this, in which one can statically check properties of individual tuples in the resulting database. That is, one often would like to find a condition that must be satisfied by a database D and an object t (typically, a record) in order for t to belong to $T(D)$. It is natural to compare the class $\mathcal{WPC}(\mathcal{L})$ with the class of transactions for which this *local* checking can be carried out. For a specification language $\mathcal{L}$, the natural definition of such a class is the set of transactions admitting *prerelations*. Informally, a transaction T admits prerelations over $\mathcal{L}$ if it is definable by $\mathcal{L}$-formulae on a set that extends the active domain of the database. That is, for every relation R of arity n there exists a $\mathcal{L}$-formula $\mathsf{pre}^T_R(x_1, \ldots, x_n)$ such that, for every n-tuple $t = (d_1, \ldots, d_n)$ and every database D:

sometimes fail to maintain some of the constraints, then modifications can be made until correctness is achieved.

If T is a transaction and α is a constraint, we say that T preserves α if for every database D:

$$D \models \alpha \text{ implies } T(D) \models \alpha.$$

We use the notation $D \models \alpha$ to mean that α is true for the database D.

According to the approach described above, we would like to be able to check transactions to see if they will always preserve α. But, as the following result tells us, one cannot hope to have an algorithm for checking this, even for simple transactions and simple constraint specification languages.

Fact A (cf. [2]) *For transactions specified as select-project-join expressions of the relational algebra, and integrity constraints specified as first-order sentences, it is undecidable to check if a given transaction preserves a given integrity constraint.*

Similar results hold for transactions that consist of simple updates and deletions. Therefore, any approach to automatic verification of transaction safety (such as the approach of [35, 37] that uses theorem-provers) is inherently limited. Instead, we adopt an approach based on modifying transactions to ensure safety.

Given a transaction T and a constraint α, we would like to find another constraint β that is satisfied by a database D *before* T is committed if and only if α is satisfied *afterwards*. In other words, for every database D:

$$D \models \beta \text{ if and only if } T(D) \models \alpha.$$

We shall call such a β a *weakest precondition* for α, with respect to T, and denote it by $\mathsf{wpc}(T, \alpha)$. We will require that $\mathsf{wpc}(T, \alpha)$ be computable for given T and α. Since we consider only transactions that do not diverge, this notion of weakest precondition is equivalent to the classical one of Dijkstra [11, 12].

The existence of weakest preconditions would not contradict Fact A, as we would not be able to determine the validity of the implication $\mathsf{wpc}(T, \alpha) \rightarrow \alpha$. At the same time it would solve the problem of statically verifying database integrity since any transaction T could be transformed to

$$\underline{\mathsf{if}}\ \mathsf{wpc}(T, \alpha)\ \underline{\mathsf{then}}\ T\ \underline{\mathsf{else}}\ \ \mathit{abort}$$

which will maintain consistency while avoiding any roll-back operations. Furthermore, we could then modify the resulting transaction by applying simplification algorithms to $\mathsf{wpc}(T, \alpha)$, thus recapturing many of the benefits of approaches based on validity-checking. This is the fundamental idea underlying many algorithms for the automatic maintenance of integrity constraints [29, 21, 22, 31, 28].

Given the attractiveness of this approach to integrity maintenance, it is important to understand the tradeoffs involved in designing transaction and specification languages with respect to our ability to express weakest preconditions. In this paper, we will concentrate on the basic principles behind developing safe transactions via preconditions, rather than specific algorithms that attempt to detect violations of integrity.

Verifiable Properties of Database Transactions

Michael Benedikt
Bell Laboratories
1000 E. Warrenville Rd
Naperville IL 60566, USA
benedikt@research.bell-labs.com

Timothy Griffin
Bell Laboratories
600 Mountain Avenue
Murray Hill NJ 07974, USA
griffin@research.bell-labs.com

Leonid Libkin*
Bell Laboratories
600 Mountain Avenue
Murray Hill NJ 07974, USA
libkin@research.bell-labs.com

Abstract

It is often necessary to ensure that database transactions preserve integrity constraints that specify valid database states. While it is possible to monitor for violations of constraints at run-time, rolling back transactions when violations are detected, it is preferable to verify correctness statically, *before* transactions are executed. This can be accomplished if we can verify transaction safety with respect to a set of constraints by means of calculating *weakest preconditions*. We study properties of weakest preconditions for a number of transaction and specification languages. We show that some simple transactions do not admit weakest preconditions over first-order logic and some of its extensions such as first-order logic with counting and monadic Σ^1_1. We also show that the class of transactions that admit weakest preconditions over first-order logic cannot be captured by any transaction language. We consider a strong local form of verifiability, and show that it is different from the general form. We define robustly verifiable transactions as those that can be statically analyzed regardless of extensions to the signature of the specification language, and we show that the class of robustly verifiable transactions over first order logic is exactly the class of transactions that admit the local form of verifiability. We discuss the implications of these results for the design of verifiable transaction languages.

1 Introduction

Databases are typically required to satisfy a collection of *integrity constraints* — sentences that specify the valid states of the database. Database transactions that update the database must preserve these constraints. One approach to the integrity maintenance problem is to defer the detection of potential violations of integrity constraints until run-time, thus preventing the violation of integrity that may be caused by transactions. However, this approach results in potentially expensive roll-back operations. In order to avoid this, one can attempt to verify *statically*, before a transaction is executed, that it will necessarily maintain all constraints. If the result of this verification shows that the transaction may

*Contact author. Address: Bell Laboratories, Room 2C-407, 600 Mountain Avenue, Murray Hill NJ 07974, USA. E-mail: libkin@research.bell-labs.com. Phone: (908)582-7647. Fax: (908)582-5857.