Preservation theorems for Tarski's relation algebra

Non Peer-reviewed author version

PRESERVATION THEOREMS FOR TARSKI'S RELATION ALGEBRA

BART BOGAERTS ● [a], BALDER TEN CATE ● [b], BRETT MCLEAN ● [c],
AND JAN VAN DEN BUSSCHE ● [d]

[a] Vrije Universiteit Brussel, Brussel, Belgium
e-mail address: bart.bogaerts@vub.be

[b] ILLC, University of Amsterdam, Amsterdam, The Netherlands
e-mail address: b.d.tencate@uva.nl

[c] Ghent University, Ghent, Belgium
e-mail address: brett.mclean@ugent.be

[d] Hasselt University, Hasselt, Belgium
e-mail address: jan.vandenbussche@uhasselt.be

ABSTRACT. We investigate a number of semantically defined fragments of Tarski's algebra of binary relations, including the function-preserving fragment. We address the question of whether they are generated by a finite set of operations. We obtain several positive and negative results along these lines. Specifically, the homomorphism-safe fragment is finitely generated (both over finite and over arbitrary structures). The function-preserving fragment is not finitely generated (and, in fact, not expressible by any finite set of guarded second-order definable function-preserving operations). Similarly, the total-function-preserving fragment is not finitely generated (and, in fact, not expressible by any finite set of guarded second-order definable total-function-preserving operations). In contrast, the forward-looking function-preserving fragment is finitely generated by composition, intersection, antidomain, and preferential union. Similarly, the forward-and-backward-looking injective-function-preserving fragment is finitely generated by composition, intersection, antidomain, inverse, and an 'injective union' operation.

1. INTRODUCTION

Just as Boolean algebra can be viewed as a language for describing operations on sets, *Tarski's relation algebra* ($\mathbb{TRA}$) is a language for describing operations on binary relations. It consists of a small, finite collection of operations on binary relations (which includes, for instance, composition and union), governed by natural equations such as $R \circ (S \cup T) = (R \circ S) \cup (R \circ T)$. The origins of $\mathbb{TRA}$ trace back to the 19th century, and, more specifically, to the work of Augustus De Morgan and Charles Peirce, but its study intensified when it was picked up by Tarski and his students in the 1940s [Tar41, Mad91, Pra92]. We can view $\mathbb{TRA}$ as a language for specifying operations on binary relations. Its expressive power, in terms of the term-definable operations, corresponds precisely to the three-variable fragment of first-order logic (FO^3) [TG87].

Many modern graph and tree query languages, such as regular path queries, SPARQL, and XPath, which describe ways of navigating through graph-structured data, can be identified with variants of $\mathbb{TRA}$, each involving a different set of allowed operations. This

		homomorphism-safe	⊆-safe	function-preserving	forward
id	identity relation	yes	yes	yes	yes
∅	empty relation	yes	yes	yes	yes
⊤	universal relation (all pairs)	yes	yes	no	no
$-(\cdot)$	complement	no	no	no	no
$(\cdot)^{\smallsmile}$	inverse	yes	yes	no	no
$\mathsf{D}(\cdot)$	domain ($\mathsf{D}(R) = \{(x,x) \mid R(x,y)\}$)	yes	yes	yes	yes
$\mathsf{R}(\cdot)$	range ($\mathsf{R}(R) = \{(y,y) \mid R(x,y)\}$)	yes	yes	yes	no
$\sim(\cdot)$	antidomain ($\sim R = \{(x,x) \mid \neg\exists y\, R(x,y)\}$)	no	no	yes	yes
$\cdot \cup \cdot$	union	yes	yes	no	yes
$\cdot \cap \cdot$	intersection	yes	yes	yes	yes
$\cdot \setminus \cdot$	relative complement	no	yes	yes	yes
$\cdot \circ \cdot$	composition	yes	yes	yes	yes
$\cdot \ltimes \cdot$	left semi-join ($R \ltimes S = \{(x,y) \in R \mid \exists z\, S(y,z)\}$)	yes	yes	yes	yes
$\cdot \sqcup \cdot$	preferential union ($R \sqcup S = R \cup \{(x,y) \in S \mid \neg\exists z\, R(x,z)\}$)	no	no	yes	yes

Table 1: Operations on binary relations

has generated an interest in systematically understanding the expressive power of fragments and extensions of $\mathbb{TRA}$ [FGL$^+$15a, FGL$^+$15b, HWGV22].

Here, we study the question whether certain semantically-defined fragments of $\mathbb{TRA}$ can be generated by a finite set of operations. One known positive result along these lines is the following, where $\mathbb{BRA}(\mathcal{O})$ denotes the binary relation algebra generated by the operations in $\mathcal{O}$ (see Table 1 for a definition of the operations).

Theorem 1.1 [Ben98]. *A $\mathbb{TRA}$-term is "bisimulation safe" if and only if it is equivalent to a $\mathbb{BRA}(\mathrm{id}, \circ, \cup, \sim)$-term.*

The precise definition of *bisimulation* and of *bisimulation safety* is not important for us here. It suffices that bisimulation is an important equivalence relation that captures behavioral equivalence of processes, and that an operation on binary relations is bisimulation safe if commutes, in a natural way, with bisimulation.

We can think of Theorem 1.1 result as analogous to a *preservation theorem* in model theory: it correlates a semantic property with expressibility in a natural, finitely-generated, syntactic fragment. The above result may suggest that various other semantically-defined fragments of $\mathbb{TRA}$ could be similarly characterised syntactically by a finite basis of operations. One particular prominent semantic fragment that arises naturally in different contexts, is the function-preserving fragment of $\mathbb{TRA}$ [McL18]. An operation on binary relations is said to be *function-preserving* if, whenever the input relations are partial functions, so is the output relation. It is a natural question, and an open problem in the community (although we could not locate an explicit reference) whether the function-preserving fragment of $\mathbb{TRA}$ is finitely generated.

Contributions. As our main contribution, we establish the following positive and negative results:

- The *homomorphism-safe* fragment of $\mathbb{TRA}$ is finitely generated (Section 3).
- The *function-preserving* fragment of $\mathbb{TRA}$ is not finitely generated (and, in fact, not expressible by any finite set of guarded second-order definable function-preserving operations). The same holds for the *total-function-preserving* fragment (Section 4).
- The *forward function-preserving* fragment and the *local injective-function-preserving* fragment are finitely generated (Section 5).

We study each of these fragments both in the general case (i.e., where the input relations may be relations over an infinite domain) and in the finite.

Naturally, there are many other semantic fragments of $\mathbb{TRA}$ for which one could ask the same finite-generatedness question. Our intention, with the above results, is to provide a sample of interesting results when it comes to the question of finite generation for semantic fragments of $\mathbb{TRA}$. In the concluding Section 6, we will further comment on directions for future work and connections to the formalisms we mentioned in our motivation above.

Related Work. Börner and Pöschel [BP91] studied whether various clones of operations on binary relations *over a fixed finite structure* are finitely generated. Their study includes the "logical clone" (which is the set of all first-order definable operations) as well as the "positive clone" (which is the set of all operations definable by positive-existential first-order formulas). Our investigation is different in that we are interested in the existence of finite bases *over all (finite) structures*. We will further comment on the relationship between our results and those by Börner and Pöschel in Section 3.

Andréka et al. [ACN85] and Börner [B̆86] consider the problem whether certain finitely generated clones of operations on binary relations are in fact generated by a single operator (analogous to the Sheffer stroke in Boolean algebra), and what is the minimum possible arity of such an operation.

There is a substantial literature on algebras of partial functions (that is, function-preserving fragments of $\mathbb{TRA}$), focusing on the axiomatisation of their first-order theories as well as computational aspects such as decidability and the finite model property. An in-depth overview of known results along these lines can be found in [McL18].

In the literature on temporal logics, there have been extensive studies concerning the existence of temporal logics generated by a finite set of operations, that are expressively complete for first-order logic in the sense of Kamp's theorem [Kam68] (see [GHR94] for an overview). One of the main differences with our setting is that, in temporal logic, the operators are typically monadic (i.e., they correspond to FO-formulas in one free variable), whereas in our case, the operators act on, and produce, binary relations (and hence correspond to FO-formulas in two free variables). Closer to our setting is Venema [Ven90], who studies expressive completeness for interval temporal logics, and showed that, on dense linear orders, no finite set of binary operations is expressively complete for FO; and the results on *conditional XPath* by Marx [Mar05], which imply that (a fragment of) $\mathbb{TRA}$ is expressively complete for FO over finite sibling-ordered trees. Both are concerned with definability of binary relations. Note however, that our objective differs from that of [Ven90, Mar05]: we are not restricted to linear orders or trees, and we are not primarily interested in expressive completeness with respect to FO, but rather expressive completeness with respect to (semantic fragments of) Tarski's relation algebra, or, equivalently, FO^3.

2. Preliminaries

First-order logic and guarded second-order logic. We restrict to structures over signatures consisting of binary relation symbols only. We write FO for first-order logic, and we denote by FO^k (for $k \geq 1$) the k-variable fragment of FO, that is, the fragment of FO consisting of formulas that use only k variables, where nested quantifiers may reuse the same variable.

We will also consider *guarded second-order logic* (GSO [GHO02], also known as MSO_2 [CE12]), which extends first-order logic with monadic second-order quantification (i.e., quantification over sets) as well as guarded second-order quantification, by which we mean quantification over subrelations of relations in the signature. Thus, for example, we can express in GSO that a pair (x, y) lies on a Hamiltonian cycle in a digraph, which is a property that cannot be expressed in MSO [Lib04].

By the *quantifier rank* of a GSO-formula ϕ we will mean the maximum nesting depth of first-order and/or second-order quantifiers. We will write $A \equiv^n_{\text{GSO}} B$ to indicate that two structures agree on all GSO-sentences of quantifier rank at most n.

Binary relation algebras. An *n-ary operation on binary relations* is a map O from first-order structures $A = (\text{dom}(A), R_1^A, \ldots, R_n^A)$ to binary relations $O(A) \subseteq \text{dom}(A)^2$ that is isomorphism invariant: for every isomorphism $h : A \cong B$, it holds that $h : O(A) \cong O(B)$. Equivalently, one may think of an n-ary operation on binary relations as mapping first-order structures $A = (\text{dom}(A), R_1^A, \ldots, R_n^A)$ to first-order structures $A' = (\text{dom}(A), O(A))$, where the domain of the structure remains unchanged. We say that O is FO-*definable* if there is an FO-formula $\phi(x, y)$ such that $O(A) = \{(a, b) \in \text{dom}(A)^2 \mid A \models \phi(a, b)\}$ for all A. A *binary relation algebra* is given by a collection $\mathcal{O}$ of operations on binary relations. We denote it by $\mathbb{BRA}(\mathcal{O})$. We say that the algebra is FO if all its operations are FO-definable.

Terms, term definable, finitely generated. Let $\mathbb{A} = \mathbb{BRA}(\mathcal{O})$ be a binary relation algebra, and fix some countable infinite set of binary relation symbols $R_1, R_2, \ldots$ By an *n-ary term* of $\mathbb{A}$ we mean a syntactic expression built up from the relation symbols $R_1, \ldots, R_n$ using the operations in $\mathcal{O}$ as function symbols. For instance, $R_1 \cup R_1^\smile$ is an example of a 1-ary $\mathbb{TRA}$-term. We denote by O_t the n-ary operation on binary relations defined by the term t. We say that two n-ary terms t and t' are *equivalent (in the finite)* if, for all (finite) structures $A = (\text{dom}(A), R_1^A, \ldots, R_n^A)$, $O_t(A) = O_{t'}(A)$. We say that an operation on binary relations is *term definable (in the finite)* in $\mathbb{A}$ if there is a term of $\mathbb{A}$ that defines it (over finite structures). Note that, if $\mathcal{O}$ consists of FO-definable operations, then every term of $\mathbb{BRA}(\mathcal{O})$ defines an FO-definable operation. In fact, if every operation in $\mathcal{O}$ is FO^k-definable (for some $k \geq 2$) then every $\mathbb{BRA}(\mathcal{O})$-term also defines an FO^k-definable operation. The same applies in the finite.

We say that a binary relation algebra $\mathbb{BRA}(\mathcal{O})$ is *finitely generated* if there is a finite subset $\mathcal{O}' \subseteq \mathcal{O}$, such that every operation in $\mathcal{O}$ is term definable in $\mathbb{BRA}(\mathcal{O}')$.

Tarski's relation algebra. Tarski's relation algebra ($\mathbb{TRA}$) is an example of an FO binary relation algebra. It can be defined as $\mathbb{TRA} := \mathbb{BRA}(\mathrm{id}, \emptyset, -, \cap, \circ, \smile)$. All operations in Table 1 are term definable in $\mathbb{TRA}$. The following two classic results on $\mathbb{TRA}$ will be relevant for us.

Theorem 2.1 [TG87, Section 3.9]. *Both in general and in the finite: an operation on binary relations is term definable in $\mathbb{TRA}$ if and only if it is FO^3-definable.*

Theorem 2.2 [Tar41, Löw15]. *Both in general and in the finite: the binary relation algebra consisting of all FO-definable operations is not finitely generated.*

Theorem 2.2 in fact follows from Theorem 2.1 together with the well-known fact in (finite) model theory that FO does not collapse to any of its finite variable fragments; cf. also [Ven90, Theorem 2.13].

Kleene Algebra is an example of a non-FO binary relation algebra, which includes the (GSO-definable) reflexive transitive closure operation. We omit the definition, as we will not study it in this paper.

3. The homomorphism-safe fragment is finitely generated

Recall that a homomorphism $h : A \to B$ is a function from the domain of A to the domain of B that preserves structure, i.e. such that $(a, b) \in R^A$ implies $(h(a), h(b)) \in R^B$. We say that an operation O on binary relations is *homomorphism safe* if, for every homomorphism $h : A \to B$ and $(a, b) \in O(A)$, $(h(a), h(b)) \in O(B)$. Equivalently, O is homomorphism safe if and only if every homomorphism $h : A \to B$ is also a homomorphism $h : (A, O(A)) \to (B, O(B))$, where $(A, O(A))$ denotes the expansion of the structure A with $O(A)$ as an additional relation, and similarly for $(B, O(B))$. Thus, intuitively, one can think of homomorphism-safe operations as *homomorphism-preserving* operations.

As indicated in Table 1, examples of homomorphism-safe operations are $\cup$, $\cap$, and $\circ$, but not $-$.

Theorem 3.1. *Both in general and in the finite: a $\mathbb{TRA}$-term is homomorphism-safe if and only if it is equivalent to a $\mathbb{BRA}(\mathrm{id}, \emptyset, \top, \circ, \cup, \cap, \smile)$-term.*

Proof. The right-to-left direction can be proved by a straightforward induction. We will focus on the more interesting left-to-right direction.

We will make use of recent results regarding homomorphism-preserved FO-formulas [BC19]. Formally, we say that an FO-formula $\phi(x_1, \ldots, x_n)$ is *homomorphism preserved* if for every homomorphism $h : A \to B$ and tuple $a_1, \ldots, a_n$, we have $A \models \phi(a_1, \ldots, a_n)$ implies $B \models \phi(h(a_1), \ldots, h(a_n))$. A classic theorem in model theory (known as the *homomorphism preservation theorem*) states that a first-order formula is homomorphism preserved if and only if it is equivalent to a positive-existential FO-formula (i.e., a formula built up from atomic formulas using only existential quantification, conjunction, and disjunction). Rossman [Ros08] proved that this holds also in the finite. Bova and Chen [BC19, Corollary 24] further refined this to finite-variable fragments (both on arbitrary structures and in the finite): they showed that every homomorphism-preserved FO^k formula is equivalent to a positive-existential FO^k-formula.

Let us now proceed with the proof of our theorem. By Theorem 2.1, it suffices to show that every FO^3-formula $\phi(x_1, x_2)$ (with two free variables) that is homomorphism preserved can be translated to the $\mathbb{TRA}$ fragment in question. Moreover, by the aforementioned results

of Bova and Chen, we may assume that $\phi(x_1, x_2)$ is a positive-existential FO^3-formula. We inductively translate $\phi(x_1, x_2)$ to a term t in the specified fragment of $\mathbb{TRA}$, such that $(a, b) \in O_t(A)$ iff A satisfies ϕ under the assignment that maps x_1 and x_2 to a and b, respectively. The base cases are straightforward. In particular, $R(x_1, x_2)$ translates to R, $R(x_2, x_1)$ translates to $R^\smile$, $R(x_1, x_1)$ translates to $(R \cap id) \circ \top$, $x = y$ translates to id, etc. Conjunction and disjunction translate to $\cap$ and $\cup$, respectively (note that, here we take advantage of the fact that our induction hypothesis was stated specifically for formulas $\phi(x_1, x_2)$). Therefore, only the case remains where $\phi(x_1, x_2)$ is of the form $\exists y \psi(x_1, x_2, y)$. It is not hard to see that ψ must, in this case, be a positive Boolean combination of formulas with at most two free variables. That is, ψ can be written as a disjunction of conjunctions of formulas with at most two free variables. Furthermore, we can pull the disjunction out from under the existential quantifier, and deal with it separately. Therefore, we can assume without loss of generality that ψ is a conjunction of formulas with two free variables. By grouping the conjuncts appropriately, we can write ψ as $\psi_1(x_1, x_2) \wedge \psi_2(x_1, y) \wedge \psi_3(x_2, y)$. By the induction hypothesis, each of these conjuncts can be translated to a $\mathbb{TRA}$-term, say, t_1, t_2, t_3. We can then translate ϕ as $t_1 \cap (t_2 \circ t_3^\smile)$. $\square$

It is worth comparing Theorem 3.1 to results by Börner and Pöschel [BP91], which state that the "logical clone" (which is defined as the binary relation algebra consisting of all FO-definable operations on binary relations) as well as the "positive clone" (the binary relation algebra consisting of all operations on binary relations definable by a positive-existential FO-formula) over any fixed finite structure are finitely generated. By Rossman [Ros08], the operations that can be defined by a positive-existential FO-formula are precisely the homomorphism-safe FO-definable operations. We see that Theorem 3.1 is incomparable to the results just mentioned. On the one hand, it is only concerned with $\mathbb{TRA}$-term-definable operations. On the other hand, it states that there there is a finite basis of operations from which all homomorphism-safe $\mathbb{TRA}$-terms are term definable *over all (finite) structures*.

One may wonder whether the approach taken in the proof of Theorem 3.1 could be used to establish a Łos–Tarski-style theorem for $\mathbb{TRA}$, characterising the fragment of $\mathbb{TRA}$ that is preserved by the $\subseteq$ relation, where, by $A \subseteq B$, we mean that A is an induced substructure of B. More precisely we say that a first-order formula $\phi(x_1, \ldots, x_n)$ is $\subseteq$-preserved if, whenever $A \subseteq B$ and $a_1, \ldots, a_n \in \mathrm{dom}(A)$ and $A \models \phi(a_1, \ldots, a_n)$, then $B \models \phi(a_1, \ldots, a_n)$. The classic Łos–Tarski preservation theorem states that, on unrestricted (i.e., possibly infinite) structures, an FO formula is $\subseteq$-preserved if and only if it is equivalent to an existential FO-formula. As it turns out, however, the Łos–Tarski theorem fails for FO^3. More precisely, it has been shown [AvBN23, Lemma 1 and 2] that there is a FO^3-sentence over a signature consisting of a single binary relation, that is $\subseteq$-preserved, but that is not equivalent, even over finite structures, to an existential FO^3-sentence.[1] This shows that the approach we used for the homomorphism-safe fragment of $\mathbb{TRA}$ will not work for the $\subseteq$-preserved fragment, which we also refer to as the $\subseteq$-*safe* fragment. So, it leaves the following question open.

Question 3.2. *Is the $\subseteq$-safe fragment of $\mathbb{TRA}$ finitely generated?*

---

[1] The result in [AvBN23] is stated in terms of preservation under taking induced substructures, and it talks about the universal fragment of FO. It is, however, equivalent by a duality argument.

4. The function-preserving fragment is not finitely generated

Let O be an n-ary operation on binary relations. We say that O is *function preserving* if the following holds for all structures $A = (\mathrm{dom}(A), R_1^A, \ldots, R_n^A)$: if each R_i^A is a partial function on $\mathrm{dom}(A)$, then $O(A)$ is a partial function on $\mathrm{dom}(A)$. Similarly, we say that O is *total-function preserving* if the following holds for all structures $A = (\mathrm{dom}(A), R_1^A, \ldots, R_n^A)$: if each R_i^A is a total function on $\mathrm{dom}(A)$, then $O(A)$ is a total function on $\mathrm{dom}(A)$.

As indicated in Table 1, the following are function preserving: id, $\emptyset$, D, R, $\sim$, $\cap$, $\setminus$, $\circ$, $\ltimes$, and $\sqcup$. Let us call the binary relation algebra consisting of these operations *function algebra* ($\mathbb{FA}$). $\mathbb{FA}$ was described in [HJM16] as "in an informal sense at least, the richest natural case" of an algebra of partial functions. In the same paper, a finite axiomatisability result was established for $\mathbb{FA}$ (see also [McL18] for a systematic study of algebras of partial functions).

Our main result in this section is the following theorem.

Theorem 4.1. *Let $\mathcal{O}$ be any finite set of function-preserving GSO-definable operations on binary relations. Then there is a function-preserving operation on binary relations O that is term definable in $\mathbb{TRA}$ but not in $\mathbb{BRA}(\mathcal{O})$, even over finite structures in which all relations are partial functions.*

The proof will make use of the following lemma (where $\uplus$ denotes the operation of *disjoint union*).

Lemma 4.2. *For all structures A, A', B, B' and $n > 0$, if $A \equiv_{\mathrm{GSO}}^n A'$ and $B \equiv_{\mathrm{GSO}}^n B'$ then $A \uplus B \equiv_{\mathrm{GSO}}^n A' \uplus B'$.*

Proof. The lemma can be derived from a (suitable adaptation to GSO of a) more general Feferman–Vaught theorem for MSO [Mak04]. However, here, we give a direct argument using an Ehrenfeucht–Fraisse-style game argument. The game we will consider is played, as usual, between two structures, C and C', and has two players, Spoiler and Duplicator. In each round, Spoiler plays first and can make two types of moves: those corresponding to first-order quantification and those corresponding to monadic or guarded second-order quantification. A move of the first type means that Spoiler picks an element of C or of C'. In this case, Duplicator must respond by picking a corresponding element of the other structure. A move of the second type means that Spoiler picks either a subset of the domain of C or C' (in which case Duplicator responds by picking a corresponding subset of the domain of the other structure) or a subrelation of one of the relations of C or C' (in which case, Duplicator responds by picking a corresponding subrelation of the same relation in the other structure). The game then continues using the same pair of structures expanded with the chosen elements/sets/relations. The game is played for a fixed number of rounds, n. Duplicator wins if, after n rounds, the resulting substructures satisfy the same quantifier-free FO formulas (with the chosen elements as parameters and chosen sets/relations as relations). It is a standard exercise to show that Duplicator has a winning strategy for the n-round game if and only if $C \equiv_{\mathrm{GSO}}^n C'$.

We now use the above game to prove the statement. Since $A \equiv_{\mathrm{GSO}}^n A'$ and $B \equiv_{\mathrm{GSO}}^n B'$, Duplicator has winning strategies in the two corresponding n-round games. Consider now the n-round game between $A \uplus B$ and $A' \uplus B'$. We will refer to A and B as the "left" part and the "right" part of $A \uplus B$ and similarly for A' and B'. Recall the two types of moves Spoiler can make in the game. A move of the first type consists of choosing an element, which must belong either the "left half" of the structure or to the "right half". In this case,

Duplicator can respond using their assumed strategy in the corresponding game. A move of the second type involves selecting either a set of elements, or a set of a tuples from a relation in the structure. In either case, the set in question can be naturally partitioned into two halves, the "left half" and the "right half". Duplicator can therefore respond to each type of move simply by using her winning strategies for the two parts of the structure. It is easy to see that this yields a winning strategy for Duplicator. □

Proof of Theorem 4.1. Let n be a number greater than the maximum quantifier rank of the GSO-formulas defining the operations in $\mathcal{O}$.

For $m \geq 0$, let C_m be the directed graph that has a vertex $a_{i,j}$ for every $i \in \{1, \ldots, m\}$ and $j \in \{1, 2, 3\}$, and that has an edge from $a_{i,j}$ to $a_{i',j'}$ whenever $i' = (i \bmod m) + 1$. In other words, C_m is a directed cycle of length m in which every vertex is replaced by three vertices. Then let $C_m^{\vee}$ be the structure over the signature $\{f, g\}$ obtained from C_m by replacing every edge by an $(f^{\smile} \circ g)$-path (using a fresh intermediate vertex each time). See Figure 1. We will refer to the vertices of the form $a_{i,j}$ as "normal nodes" and the added intermediate vertices as "auxiliary nodes". In addition, by a "cluster of auxiliary nodes" we mean the family of nine auxiliary nodes added between the points $a_{i,j}$ and $a_{(i+1 \bmod m),j'}$ for $j, j' \in \{1, 2, 3\}$, for some $i \in \{1, \ldots, m\}$.

Claim 1: There are $m \neq m'$ such that, in the structure $C := C_m^{\vee} \uplus C_{m'}^{\vee}$, all normal nodes satisfy the same GSO-formulas $\phi(x)$ of quantifier depth n and likewise for the auxiliary nodes.

Proof of Claim 1: Since there are (up to equivalence) only finitely many GSO-sentences of quantifier rank at most $n + 1$, by the pigeonhole principle, there exist $m \neq m'$ such that $C_m^{\vee} \equiv_{\text{GSO}}^{n+1} C_{m'}^{\vee}$. Therefore, by Lemma 4.2, $C_m^{\vee} \uplus C_{m'}^{\vee} \equiv_{\text{GSO}}^{n+1} C_m^{\vee} \uplus C_m^{\vee}$. It follows by invariance under isomorphism that every normal node in $C_m^{\vee} \uplus C_m^{\vee}$ satisfies the same GSO-formulas $\phi(x)$, and similarly for the auxiliary nodes. In other words, for all GSO-formulas $\phi(x)$, we have that

$$C_m^{\vee} \uplus C_m^{\vee} \models \forall x (\text{normal}(x) \to \phi(x)) \vee \forall x (\text{normal}(x) \to \neg\phi(x))$$

and

$$C_m^{\vee} \uplus C_m^{\vee} \models \forall x (\text{auxiliary}(x) \to \phi(x)) \vee \forall x (\text{auxiliary}(x) \to \neg\phi(x))$$

where $\text{normal}(x)$ is a shorthand for $\exists y f(y, x)$ and $\text{auxiliary}(x)$ is a shorthand for $\exists y f(x, y)$. Since $C_m^{\vee} \uplus C_{m'}^{\vee} \equiv_{\text{GSO}}^{n+1} C_m^{\vee} \uplus C_m^{\vee}$, the same holds in the structure $C_m^{\vee} \uplus C_{m'}^{\vee}$ for ϕ of quantifier rank at most n. This concludes the proof of Claim 1.

Note that the signature of C is $\{f, g\}$ and that f and g are partial functions. Let X be the set consisting of the following partial functions over the domain of C:

- f,
- g,
- the identity function id,
- id_1 which is id restricted to the auxiliary nodes,
- id_2 which is id restricted to the normal nodes,
- $f \cup \text{id}_2$,
- $g \cup \text{id}_2$,
- the empty relation $\emptyset$.

Each of the partial functions in X is $\mathbb{TRA}$-term definable in C, and it will be convenient to expand C with these partial functions. That is, we will treat C as a structure over a signature consisting of these eight partial functions.

Figure 1: Structure $C_m^\vee$

Claim 2: Let $\phi(x,y)$ be any GSO-formula that is function-preserving. Then $C \models \phi(a,b)$ implies that (a,b) belongs to $f \cup g \cup \mathrm{id}$. In other words, ϕ defines a subrelation of $f \cup g \cup \mathrm{id}$ in C.

Proof of Claim 2: This can be shown using an automorphism argument: suppose that $C \models \phi(a,b)$, and suppose, for the sake of a contradiction, that b is not equal to $f(a)$, $g(a)$, or a itself. We will show that, then, there exists some $b' \neq b$ such that $(C,a,b) \cong (C,a,b')$, and therefore $C \models \phi(a,b')$, contradicting the assumption that $\phi(x,y)$ was function preserving. We argue by cases. First, suppose that a is a normal node. We may assume without loss of generality that $a = a_{1,1}$. Recall that $(a,b) \notin \mathrm{id}$. If $b = a_{1,2}$ or $b = a_{1,3}$, then we can pick b' to be $a_{1,3}$, respectively, $a_{1,2}$. It then follows from the construction of the structure C that $(C,a,b) \cong (C,a,b')$. Similarly, if $b = a_{i,j}$ with $i \neq 1$, then it follows from the construction of the structure C that $(C,a,b) \cong (C,a,b')$ for all $b' = a_{i,j'}$. Finally, if b is an auxiliary node, then it follows from the construction of the structure C that $(C,a,b) \cong (C,a,b')$ for some auxiliary node $b' \neq b$ from the same cluster. This concludes the case where a is a normal node. Next, suppose that a is an auxiliary node, and recall that $(a,b) \notin f \cup g \cup \mathrm{id}$. Regardless whether b is a normal or a special node, it follows that a and b do not co-occur in any fact (i.e., tuple in a relation) of C. It easy to see that, then, $(C,a,b) \cong (C,a,b')$ must be satisfied if we choose $b' \neq b$ to be another node from the same cluster as b. This concludes the proof of Claim 2.

Claim 3: Let $\phi(x,y)$ be any GSO-formula of quantifier rank less than n that is function-preserving. If $C \models \phi(a,b)$ and $f(a) = b$, then for all a' and b' with $f(a') = b'$ we have that $C \models \phi(a',b')$. Likewise for the functions g, id_1, and id_2.

Proof of Claim 3: We will discuss the proof for the case for f. The same argument applies to g, while the cases for id_1 and id_2 follow immediately from Claim 1. Assume $C,a,b \models \phi(x,y)$. Then $C,a \models \exists y(f(x,y) \wedge \phi(x,y))$. Therefore, by Claim 1, we have $C,a' \models \exists y(f(x,y) \wedge \phi(x,y))$, and therefore, since f is a partial function and $f(a') = b'$, we have $C,a',b' \models \phi(x,y)$. This concludes the proof of Claim 3.

The next claim follows from Claim 2 and 3.

Claim 4: If $\phi(x, y)$ is any GSO-formula of quantifier rank less than n that is function-preserving, then the partial function defined by $\phi(x, y)$ in C belongs to X.

Proof of Claim 4: By Claim 2, the relation $R = \{(c, d) \mid C, c, d \models \phi(x, y)\}$ is contained in $f \cup g \cup \mathrm{id}_1 \cup \mathrm{id}_2$, while by Claim 3, $R \cap f \neq \emptyset$ implies $f \subseteq R$, and likewise for g, id_1 and id_2. It follows that R must be equal to the union of a subset of the relations $f, g, \mathrm{id}_1, \mathrm{id}_2$. In other words, R belongs to X. This concludes the proof of Claim 4.

Claim 4 tell us that no function-preserving GSO-operation with quantifier rank smaller than n can take us outside of the set X. Since each operation in $\mathcal{O}$ is defined by a GSO-formula of quantifier rank less than n, and is function preserving, this implies, by induction, that every term of $\mathbb{BRA}(\mathcal{O})$ denotes one of the relations in X in C.

This implies the theorem: consider the $\mathbb{TRA}$-term $(f^{\smile} \circ g)^m \cap \mathrm{id}$, where $(\,\cdot\,)^m$ stands for an m-fold composition. This term denotes the identity relation restricted to the normal nodes of C_m only; this relation does not belong to X. Therefore, this term cannot be equivalent to any term of $\mathbb{BRA}(\mathcal{O})$. Nevertheless it is function preserving, simply because its interpretation always consists only of reflexive edges. $\qquad\square$

With some minor modifications, the same argument applies to total-function-preserving operations:

Theorem 4.3. *Let $\mathcal{O}$ be a finite set of total-function-preserving GSO-definable operations on binary relations. Then there is a total-function-preserving operation O that is term definable in $\mathbb{TRA}$ but not in $\mathbb{BRA}(\mathcal{O})$, even over finite structures in which every relation is a total function.*

Proof. (sketch) We use the same construction as before, except that we extend the structure C with an additional "sink node" s and an additional function $\hat{\emptyset}$ where $\hat{\emptyset}(c) = s$ for all nodes c (including s itself). Observe that $\hat{\emptyset}$ is a total function. We also extend the partial functions f and g to total functions $\hat{f}$ and $\hat{g}$, by setting $\hat{f}(c) = \hat{g}(c) = s$ for every normal node c and $\hat{f}(s) = \hat{g}(s) = s$. Note that the old partial functions f and g are $\mathbb{TRA}$-term definable from the new ones, namely as $f = \hat{f} - (\top \circ \hat{\emptyset})$ and $g = \hat{g} - (\top \circ \hat{\emptyset})$. Now the same argument as before shows that the $\mathbb{TRA}$-term

$$\left((f^{\smile} \circ g)^m \cap \mathrm{id}\right) \sqcup \hat{\emptyset}$$

(where f and g are now shorthand for the aforementioned terms, and where $\sqcup$ is the preferential union operator) defines a total-function-preserving operation that is not term definable in $\mathbb{BRA}(\mathcal{O})$. $\qquad\square$

As a consequence of Theorem 4.1, we obtain the following.

Corollary 4.4. *Both in general and in the finite:*

(1) *The function-preserving fragment of $\mathbb{TRA}$ is not finitely generated. In particular, not every function-preserving $\mathbb{TRA}$-term is term definable in $\mathbb{FA}$.*
(2) *The homomorphism-safe function-preserving fragment of $\mathbb{TRA}$ is not finitely generated.*
(3) *The $\subseteq$-safe function-preserving fragment of $\mathbb{TRA}$ is not finitely generated.*

Proof. The first item follows immediately from Theorem 4.1. The other items follow from its proof. This is because the $\mathbb{TRA}$-term used as counterexample in the proof, i.e., $(f^{\smile} \circ g)^m \cap \mathrm{id}$, uses only operations that are homomorphism safe and $\subseteq$-safe. (Note that the same does *not* hold in the total-function-preserving case because there we used preferential union.) $\qquad\square$

Question 4.5. *Is the homomorphism-safe total-function-preserving fragment of $\mathbb{TRA}$ finitely generated?*

Given that the function-preserving fragment of $\mathbb{TRA}$ is not finitely generated, one may ask if it is at least generated by a *recursive* set of operations. This is indeed the case, for a trivial reason: for any $\mathbb{TRA}$ term t, consider the term $t' = t \setminus (t \circ (\top \setminus \mathrm{id}))$. By construction t' always outputs a partial function. Furthermore, on any input where t produces a partial function, t' produces the same output as t. Therefore, the function-preserving fragment of $\mathbb{TRA}$ is generated by the (recursive) set of all $\mathbb{TRA}$-terms of the form $t \setminus (t \circ (\top \setminus \mathrm{id}))$.

Another question left open by the above results is whether $\mathbb{FA}$, although it is not the function-preserving fragment of $\mathbb{TRA}$, can still be characterized as a natural fragment of $\mathbb{TRA}$.

Question 4.6. *Can $\mathbb{FA}$ be characterised as a fragment of $\mathbb{TRA}$ using additional properties besides function preserving (or using a strengthening of the notion of "function preserving")?*

5. The forward function-preserving fragment is finitely generated

In our proof of Theorem 4.1, we implicitly made use of the fact that any binary relation can be represented as a composition $f^{\smile} \circ g$, where f, g are partial functions. That is, we crucially made use of the inverse operation. This is indeed essential to the proof: as we will now show, if we restrict attention to direction-preserving operations (*forward* operations, as we will call them below), then we do get a binary relation algebra that is finitely generated.

Formally, we say that an n-ary operation O on binary relations is *forward* if for all structures A over signature $\sigma = \{R_1, \ldots, R_n\}$ and for all pairs $(a, b) \in \mathrm{dom}(A)$, we have that $(a, b) \in O(A)$ if and only if $(a, b) \in O(A_a)$ where A_a is the substructure of A generated by a, i.e., the induced substructure of A whose domain consists of all elements reachable from a by a finite directed path along the relations $R_1^A, \ldots, R_n^A$. In particular, this implies that, whenever $(a, b) \in O(A)$ then b must belong to A_a. We say that O is *forward over a class of structures* K if the above holds for all structures $A \in K$.[2]

Lemma 5.1. *Let K be any FO-definable class of structures, and let O be any FO-definable operation on binary relations that is forward over K. Then there is a natural number m such that, for all structures $A \in K$ and $a, b \in \mathrm{dom}(A)$, whether (a, b) belongs to $O(A)$ depends only on the substructure of A consisting of the elements reachable from a by a directed path of length at most m.*

Proof. This can be shown using a simple compactness argument [Ben07]: let χ be the FO-sentence defining K, and let n be the arity of the operation O. By assumption, O is defined by a first-order formula $\phi(x, y)$ over the signature consisting of the relation symbols $R_1, \ldots, R_n$. Let P be a fresh unary relation symbol, let ϕ^P be the result of relativising all quantifiers in ϕ by P (i.e., replacing $\exists z$ by $\exists z(P(z) \wedge \ldots)$ and replacing $\forall z$ by $\forall z(P(z) \to \ldots)$). Furthermore, for every natural number k, let $\psi_k(x)$ be the FO-formula expressing that all elements reachable from x by a directed path of length at most k satisfy P. Then $\{\chi, \psi_k(x) \mid k \geq 0\} \models \forall y(\phi(x, y) \leftrightarrow (P(y) \wedge \phi^P(x, y)))$. It follows by compactness

---

[2]Note that being forward is a stronger requirement than requiring that $b \in A_a$ for all $(a, b) \in O(A)$. Indeed, the operation defined by the $\mathbb{TRA}$-expression $R \cap (R^{\smile} \circ R)$ satisfies the latter requirement but is not forward.

that, for some m, $\{\chi, \psi_k(x) \mid 0 \leq k \leq m\} \models \forall y(\phi(x,y) \leftrightarrow (P(y) \wedge \phi^P(x,y)))$. This proves the lemma. $\qquad \square$

Theorem 5.2. *Let K_{pf} be the class of structures in which each relation is a partial function, and let O be any FO operation on binary relations. The following are equivalent:*

(1) *O is function preserving and forward over K_{pf},*
(2) *O is term-definable in $\mathbb{BRA}(\circ, \sim, \cap, \sqcup)$ over K_{pf}.*

Proof. The direction from 2 to 1 is straightforward. For the direction from 1 to 2: let O be any n-ary FO operation that is function preserving and forward over K_{pf}. From the fact that O is forward over K_{pf}, it follows by Lemma 5.1 that there exists a constant $m > 0$ (depending on O) such that whether a pair (a, b) belongs to $O(A)$, for $A \in K_{\mathrm{pf}}$, depends only on the substructure $B \subseteq A$ consisting of the elements reachable from a by a directed path of length at most m. For $A \in K_{\mathrm{pf}}$, such a substructure B can be of size at most $(n+1)^m$. There are only finitely many isomorphism types of such structures B. Furthermore, for each such B, the structure (B, a) can be characterised up to isomorphism by an intersection $\chi_{B,a}$ of terms of the following forms:

- $\sim(f_1 \circ \cdots \circ f_k)$
  "there is no outgoing $f_1 \circ \cdots \circ f_k$ path"
- $\sim\sim(f_1 \circ \cdots \circ f_k)$
  "there is an outgoing $f_1 \circ \cdots \circ f_k$ path"
- $\sim(f_1 \circ \cdots \circ f_k \cap g_1 \circ \cdots \circ g_l)$
  "the outgoing $f_1 \circ \cdots \circ f_k$ path and the outgoing $g_1 \circ \cdots \circ g_l$ path do not lead to the same node"
- $\sim\sim(f_1 \circ \cdots \circ f_k \cap g_1 \circ \cdots \circ g_l)$
  "the outgoing $f_1 \circ \cdots \circ f_k$ path and the outgoing $g_1 \circ \cdots \circ g_l$ path do lead to the same node"

Note that here we implicitly use id (which is definable as $\sim(\sim f \circ f)$) for the case where $k = 0$ or $l = 0$. Finally, we can take our term to be $\chi_{B,a} \circ (f_1 \circ \cdots \circ f_k)$ where $f_1, \ldots, f_k$ describes an arbitrary directed path from a to b (or simply $\chi_{B,a}$ if the path is empty). Doing this for each isomorphism type of structure $B \models \phi(a, b)$, we obtain finitely many terms (defining relations that are guaranteed to be pairwise disjoint from each other) we then combine using the preferential union operator (in arbitrary order, since they are pairwise disjoint). In the special case where there is no $B \models \phi(a, b)$, we may choose as our term $\emptyset$ (which is definable as $\sim f \circ f$). $\qquad \square$

The collection $\{\circ, \sim, \cap, \sqcup\}$ of operations identified in Theorem 5.2 is one that has already been investigated in the literature. Specifically, Jackson and Stokes [JS11] give a finite equational axiomatisation of the class of algebras isomorphic to a set of partial functions equipped with these operations. The equational theory of these algebras is coNP-complete [HJM16].

Question 5.3. *Does Theorem 5.2 hold in the finite?*

Although we do not know the answer to this question, we can show that Lemma 5.1 fails in the finite, and therefore, a different approach is required.

Proposition 5.4. *Lemma 5.1 fails when K is the class of all finite structures (which is not FO-definable).*

Proof. Let $\phi(u)$ be the conjunction of the following FO-formulas:

Figure 2: Structure satisfying $\psi(a,a)$.

- $R_3(u,u)$
- $\forall v(R_3(u,v) \rightarrow \exists w(R_3(v,w) \wedge R_2(w,v)))$
- $\forall vw(R_3(u,v) \wedge R_3(v,w) \rightarrow R_3(u,w))$
- $\neg \exists v R_2(u,v)$
- $\forall vw(R_3(u,v) \wedge (\exists^{\geq 2} s R_2(v,s)) \wedge R_1(u,w) \rightarrow R_4(w,v))$

It follows from the fact that every quantifier is bounded by a forward-oriented atom, that $\phi(u)$ is invariant for generated substructures [Fef68]. That is, for all structures A and elements a, we have $A \models \phi(a)$ if and only if $A_a \models \phi(a)$.

Next, let $\psi(x,y) := (x = y) \wedge \exists u(R_1(u,x) \wedge \phi(u))$.

It follows immediately from the presence of the equality conjunct that $\psi(x,y)$ is function preserving.

Claim 1: $\psi(x,y)$ defines a forward operation, i.e., for all finite structures A, we have $A \models \psi(a,a)$ if and only if $A_a \models \psi(a,a)$.

The right-to-left direction is easy (and does not depend on the restriction to finite structures). For the other direction, suppose that $A \models \psi(a,a)$. It follows, by the construction of ψ and the finiteness of the structure A, that there exist elements connected as in Figure 2. (In fact, further facts hold that have not been drawn in the figure to avoid cluttering. Specifically, $R_3(b_i, b_j)$ holds for all $i < j$.)

It follows that all the depicted elements belong to A_a. In particular, b_0 belongs to A_a. From this, it follows that $A_a \models \psi(a,a)$. This concludes the proof of Claim 1.

Now let m be any natural number. Let A be the structure drawn above, with $n = m+1$. Let B be the identical structure but with the node b_0 removed. Clearly, $A \models \psi(a,a)$ and $B \not\models \psi(a,a)$ (because B lacks a reflexive R_3-edge). However, the induced substructures consisting of nodes reachable from a by a directed path of length at most m are identical. □

Proposition 5.4, incidentally, also resolves in the negative an open question about hybrid logic posed in [AM22, Section 7], namely whether the technique used [AM22] for proving a preservation theorem for hybrid temporal logic in the finite could be extended to prove a similar result for the case without backward modalities. It follows from Proposition 5.4 that the corresponding preservation theorem in the finite in fact fails for hybrid logic without backward modalities.

We can adapt the proof of Theorem 5.2 to obtain a similar, but undirected, result for *injective* partial functions. For this, we say that O is *injective-function preserving* if the following holds for all structures $A = (\text{dom}(A), R_1^A, \ldots, R_n^A)$: if each R_i^A is an injective partial function on $\text{dom}(A)$, then $O(A)$ is an injective partial function on $\text{dom}(A)$. Let us also say that that an n-ary operation O on binary relations is *local* if for all structures A over signature $\sigma = \{R_1, \ldots, R_n\}$ and for all pairs $(a, b) \in \text{dom}(A)$, we have that $(a, b) \in O(A)$ if and only if $(a, b) \in O(A_a^\leftrightarrow)$ where $A_a^\leftrightarrow$ is the induced substructure of A whose domain consists of all elements reachable from a by a finite undirected path along the relations $R_1^A, \ldots, R_n^A$. As before, this implies that, whenever $(a, b) \in O(A)$ then b must belong to $A_a^\leftrightarrow$.

To state the result, we first define a variant of preferential union that is injective-function preserving. We call this new operation *injective union* and use $\stackrel{\perp}{\sqcup}$ to denote it. The operation adds to its first argument any pairs from its second argument whose addition does not violate functionality or injectivity. One possible term definition of injective union is as follows.

$$f \stackrel{\perp}{\sqcup} g := (f \sqcup g) \cap (f^\smile \sqcup g^\smile)^\smile$$

Theorem 5.5. *Let K_{ipf} be the class of structures in which each relation is an injective partial function, and let O be any FO operation on binary relations. The following are equivalent:*

(1) *O is injective-function preserving and local over K_{ipf}.*
(2) *O is term-definable in $\mathbb{BRA}(\circ, \sim, \cap, \smile, \stackrel{\perp}{\sqcup})$ over K_{ipf}.*

Proof. (sketch) First note that we can obtain an undirected analog of Lemma 5.1 using a similar proof. That is, if an FO-definable operation is local over K_{ipf}, then there is a natural number m such that, for all structures $A \in K_{\text{ipf}}$ and $a, b \in \text{dom}(A)$, whether (a, b) belongs to $O(A)$ depends only on the substructure of A consisting of the elements reachable from a by an undirected path of length at most m.

Next, the same proof used for Theorem 5.2 works if we replace every instance of 'directed path' by 'oriented path' (i.e., sequence of possibly reverse-oriented edges), use $\smile$ to express reverse-oriented edges in such paths, and use $\stackrel{\perp}{\sqcup}$ in place of $\sqcup$. □

The collection $\{\circ, \sim, \cap, \smile, \stackrel{\perp}{\sqcup}\}$ of operations identified in Theorem 5.5 is one that has been considered in the literature on inverse semigroups. Any set of injective partial functions closed under these operations forms a Boolean inverse monoid in the sense of Lawson [Law10]; indeed these are the canonical examples of Boolean inverse monoids.[3] Conversely, from the results of Lawson it can be seen that any Boolean inverse monoid is isomorphic to one of these algebras of injective partial functions [Law10, Proposition 2.23(2)]. Thus Theorem 5.5

---

[3]In the definition of Boolean inverse monoids, only joins of pairs of *orthogonal* elements are required, that is, elements a and b such that $a^\smile \circ b$ and $a \circ b^\smile$ are both zero. This is equivalent to the presence of $\stackrel{\perp}{\sqcup}$, since clearly if a collection of injective partial functions is closed under $\stackrel{\perp}{\sqcup}$ then it is closed under orthogonal joins, and conversely, in the presence of the other operations, $\stackrel{\perp}{\sqcup}$ is expressible as the orthogonal join of $\sim b \circ a \circ \sim(b^\smile)$, $\sim a \circ b \circ \sim(a^\smile)$, and $a \cap b$.

demonstrates that within the program of studying enrichments of inverse semigroups, the Boolean inverse monoids are in a sense the *fully enriched* instances.

6. Conclusion

In summary, our results show that certain semantic fragments of Tarski's relation algebra, such as the homomorphism-safe fragment, admit a syntactic characterisation in terms of a finite set of operations, while others, such as the function-preserving fragment, do not. We hope that these results show that the study of preservation theorems in the context of algebras of binary relations is an interesting topic. We conclude by listing a few directions that deserve further exploration.

Firstly, one could explore the same questions for other semantic properties of operations on binary relations (e.g., $\subseteq$-safety, as mentioned in Section 3, as well as additivity [BOP+19]). Secondly, our results concern fragments of $\mathbb{TRA}$, but the same questions can be asked for other binary relation algebras, including ones that contain the transitive closure operator. In particular, our results leave open the question whether the function-preserving fragment of Kleene Algebra with Tests (KAT) is finitely generated.

Finally, various applications of $\mathbb{TRA}$ in computer science and elsewhere are concerned with a restricted class of structures, such as finite trees (e.g., XPath), linear orders (e.g., interval temporal logics), or variable-assignment spaces (e.g., dynamic predicate logic [GS91] and the Logic of Information Flows (LIF) [Ter19, Moh23]). It is therefore meaningful to ask whether our results hold also over these restricted classes of structures.

References

[ACN85]   Hajnal Andréka, Stephen D. Comer, and Istvan Németi. Clones of operations on relations. In Stephen D. Comer, editor, *Universal Algebra and Lattice Theory*, pages 17–21, Berlin, Heidelberg, 1985. Springer. `doi:10.1007/BFb0098452`.

[AM22]   Samson Abramsky and Dan Marsden. Comonadic semantics for hybrid logic. In *Proceedings of MFCS 2022*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:14, Dagstuhl, Germany, 2022. `doi:10.4230/LIPIcs.MFCS.2022.7`.

[AvBN23] Hajnal Andréka, Johan van Benthem, and Istvan Németi. A note on the submodel preservation property in fragments of first-order logic. *CoRR*, abs/2303.13222, 2023. URL: `http://arxiv.org/abs/2303.13222`.

[B̈86]   Ferdinand Börner. One-generated clones of operations on binary relations. *Beiträge zur Algebra und Geometrie = Contributions to algebra and geometry*, 23:73–84, 1986. URL: `http://eudml.org/doc/138351`.

[BC19]   Simone Bova and Hubie Chen. How many variables are needed to express an existential positive query? *Theory Comput. Syst.*, 63:1573–1594, 2019. `doi:10.1007/s00224-018-9884-z`.

[Ben98]   Johan van Benthem. Program constructions that are safe for bisimulation. *Studia Logica*, 60(2):311–330, 1998. `doi:10.1023/A:1005072201319`.

[Ben07]   Johan van Benthem. A new modal Lindström theorem. *Logica Universalis*, 1(1):125–138, 2007. `doi:10.1007/s11787-006-0006-3`.

[BOP+19] Gerald Berger, Martin Otto, Andreas Pieris, Dimitri Surinx, and Jan Van den Bussche. Additive first-order queries. In P. Barcelo and M. Calautti, editors, *Proceedings 22nd International Conference on Database Theory (ICDT 2019)*, volume 127 of *Leibniz International Proceedings in Informatics*, pages 19:1–19:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019. `doi:10.4230/LIPIcs.ICDT.2019.19`.

[BP91]   Ferdinand Börner and Reinhard Pöschel. Clones of operations on binary relations. In *Contributions to general algebra 7. Proceedings of the Vienna conference, June 14-17, 1990, Vienna, Austria*, pages 50–70, 1991. URL: `http://eudml.org/doc/138351`.

[CE12]   Bruno Courcelle and Joeri Engelfriet. *Graph Structure and Monadic Second-Order Logic*. Cambridge University Press, 2012.

[Fef68]   Solomon Feferman. Persistent and invariant formulas for outer extensions. *Compositio Mathematica*, 20:29–52, 1968. URL: `http://eudml.org/doc/88983`.

[FGL+15a]   George H.L. Fletcher, Marc Gyssens, Dirk Leinders, Dimitri Surinx, Jan Van den Bussche, Dirk Van Gucht, Stijn Vansummeren, and Yuqing Wu. Relative expressive power of navigational querying on graphs. *Information Sciences*, 298:390–406, 2015. `doi:10.1016/j.ins.2014.11.031`.

[FGL+15b]   George H.L. Fletcher, Marc Gyssens, Dirk Leinders, Dimitri Surinx, Jan Van den Bussche, Dirk Van Gucht, Stijn Vansummeren, and Yuqing Wu. Relative expressive power of navigational querying on graphs using transitive closure. *Logic Journal of the IGPL*, 23(5):759–788, 2015. `doi:10.1016/j.ins.2014.11.031`.

[GHO02]   Erich Grädel, Colin Hirsch, and Martin Otto. Back and forth between guarded and modal logics. *ACM Transactions on Computational Logic*, 3(3):418–463, 2002. `doi:10.1145/507382.507388`.

[GHR94]   Dov Gabbay, Ian Hodkinson, and Mark Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects, Vol. 1*. Clarendon Press, Oxford, 1994.

[GS91]   Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100, 1991. `doi:10.1007/BF00628304`.

[HJM16]   Robin Hirsch, Marcel Jackson, and Szabolcs Mikulás. The algebra of functions with antidomain and range. *Journal of Pure and Applied Algebra*, 220(6):2214–2239, 2016. `doi:10.1016/j.jpaa.2015.11.003`.

[HWGV22]   Jelle Hellings, Yuqing Wu, Marc Gyssens, and Dirk Van Gucht. The power of Tarski's relation algebra on trees. *Journal of Logical and Algebraic Methods in Programming*, 126:100748, 2022. `doi:10.1016/j.jlamp.2022.100748`.

[JS11]   Marcel Jackson and Tim Stokes. Modal restriction semigroups: towards an algebra of functions. *International Journal of Algebra and Computation*, 21(7):1053–1095, 2011. `doi:10.1142/S0218196711006844`.

[Kam68]   Hans Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, UCLA, 1968.

[Law10]   Mark V. Lawson. A noncommutative generalization of Stone duality. *Journal of the Australian Mathematical Society*, 88(3):385–404, 2010. `doi:10.1017/s1446788710000145`.

[Lib04]   Leonid Libkin. *Elements of Finite Model Theory*. Springer, August 2004.

[Löw15]   Leopold Löwenheim. Über möglichkeiten im relativkalkül. *Mathematische Annalen*, 76:447–470, 1915.

[Mad91]   Roger D. Maddux. The origin of relation algebras in the development and axiomatization of the calculus of relations. *Studia Logica*, 50(3/4):421–455, 1991. `doi:10.1007/BF00370681`.

[Mak04]   Johann A. Makowsky. Algorithmic uses of the Feferman–Vaught theorem. *Annals of Pure and Applied Logic*, 126(1):159–213, 2004. Provinces of logic determined. Essays in the memory of Alfred Tarski. Parts I, II and III. `doi:10.1016/j.apal.2003.11.002`.

[Mar05]   Maarten Marx. Conditional XPath. *ACM Transactions on Database Systems*, 30(4):929–959, dec 2005. `doi:10.1145/1114244.1114247`.

[McL18]   Brett McLean. *Algebras of partial functions*. PhD thesis, University College London, 2018.

[Moh23]   Heba Mohamed. *Logical Analysis of Input and Output Sensitivity in the Logic of Information Flows*. PhD thesis, Hasselt University, 2023.

[Pra92]   Vaughan Pratt. Origins of the calculus of binary relations. In *Proceedings 7th Annual IEEE Symposium on Logic in Computer Science (LICS 1992)*, pages 248–254, 1992. `doi:10.1109/LICS.1992.185537`.

[Ros08]   Benjamin Rossman. Homomorphism preservation theorems. *Journal of the ACM*, 55(3):1–53, 2008. `doi:10.1145/1379759.1379763`.

[Tar41]   Alfred Tarski. On the calculus of relations. *The Journal of Symbolic Logic*, 6(3):73–89, 1941. URL: `http://www.jstor.org/stable/2268577`.

[Ter19]   Eva Ternovska. An algebra of modular systems: static and dynamic perspectives. In A. Herzig and A. Popescu, editors, *Frontiers of Combining Systems: Proceedings 12th FroCos*, volume 11715 of *Lecture Notes in Artificial Intelligence*, pages 94–111. Springer, 2019. `doi:10.1007/978-3-030-29007-8_6`.

[TG87]   Alfred Tarski and Steven Givant. *A Formalization Of Set Theory Without Variables*. AMS, 1987.

[Ven90]   Yde Venema. Expressiveness and completeness of an interval tense logic. *Notre Dame Journal of Formal Logic*, 31:529–547, 1990. `doi:10.1305/NDJFL/1093635589`.