The complexity of positive first-order logic without equality II: The four-element case

Barnaby Martin[1] and Jos Martin[2]

[1] Department of Computer Science, University of Durham,
Science Labs, South Road, Durham DH1 3LE, U.K.
[2] The MathWorks, Matrix House, Cambridge, CB4 0HH, U.K.

Abstract. We study the complexity of evaluating positive equality-free sentences of first-order (FO) logic over fixed, finite structures $\mathcal{B}$. This may be seen as a natural generalisation of the non-uniform quantified constraint satisfaction problem QCSP($\mathcal{B}$). Extending the algebraic methods of a previous paper, we derive a complete complexity classification for these problems as $\mathcal{B}$ ranges over structures of domain size 4. Specifically, each problem is either in L, is NP-complete, is co-NP-complete or is Pspace-complete.

1 Introduction

We continue the study of the evaluation problem for positive equality-free first-order logic, on a fixed finite structure $\mathcal{B}$, denoted $\{\exists, \forall, \land, \lor\}$-FO($\mathcal{B}$), started in [5]. This problem is a close relative of the *constraint satisfaction problem*, CSP($\mathcal{B}$), and an even closer relative of the *quantified CSP*, QCSP($\mathcal{B}$). In fact, it is noted in [5] that among a wide family of problems, the only interesting case, other than the CSP and QCSP, is the one addressed in this paper. The bulk of the theoretical research into CSPs concerns the so-called *dichotomy conjecture*: that the complexity of the problem of evaluating a primitive positive sentence on a fixed finite $\mathcal{B}$, CSP($\mathcal{B}$), is either in P or is NP-complete. This was solved for structures with two elements domains in [8] and improved to encompass structures with three element domains in [2]. The most successful approach to date, and the method used in [2], has been the so-called algebraic method, in which the problem of classification reverts to classes of functions under which the relevant relational systems are invariant. A similar algebraic approach has been successful in the study of the evaluation problem for positive Horn sentences, QCSP($\mathcal{B}$), and, while no formal trichotomy has there been conjectured, the only known attainable complexities are P, NP-complete and Pspace-complete.

In [5], the complexity of $\{\exists, \forall, \land, \lor\}$-FO($\mathcal{B}$) is studied through an analogous algebraic method to that used for CSP($\mathcal{B}$). The paper culminates in a full classification – a tetrachotomy – as $\mathcal{B}$ ranges over structures with three-element domains. Specifically, the problems $\{\exists, \forall, \land, \lor\}$-FO($\mathcal{B}$) are either Pspace-complete, NP-complete, co-NP-complete or in L. In this paper we extend this work to deriving a similar tetrachotomy for structures with a four element domain. This is

the largest domain size for which a comparable result is known. However, our task is somewhat easier since our lattices are finite for all finite domain sizes, while the clone lattice even for three-element domains – that is used for the study of $\mathrm{CSP}(\mathcal{B})$ – is not only infinite, but actually uncountable.

We derive our result by studying sets of surjective hyper-operations (shops) under which our relational $\mathcal{B}$ may be invariant. These sets, always containing the identity and closed under composition and sub-shops, are known as down shop-monoids (DSMs). For membership of L, NP and co-NP, it turns out that it is sufficient to have certain special shops as surjective hyper-endomorphisms (shes). In the four-element case, we are able to prove that these are necessary, too. We do this by isolating certain maximal DSMs and proving the relevant properties of them. Many of our proofs contain parts where the case-verification is undertaken by computer. There are 41503 shops on a four element domain, so the number of DSMs is bound by 2^{41503}. This is far too large a search space to simply be scanned, therefore we have come up with a novel, computational inductive way to demonstrate that certain small collections of DSMs are exactly those that are in some sense "maximal". A trade-off between size of search space and number of putative solutions ensures that our computational procedures are not of too high complexity.

Having isolated the maximal DSMs, we find that we can use the methods of [5] to classify most of them, but there is one completely new class in the four-element case that is unlike anything in the three-element case. For this class, a separate proof of Pspace-hardness is required.

The paper is organised as follows. In Section 2, we detail the necessary preliminarlies and necessary results of [5]. In Section 3, we prove the completeness of our Pspace-complete classes. Section 4, we prove the completeness of the NP-complete classes; Section 5 does likewise with the co-NP-complete classes. In Section 6, we conclude with the *coup de grâce* and some final remarks. The code used for the computational verification procedures is available at `http://www.dur.ac.uk/barnaby.martin/`.

2   Preliminaries

Throughout, let $\mathcal{B}$ be a finite structure, with domain B, over the finite relational signature σ. Let $\{\exists, \forall, \wedge, \vee\}$-FO be the positive fragment of first-order (FO) logic without equality. An *extensional* relation is one that appears in the signature σ. We will usually denote extensional relations of $\mathcal{B}$ by R and other relations by S (or by some formula that defines them). In $\{\exists, \forall, \wedge, \vee\}$-FO the atomic formulae are exactly substitution instances of extensional relations. The problem $\{\exists, \forall, \wedge, \vee\}$-FO$(\mathcal{B})$ has:

 – Input: a sentence $\varphi \in \{\exists, \forall, \wedge, \vee\}$-FO.
 – Question: does $\mathcal{B} \models \varphi$?

When $\mathcal{B}$ is of size one, the evaluation of any FO sentence may be accomplished in L (essentially, the quantifiers are irrelevant and the problem amounts

to the *boolean sentence value problem*, see [4]). In this case, it follows that $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in L. Furthermore, by inward evaluation of the quantifiers, $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is readily seen to always be in Pspace.

Consider the set B and its power set $\mathfrak{P}(B)$. A *hyper-operation* on B is a function $f : B \to \mathfrak{P}(B) \setminus \{\emptyset\}$ (that the image may not be the empty set corresponds to the hyper-operation being *total*, in the parlance of [1]). If the hyper-operation f has the additional property that

– for all $y \in B$, there exists $x \in B$ such that $y \in f(x)$,

then we designate (somewhat abusing terminology) f *surjective*. A surjective hyper-operation (shop) in which each element is mapped to a singleton set is identified with a *permutation* (bijection). A *surjective hyper-endomorphism* (she) of $\mathcal{B}$ is a shop f on B that satisfies, for all extensional relations R of $\mathcal{B}$,

– if $\mathcal{B} \models R(x_1, \ldots, x_i)$ then, for all $y_1 \in f(x_1), \ldots, y_i \in f(x_i)$, $\mathcal{B} \models R(y_1, \ldots, y_i)$.

More generally, for $r_1, \ldots, r_k \in B$, we say f is *a she from* $(\mathcal{B}, r_1, \ldots, r_k)$ to $(\mathcal{B}, r_1', \ldots, r_k')$ if f is a she of $\mathcal{B}$ and $r_1' \in f(r_1), \ldots, r_k' \in f(r_k)$. A she may be identified with a *surjective endomorphism* if each element is mapped to a singleton set. On finite structures surjective endomorphisms are necessarily automorphisms.

2.1 Galois Connections

For a set F of shops on the finite domain B, let $\mathsf{Inv}(F)$ be the set of relations on B of which each $f \in F$ is a she (when these relations are viewed as a structure over B). We say that $S \in \mathsf{Inv}(F)$ is invariant or *preserved* by (the shops in) F. Let $\mathsf{shE}(\mathcal{B})$ be the set of shes of $\mathcal{B}$. Let $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-}\mathsf{FO}}$ be the set of relations that may be defined on $\mathcal{B}$ in $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}$. We recall the following results from [5].

Theorem 1 ([5]). *For a finite structure $\mathcal{B}$ we have*

(i). $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-}\mathsf{FO}} = \mathsf{Inv}(\mathsf{shE}(\mathcal{B}))$*, and*
(ii). *If $\mathsf{shE}(\mathcal{B}) \subseteq \mathsf{shE}(\mathcal{B}')$ then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B}') \leq_{\mathsf{L}} \{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$.*

2.2 Down-she-monoids

Consider a finite domain B. The *identity* shop id_B is defined by $x \mapsto \{x\}$. Given shops f and g, define the *composition* $g \circ f$ by $x \mapsto \{z : \exists y \; z \in g(y) \wedge y \in f(x)\}$. Finally, a shop f is a *sub-shop* of g – denoted $f \subseteq g$ – if $f(x) \subseteq g(x)$, for all x. A set of shops on a finite set B is a *down-she-monoid* (DSM), if it contains id_B, and is closed under composition and sub-shops (of course, not all sub-hyper-operations of a surjective hyper-operation are surjective – we are only concerned with those that are). id_B is a she of all structures, and, if f and g are shes of $\mathcal{B}$, then so is $g \circ f$. Further, if g is a she of $\mathcal{B}$, then so is f for all (surjective) $f \subseteq g$. It follows that $\mathsf{shE}(\mathcal{B})$ is always a DSM. The DSMs of B form a lattice under (set-theoretic) inclusion and, as per the Galois connection of the previous

3

section, classify the complexities of $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$. If F is a set of shops on B, then let $\langle F \rangle$ denote the minimal DSM containing the operations of F. If F is the singleton $\{f\}$, then, by abuse of notation, we write $\langle f \rangle$ instead of $\langle \{f\} \rangle$

Lemma 1 ([5]). *Let $\mathcal{B}$, with $|B| \geq 2$, be a structure s.t. $\mathsf{shE}(\mathcal{B})$ is a permutation subgroup. Then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is Pspace-complete.*

Proof. Let $\mathcal{B}_{NAE}$ be the structure on B with a single ternary relation $R_{NAE} := B^3 \setminus \{(b, b, b) : b \in B\}$. $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B}_{NAE})$ is a generalisation of the problem $\mathrm{QCSP}(\mathcal{B}_{NAE})$, well-known to be Pspace-complete (in the case $|B| = 2$, this is *quantified not-all-equal 3-satisfiability*, see, e.g., [7]). $\mathsf{shE}(\mathcal{B}_{NAE})$ is the symmetric group $S_{|B|}$. The statement of the theorem now follows from Theorem 1, since $\mathsf{shE}(\mathcal{B}) \subseteq \mathsf{shE}(\mathcal{B}_{NAE})$.

The following is a generalisation of Lemma 1.

Lemma 2 ([5]). *Let $\mathcal{B}$ be a structure whose universe admits the partition $B_1, \ldots, B_l$ ($l \geq 2$). If all shes of $\mathcal{B}$ are sub-hyper-operations of some f of the form $f(x) := B_i$ iff $x \in B_{\pi(i)}$, for π a permutation on the set $\{1, \ldots, l\}$, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is Pspace-complete.*

Proof. Let $\mathcal{K}_{|B_1|, \ldots, |B_l|}$ be the complete l-partite graph with partitions of size $|B_1|, \ldots, |B_l|$. It may easily be verified that the shes of $\mathcal{K}_{|B_1|, \ldots, |B_l|}$ are of the form of the lemma. Furthermore, $\mathcal{K}_{|B_1|, \ldots, |B_l|}$ agrees with the antireflexive l-clique $\mathcal{K}_l$ on all sentences of equality-free FO logic (for more detail on why this is, see, e.g., the Homomorphism Theorem of [3]), and certainly $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}$. Pspace-hardness of $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{K}_l)$ follows from the Lemma 1, and so Pspace-hardness of $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{K}_{|B_1|, \ldots, |B_l|})$ follows a fortiori. Finally, Pspace-hardness of $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ now follows from Theorem 1, since $\mathsf{shE}(\mathcal{B}) \subseteq \mathsf{shE}(\mathcal{K}_{|B_1|, \ldots, |B_l|})$.

The following is a generalisation of a result from [5] that appears in the journal version of that paper, [6].

Theorem 2. *If $\mathcal{B}$ has an A-shop as a she then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in NP. If $\mathcal{B}$ has an E-shop as a she then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in $\mathsf{co\text{-}NP}$. If $\mathcal{B}$ has a both an A-shop and an E-shop as a she then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in L.*

The following is the main result of [5].

Theorem 3 (Tetrachotomy [5]). *Let $\mathcal{B}$ be s.t. $|B| \leq 3$.*

I. *If $\mathsf{shE}(\mathcal{A})$ contains both an A-shop and an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in L.*
II. *If $\mathsf{shE}(\mathcal{A})$ contains an A-shop but no E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in NP-complete.*
III. *If $\mathsf{shE}(\mathcal{A})$ contains an E-shop but no A-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in $\mathsf{co\text{-}NP}$-complete.*
IV. *If $\mathsf{shE}(\mathcal{A})$ contains neither an A-shop nor an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in Pspace-complete.*

3 The Pspace-case

This section will progress in the pursuit of the following.

Proposition 1. *If* shE($\mathcal{B}$)*, where* $|B| = 4$*, contains neither an* A*-shop nor an* E*-shop, then* $\{\exists, \forall, \wedge, \vee\}$*-*FO($\mathcal{B}$) *is* Pspace*-complete.*

This will follow from the classification of all maximal DSMs which contain neither an A-shop nor an E-shop. Call a DSM M *maximally* Pspace*-hard* if M contains neither an A-shop nor an E-shop, but for any shop $f \notin M$, $\langle M \cup \{f\} \rangle$ contains either an A-shop or an E-shop. It turns out that there are twenty such DSMs, in five classes.

Class I   Class II Class III Class IV   Class V

Class V contains only the symmetric group S_4. Class IV contains variants of S_3 formed from a four-element domain by blurring two elements to one. Classes III and II are each variants of S_2, in III there are three elements blurred to one, while in IV there are twice two elements blurred to one. Class I, however, is of a different form altegether. It is not based around a symmetric group, and provides sport in the four-element case that was absent from the three-element case – where all maximally Pspace-hard DSMs were variants of S_2 or S_3. We know from [5] how to classify the complexity for Classes II – V, but we will need a new proof for Class I.

Note, that although one may verify that the single DSM of Class V is maximally Pspace-hard, the union of the DSMs of Classes I-IV in fact equals the union of the DSMs of Classes I-V.

Lemma 3. *Let* shE($\mathcal{B}$) *be any of the DSMs in Classes II – V. Then* $\{\exists, \forall, \wedge, \vee\}$*-*FO($\mathcal{B}$) *is* Pspace*-complete.*

Proof. A canonical representative $\mathcal{B}$ for each of the Classes II – V is drawn in Figure 1. That is, a $\mathcal{B}$ is drawn s.t. shE($\mathcal{B}$) ranges over the DSMs of Classes II – V. The result follows from Lemmas 1 and 2.
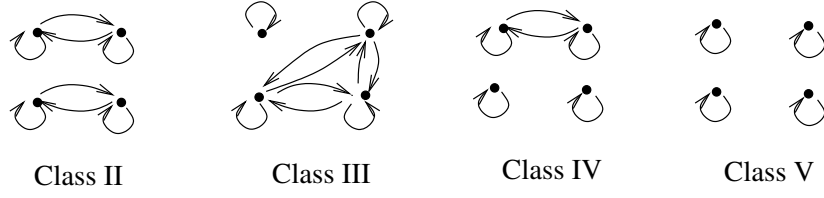
Class II   Class III   Class IV   Class V

Fig. 1. Canonical representatives

Lemma 4. *Let* shE($\mathcal{B}$) *be any of the the DSMs in Class I. Then* $\{\exists, \forall, \wedge, \vee\}$-FO($\mathcal{B}$) *is* Pspace-*complete.*

Proof. Let us begin by introducing the following graph $\mathcal{H}$. Note that shE($\mathcal{H}$) is in fact $\langle \begin{smallmatrix} 0|123 \\ \hline 1|2 \\ \hline 2|1 \\ \hline 3|012 \end{smallmatrix} \rangle \subset \langle \begin{smallmatrix} 0|123 \\ \hline 1|1 \\ \hline 2|2 \\ \hline 3|012 \end{smallmatrix}, \begin{smallmatrix} 0|012 \\ \hline 1|2 \\ \hline 2|1 \\ \hline 3|123 \end{smallmatrix} \rangle$.

We will first prove that $\{\exists, \forall, \wedge, \vee\}$-FO($\mathcal{H}$) is Pspace-complete, by reduction from QCSP($\mathcal{B}_{NAE}$). Define $\Theta(x, v^\top, v^\bot, y)$ to be

$$E(x, v^\top) \wedge E(v^\top, v^\top) \wedge E(v^\top, v^\bot) \wedge E(v^\bot, v^\bot) \wedge E(v^\bot, y).$$

From an instance of QCSP($\mathcal{B}_{NAE}$), we build a prenex instance of $\{\exists, \forall, \wedge, \vee\}$-FO($\mathcal{H}$) as follows. Firstly,

- for each clause $C_j := (v_\alpha \vee v_\beta \vee v_\gamma)$ we introduce to the quantifier prefix $\forall s_j, t_j$, and to the body

$$\Theta(s_j, v_\alpha^\top, v_\alpha^\bot, t_j) \vee \Theta(s_j, v_\beta^\top, v_\beta^\bot, t_j) \vee \Theta(s_j, v_\gamma^\top, v_\gamma^\bot, t_j).$$

Thereafter, outside of the clause quantifiers, and following the prefix order of the instance of QCSP($\mathcal{B}_{NAE}$),

- for each existentially quantified variable v_i, we add $\exists v_i^\top, v_i^\bot \forall x_i, y_i$ to the quantifier prefix and $\Theta(x_i, v_i^\top, v_i^\bot, y_i) \vee \Theta(x_i, v_i^\bot, v_i^\top, y_i)$ to the body.
- For each universally quantified variable v_i, we add $\forall z_i \exists v_i^\top, v_i^\bot \forall x_i, y_i$ to the quantifier prefix and $E(z_i, v_i^\top) \wedge (\Theta(x_i, v_i^\top, v_i^\bot, y_i) \vee \Theta(x_i, v_i^\bot, v_i^\top, y_i))$ to the body.

Suppose we are given a yes-instance QCSP($\mathcal{B}_{NAE}$). No matter how the z_i are evaluated, there is the possibility of placing each $(v_i^\top, v_i^\bot)$ on one of $(1, 2)$ or $(2, 1)$, corresponding to the variable v_i being true or false, respectively. Finally, we check with the clause variables s_j and t_j that the assignment to each clause is not-all-equal. It follows that we have a yes-instance of $\{\exists, \forall, \wedge, \vee\}$-FO($\mathcal{H}$). Conversely, a

yes-instance of $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{H})$ verifies all mappings to the universal variables v_i, through the choice of z_i, since each $(v_i^\top, v_i^\perp)$ is forced to map to one of $(1,2)$ or $(2,1)$. A valuation for the variables of $\mathrm{QCSP}(\mathcal{B}_{NAE})$ may be read from this.

Let Σ be the set of pairs of the edge set of $\mathcal{H}$, i.e. $\{(0,1),(1,0),(1,1),(1,2),$ $(2,1),(2,2),(2,3),(3,2)\}$. Let

$$\Sigma' := \{(0,2),(2,0),(2,2),(2,1),(1,2),(1,1),(1,3),(3,1)\}.$$

Note that the digraph given by Σ' is isomorphic to $\mathcal{H}$. We now introduce the structure $\mathcal{H}'$, on domain $\{0,1,2,3\}$, with a single 4-ary relation R given by

$$
\begin{array}{ll}
(1,3) \times \Sigma & (1,3) \times \Sigma' \\
(2,0) \times \Sigma & (2,3) \times \Sigma' \\
(1,1) \times \Sigma & (1,1) \times \Sigma' \\
(1,2) \times \Sigma & (1,2) \times \Sigma' \\
(2,1) \times \Sigma & (2,1) \times \Sigma' \\
(2,2) \times \Sigma & (2,2) \times \Sigma'
\end{array}
$$

Now, $\mathsf{shE}(\mathcal{H}')$ is $\langle\; \begin{array}{c|c} 0 & 123 \\ \hline 1 & 1 \\ 2 & 2 \\ \hline 3 & 012 \end{array}\;,\; \begin{array}{c|c} 0 & 012 \\ \hline 1 & 2 \\ 2 & 1 \\ \hline 3 & 123 \end{array}\;\rangle$. Further, we can prove $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{H}')$ is Pspace-complete exacly as we did $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{H})$, by adding $\forall q \exists p\; R(p,q,p,p)$ to the sentence and turning all instances of $E(u,v)$ to $R(p,q,u,v)$.

It remains for us to demonstrate that all DSMs that contain neither an A-shop nor an E-shop are sub-DSMs of one of the twenty maximally Pspace-hard DSMs of Classes I-V.

Lemma 5. *The shops f s.t. $\langle f \rangle$ contains neither an A-shop nor an E-shop number 1478 and are exactly those from the twenty DSMs in Classes I-V.*

Proof. Verification by computer. For any shop f, if none of f, f^2, f^4, f^8 is an A-shop or E-shop then check that f is in one of the DSMs $M_1, \ldots, M_{20}$.

It follows that DSMs which contain neither an A-shop nor an E-shop must draw their members exclusively from among $M_1, \ldots, M_{20}$.

Lemma 6. *The maximally Pspace-hard DSMs are precisely those of $M_1, \ldots, M_{20}$.*

Proof. That for no $i, j \in \{1, \ldots, 20\}$ is $M_i \subseteq M_j$ may readily be seen. We will prove that there is no further maximally Pspace-hard DSM.

Let M be a maximally Pspace-hard DSM that is not among $M_1, \ldots, M_{20}$. We will aim for a contradiction which will follow immediately from this claim.

(*) if M contains some k elements other than the identity, then these k elements must all be within one of $M_1, \ldots, M_{20}$.

Proof of claim, by induction. It is clearly true for $k = 1$. Suppose it is true for k. Consider a set of shops $\{f_1, \ldots, f_{k+1}\} \subseteq M$. If $\{f_1, \ldots, f_{k+1}\}$ is not already contained in one of $M_1, \ldots, M_{20}$, then by closure and inductive hypothesis, each

of the $k+1$ sets $S_i := \{f_1, \ldots, f_{k+1}\} \setminus \{f_i\}$ must be contained within a *different* DSM among $M_1, \ldots, M_{20}$. That is, there exist pairwise distinct $a_1, \ldots, a_{k+1}$ s.t. $S_i \subseteq M_{a_i}$. But then we may verify that, either $\langle \{f_1, \ldots, f_{k+1}\} \rangle$ contains an A-shop or an E-shop, or $\langle \{f_1, \ldots, f_{m+1}\} \rangle$ is contained in one of $M_1, \ldots, M_{20}$. The step of verification is accomplished by computer. In the case of $k < 5$, this verification is done by considering all $k + 1$-tuples such that every sub-k-tuple is drawn from a distinct DSM among $M_1, \ldots, M_{20}$. Either the $k+1$-tuple is shown to generate an A-shop or an E-shop, or to be contained in one among $M_1, \ldots, M_{20}$. The key idea is that the number of cases to verify is maintained low by the combinatorial stipulation that the shops come from different DSMs. The cases $k \geq 5$ turn out to be trivial, as there are not k distinct DSMs containing any shop other than the identity.

Proof (Proof of Proposition 1). From Lemmas 6, 3 and 4, in light of Theorem 1.

4 The NP-case

This section will culminate in a proof of the following.

Proposition 2. *If* $\mathsf{shE}(\mathcal{B})$, *where* $|B| = 4$, *contains an* A-*shop but no* E-*shop, then* $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ *is* NP-*complete.*

This will ultimately follow from the classification of all (maximal) DSMs which contain no A-shop. Call a DSM M *maximally* NP-*hard* if M contains no E-shop, but for any shop $f \notin M$, $\langle M \cup \{f\} \rangle$ contains an E-shop. The twenty DSMs of Classes I -V are clearly such; it turns out that there are thirty-four more DSMs, in four further classes.

Let $N_1, \ldots, N_{32}$ enumerate the DSMs of Classes VI – IX.
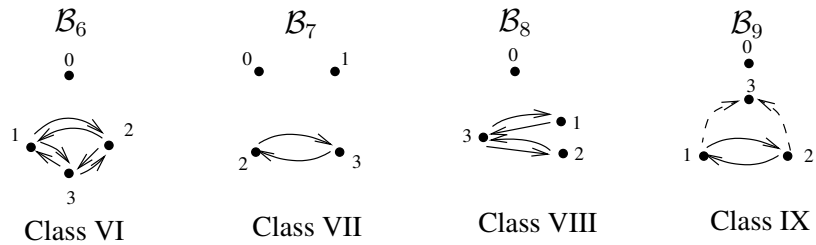

Fig. 2. Canonical representatives. The dotted lines in $\mathcal{B}_9$ are a second binary relation.

Lemma 7. *Let* $\mathsf{shE}(\mathcal{B})$ *be any of the DSMs in Classes VI – IX. Then* $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ *is* NP-*complete.*

Proof. Membership in NP follows from Proposition 2. Consider the structures depicted for each of these classes in Figure 2. When $\mathcal{B}$ is any of these structures we must demonstrate that $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is NP-hard. Let $\mathcal{B}_9^E$ be the restriction of $\mathcal{B}_9$ to the binary relation depicted as a solid line. Our first observation is that $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B}_7) = \{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B}_8) = \{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}\ (\mathcal{B}_9^E) = \{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{K}_2 \uplus \mathcal{K}_1)$, where $\mathcal{K}_2 \uplus \mathcal{K}_1$ is the digraph on vertices $\{0, 1, 2\}$ with edge set $\{(1, 2), (2, 1)\}$. This is because $\mathcal{B}_7$, $\mathcal{B}_8$, $\mathcal{B}_7^E$ and $\mathcal{K}_2 \uplus \mathcal{K}_1$ agree on all equality-free FO sentences. Note that $\mathcal{B}_6 := \mathcal{K}_3 \uplus \mathcal{K}_1$. The problem of evaluating primitive positive ($\{\exists, \wedge\}$-) sentences on $\mathcal{K}_3 \uplus \mathcal{K}_1$ is equivalent to that of evaluating primitive positive sentences on $\mathcal{K}_3$, and this is equivalent to the well-known NP-complete problem of *graph 3-colourability*. It follows that $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B}_7)$ is NP-hard. Similarly, the problem of evaluating existential positive ($\{\exists, \wedge, \vee\}$-) sentences on $\mathcal{K}_2 \uplus \mathcal{K}_1$) is equivalent to that of evaluating existential positive sentences on $\mathcal{K}_2$, and this gives an easy encoding of the NP-complete *not-all-equal 3-colourability* (encode $R_{NAE}(x, y, z) := E(x, y) \vee E(y, z) \vee E(x, z)$ – see [5]). It follows that $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{K}_2 \uplus \mathcal{K}_1))$ is NP-hard.

It remains for us to demonstrate that all DSMs that contain no E-shops are among the twenty plus thirty-two (NP-hard) DSMs of Classes I-IX.

Lemma 8. *The shops f s.t. $\langle f \rangle$ contains no E-shop number* 2678 *and are exactly those from* $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$.

Proof. Verification by computer. For any shop f, if none of f, f^2, f^4, f^8 is an E-shop then check that f is in one of the DSMs $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$.

It follows that DSMs which contain no E-shop must draw their members exclusively from $M_1, \ldots, M_{20}, M'_1, \ldots, M'_{32}$.

Lemma 9. *The* NP-*hard DSMs are precisely those of* $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$.

Proof. That for no distinct $M, M' \in \{M_1, \ldots, M_{20}, N_1, \ldots, N_{32}\}$ is $M \subseteq M'$ may readily be seen. We will prove that there is no further NP-hard DSM.

Let M be a maximally hard DSM that is not among $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$. We will aim for a contradiction which will follow immediately from this claim.

(*) if M contains some k elements other than the identity, then these k elements must all be within one of $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$.

Proof of claim, by induction. It is clearly true for $k = 1$. Suppose it is true for k. Consider a set of shops $\{f_1, \ldots, f_{k+1}\} \subseteq M$. If $\{f_1, \ldots, f_{k+1}\}$ is not already contained in one of $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$, then by closure and inductive hypothesis, each of the $k + 1$ sets $S_i := \{f_1, \ldots, f_{k+1}\} \setminus \{f_i\}$ must be contained within a *different* DSM among $M_1, \ldots, M_{20}$. But then we may verify that, either $\langle \{f_1, \ldots, f_{k+1}\} \rangle$ contains an E-shop, or $\langle \{f_1, \ldots, f_{m+1}\} \rangle$ is contained in one of $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$. The step of verification is accomplished by

computer. In the case of $k < 8$, this verification is done by considering all $k + 1$-tuples such that every sub-k-tuple is drawn from a distinct DSM among $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$. Either the $k+1$-tuple is shown to generate an E-shop, or to be contained in one among $M_1, \ldots, M_{20}, N_1, \ldots, N_{32}$. The cases $k \geq 8$ turn out to be trivial, as there are not k distinct DSMs containing any shop other than the identity.

Proof (Proof of Proposition 2). Membership in NP follows from Theorem 2. NP-completeness follws from Lemmas 9, 7, in light of Theorem 1.

5 The co-NP-case

This case is perfectly dual to the NP-case. We state the results nonetheless.[3]

Proposition 3. *If* $\mathsf{shE}(\mathcal{B})$, *where* $|B| = 4$, *contains an* E-*shop but no* A-*shop, then* $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ *is* co-NP-*complete.*

Membership in co-NP follows from Theorem 2. Completeness will follow from the classification of all (maximal) DSMs which contain no E-shop. Call a DSM M *maximally* co-NP-*hard* if M contains no A-shop, but for any shop $f \notin M$, $\langle M \cup \{f\} \rangle$ contains an A-shop. The twenty DSMs of Classes I -V are clearly such; it turns out that there are thirty-four more DSMs, in four further classes.

$$
\begin{array}{cccc}
\text{Class VI'} & \text{Class VII'} & \text{Class VIII'} & \text{Class IX'} \\[4pt]
\left\langle
\begin{array}{c|c} 0 & 0 \\ 1 & 02 \\ 2 & 01 \\ 3 & 03 \end{array},
\begin{array}{c|c} 0 & 0 \\ 1 & 01 \\ 2 & 03 \\ 3 & 02 \end{array}
\right\rangle
&
\left\langle
\begin{array}{c|c} 0 & 01 \\ 1 & 01 \\ 2 & 013 \\ 3 & 012 \end{array}
\right\rangle
&
\left\langle
\begin{array}{c|c} 0 & 0 \\ 1 & 03 \\ 2 & 03 \\ 3 & 012 \end{array}
\right\rangle
&
\left\langle
\begin{array}{c|c} 0 & 0 \\ 1 & 02 \\ 2 & 01 \\ 3 & 0 \end{array}
\right\rangle
\\[6pt]
+\ 3 \text{ others} & +\ 5 \text{ others} & +\ 11 \text{ others} & +\ 11 \text{ others}
\end{array}
$$

Let $N'_1, \ldots, N'_{32}$ enumerate the DSMs of Classes VI – IX.

Lemma 10. *Let* $\mathsf{shE}(\mathcal{B})$ *be any of the DSMs in Classes VI' – IX'. Then* $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ *is* co-NP-*complete.*

Indeed, for the proof of the lemma, we may use the structures that are the set-theoretic complements of those depicted in Figure 2. It remains for us to demonstrate that all DSMs that contain no A-shops are among the twenty plus thirty-two (NP-hard) DSMs of Classes I-IX.

Lemma 11. *The shops* f *s.t.* $\langle f \rangle$ *contains no* E-*shop number* 2678 *and are exactly those from* $M_1, \ldots, M_{20}, N'_1, \ldots, N'_{32}$.

Lemma 12. *The* co-NP-*hard DSMs are precisely those of* $M_1, \ldots, M_{20}, N'_1, \ldots, N'_{32}$.

---

[3] As we also undertook the respective computations, by way of a checksum.

6 Final remarks

We have now derived the following.

Theorem 4 (Tetrachotomy [5]). *Let $\mathcal{B}$ be s.t. $|B| \leq 4$.*

   I. *If $\mathsf{shE}(\mathcal{A})$ contains both an A-shop and an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in L.*

   II. *If $\mathsf{shE}(\mathcal{A})$ contains an A-shop but no E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in NP-complete.*

   III. *If $\mathsf{shE}(\mathcal{A})$ contains an E-shop but no A-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in $\mathsf{co}\text{-}\mathsf{NP}$-complete.*

   IV. *If $\mathsf{shE}(\mathcal{A})$ contains neither an A-shop nor an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in Pspace-complete.*

Proof. The case $|B| \leq 3$ was taken care of in Theorem 4. The case $|B| = 4$ follows from Propositions 1, 2 and 3.

There are two reasons why our work is a substantial addition to that which appeared in [5]. Firstly, we have encountered cases of a distinctly different nature. This is especially true for the DSMs of Class I, analogs of which do not appear in the case of three-element domains, and whose associated structures needed a fundamentally new proof for Pspace-hardness. Secondly, the search space of DSMs is too large for four-element domains to simply compute it or its significant components. Therefore, an intelligent computational procedure was required to prove that the requisite DSMs were indeed exactly those that are maximal. It is not clear that these methods even extend to five-element domains (at least on modest computers); certainly we can not use these methods to settle the conjecture of [5] here reiterated.

Conjecture 1. For any finite $\mathcal{B}$.

   I. If $\mathsf{shE}(\mathcal{A})$ contains both an A-shop and an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in L.

   II. If $\mathsf{shE}(\mathcal{A})$ contains an A-shop but no E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in NP-complete.

   III. If $\mathsf{shE}(\mathcal{A})$ contains an E-shop but no A-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in $\mathsf{co}\text{-}\mathsf{NP}$-complete.

   IV. If $\mathsf{shE}(\mathcal{A})$ contains neither an A-shop nor an E-shop, then $\{\exists, \forall, \wedge, \vee\}$-$\mathsf{FO}(\mathcal{B})$ is in Pspace-complete.

It seems as if the resolution of this conjecture will require different methods, and perhaps a greater algebraic understanding of the properties of DSMs on arbitrary finite domains.

References

1. Börner, F. Total multifunctions and relations. In *AAA60: Workshop on General Algebra, Dresden, Germany* (2000).
2. Bulatov, A. A. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM 53*, 1 (2006), 66–120.
3. Enderton, H. B. *A Mathematical Introduction to Logic*. Academic Press, 1972.
4. Lynch, N. Log space recognition and translation of parenthesis languages. *J. ACM 24* (1977), 583–590.
5. Madelaine, F., and Martin, B. The complexity of positive first-order logic without equality. *Logic in Computer Science, Symposium on* (2009), 429–438.
6. Madelaine, F., and Martin, B. The complexity of positive first-order logic without equality. http://www.logique.jussieu.fr/∼barnabym/, 2009. Manuscript.
7. Papadimitriou, C. *Computational Complexity*. Addison-Wesley, 1994.
8. Schaefer, T. The complexity of satisfiability problems. In *STOC* (1978).