

The complexity of positive first-order logic without equality

FLORENT MADELAINE

Clermont Universit, Universit d'Auvergne, LIMOS.
CNRS, UMR 6158, LIMOS.

and

BARNABY MARTIN

Durham University.

We study the complexity of evaluating positive equality-free sentences of first-order (FO) logic over a fixed, finite structure \mathcal{B} . This may be seen as a natural generalisation of the non-uniform quantified constraint satisfaction problem $\text{QCSP}(\mathcal{B})$. We introduce surjective hyper-endomorphisms and use them in proving a Galois connection that characterises definability in positive equality-free FO. Through an algebraic method, we derive a complete complexity classification for our problems as \mathcal{B} ranges over structures of size at most three. Specifically, each problem is either in L , is NP-complete, is co-NP-complete or is Pspace-complete.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Computational Logic

General Terms: Languages, Theory

Additional Key Words and Phrases: Quantified Constraints, Equality-free Logics, Galois Connection

1. INTRODUCTION

The evaluation problem under a logic \mathcal{L} – here always a fragment of first-order logic (FO) – takes as input a structure (model) \mathcal{B} and a sentence φ of \mathcal{L} , and asks whether $\mathcal{B} \models \varphi$.¹ When \mathcal{L} is the *primitive positive* fragment of FO, $\{\exists, \wedge\}$ -FO, the evaluation problem is equivalent to the much-studied *constraint satisfaction problem* (CSP). Similarly, when \mathcal{L} is the *positive Horn* fragment of FO, $\{\exists, \forall, \wedge\}$ -FO, the evaluation problem is equivalent to the well-studied *quantified constraint satis-*

¹We resist the better-established terminology of ‘model checking problem’ because in the majority of this paper we consider the structure \mathcal{B} to be fixed.

Author’s addresses: Florent Madelaine, Dpartement Informatique, Institut Universitaire de Technologie, Campus des Czeaux B.P. 86, 63172 AUBIERE cedex, FRANCE. florent.madelaine@u-clermont1.fr. Barnaby Martin, School of Engineering and Computing Sciences, Durham University, Durham DH1 3LE, U.K. barnabymartin@gmail.com. Second author supported by EPSRC grant EP/G020604/1.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1529-3785/20YY/0700-0111 \$5.00

fraction problem (QCSP). In this manner, the QCSP is the generalisation of the CSP in which universal quantification is restored to the mix. In both cases it is essentially irrelevant whether or not equality is permitted in the sentences, as it may be propagated out by substitution. Much work has been done on the parameterisation of these problems by the structure \mathcal{B} – that is, where \mathcal{B} is fixed and only the sentence is input. It is conjectured [8] that the ensuing problems $\text{CSP}(\mathcal{B})$ attain only the complexities P and NP-complete. This may appear surprising given that 1.) so many natural NP problems may be expressed as CSPs (see, e.g., myriad examples in [10]) and 2.) NP itself does not have this ‘dichotomy’ property (assuming $\text{P} \neq \text{NP}$) [12]. While this *dichotomy conjecture* remains open, it has been proved for certain classes of \mathcal{B} (e.g., for structures of size at most three [4] and for undirected graphs [9]). The like parameterisation of the QCSP is also well-studied, and while no overarching polychotomy has been conjectured, only the complexities P, NP-complete and Pspace-complete are known to be attainable (for trichotomy results on certain classes see [3; 20], as well as the dichotomy for Boolean structures, e.g., in [6]).

In previous work, [17; 16], we have studied the evaluation problem, parameterised by the structure, under various fragments of FO obtained by restrictions on which of the symbols of $\{\exists, \forall, \wedge, \vee, \neg, =, \neq\}$ is permitted. Of course, many of the ostensibly 2^7 such fragments may be discarded as totally trivial or as repetitions through de Morgan duality. There are four fragments each equivalent to the CSP and QCSP: these are $\{\exists, \wedge\}$ -FO, $\{\exists, \wedge, =\}$ -FO, $\{\forall, \vee\}$ -FO, $\{\forall, \vee, \neq\}$ -FO and $\{\exists, \forall, \wedge\}$ -FO, $\{\exists, \forall, \wedge, =\}$ -FO, $\{\exists, \forall, \vee\}$ -FO, $\{\exists, \forall, \vee, \neq\}$ -FO, respectively. Here, equivalent means that a complexity classification for one yields a complexity classification for the other; but, the complexity classes need not be the same. For example, the class of problems given by fixing the structure under $\{\exists, \wedge\}$ -FO would display dichotomy between P and NP-complete iff the like class of problems under $\{\forall, \vee\}$ -FO displays dichotomy between P and co-NP-complete. Various complexity classifications are obtained in [17; 16] and it is observed that the only interesting fragment, other than the eight associated with CSP and QCSP, is $\{\exists, \forall, \wedge, \vee\}$ -FO.² The evaluation problem over $\{\exists, \forall, \wedge, \vee\}$ -FO may be seen as the generalisation of the QCSP in which disjunction is returned to the mix. Note that the absence of equality is here important, as there is no general method for its being propagated out by substitution. Indeed, we will see that evaluating the related fragment $\{\exists, \forall, \wedge, \vee, =\}$ -FO is Pspace-complete on any structure \mathcal{B} of size at least two.

In this paper we initiate a study of the evaluation problem for the fragment $\{\exists, \forall, \wedge, \vee\}$ -FO over a fixed relational \mathcal{B} – the problem we denote $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}). We demonstrate at least that this class displays a complexity-theoretic richness absent from those other fragments that are not associated with the CSP or QCSP.

²For many of the other fragments the complexity classification is nearly trivial. For example, this is true for $\{\exists, \wedge, \vee\}$ -FO, $\{\forall, \wedge, \vee\}$ -FO and $\{\exists, \forall, \wedge, \vee, \neg\}$ -FO (also for these classes with $=$ or \neq). For others the classification may be read through the Schaefer classification for Boolean CSP and QCSP, because computational hardness is clear over fixed structures of size at least three. For example, this is the case for $\{\exists, \wedge, \neq\}$ -FO, $\{\forall, \vee, =\}$ -FO and $\{\exists, \forall, \wedge, \neq\}$ -FO, $\{\exists, \forall, \vee, =\}$ -FO. Note that the consideration of \neq is not explicit in [17; 16]. Similarly, fragments involving both quantifiers and $=$ or \neq are not explicitly considered. In both cases, the results may be read off from de Morgan duality together with standard Schaefer class results (for which we refer to [6]).

It is possibly to be hoped, however, that a full classification for this class is not as resistant as that for the CSP or QCSP. We undertake our study through the algebraic method that has been so fruitful in the study of the CSP and QCSP (see [11; 4; 3; 5]). To this end, we define surjective hyper-endomorphisms and use them to define a new Galois connection that characterises definability under $\{\exists, \forall, \wedge, \vee\}$ -FO. While this Galois connection appears here for the first time, it does follow a general recipe as outlined, e.g., in [2]. It is not clear that the many different Galois connections associated with fragments of FO can be proved in a straightforwardly uniform manner. We are able to prove a complete complexity classification for $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) when \mathcal{B} ranges over structures of size at most three. On the class of Boolean structures we see dichotomy between L and Pspace-complete. On the class of structures of size three we see tetrachotomy between L, NP-complete, co-NP-complete and Pspace-complete. Some of the results that appear in this paper had been obtained through adhoc methods in [18] – although there the tetrachotomy extends only to digraphs and not arbitrary relational structures. Also, little insight was provided as to the underlying properties of the classification. It is a pleasing consequence of our algebraic approach that we can give quite simple explanation to the delineation of our subclasses.

The paper is organised as follows. In Section 2, we introduce the preliminaries, including the relevant Galois connection together with the central notions of surjective hyper-operation (shop), surjective hyper-endomorphism (she) and downshop-monoid. In Section 3, we outline conditions under which $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) either drops from or attains maximal complexity. In Section 4, we classify the complexity of the problems $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}), when \mathcal{B} ranges over, firstly, Boolean structures and, secondly, structures of size three. In the first instance a dichotomy – between L and Pspace-complete – is obtained; in the second instance a tetrachotomy – between L, NP-complete, co-NP-complete and Pspace-complete – is obtained. We conclude, in Section 5, with some final remarks.

An extended abstract of this paper has appeared as [14].

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 and $\{\exists, \forall, \wedge, \vee, =\}$ -FO be the positive fragments of first-order (FO) logic, without and with equality, respectively. 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 $\mathcal{B} \models \varphi$?

The related problem $\{\exists, \forall, \wedge, \vee, =\}$ -FO(\mathcal{B}) permits sentences φ that may involve equalities, in the obvious way. 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 [13]). In this case, it follows that both $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) and $\{\exists, \forall, \wedge, \vee, =\}$ -FO(\mathcal{B}) are also in L.

Consider the set B and its power set $\mathfrak{P}(B)$. A *hyper-operation* on B is a function f from 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 call (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 surjective hyper-operation f on B that satisfies, for all extensional relations R of \mathcal{B} ,

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

More generally, for $r_1, \dots, r_k \in B$, we say f is a *she from* $(\mathcal{B}, r_1, \dots, r_k)$ to $(\mathcal{B}, r'_1, \dots, r'_k)$ if f is a she of \mathcal{B} and $r'_1 \in f(r_1), \dots, r'_k \in f(r_k)$. 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.

For $b_1, \dots, b_{|B|}$ an enumeration of the elements of \mathcal{B} , let the quantifier-free formula $\Phi_{\mathcal{B}}(v_1, \dots, v_{|B|})$ be a conjunction of the positive facts of \mathcal{B} , where the variables $v_1, \dots, v_{|B|}$ correspond to the elements $b_1, \dots, b_{|B|}$. That is, for R an extensional relation of \mathcal{B} , $R(v_{\lambda_1}, \dots, v_{\lambda_i})$ appears as an atom in $\Phi_{\mathcal{B}}$ iff $\mathcal{B} \models R(b_{\lambda_1}, \dots, b_{\lambda_i})$. For example, let \mathcal{K}_3 be the antireflexive 3-clique, that is the structure with domain $\{0, 1, 2\}$ and single binary relation

$$E := \{(0, 1), (1, 0), (1, 2), (2, 1), (2, 0), (0, 2)\}.$$

Then

$$\Phi_{\mathcal{K}_3}(v_0, v_1, v_2) := E(v_0, v_1) \wedge E(v_1, v_0) \wedge E(v_1, v_2) \wedge E(v_2, v_1) \wedge E(v_2, v_0) \wedge E(v_0, v_2).$$

The existential sentence $\exists v_1, \dots, v_{|B|} \Phi_{\mathcal{B}}(v_1, \dots, v_{|B|})$ is known as the *canonical query* of \mathcal{B} . More generally, for a (not necessarily distinct) l -tuple of elements $\mathbf{r} := (r_1, \dots, r_l) \in B^l$, define the quantifier-free $\Phi_{\mathcal{B}(\mathbf{r})}(v_1, \dots, v_l)$ to be the conjunction of the positive facts of \mathbf{r} , where the variables v_1, \dots, v_l correspond to the elements r_1, \dots, r_l . That is, $R(v_{\lambda_1}, \dots, v_{\lambda_i})$ appears as an atom in $\Phi_{\mathcal{B}(\mathbf{r})}$ iff $\mathcal{B} \models R(r_{\lambda_1}, \dots, r_{\lambda_i})$. For example,

$$\Phi_{\mathcal{K}_3(0,0,2)}(v_0, v_1, v_2) := E(v_0, v_2) \wedge E(v_2, v_0) \wedge E(v_1, v_2) \wedge E(v_2, v_1).$$

We refer to elements in \mathcal{B} as r, s, t (also x, y), or $b_1, \dots, b_{|B|}$ when this is an enumeration. We reserve u, v, w to refer to variables in FO formulae.

2.1 Galois Connections

For a set F of shops on the finite domain B , let $\text{Inv}(F)$ 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 $S \in \text{Inv}(F)$ is *invariant* or *preserved* by (the shops in) F . Let $\text{shE}(\mathcal{B})$ be the set of shes of \mathcal{B} . Let $\text{Aut}(\mathcal{B})$ be the set of automorphisms of \mathcal{B} .

Let $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}}$ and $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee, =\}\text{-FO}}$ be the sets of relations that may be defined on \mathcal{B} in $\{\exists, \forall, \wedge, \vee\}$ -FO and $\{\exists, \forall, \wedge, \vee, =\}$ -FO, respectively.

LEMMA 2.1. *Let $\mathbf{r} := (r_1, \dots, r_k)$ be a k -tuple of elements of \mathcal{B} . There exists:*

- (i). a formula $\theta_{\mathbf{r}}(u_1, \dots, u_k) \in \{\exists, \forall, \wedge, \vee, =\}$ -FO s.t. $(\mathcal{B}, r'_1, \dots, r'_k) \models \theta_{\mathbf{r}}(u_1, \dots, u_k)$ iff there is an automorphism from $(\mathcal{B}, r_1, \dots, r_k)$ to $(\mathcal{B}, r'_1, \dots, r'_k)$.
- (ii). a formula $\theta_{\mathbf{r}}(u_1, \dots, u_k) \in \{\exists, \forall, \wedge, \vee\}$ -FO s.t. $(\mathcal{B}, r'_1, \dots, r'_k) \models \theta_{\mathbf{r}}(u_1, \dots, u_k)$ iff there is a she from $(\mathcal{B}, r_1, \dots, r_k)$ to $(\mathcal{B}, r'_1, \dots, r'_k)$.

PROOF. For Part (i), let $b_1, \dots, b_{|B|}$ an enumeration of the elements of \mathcal{B} and $\Phi_{\mathcal{B}}(v_1, \dots, v_{|B|})$ be the associated conjunction of positive facts. Set $\theta_{\mathbf{r}}(u_1, \dots, u_k) := \exists v_1, \dots, v_{|B|} \Phi_{\mathcal{B}}(v_1, \dots, v_{|B|}) \wedge \forall v (v = v_1 \vee \dots \vee v = v_{|B|}) \wedge u_1 = v_{\lambda_1} \wedge \dots \wedge u_k = v_{\lambda_k}$, where $r_1 = b_{\lambda_1}, \dots, r_k = b_{\lambda_k}$. The forward direction follows since \mathcal{B} is finite, so any surjective endomorphism is necessarily an automorphism. The backward direction follows since all first-order formulae are preserved by automorphism.

[Part (ii).] This will require greater dexterity. Let $\mathbf{r} \in B^k$, $\mathbf{s} := (b_1, \dots, b_{|B|})$ be an enumeration of B and $\mathbf{t} \in B^{|B|}$. Recall that $\Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s})}(u_1, \dots, u_k, v_1, \dots, v_{|B|})$ is a conjunction of the positive facts of (\mathbf{r}, \mathbf{s}) , where the variables (\mathbf{u}, \mathbf{v}) correspond to the elements (\mathbf{r}, \mathbf{s}) . Similarly, $\Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(u_1, \dots, u_k, v_1, \dots, v_{|B|}, w_1, \dots, w_{|B|})$ is the conjunction of the positive facts of $(\mathbf{r}, \mathbf{s}, \mathbf{t})$, where the variables $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ correspond to the elements $(\mathbf{r}, \mathbf{s}, \mathbf{t})$. Set $\theta_{\mathbf{r}}(u_1, \dots, u_k) :=$

$$\exists v_1, \dots, v_{|B|} \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s})}(u_1, \dots, u_k, v_1, \dots, v_{|B|}) \wedge \forall w_1 \dots w_{|B|} \bigvee_{\mathbf{t} \in B^{|B|}} \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(u_1, \dots, u_k, v_1, \dots, v_{|B|}, w_1, \dots, w_{|B|}).$$

[Part (ii), backwards.] Suppose f is a she from $(\mathcal{B}, r_1, \dots, r_k)$ to $(\mathcal{B}', r'_1, \dots, r'_k)$, where $\mathcal{B}' := \mathcal{B}$ (we will wish to differentiate the two occurrences of \mathcal{B}). We aim to prove that $\mathcal{B}' \models \theta_{\mathbf{r}}(r'_1, \dots, r'_k)$. Choose arbitrary $s'_1 \in f(b_1), \dots, s'_{|B|} \in f(b_{|B|})$ as witnesses for $v_1, \dots, v_{|B|}$. Let $\mathbf{t}' := (t'_1, \dots, t'_{|B|}) \in B'^{|B|}$ be any valuation of $w_1, \dots, w_{|B|}$ and take arbitrary $t_1, \dots, t_{|B|}$ s.t. $t'_1 \in f(t_1), \dots, t'_{|B|} \in f(t_{|B|})$ (here we use surjectivity). Let $\mathbf{t} := (t_1, \dots, t_{|B|})$. It follows from the definition of she that

$$\mathcal{B}' \models \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s})}(r'_1, \dots, r'_k, s'_1, \dots, s'_{|B|}) \wedge \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(r'_1, \dots, r'_k, s'_1, \dots, s'_{|B|}, t'_1, \dots, t'_{|B|}).$$

[Part (ii), forwards.] Assume that $\mathcal{B}' \models \theta_{\mathbf{r}}(r'_1, \dots, r'_k)$, where $\mathcal{B}' := \mathcal{B}$. Let $b'_1, \dots, b'_{|B|}$ be an enumeration of $\mathcal{B}' := B$.³ Choose some witness elements $s'_1, \dots, s'_{|B|}$ for $v_1, \dots, v_{|B|}$ and a witness tuple $\mathbf{t} := (t_1, \dots, t_{|B|}) \in B^{|B|}$ s.t.

$$(\dagger) \mathcal{B}' \models \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s})}(r'_1, \dots, r'_k, s'_1, \dots, s'_{|B|}) \wedge \Phi_{\mathcal{B}(\mathbf{r}, \mathbf{s}, \mathbf{t})}(r'_1, \dots, r'_k, s'_1, \dots, s'_{|B|}, b'_1, \dots, b'_{|B|}).$$

Consider the following partial hyper-operations from $B \rightarrow \mathfrak{P}(B') \setminus \{\emptyset\}$.

1. $f_{\mathbf{r}}$ given by $f_{\mathbf{r}}(r_i) := \{r'_i\}$, for $1 \leq i \leq k$.
2. $f_{\mathbf{s}}$ given by $f_{\mathbf{s}}(b_i) = \{s'_i\}$, for $1 \leq i \leq |B|$. (Totality.)
3. $f_{\mathbf{t}}$ given by $b'_i \in f_{\mathbf{t}}(b_j)$ iff $t_i = b_j$, for $1 \leq i, j \leq |B|$. (Surjectivity.)

³One may imagine $b_1, \dots, b_{|B|}$ and $b'_1, \dots, b'_{|B|}$ to be the same enumeration, but this is not essential. In any case, we will wish to keep the dashes on the latter set to remind us they are in \mathcal{B}' and not \mathcal{B} .

Let $f := f_{\mathbf{r}} \cup f_{\mathbf{s}} \cup f_{\mathbf{t}}$; f is a hyper-operation whose surjectivity is guaranteed by $f_{\mathbf{t}}$ (note that totality is guaranteed by $f_{\mathbf{s}}$). That f is a she follows from the right-hand conjunct of (†). \square

THEOREM 2.2. *For a finite structure \mathcal{B} we have*

- (i). $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee, =\}\text{-FO}} = \text{Inv}(\text{Aut}(\mathcal{B}))$ and
- (ii). $\langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}} = \text{Inv}(\text{shE}(\mathcal{B}))$.

PROOF. Part (i) is well-known and may be proved in a similar, albeit simpler, manner to Part (ii), which we now prove.

[$\varphi(\mathbf{v}) \in \langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}} \Rightarrow \varphi(\mathbf{v}) \in \text{Inv}(\text{shE}(\mathcal{B}))$.] This is proved by induction on the complexity of $\varphi(\mathbf{v})$.

(Base Case.) $\varphi(\mathbf{v}) := R(\mathbf{v})$.⁴ Follows from the definition of she.

(Inductive Step.) There are four subcases. We progress through them in a workmanlike fashion. Take $f \in \text{shE}(\mathcal{B})$.

$\varphi(\mathbf{v}) := \psi(\mathbf{v}) \wedge \psi'(\mathbf{v})$.⁵ Let $\mathbf{v} := (v_1, \dots, v_l)$. Suppose $\mathcal{B} \models \varphi(x_1, \dots, x_l)$; then both $\mathcal{B} \models \psi(x_1, \dots, x_l)$ and $\mathcal{B} \models \psi'(x_1, \dots, x_l)$. By Inductive Hypothesis (IH), for any $y_1 \in f(x_1), \dots, y_l \in f(x_l)$, both $\mathcal{B} \models \psi(y_1, \dots, y_l)$ and $\mathcal{B} \models \psi'(y_1, \dots, y_l)$, whence $\mathcal{B} \models \varphi(y_1, \dots, y_l)$.

$\varphi(\mathbf{v}) := \psi(\mathbf{v}) \vee \psi'(\mathbf{v})$. Let $\mathbf{v} := (v_1, \dots, v_l)$. Suppose $\mathcal{B} \models \varphi(x_1, \dots, x_l)$; then one of $\mathcal{B} \models \psi(x_1, \dots, x_l)$ or $\mathcal{B} \models \psi'(x_1, \dots, x_l)$; w.l.o.g. the former. By IH, for any $y_1 \in f(x_1), \dots, y_l \in f(x_l)$, $\mathcal{B} \models \psi(y_1, \dots, y_l)$, whence $\mathcal{B} \models \varphi(y_1, \dots, y_l)$.

$\varphi(\mathbf{v}) := \forall w \psi(\mathbf{v}, w)$. Let $\mathbf{v} := (v_1, \dots, v_l)$. Suppose $\mathcal{B} \models \forall w \psi(x_1, \dots, x_l, w)$; then for each x' , $\mathcal{B} \models \psi(x_1, \dots, x_l, x')$. By IH, for any $y_1 \in f(x_1), \dots, y_l \in f(x_l)$, we have for all y' (remember f is surjective), $\mathcal{B} \models \psi(y_1, \dots, y_l, y')$, whereupon $\mathcal{B} \models \forall w \psi(y_1, \dots, y_l, w)$.

$\varphi(\mathbf{v}) := \exists w \psi(\mathbf{v}, w)$. Let $\mathbf{v} := (v_1, \dots, v_l)$. Suppose $\mathcal{B} \models \exists w \psi(x_1, \dots, x_l, w)$; then for some x' , $\mathcal{B} \models \psi(x_1, \dots, x_l, x')$. By IH, for any $y_1 \in f(x_1), \dots, y_l \in f(x_l)$, $y' \in f(x')$ (remember $f(x')$ can not be empty), $\mathcal{B} \models \psi(y_1, \dots, y_l, y')$, whereupon $\mathcal{B} \models \exists w \psi(y_1, \dots, y_l, w)$.

[$S \in \text{Inv}(\text{shE}(\mathcal{B})) \Rightarrow S \in \langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}}$.] Consider the k -ary relation $S \in \text{Inv}(\text{shE}(\mathcal{B}))$. Let $\mathbf{r}_1, \dots, \mathbf{r}_m$ be the tuples of S . Set

$$\theta_S(u_1, \dots, u_k) := \theta_{\mathbf{r}_1}(u_1, \dots, u_k) \vee \dots \vee \theta_{\mathbf{r}_m}(u_1, \dots, u_k).$$

Obviously, $\theta_S(u_1, \dots, u_k) \in \{\exists, \forall, \wedge, \vee\}\text{-FO}$. For $\mathbf{r}_i := (r_{i1}, \dots, r_{ik})$, note that $(\mathcal{B}, r_{i1}, \dots, r_{ik}) \models \theta_{\mathbf{r}_i}(u_1, \dots, u_k)$ (the ‘identity’ she will be formally introduced in the next section). That $\theta_S(u_1, \dots, u_k) = S$ now follows from Part (ii) of Lemma 2.1, since $S \in \text{Inv}(\text{shE}(\mathcal{B}))$. \square

Let \leq_{L} indicate the existence of a logspace many-to-one reduction. The following theorem is our counterpart to Corollary 4.11 of [10] (for CSP) and Theorem 3.1 of [3] (for QCSP).

THEOREM 2.3. *Let \mathcal{B} and \mathcal{B}' be finite structures over the same domain B .*

⁴The variables \mathbf{v} may appear multiply in R and in any order. Thus R is an instance of an extensional relation under substitution and permutation of positions.

⁵The presence of, e.g., \mathbf{v} in $\psi(\mathbf{v}) \wedge \psi'(\mathbf{v})$ should not be taken as indication that *all* \mathbf{v} appear free in both ψ and ψ' .

- (i). If $\text{Aut}(\mathcal{B}) \subseteq \text{Aut}(\mathcal{B}')$ then $\{\exists, \forall, \wedge, \vee, =\}\text{-FO}(\mathcal{B}') \leq_L \{\exists, \forall, \wedge, \vee, =\}\text{-FO}(\mathcal{B})$.
- (ii). If $\text{shE}(\mathcal{B}) \subseteq \text{shE}(\mathcal{B}')$ then $\{\exists, \forall, \wedge, \vee\}\text{-FO}(\mathcal{B}') \leq_L \{\exists, \forall, \wedge, \vee\}\text{-FO}(\mathcal{B})$.

PROOF. Again, Part (i) is well-known and the proof is similar to that of Part (ii), which we give. If $\text{shE}(\mathcal{B}) \subseteq \text{shE}(\mathcal{B}')$, then $\text{Inv}(\text{shE}(\mathcal{B}')) \subseteq \text{Inv}(\text{shE}(\mathcal{B}))$. From Theorem 2.2, it follows that $\langle \mathcal{B}' \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}} \subseteq \langle \mathcal{B} \rangle_{\{\exists, \forall, \wedge, \vee\}\text{-FO}}$. Recalling that \mathcal{B}' contains only a finite number of extensional relations, we may therefore effect a Logspace reduction from $\{\exists, \forall, \wedge, \vee\}\text{-FO}(\mathcal{B}')$ to $\{\exists, \forall, \wedge, \vee\}\text{-FO}(\mathcal{B})$ by straightforward substitution of predicates. \square

2.2 Down-shop-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 surjective shops on a finite set B is a *down-shop-monoid* (DSM), if it contains id_B , and is closed under composition and sub-shops (of course, not all sub-hyperoperations of a shop 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 $\text{shE}(\mathcal{B})$ is always a DSM. The DSMs over B form a lattice under (set-theoretic) inclusion and, as per the Galois connection of the previous section, classify the complexities of $\{\exists, \forall, \wedge, \vee\}\text{-FO}(\mathcal{B})$. 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 $\langle \{f\} \rangle$

For a shop f , define its inverse f^{-1} by $x \mapsto \{y : x \in f(y)\}$. Note that f^{-1} is also a shop and $(f^{-1})^{-1} = f$, though $f \circ f^{-1} = id_B$ only if f is a permutation. For a set of shops F , let $F^{-1} := \{f^{-1} : f \in F\}$. If F is a DSM then so is F^{-1} . We will see this algebraic duality resonates with the de Morgan duality of \exists and \forall , and the complexity-theoretic duality of NP and co-NP.⁶

A *permutation subgroup* on a finite set B is a set of permutations of B closed under composition. It may easily be verified that such a set contains the identity and is closed under inverse. A permutation subgroup may be identified with a particular type of DSM in which all shops have only singleton sets in their range. The permutation subgroups form a lattice under inclusion whose minimal element contains just the identity and whose maximal element is the symmetric group $S_{|B|}$. As per the Galois connection of the previous section, this lattice classifies the complexities of $\{\exists, \forall, \wedge, \vee, =\}\text{-FO}(\mathcal{B})$ – although we shall see this complexity classification is relatively uninteresting.

In the lattice of DSMs, the minimal element still contains just id_B , but the maximal element contains all shops. However, the lattice of permutation subgroups always appears as a sub-lattice within the lattice of DSMs.

⁶We resist discussing this further as it plays no direct role in the derivation of our results. Since this paper was submitted, the duality born of inverse has been studied further in [19].

3. CLASSIFICATION METHODS

We are now in a position to study the interplay between the shes of a structure \mathcal{B} and the complexity of the problem $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}).

3.1 Shes inducing lower complexity

We begin by studying three classes of shes, the presence of any of which reduces the complexity of the problem $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}). Let \mathcal{B} be a finite structure, with distinct elements b, b' . We define the following shops from B to $\mathfrak{P}(B) \setminus \{\emptyset\}$.

$$\forall_b(x) := \begin{cases} B & \text{if } x = b \\ \{x\} & \text{otherwise.} \end{cases}$$

$$\exists_b(x) := \{x, b\}$$

$$\forall_b \exists_{b'}(x) := \begin{cases} B & \text{if } x = b \\ \{b'\} & \text{otherwise.} \end{cases}$$

We call their classes \forall -, \exists - and $\forall\exists$ -shops, respectively. In Figure 1, four digraphs \mathcal{G}_1 –




Fig. 1. Sample digraphs admitting \forall -, \exists - and $\forall\exists$ -hyper-operations as shes.

\mathcal{G}_4 are drawn. For typographic reasons we will mark-up, e.g., the shop $0 \mapsto \{0, 1\}$, $1 \mapsto \{1\}$ and $2 \mapsto \{1, 2\}$ as $\frac{0|01}{1|1}$. It may easily be verified that the DSMs $\text{shE}(\mathcal{G}_1)$ – $\text{shE}(\mathcal{G}_4)$ are as follows.

$$\begin{array}{cccc} \text{shE}(\mathcal{G}_1) & \text{shE}(\mathcal{G}_2) & \text{shE}(\mathcal{G}_3) & \text{shE}(\mathcal{G}_4) \\ \langle \frac{0|01}{1|1} \rangle & \langle \frac{0|0}{1|012} \rangle & \langle \frac{0|2}{1|012}, \frac{0|0}{1|01} \rangle & \langle \frac{0|012}{1|1} \rangle \\ & & \langle \frac{0|0}{2|2} \rangle & \langle \frac{0|012}{2|012} \rangle \end{array}$$

We see that \mathcal{G}_1 , \mathcal{G}_2 and \mathcal{G}_3 admit the shes \exists_1 , \forall_1 and $\forall_1 \exists_2$, respectively. \mathcal{G}_4 admits each of the shes \forall_0 , \forall_2 , \exists_1 , $\forall_0 \exists_1$ and $\forall_2 \exists_1$.

Remarks 3.1. We have not considered shes $\forall_b \exists_b$, defined as above but with $b' := b$. The DSM $\langle \forall_b \exists_b \rangle$ is easily seen to contain all shops. It follows that any structure \mathcal{B} that has $\forall_b \exists_b$ as a she already has all shes of the form $\forall_{b'} \exists_{b''}$ with $b' \neq b''$.

Note that the DSMs $\langle \forall_b \exists_{b'} \rangle$ and $\langle \{\forall_b, \exists_{b'}\} \rangle = \langle \forall_b \circ \exists_{b'} \rangle = \langle \exists_{b'} \circ \forall_b \rangle$ do not in general coincide, though the first is always a subset of the following three. Also, we note the identities $\exists_b^{-1} = \forall_b$, $\forall_b^{-1} = \exists_b$ and $(\forall_b \exists_{b'})^{-1} = \forall_{b'} \exists_b$.

We now give a series of three lemmata, one associated with each of the shops \forall_b , \exists_b and $\forall_b \exists_{b'}$. They will ultimately be used in a form of quantifier elimination that will diminish the complexity of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}), if \mathcal{B} has one of these as a she.

LEMMA 3.2. *Let $\varphi(u, \mathbf{v})$ be a formula of $\{\exists, \forall, \wedge, \vee\}$ -FO. Let \mathcal{B} be a finite structure with \forall_b as a she. Then*

$$\mathcal{B} \models \forall u \varphi(u, \mathbf{v}) \iff \mathcal{B} \models \varphi(b, \mathbf{v}).$$

PROOF. The forward direction is trivial; we prove the backward. Consider the relation defined by the formula $\varphi(u, \mathbf{v})$, where $\mathbf{v} := (v_1, \dots, v_k)$, of $\{\exists, \forall, \wedge, \vee\}$ -FO on \mathcal{B} . By Theorem 2.2, it is invariant under $\forall_b \in \text{shE}(\mathcal{B})$. For any $x_1, \dots, x_k \in B$, assume $\mathcal{B} \models \varphi(b, x_1, \dots, x_k)$. Taking an arbitrary $c \in B$, and noting each $x_i \in \forall_b(x_i)$ and $c \in \forall_b(b)$, we derive $\mathcal{B} \models \varphi(c, x_1, \dots, x_k)$. The result follows. \square

LEMMA 3.3. *Let $\varphi(u, \mathbf{v})$ be a formula of $\{\exists, \forall, \wedge, \vee\}$ -FO. Let \mathcal{B} be a finite structure with \exists_b as a she. Then*

$$\mathcal{B} \models \exists u \varphi(u, \mathbf{v}) \iff \mathcal{B} \models \varphi(b, \mathbf{v}).$$

PROOF. The backward direction is trivial; we prove the forward. Consider the relation defined by the formula $\varphi(u, \mathbf{v})$, where $\mathbf{v} := (v_1, \dots, v_k)$, of $\{\exists, \forall, \wedge, \vee\}$ -FO on \mathcal{B} . By Theorem 2.2, it is invariant under $\exists_b \in \text{shE}(\mathcal{B})$. For any $x_1, \dots, x_k \in B$, and some $c \in B$, assume $\mathcal{B} \models \varphi(c, x_1, \dots, x_k)$. Noting each $x_i \in \exists_b(x_i)$ and $b \in \exists_b(c)$, we derive $\mathcal{B} \models \varphi(b, x_1, \dots, x_k)$. The result follows. \square

LEMMA 3.4. *Let $\varphi(u, \mathbf{v})$ be a formula of $\{\exists, \forall, \wedge, \vee\}$ -FO, where the arity of \mathbf{v} is k . Let \mathcal{B} be a finite structure with $\forall_b \exists_{b'}$ as a she. For all $c \in B$ and $\mathbf{x} := (x_1, \dots, x_k) \in \{b, b'\}^k$,*

$$\mathcal{B} \models \varphi(b, \mathbf{x}) \xrightarrow{(I)} \mathcal{B} \models \varphi(c, \mathbf{x}) \xrightarrow{(II)} \mathcal{B} \models \varphi(b', \mathbf{x}).$$

PROOF. Consider the relation defined by the formula $\varphi(u, \mathbf{v})$ of $\{\exists, \forall, \wedge, \vee\}$ -FO on \mathcal{B} . By Theorem 2.2, it is invariant under $\forall_b \exists_{b'} \in \text{shE}(\mathcal{B})$. Take arbitrary $c \in B$. Noting that x_i is from $\{b, b'\}$, we have $x_i \in \forall_b \exists_{b'}(x_i)$ and $c \in \forall_b \exists_{b'}(b)$. Part (I) follows. Now noting that $b' \in \forall_b \exists_{b'}(c)$, Part (II) follows. \square

We are now ready to state how the presence of \forall -, \exists - or $\forall\exists$ -shops as shes of \mathcal{B} can diminish the complexity of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}). In each case we proceed by quantifier elimination.

THEOREM 3.5. *If \mathcal{B} has a \forall -shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in NP. If \mathcal{B} has an \exists -shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in co-NP. If \mathcal{B} has a $\forall\exists$ -shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in L.*

PROOF. Let φ be a sentence of $\{\exists, \forall, \wedge, \vee\}$ -FO, and let $\varphi_{[\forall/b]}$ (resp., $\varphi_{[\exists/b]}$ and $\varphi_{[\forall/b, \exists/b']}$) be φ with all universal variables substituted by b (resp., existential variables substituted by b and universal variables substituted by b and existential variables substituted by b').

If \mathcal{B} has a she \forall_b , then consider a sentence $\varphi \in \{\exists, \forall, \wedge, \vee\}$ -FO, w.l.o.g. in prenex form. It follows by repeated application of Lemma 3.2 on φ – either from the outermost quantifier in, or from the innermost quantifier out – that $\mathcal{B} \models \varphi$ iff $\mathcal{B} \models \varphi_{[\forall/b]}$. Similarly, if \mathcal{B} has a she $\exists_{b'}$, then it follows by repeated application of Lemma 3.3 that $\mathcal{B} \models \varphi$ iff $\mathcal{B} \models \varphi_{[\exists/b']}$.

If \mathcal{B} has a she $\forall_b \exists_{b'}$, then, again, assume the sentence $\varphi \in \{\exists, \forall, \wedge, \vee\}$ -FO to be in prenex form. It follows by repeated application of Lemma 3.4 – from the outermost

quantifier in $-$ that $\mathcal{B} \models \varphi$ iff $\mathcal{B} \models \varphi_{[\forall/b, \exists/b']}$. Note that, in this case, one can not move from the innermost quantifier out because this may involve the possibility of free variables taking values from outside the set $\{b, b'\}$. The result now follows since evaluating $\varphi_{[\forall/b, \exists/b']}$ on \mathcal{B} is equivalent to a Boolean sentence value problem, known to be in L [13]. \square

Returning to the examples of Figure 1, we see that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_1) is in co-NP, $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_2) is in NP and both $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_3) and $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_4) are in L.

3.2 Extending the method

Call a shop $f : B \rightarrow \mathfrak{P}(B) \setminus \{\emptyset\}$ an *A-shop* if there exists $b \in B$ s.t. $f(b) = B$. Call f an *E-shop* if there exists $b \in B$ s.t. $b \in f(x)$, for all $x \in B$. Clearly, \forall -shops are A-shops and \exists -shops are E-shops. Let f^r denote f composed with itself r times. It is clear that, if f is an A-shop (resp., an E-shop) as just defined, then $f^r(b) = B$ (resp., $b \in f^r(x)$, for all $x \in B$), for all r . With an arbitrary shop f on B , we may associate the digraph \mathcal{G}_f on B in which there is an edge (x, y) if $f(x) \ni y$. In a digraph, a *source* is a vertex of in-degree zero and a *sink* is a vertex of out-degree zero (a self-loop is neither). The condition of totality ensures \mathcal{G}_f has no sinks and the condition of surjectivity ensures \mathcal{G}_f has no sources.

LEMMA 3.6. *Let f be an A-shop on a set B , with $|B| \geq 2$. Then $\langle f \rangle$ contains an A-shop g with a tripartition $\{b\}; B'; B''$ of B (B' non-empty) s.t.*

- $g(b) = B$
- for all $x \in B'$, $g(x) = \{x\}$
- for all $x \in B''$, there exists $y \in B'$, $g(x) = \{y\}$.

PROOF. Let b be s.t. $f(b) = B$. It is possible that there is $x \in B$ s.t. $f(x) = \{b\}$. However, by considering paths in \mathcal{G}_f , it is easy to see that

$$(3.6 \text{ *}) \quad \text{for no } x \text{ does } f^{|B|}(x) = \{b\}.$$

Consider now $\mathcal{G}'_{f|B|}$ to be the graph $\mathcal{G}_{f|B|}$ with the vertex b removed. Owing to (3.6 *), $\mathcal{G}'_{f|B|}$ will still have no sinks, but it may now have sources. Build $\mathcal{G}''_{f|B|}$ from $\mathcal{G}'_{f|B|}$ by recursively removing sources from $\mathcal{G}'_{f|B|}$ until none is left. Let us say this takes d steps. $\mathcal{G}''_{f|B|}$ is therefore the disjoint union of strongly-connected components. For each of its strongly-connected components C_1, \dots, C_k pick a cycle (not necessarily Hamiltonian) that visits each vertex in the component at least once. Let the lengths of these cycles be c_1, \dots, c_k and let c be the least common multiple of $\{c_1, \dots, c_k, d\}$. It is not hard to see that some sub-shop g of $(f^{|B|})^c$ has the desired properties, with B' being those vertices that remain in $\mathcal{G}''_{f|B|}$ (non-emptiness of B' follows from $\mathcal{G}'_{f|B|}$ having no sinks). \square

LEMMA 3.7. *Let f be an E-shop on a set B , with $|B| \geq 2$, where $b \in f(x)$ for all x . Then $\langle f \rangle$ contains an E-shop g with a bipartition $B'; B''$ of B (B' non-empty) s.t.*

- for all $x \in B'$, $g(x) \supseteq \{x, b\}$
- for all $x \in B''$, $g(x) \supseteq \{b\}$

—for all $y \in B$, exists $x \in B'$, $y \in g(x)$.

PROOF. It is possible that there is $x \in B$ s.t.

$$x \in f(b) \text{ but for all } y \in B \setminus \{b\}, x \notin f(y).$$

However, by considering paths in \mathcal{G}_f , it is easy to see that

$$(3.7^*) \text{ there is no } x \in B \text{ s.t. } x \in f^{|B|}(b) \text{ but for all } y \in B \setminus \{b\}, x \notin f^{|B|}(y).$$

Consider now $\mathcal{G}'_{f|B|}$ to be the graph $\mathcal{G}_{f|B|}$ with the vertex b removed. Owing to (3.7*), $\mathcal{G}'_{f|B|}$ will still have no sources, but it may now have sinks. Build $\mathcal{G}''_{f|B|}$ from $\mathcal{G}'_{f|B|}$ by recursively removing sinks from $\mathcal{G}'_{f|B|}$ until none is left. Let us say this takes d steps. $\mathcal{G}''_{f|B|}$ is therefore the disjoint union of strongly-connected components. For each of its strongly-connected components C_1, \dots, C_k pick a cycle (not necessarily Hamiltonian) that visits each vertex in the component at least once. Let the lengths of these cycles be c_1, \dots, c_k and let c be the least common multiple of $\{c_1, \dots, c_k, d\}$. It is not hard to see that some sub-shop g of $(f^{|B|})^c$ has the desired properties, with B' being those vertices that remain in $\mathcal{G}''_{f|B|}$ (non-emptiness of B' follows from $\mathcal{G}'_{f|B|}$ having no sources). \square

We give the following examples of shops g of the given forms.

Lemma 3.6 Lemma 3.7

$$\begin{array}{c|c} 0|0123 & 0|0 \\ \hline 1|1 & 1|0 \\ \hline 2|1 & 2|012 \\ \hline 3|3 & 3|03 \end{array}$$

LEMMA 3.8. Let $\varphi(u, \mathbf{v})$ be a formula of $\{\exists, \forall, \wedge, \vee\}$ -FO, where the arity of \mathbf{v} is k . Let \mathcal{B} be a finite structure with an **A**-shop g , satisfying the conditions of Lemma 3.6, as a she. For all $\mathbf{x} := (x_1, \dots, x_k) \in (B' \cup \{b\})^k$,

$$\mathcal{B} \models \varphi(b, \mathbf{x}) \xrightarrow{(I)} \mathcal{B} \models \forall u \varphi(u, \mathbf{x})$$

$$\mathcal{B} \models \exists u \varphi(u, \mathbf{x}) \xrightarrow{(II)} \mathcal{B} \models \exists u \in B' \varphi(u, \mathbf{x}).$$

PROOF. Consider the relation defined by the formula $\varphi(u, \mathbf{v})$ of $\{\exists, \forall, \wedge, \vee\}$ -FO on \mathcal{B} . By Theorem 2.2, it is invariant under $g \in \text{shE}(\mathcal{B})$. Take arbitrary $u \in B$. Noting that x_i is from $B' \cup \{b\}$, we have $x_i \in g(x_i)$ and $u \in g(b)$. Part (I) follows. Now noting that, for each $u \in B$ there is some $u' \in B'$ s.t. $u' \in g(u)$, Part (II) follows. \square

LEMMA 3.9. Let $\varphi(u, \mathbf{v})$ be a formula of $\{\exists, \forall, \wedge, \vee\}$ -FO, where the arity of \mathbf{v} is k . Let \mathcal{B} be a finite structure with an **E**-shop g , satisfying the conditions of Lemma 3.7, as a she. For all $\mathbf{x} := (x_1, \dots, x_k) \in (B' \cup \{b\})^k$,

$$\mathcal{B} \models \forall u \in B' \varphi(u, \mathbf{x}) \xrightarrow{(I)} \mathcal{B} \models \forall u \varphi(u, \mathbf{x})$$

$$\mathcal{B} \models \exists u \varphi(u, \mathbf{x}) \xrightarrow{(II)} \mathcal{B} \models \varphi(b, \mathbf{x}).$$

PROOF. Consider the relation defined by the formula $\varphi(u, \mathbf{v})$ of $\{\exists, \forall, \wedge, \vee\}$ -FO on \mathcal{B} . By Theorem 2.2, it is invariant under $g \in \text{shE}(\mathcal{B})$. Take arbitrary $u \in B$.

Noting that x_i is from $B' \cup \{b\}$, and $x_i \in g(x_i)$, and that there is some $u' \in B'$ s.t. $u \in g(u')$, Part (I) follows. Now noting that, for each $u \in B$ we have $b \in g(u)$, Part (II) follows. \square

THEOREM 3.10. *If \mathcal{B} has an A-shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in NP. If \mathcal{B} has an E-shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in co-NP. If \mathcal{B} has a both an A-shop and an E-shop as a she then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in L.*

PROOF. Let φ be a sentence of $\{\exists, \forall, \wedge, \vee\}$ -FO, and let $\varphi_{[\forall/b, \exists/B']}$ (respectively, $\varphi_{[\exists/b, \forall/B']}$) be φ with all universal variables substituted by b and existential variables restricted to B' (respectively, existential variables substituted by b and universal variables restricted to B').

If \mathcal{B} has an A-shop f as a she, then let $g \in \langle f \rangle$ and B' be as in Lemma 3.6. Consider a sentence $\varphi \in \{\exists, \forall, \wedge, \vee\}$ -FO, w.l.o.g. in prenex form. It follows by repeated application of Lemma 3.8 on φ – from the outermost quantifier in – that $\mathcal{B} \models \varphi$ iff $\mathcal{B} \models \varphi_{[\forall/b, \exists/B']}$. This may now be checked by an NP machine simply guessing the witnesses in B' . Similarly, if \mathcal{B} has an E-shop as a she then it follows by repeated application of Lemma 3.9 that $\mathcal{B} \models \varphi$ iff $\mathcal{B} \models \varphi_{[\exists/b, \forall/B']}$.

If \mathcal{B} has both an A-shop f_A and an E-shop f_E as a she, then it follows that their composition $f_A \circ f_E$ (also $f_E \circ f_A$) is a she and an $\exists\forall$ -shop. The result follows from Theorem 3.5. \square

3.3 Reduction to simpler cases

For a shop f , we recall the associated digraph \mathcal{G}_f defined in the previous section. We say that f is an equivalence relation if \mathcal{G}_f is the digraph of an equivalence relation (f maps each element to its equivalence class). For a structure \mathcal{B} and an equivalence relation f on B , we define the structure $\mathcal{B}_{/f}$ as follows. The elements of $\mathcal{B}_{/f}$ are the equivalence classes of f and the relation $R^{\mathcal{B}_{/f}}(\tilde{b}_1, \dots, \tilde{b}_r)$ holds if for some representatives b_1, \dots, b_r of the classes $\tilde{b}_1, \dots, \tilde{b}_r$ the relation $R^{\mathcal{B}}(b_1, \dots, b_r)$ holds.

LEMMA 3.11. *Let \mathcal{B} be a structure with an equivalence relation f as a she. Then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) = $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{B}_{/f}$).*

PROOF. In fact, it is easy to see that \mathcal{B} and $\mathcal{B}_{/f}$ agree on all sentences of equality-free FO logic. This is because the she f guarantees that all elements in an equivalence class partake in exactly the same relations as one another. Indeed, there is a strong homomorphism from \mathcal{B} to $\mathcal{B}_{/f}$ (for more details, see, e.g., the Homomorphism Theorem in [7]). \square

3.4 Down-shop-monoids of high complexity

LEMMA 3.12. *Let \mathcal{B} , with $|B| \geq 2$, be a structure s.t. $\text{shE}(\mathcal{B})$ is a permutation subgroup. Then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is Pspace-complete.*

PROOF. Let \mathcal{B}_{NAE} be the structure on B with a single ternary relation $R_{\text{NAE}} := B^3 \setminus \{(b, b, b) : b \in B\}$. $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}_{NAE}) is a generalisation of the problem 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., [22]). $\text{shE}(\mathcal{B}_{\text{NAE}})$ is the symmetric group $S_{|B|}$. The statement of the theorem now follows from Theorem 2.3, since $\text{shE}(\mathcal{B}) \subseteq \text{shE}(\mathcal{B}_{\text{NAE}})$. \square




Fig. 2. Further sample digraphs.

COROLLARY 3.13. *For all \mathcal{B} s.t. $|B| \geq 2$, $\{\exists, \forall, \wedge, \vee, =\}$ -FO(\mathcal{B}) is Pspace-complete.*

PROOF. $\{\exists, \forall, \wedge, \vee, =\}$ -FO(\mathcal{B}) may be given as the problem $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}'), where \mathcal{B}' is the structure \mathcal{B} expanded with the graph of equality. Owing to the presence of the graph of equality, $\text{shE}(\mathcal{B}')$ must be a permutation subgroup, and the result follows from the previous lemma. \square

The following is a generalisation of Lemma 3.12.

LEMMA 3.14. *Let \mathcal{B} be a structure whose universe admits the partition B_1, \dots, B_l ($l \geq 2$). If all shes of \mathcal{B} are sub-shops of some f of the form $f(x) := B_i$ iff $x \in B_{\pi(i)}$, for π a permutation on the set $\{1, \dots, l\}$, then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is Pspace-complete.*

PROOF. Let $\mathcal{K}_{|B_1|, \dots, |B_l|}$ be the complete l -partite graph with partitions of size $|B_1|, \dots, |B_l|$. It may easily be verified that $\text{shE}(\mathcal{B}) \subseteq \text{shE}(\mathcal{K}_{|B_1|, \dots, |B_l|})$. $\mathcal{K}_{|B_1|, \dots, |B_l|}$ clearly has the equivalence relation g that maps each element in B_i to B_i ($1 \leq i \leq l$) as a she. Hence, $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{K}_{|B_1|, \dots, |B_l|}$) = $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{K}_l) by Lemma 3.11, and the result follows from Theorem 2.3 and Lemma 3.12. \square

In Figure 2, four more digraphs \mathcal{G}_5 – \mathcal{G}_8 are drawn. It may easily be verified that $\text{shE}(\mathcal{G}_5)$ – $\text{shE}(\mathcal{G}_8)$ are as follows.

$$\begin{array}{cccc} \text{shE}(\mathcal{G}_5) & \text{shE}(\mathcal{G}_6) & \text{shE}(\mathcal{G}_7) & \text{shE}(\mathcal{G}_8) \\ \langle \begin{array}{c} 0|0 \\ 1|1 \\ 2|2 \end{array} \rangle & \langle \begin{array}{c} 0|0 \\ 1|1 \\ 2|2 \end{array} \rangle & \langle \begin{array}{c} 0|0 \\ 1|2 \\ 2|1 \end{array} \rangle & \langle \begin{array}{c} 0|02 \\ 1|1 \\ 2|02 \end{array} \rangle \end{array}$$

It follows from Lemmata 3.12 and 3.14 that each of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_5), \dots , $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_8) is Pspace-complete.

4. CLASSIFICATION RESULTS

We are now in a position to use the methods of the previous section to classify the complexities of $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) as \mathcal{B} ranges over, firstly, Boolean structures, and then structures of size three.

4.1 The Boolean case

We consider the case $|B| = 2$, with the normalised domain $B := \{0, 1\}$. It may easily be verified that there are five DSMS in this case, depicted as a lattice in Figure 3. The two elements of this lattice that represent the two subgroups of S_2 are drawn in the middle and bottom.




Fig. 3. The Boolean lattice of DSMs with their associated complexity.

THEOREM 4.1 DICHOTOMY. *Let \mathcal{B} be a Boolean structure.*

- I. *If either $\forall_0\exists_1$ or $\forall_1\exists_0$ (i.e., $\frac{0|01}{1|1}$ or $\frac{0|0}{1|01}$) is a she of \mathcal{B} , then $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in L.*
- II. *Otherwise, $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is Pspace-complete.*

PROOF. $\text{shE}(\mathcal{B})$ must be one of the five DSMs depicted in Figure 3. If $\text{shE}(\mathcal{B})$ contains one of $\forall_0\exists_1$ or $\forall_1\exists_0$, then L membership follows from Theorem 3.5. Otherwise $\text{shE}(\mathcal{B})$ is either $\langle \frac{0|0}{1|1} \rangle$ or $\langle \frac{0|1}{1|0} \rangle$; in both cases the hardness result follows from Lemma 3.12. \square

Remark 4.2. In the Boolean case, $\langle \forall_1\exists_0 \rangle = \langle \{\forall_1, \exists_0\} \rangle$ ($\forall_1 = \exists_0$) and $\langle \forall_0\exists_1 \rangle = \langle \{\forall_0, \exists_1\} \rangle$ ($\forall_0 = \exists_1$).

4.2 The three-element case

We consider the case $|B| = 3$, with the normalised domain $B := \{0, 1, 2\}$. We will move straight to the classification theorem.

THEOREM 4.3 TETRACHOTOMY. *Let \mathcal{B} be a three-element structure.*

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

PROOF.

I. $\text{shE}(\mathcal{B})$ contains both an A-shop and an E-shop. If $\text{shE}(\mathcal{B})$ contains both an A-shop f and an E-shop g , then both $f \circ g$ and $g \circ f$ are $\forall\exists$ -shops, and it follows from Theorem 3.5 that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) is in L.

II. $\text{shE}(\mathcal{B})$ contains an A-shop but no E-shop. Removing a symmetry, and without loss of generality, we assume that $\text{shE}(\mathcal{B})$ contains a shop f s.t. $f(0) = \{0, 1, 2\}$. If either 1 or 2 were in both $f(1)$ and $f(2)$ then f would be a $\forall\exists$ -shop. If either $f(1)$ or $f(2)$ were $\{1, 2\}$ then either f would be a $\forall\exists$ -shop or the other of $f(1)$ or $f(2)$ would contain 0. And, if either of $f(1)$ or $f(2)$ contained 0, then f^2 would be a $\forall\exists$ -shop. It follows that f must be either $\frac{0|012}{1|1} \frac{1}{2}$ or $\frac{0|012}{1|2} \frac{1}{1}$, and that $\frac{0|012}{1|1} \frac{1}{2}$ necessarily appears in $\text{shE}(\mathcal{B})$.

Which other shops may also appear in $\text{shE}(\mathcal{B})$? Again no g with any of the properties

- (1) either 1 or 2 are in both $g(1)$ and $g(2)$,
- (2) either $g(1)$ or $g(2)$ are $\{1, 2\}$,
- (3) either $g(1)$ or $g(2)$ contain 0,

since then $g \circ \frac{0|012}{1|1} \frac{1}{2}$ would be forbidden as in the previous paragraph. It follows

that only shops among $\langle \frac{0|012}{1|2} \frac{1}{1} \rangle$ may also be in $\text{shE}(\mathcal{B})$. Reconsidering symmetries, it

follows that $\text{shE}(\mathcal{B})$ is one of the six DSMs, $\langle \frac{0|012}{1|1} \frac{1}{2} \rangle$, $\langle \frac{0|012}{1|2} \frac{1}{1} \rangle$, $\langle \frac{0|0}{1|012} \frac{1}{2} \rangle$, $\langle \frac{0|2}{1|012} \frac{1}{0} \rangle$, $\langle \frac{0|0}{1|1} \frac{1}{012} \rangle$

and $\langle \frac{0|1}{1|0} \frac{1}{012} \rangle$. Membership in NP follows from either of Theorems 3.5 or 3.10. For NP-

hardness, consider the disjoint union $\mathcal{K}_2 \uplus \mathcal{K}_1$ of the antireflexive 2- and 1-cliques, as drawn in Figure 4.2. $\text{shE}(\mathcal{K}_2 \uplus \mathcal{K}_1)$ is either $\langle \frac{0|012}{1|2} \frac{1}{1} \rangle$, $\langle \frac{0|2}{1|012} \frac{1}{0} \rangle$ or $\langle \frac{0|1}{1|0} \frac{1}{012} \rangle$, depending on

the vertex labelling. $\{\exists, \wedge, \vee\}$ -FO(\mathcal{K}_2) is NP-complete (by reduction from 3-*not-all-equal satisfiability*, set $R_{\text{NAE}}(u, v, w) := E(u, v) \vee E(v, w)$), and $\mathcal{K}_2 \uplus \mathcal{K}_1$ agrees with \mathcal{K}_2 on all sentences of $\{\exists, \wedge, \vee\}$ -FO (see [17]). It follows that $\{\exists, \wedge, \vee\}$ -FO($\mathcal{K}_2 \uplus \mathcal{K}_1$) is NP-complete and that $\{\exists, \forall, \wedge, \vee\}$ -FO($\mathcal{K}_2 \uplus \mathcal{K}_1$) is NP-hard. The result for NP-hardness now follows from Theorem 2.3, since $\text{shE}(\mathcal{B}) \subseteq \text{shE}(\mathcal{K}_2 \uplus \mathcal{K}_1)$, for one of the three vertex labellings.

III. $\text{shE}(\mathcal{B})$ contains an E-shop but no A-shop. Removing a symmetry, and without loss of generality, we assume that $\text{shE}(\mathcal{B})$ contains a shop f s.t. $0 \in f(0), f(1), f(2)$. If either $f(1)$ or $f(2)$ also contained $\{1, 2\}$, then f would be a $\forall\exists$ -shop. If either 1 or 2 were in both $f(1)$ and $f(2)$ then either f would be a $\forall\exists$ -shop or the other of 1 or 2 would be in $f(0)$. And, if $f(0)$ contained either 1 or 2 (e.g. f looked like $\frac{0|01\dots}{1|0\dots} \frac{1}{0\dots}$), then f^2 would be a $\forall\exists$ -shop. It follows that f must be either

$\frac{0|0}{1|01} \frac{1}{02}$ or $\frac{0|0}{1|02} \frac{1}{01}$, and that $\frac{0|0}{1|01} \frac{1}{02}$ necessarily appears in $\text{shE}(\mathcal{B})$.

Which other shops may also appear in $\text{shE}(\mathcal{B})$? Again no g with any of the properties

- (1) either $g(1)$ or $g(2)$ contains $\{1, 2\}$,
- (2) either 1 or 2 in both $g(1)$ and $g(2)$,
- (3) $g(0)$ contains either 1 or 2,

since then $g \circ \frac{0|0}{1|01}$ would be forbidden as in the previous paragraph. It follows that only shops among $\langle \frac{0|0}{1|02} \rangle$ may also be in $\text{shE}(\mathcal{B})$. Reconsidering symmetries, it follows that $\text{shE}(\mathcal{B})$ is one of the six DSMs, $\langle \frac{0|0}{1|01} \rangle$, $\langle \frac{0|0}{1|02} \rangle$, $\langle \frac{0|12}{1|1} \rangle$, $\langle \frac{0|01}{1|1} \rangle$, $\langle \frac{0|12}{1|2} \rangle$ and $\langle \frac{0|02}{1|2} \rangle$. Membership in co-NP follows from either of Theorems 3.5 or 3.10. For co-NP-hardness, consider now the complement graph $\overline{\mathcal{K}_2 \uplus \mathcal{K}_1}$ (for a graph \mathcal{G} , define its complement $\overline{\mathcal{G}}$ over the same vertex set to have the complementary edge set – i.e. $\overline{\mathcal{G}} \models E(x, y)$ iff $\mathcal{G} \not\models E(x, y)$). It is a simple application of de Morgan duality that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}) is in NP (resp., is NP-complete) iff $\{\exists, \forall, \wedge, \vee\}$ -FO($\overline{\mathcal{G}}$) is in co-NP (resp., is co-NP-complete) – see [17]. A similar argument to that for Class II, but with $\overline{\mathcal{K}_2 \uplus \mathcal{K}_1}$ yields the co-NP-hardness result for Class III.

IV. $\text{shE}(\mathcal{B})$ contains neither an E-shop nor an A-shop. If $\text{shE}(\mathcal{B})$ is not a sub-DSM of the DSM associated with the symmetric group $S_3 = \langle \frac{0|1}{1|2}, \frac{0|0}{1|2} \rangle$, then we may assume that $\text{shE}(\mathcal{B})$ contains a shop f where some element is mapped to exactly two elements. There are two possibilities, either the element is included among the two it is mapped to, or it is not. Without loss of generality, we consider these two cases separately as follows.

f is of the form $\frac{0|01}{1|\dots}$. 2 must appear somewhere among $f(0)$, $f(1)$ and $f(2)$, but may not appear among $f(0)$ or $f(1)$ as then f^2 would be an A-shop. It follows that $2 \in f(2)$, whereupon neither may 0 be in $f(2)$ (as then f^2 would be an A-shop) nor may 1 be in $f(2)$ (as then either f or f^2 would be an E-shop – depending on which of 0, 1 $\in f(1)$). It follows that $f(2) = \{2\}$. In fact, all of the remaining possibilities are valid, i.e. f may be any of $\frac{0|01}{1|0}$, $\frac{0|01}{1|1}$ and $\frac{0|01}{1|2}$.

f is of the form $\frac{0|12}{1|\dots}$. W.l.o.g. we may assume that $0 \in f(1)$, whereupon $1 \notin f(1)$ for otherwise f^2 would be an A-shop. It follows that $f(1) =$ either $\{0\}$ or $\{0, 2\}$. In fact the second of these is not possible as, if it were so, then: if $2 \in f(2)$, f would be an E-shop, and if 0 or $1 \in f(2)$, f^2 would be an A-shop. Thus, we have $f(1) = \{0\}$. If $1 \in f(2)$ then f^3 would be an A-shop and if $2 \in f(2)$ then f^2 would be an E-shop. It follows that f is $\frac{0|12}{1|0}$.

Recalling symmetries, it follows that f is some shop from among $\langle \frac{0|1}{1|2}, \frac{0|0}{1|2} \rangle$, $\langle \frac{0|12}{1|0}, \frac{0|12}{1|1} \rangle$, $\langle \frac{0|1}{1|02} \rangle$ and $\langle \frac{0|2}{1|01} \rangle$. We have in fact demonstrated that if $\text{shE}(\mathcal{B})$ contains neither an E-shop nor an A-shop then all shops of $\text{shE}(\mathcal{B})$ are among the given list. There are several DSMs that can be formed in this manner but we will demonstrate that all of them are sub-DSMs of one of $\langle \frac{0|1}{1|2}, \frac{0|0}{1|2} \rangle$, $\langle \frac{0|12}{1|0}, \frac{0|12}{1|1} \rangle$ and $\langle \frac{0|2}{1|01} \rangle$.

Arguing via symmetries, this is clear when we consider that $\frac{0|0}{1|1}$ can not combine with any element of $S_3 = \langle \frac{0|1}{1|2}, \frac{0|0}{1|2} \rangle$ apart from $\frac{0|0}{1|2}$ (and the identity). For $\frac{0|0}{1|1} \circ$




Fig. 4. The digraphs involved in Classes II and III of Theorem 4.3.

$\frac{0|1}{1|2} \circ \frac{0|0}{1|1}$ is an A-shop; $\frac{0|0}{1|1} \circ \frac{0|1}{1|0} \circ \frac{0|0}{1|1}$ is an A-shop; and $\frac{0|0}{1|1} \circ \frac{0|2}{1|1} \circ \frac{0|0}{1|1}$ is an E-shop.

Pspace-hardness when $\text{shE}(\mathcal{B}) \subseteq \langle \frac{0|1}{1|2}, \frac{0|0}{1|1} \rangle$ follows from Lemma 3.12, while Pspace-hardness when $\text{shE}(\mathcal{B}) \subseteq$ any of $\langle \frac{0|12}{1|0} \rangle, \langle \frac{0|1}{1|02} \rangle$ follows from Lemma 3.14. \square

Casting our mind back to the digraphs \mathcal{G}_1 and \mathcal{G}_2 of Figure 1, we can read from the previous theorem that $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_1) and $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{G}_2) are co-NP-complete and NP-complete, respectively.

5. FINAL REMARKS

We have introduced the class of problems $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) as well as an algebraic framework in which to study their complexity. We hope that we have adequately demonstrated that this class of problems displays complexity-theoretic richness, while not being too resistant to a full classification. The algebraic method used in our classification for the three-element case gives simple explanation where there previously was none – if one were to look at the examples of Figures 1 and 2, there is little obvious in their immediate structure that betrays their position in the classification.

We note that our positive algorithms, for membership of NP, co-NP and – especially – L, are uniform, and are based on simple quantifier elimination. Perhaps it is to be hoped that a full classification for the problems $\{\exists, \forall, \wedge, \vee\}$ -FO(\mathcal{B}) would make use only of versions of quantifier elimination. In any case, we conjecture that the tetrachotomy of Theorem 4.3 extends to all structures \mathcal{B} ; though we know we would need more sophisticated classes of shes than those of Section 3.1 to prove this.

CONJECTURE 5.1 TETRACHOTOMY. *Let \mathcal{B} be any finite structure.*

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

We note also that, unlike the situation with clones and the CSP, the down-shop-monoids associated with a finite domain are always finite. This means that their lattice should be effectively computable for low domain sizes like four or five.⁷ Very recently, and since this paper was submitted, a proof of Conjecture 5.1 has been submitted to conference [15]. It remains independently unverified.

REFERENCES

- BÖRNER, F. Total multifunctions and relations. In *AAA60: Workshop on General Algebra, Dresden, Germany* (2000).
- BÖRNER, F. Some other Galois connections, 2006. Complexity of Constraints, Dagstuhl, Germany.
- BÖRNER, F., KROKHIN, A., BULATOV, A., AND JEAVONS, P. Quantified constraints and surjective polymorphisms. Tech. Rep. PRG-RR-02-11, Oxford University, 2002.
- BULATOV, A. A. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM* 53, 1 (2006), 66–120.
- CHEN, H. The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case. *SIAM J. Comput.* 37, 5 (2008), 1674–1701.
- CREIGNOU, N., KHANNA, S., AND SUDAN, M. *Complexity classifications of Boolean Constraint Satisfaction Problems*. SIAM Monographs. 2001.
- ENDERTON, H. B. *A Mathematical Introduction to Logic*. Academic Press, 1972.
- FEDER, T., AND VARDI, M. Y. The computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory. *SIAM J. Comput.* 28 (1999).
- HELL, P., AND NEŠETŘIL, J. On the complexity of H-coloring. *J. Combin. Theory Ser. B* 48 (1990).
- JEAVONS, P. On the algebraic structure of combinatorial problems. *Theoretical Computer Science* 200, 1–2 (1998), 185–204.
- JEAVONS, P., COHEN, D., AND GYSSENS, M. Closure properties of constraints. *J. ACM* 44, 4 (1997), 527–548.
- LADNER, R. E. On the structure of polynomial time reducibility. *J. ACM* 22, 1 (1975), 155–171.
- LYNCH, N. Log space recognition and translation of parenthesis languages. *J. ACM* 24 (1977), 583–590.
- MADELAINE, F., AND MARTIN, B. The complexity of positive first-order logic without equality. *Logic in Computer Science, Symposium on* (2009), 429–438.
- MADELAINE, F., AND MARTIN, B. A tetrachotomy for positive first-order logic without equality, 2011. Manuscript.
- MARTIN, B. Dichotomies and duality in first-order model checking problems. *CoRR abs/cs/0609022* (2006).
- MARTIN, B. First order model checking problems parameterized by the model. In *CiE 2008, LNCS 5028* (2008), pp. 417–427.
- MARTIN, B. Model checking positive equality-free FO: Boolean structures and digraphs of size three. *CoRR abs/0808.0647* (2008).
- MARTIN, B. The lattice structure of sets of surjective hyper-operations. In *Principles and Practice of Constraint Programming - CP 2010 - 16th International Conference, CP 2010* (2010), pp. 368–382.
- MARTIN, B., AND MADELAINE, F. R. Towards a trichotomy for quantified H-coloring. In *CiE 2006, LNCS 3988* (2006), pp. 342–352.

⁷Since this paper was submitted, Conjecture 5.1 has been settled for four-element domains in [21]. Computer verification was used, but the lattice of DSMs itself was not computed, being too large for the available computation.

MARTIN, B., AND MARTIN, J. The complexity of positive first-order logic without equality II: The four-element case. In *Computer Science Logic, 24th International Workshop, CSL 2010* (2010), pp. 426–438.

PAPADIMITRIOU, C. *Computational Complexity*. Addison-Wesley, 1994.

Received March 2010; accepted November 2010