

Combining Theories with Shared Set Operations

Thomas Wies, Ruzica Piskac, and Viktor Kuncak

EPFL School of Computer and Communication Sciences, Switzerland

Abstract. Motivated by applications in software verification, we explore automated reasoning about the non-disjoint combination of theories of infinitely many finite structures, where the theories share set variables and set operations. We prove a combination theorem and apply it to show the decidability of the satisfiability problem for a class of formulas obtained by applying propositional connectives to formulas belonging to: 1) Boolean Algebra with Presburger Arithmetic (with quantifiers over sets and integers), 2) weak monadic second-order logic over trees (with monadic second-order quantifiers), 3) two-variable logic with counting quantifiers (ranging over elements), 4) the Bernays-Schönfinkel-Ramsey class of first-order logic with equality (with $\exists^*\forall^*$ quantifier prefix), and 5) the quantifier-free logic of multisets with cardinality constraints.

1 Introduction

Constraint solvers based on satisfiability modulo theories (SMT) [4, 8, 13] are a key enabling technique in software and hardware verification systems [2, 3]. The range of problems amenable to such approaches depends on the expressive power of the logics supported by the SMT solvers. Current SMT solvers implement the combination of quantifier-free stably infinite theories with disjoint signatures, in essence following the approach pioneered by Nelson and Oppen [24]. Such solvers serve as decision procedures for quantifier-free formulas, typically containing uninterpreted function symbols, linear arithmetic, and bit vectors. The limited expressiveness of SMT prover logics translates into a limited class of properties that automated verification tools can handle.

To support a broader set of applications, this paper considers decision procedures for the combination of *possibly quantified* formulas in *non-disjoint* theories. The idea of combining rich theories within an expressive language has been explored in interactive provers [5, 7, 23, 25]. Such integration efforts are very useful, but do not result in complete decision procedures for the combined logics. The study of completeness for non-disjoint combination is relatively recent [32] and provides foundations for the general problem. Under certain conditions, such as local finiteness, decidability results have been obtained even for non-disjoint theories [14]. Our paper considers a case of combination of non-disjoint theories sharing operations on *sets of uninterpreted elements*, a case that was not considered before. The theories that we consider have the property that the tuples of cardinalities of Venn regions over shared set variables in the models of a formula are a semilinear set (i.e., expressible in Presburger arithmetic).

Reasoning about combinations of decidable logics. The idea of deciding a combination of logics is to check the satisfiability of a conjunction of formulas $A \wedge B$ by using one decision procedure, D_A , for A , and another decision procedure, D_B , for B . To obtain a complete decision procedure, D_A and D_B must communicate to ensure that a model found by D_A and a model found by D_B can be merged into a model for $A \wedge B$.

Reduction-based decision procedure. We follow a reduction approach to decision procedures. The first decision procedure, D_A , computes a *projection*, S_A , of A onto *shared* set variables, which are free in both A and B . This projection is semantically equivalent to existentially quantifying over predicates and variables that are free in A but not in B ; it is the strongest consequence of A expressible only using the shared set variables. D_B similarly computes the projection S_B of B . This reduces the satisfiability of $A \wedge B$ to satisfiability of the formula $S_A \wedge S_B$, which contains only set variables.

A logic for shared constraints on sets. A key parameter of our combination approach is the logic of sets used to express the projections S_A and S_B . A suitable logic depends on the logics of formulas A and B . We consider as the logics for A, B several expressive logics we consider useful based on our experience with the Jahob verification system [34, 36]. Remarkably, the smallest logic needed to express the projection formulas in these logics has the expressive power of Boolean Algebra with Presburger Arithmetic (BAPA), described in [21] and in Fig. 3. We show that the decision procedures for these four logics can be naturally extended to a reduction to BAPA that captures precisely the constraints on set variables. The existence of these reductions, along with quantifier elimination [20] and NP membership of the quantifier-free fragment [21], make BAPA an appealing reduction target for expressive logics.

Contribution summary. We present a technique for showing decidability of theories that share sets of elements. Furthermore, we show that the logics

1. Boolean Algebra with Presburger Arithmetic [9, 20, 21],
2. weak monadic second-order logic of two successors WS2S [31],
3. two-variable logic with counting C^2 [29],
4. Bernays-Schönfinkel-Ramsey class [30], and
5. quantifier-free multisets with cardinality constraints [27, 28]

all meet the conditions of our combination technique. Consequently, we obtain the decidability of quantifier-free combination of formulas in these logics.¹

2 Example: Proving a Verification Condition

Our example shows a verification condition formula generated when verifying an unbounded linked data structure. The formula belongs to our new decidable class obtained by combining several decidable logics.

¹ Further details are provided in [35].

$$\begin{aligned}
& \text{tree}[\text{left}, \text{right}] \wedge \text{left } p = \text{null} \wedge p \in \text{nodes} \wedge \\
& \text{nodes} = \{x. (\text{root}, x) \in \{(x, y). \text{left } x = y \mid \text{right } x = y\}^*\} \wedge \\
& \text{content} = \{x. \exists n. n \neq \text{null} \wedge n \in \text{nodes} \wedge \text{data } n = x\} \wedge \\
& e \notin \text{content} \wedge \text{nodes} \subseteq \text{alloc} \wedge \\
& \text{tmp} \notin \text{alloc} \wedge \text{left } \text{tmp} = \text{null} \wedge \text{right } \text{tmp} = \text{null} \wedge \\
& \text{data } \text{tmp} = \text{null} \wedge (\forall y. \text{data } y \neq \text{tmp}) \wedge \\
& \text{nodes1} = \{x. (\text{root}, x) \in \{(x, y). (\text{left } (p := \text{tmp})) x = y \mid \text{right } x = y\} \wedge \\
& \text{content1} = \{x. \exists n. n \neq \text{null} \wedge n \in \text{nodes1} \wedge (\text{data}(\text{tmp} := e)) n = x\} \rightarrow \\
& \quad \text{card } \text{content1} = \text{card } \text{content} + 1
\end{aligned}$$
Fig. 1. Verification condition

SHARED SETS: nodes, nodes1, content, content1, {e}, {tmp}

WS2S FRAGMENT: tree[left,right] \wedge left p = null \wedge p \in nodes \wedge left tmp = null \wedge right tmp = null \wedge nodes = {x. (root,x) \in {(x,y). left x = y | right x = y}^{*}} \wedge nodes1 = {x. (root,x) \in {(x,y). (left (p:=tmp)) x = y | right x = y}

CONSEQUENCE: nodes1 = nodes \cup {tmp}

C2 FRAGMENT: data tmp = null \wedge ($\forall y. \text{data } y \neq \text{tmp}$) \wedge tmp \notin alloc \wedge nodes \subseteq alloc \wedge content = {x. $\exists n. n \neq \text{null} \wedge n \in \text{nodes} \wedge \text{data } n = x$ } \wedge content1 = {x. $\exists n. n \neq \text{null} \wedge n \in \text{nodes1} \wedge (\text{data}(\text{tmp} := e)) n = x$ }

CONSEQUENCE: nodes1 \neq nodes \cup {tmp} \vee content1 = content \cup {e}

BAPA FRAGMENT: e \notin content \wedge card content1 \neq card content + 1

CONSEQUENCE: e \notin content \wedge card content1 \neq card content + 1

Fig. 2. Negation of Fig. 1, and consequences on shared sets

Decidability of the verification condition. Fig. 1 shows the verification condition formula for a method (`insertAt`) that inserts a node into a linked list. The validity of this formula implies that invoking a method in a state satisfying the precondition results in a state that satisfies the postcondition of `insertAt`. The formula contains the transitive closure operator, quantifiers, set comprehensions, and the cardinality operator. Nevertheless, there is a (syntactically defined) decidable class of formulas that contains the verification condition in Fig. 1. This decidable class is a set-sharing combination of three decidable logics, and can be decided using the method we present in this paper.

To understand the method for proving the formula in Fig. 1, consider the problem of showing the unsatisfiability of the negation of the formula. Fig. 2 shows the conjuncts of the negation, grouped according to three decidable logics to which the conjuncts belong: 1) weak monadic second-order logic of two successors (WS2S) 2) two-variable logic with counting C^2 3) Boolean Algebra with Presburger Arithmetic (BAPA). For the formula in each of the fragments, Fig. 2 also shows a consequence formula that contains only shared sets and statements about their cardinalities. (We represent elements as singleton sets, so we admit formulas sharing elements as well.)

A decision procedure. Note that the conjunction of the consequences of three formula fragments is an unsatisfiable formula. This shows that the original

verification condition is valid. In general, our decidability result shows that the decision procedures of logics such as WS2S and C^2 can be naturally extended to compute strongest consequences of formulas involving given shared sets. These consequences are all expressed in BAPA, which is decidable. In summary, the following is a decision procedure for satisfiability of combined formulas: 1) split the formula into fragments (belonging to WS2S, C^2 , or BAPA); 2) for each fragment compute its strongest BAPA consequence; 3) check the satisfiability of the conjunction of consequences.

3 Syntax and Semantics of Formulas

Higher-order logic. We present our problem in a fragment of classical higher-order logic [1, Chapter 5] with a particular set of types, which we call sorts. We assume that formulas are well-formed according to sorts of variables and logical symbols. Each variable and each logical symbol have an associated sort. The primitive sorts we consider are 1) `bool`, interpreted as the two-element set $\{\text{true}, \text{false}\}$ of booleans; 2) `int`, interpreted as the set of integers \mathbb{Z} ; and 3) `obj`, interpreted as a non-empty set of elements. The only sort constructors is the binary function space constructor ‘ \rightarrow ’. We represent a function mapping elements of sorts s_1, \dots, s_n into an element of sort s_0 as a term of sort $s_1 \times \dots \times s_n \rightarrow s_0$ where $s_1 \times s_2 \times \dots \times s_n \rightarrow s_0$ is a shorthand for $s_1 \rightarrow (s_2 \rightarrow \dots (s_n \rightarrow s_0))$. When s_1, \dots, s_n are all the same sort s , we abbreviate $s_1 \times \dots \times s_n \rightarrow s_0$ as $s^n \rightarrow s_0$. We represent a relation between elements of sorts s_1, \dots, s_n as a function $s_1 \times \dots \times s_n \rightarrow \text{bool}$. We use `set` as an abbreviation for the sort `obj` \rightarrow `bool`. We call variables of sort `set` *set variables*. The equality symbol applies only to terms of the same sort. We assume to have a distinct equality symbol for each sort of interest, but we use the same notation to denote all of them. Propositional operations connect terms of sort `bool`. We write $\forall x:s.F$ to denote a universally quantified formula where the quantified variable has sort s (analogously for $\exists x:s.F$ and $\exists x:s^K.F$ for counting quantifiers of Section 5.2). We denote by $\text{FV}(F)$ the set of all free variables that occur free in F . We write $\text{FV}_s(F)$ for the free variables of sort s . Note that the variables can be higher-order (we will see, however, that the shared variables are of sort `set`). A *theory* is simply a set of formulas, possibly with free variables.

Structures. A structure α specifies a finite set, which is also the meaning of `obj`, and we denote it $\alpha(\text{obj})$.² When α is understood we use $\llbracket X \rrbracket$ to denote $\alpha(X)$, where X denotes a sort, a term, a formula, or a set of formulas. If S is a set of formulas then $\alpha(S) = \text{true}$ means $\alpha(F) = \text{true}$ for each $F \in S$. In every structure we let $\llbracket \text{bool} \rrbracket = \{\text{false}, \text{true}\}$. Instead of $\alpha(F) = \text{true}$ we often write simply $\alpha(F)$. We interpret terms of the sort $s_1 \times \dots \times s_n \rightarrow s_0$ as total functions $\llbracket s_1 \rrbracket \times \dots \times \llbracket s_n \rrbracket \rightarrow \llbracket s_0 \rrbracket$. For a set A , we identify a function $f : A \rightarrow \{\text{false}, \text{true}\}$

² We focus on the case of finite $\alpha(\text{obj})$ primarily for simplicity; we believe the extension to the case where domains are either finite or countable is possible and can be done using results from [20, Section 8.1], [29, Section 5], [31].

with the subset $\{x \in A \mid f(x) = \mathbf{true}\}$. We thus interpret variables of the sort $\mathbf{obj}^n \rightarrow \mathbf{bool}$ as subsets of $\llbracket \mathbf{obj} \rrbracket^n$. If s is a sort then $\alpha(s)$ depends only on $\alpha(\mathbf{obj})$ and we denote it also by $\llbracket s \rrbracket$. We interpret propositional operations \wedge, \vee, \neg as usual in classical logic. A quantified variable of sort s ranges over all elements of $\llbracket s \rrbracket$. (Thus, as in standard model of HOL [1, Section 54], quantification over variables of sort $s_1 \rightarrow s_2$ is quantification over all total functions $\llbracket s_1 \rrbracket \rightarrow \llbracket s_2 \rrbracket$.)

3.1 Boolean Algebra with Presburger Arithmetic

$$\begin{aligned}
F &::= A \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid \neg F \mid \forall x:s.F \mid \exists x:s.F \\
s &::= \mathbf{int} \mid \mathbf{obj} \mid \mathbf{set} \\
A &::= B_1 = B_2 \mid B_1 \subseteq B_2 \mid T_1 = T_2 \mid T_1 < T_2 \mid K \mathit{dvd} T \\
B &::= x \mid \emptyset \mid \mathbf{Univ} \mid \{x\} \mid B_1 \cup B_2 \mid B_1 \cap B_2 \mid B^c \\
T &::= x \mid K \mid \mathbf{CardUniv} \mid T_1 + T_2 \mid K \cdot T \mid \mathit{card}(B) \\
K &::= \dots -2 \mid -1 \mid 0 \mid 1 \mid 2 \dots
\end{aligned}$$

Fig. 3. Boolean Algebra with Presburger Arithmetic (BAPA)

It will be convenient to enrich the language of our formulas with operations on integers, sets, and cardinality operations. These operations could be given by a theory or defined in HOL, but we choose to simply treat them as built-in logical symbols, whose meaning must be respected by all structures α we consider. Fig. 3 shows the syntax of Boolean Algebra with Presburger Arithmetic (BAPA) [9, 20]. The sorts of symbols appearing in BAPA formulas are as expected, e.g., $\subseteq : \mathbf{set}^2 \rightarrow \mathbf{bool}$, $< : \mathbf{int}^2 \rightarrow \mathbf{bool}$, $\mathit{dvd}_K : \mathbf{int} \rightarrow \mathbf{bool}$ for each integer constant K , $\mathit{singleton} : \mathbf{obj} \rightarrow \mathbf{set}$ (with $\mathit{singleton}(x)$ denoted as $\{x\}$), and $\mathit{complement} : \mathbf{set} \rightarrow \mathbf{set}$ (with $\mathit{complement}(A)$ denoted by A^c).

We sketch the meaning of the less common among the symbols in Fig. 3. \mathbf{Univ} denotes the universal set, that is, $\llbracket \mathbf{Univ} \rrbracket = \llbracket \mathbf{obj} \rrbracket$. $\mathit{card}(A)$ denotes the cardinality of the set A . $\mathbf{CardUniv}$ is interpreted as $\mathit{card}(\mathbf{Univ})$. The formula $K \mathit{dvd} t$ denotes that the integer constant K divides the integer t . We note that the condition $x \in A$ can be written in this language as $\{x\} \subseteq A$. Note that BAPA properly extends the first-order theory of Boolean Algebras over finite structures, which in turn subsumes the first-order logic with unary predicates and no function symbols, because e.g. $\exists x:\mathbf{obj}.F(x)$ can be written as $\exists X:\mathbf{set}.\mathit{card}(X)=1 \wedge F'(X)$ where in F' e.g. $P(x)$ is replaced by $X \subseteq P$.

BAPA-definable relations between sets. A *semilinear set* is a finite union of *linear sets*. A linear set is a set of the form $\{\mathbf{a} + k_1\mathbf{b}_1 + \dots + k_n\mathbf{b}_n \mid k_1, \dots, k_n \in \{0, 1, 2, \dots\}\}$ where $\mathbf{a}, \mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{Z}^M$ (\mathbf{a} is a *base* and \mathbf{b}_i are *step* vectors). We represent a linear set by its generating vectors $\mathbf{a}, \mathbf{b}_1, \dots, \mathbf{b}_n$, and a semilinear set by the finite set of representations of its linear sets. It was shown in [15] that a set of integer vectors $S \subseteq \mathbb{Z}^M$ is a solution set of a Presburger arithmetic formula P i.e. $S = \{(v_1, \dots, v_n).P\}$ iff S is a semilinear set. We then have the

following characterization of relationships between sets expressible in BAPA, which follows from [20].

Lemma 1 (BAPA-expressible means Venn-cardinality-semilinear).

Given a finite set U and a relation $\rho \subseteq (2^U)^p$ the following are equivalent:

1. there exists a BAPA formula F whose free variables are A_1, \dots, A_p , and have the sort set, such that $\rho = \{(s_1, \dots, s_p) \mid \{A_1 \mapsto s_1, \dots, A_p \mapsto s_p\}(F)\}$;
2. the following subset of \mathbb{Z}^M for $M = 2^p$ is semilinear:
 $\{(|s_1^c \cap s_2^c \cap \dots \cap s_p^c|, |s_1 \cap s_2^c \cap \dots \cap s_p^c|, \dots, |s_1 \cap s_2 \cap \dots \cap s_p|) \mid (s_1, \dots, s_p) \in \rho\}$.

Structures of interest in this paper. In the rest of this paper we consider structures that interpret the BAPA symbols as defined above. Because the meaning of BAPA-specific symbols is fixed, a structure α that interprets a set of formulas is determined by a finite set $\alpha(\text{obj})$ as well as the values $\alpha(x)$ for each variable x free in the set of formulas. Let $\{\text{obj} \mapsto u, x_1 \mapsto v_1, \dots, x_n \mapsto v_n\}$ denote the structure α with domain u that interprets each variable x_i as v_i .

4 Combination by Reduction to BAPA

The Satisfiability Problem. We are interested in an algorithm to determine whether there exists a structure $\alpha \in \mathcal{M}$ in which the following formula is true

$$B(F_1, \dots, F_n) \tag{1}$$

where

1. F_1, \dots, F_n are formulas with $\text{FV}(F_i) \subseteq \{A_1, \dots, A_p, x_1, \dots, x_q\}$.
2. $V_S = \{A_1, \dots, A_p\}$ are variables of sort **set**, whereas x_1, \dots, x_q are the remaining variables.³
3. Each formula F_i belongs to a given class of formulas, \mathcal{F}_i . For each \mathcal{F}_i we assume that there is a corresponding theory $\mathcal{T}_i \subseteq \mathcal{F}_i$.
4. $B(F_1, \dots, F_n)$ denotes a formula built from F_1, \dots, F_n using the propositional operations \wedge, \vee .⁴
5. As the set of structures \mathcal{M} we consider all structures α of interest (with finite $\llbracket \text{obj} \rrbracket$, interpreting BAPA symbols in the standard way) for which $\alpha(\cup_{i=1}^n \mathcal{T}_i)$.
6. (Set Sharing Condition) If $i \neq j$, then $\text{FV}(\{F_i\} \cup \mathcal{T}_i) \cap \text{FV}(\{F_j\} \cup \mathcal{T}_j) \subseteq V_S$.

Note that, as a special case, if we embed a class of first-order formulas into our framework, we obtain a framework that supports sharing unary predicates, but not e.g. binary predicates.

Combination Theorem. The formula B in (1) is satisfiable iff one of the disjuncts in its disjunctive normal form is satisfiable. Consider a disjunct $F_1 \wedge$

³ For notational simplicity we do not consider variables of sort **obj** because they can be represented as singleton sets, of sort **set**.

⁴ The absence of negation is usually not a loss of generality because most \mathcal{F}_i are closed under negation so B is the negation-normal form of a quantifier-free combination.

$\dots \wedge F_m$ for $m \leq n$. By definition of the satisfiability problem (1), $F_1 \wedge \dots \wedge F_m$ is satisfiable iff there exists a structure α such that for each $1 \leq i \leq m$, for each $G \in \{F_i\} \cup \mathcal{T}_i$, we have $\alpha(G) = \text{true}$. Let each variable x_i have some sort s_i (such as $\text{obj}^2 \rightarrow \text{bool}$). Then the satisfiability of $F_1 \wedge \dots \wedge F_m$ is equivalent to the following condition:

$$\exists \text{ finite set } u. \exists a_1, \dots, a_p \subseteq u. \exists v_1 \in \llbracket s_1 \rrbracket^u \dots \exists v_q \in \llbracket s_q \rrbracket^u. \bigwedge_{i=1}^m \{\text{obj} \rightarrow u, A_1 \mapsto a_1, \dots, A_p \mapsto a_p, x_1 \mapsto v_1, \dots, x_q \mapsto v_q\}(\{F_i\} \cup \mathcal{T}_i) \quad (2)$$

By the set sharing condition, each of the variables x_1, \dots, x_q appears only in one conjunct and can be moved inwards from the top level to this conjunct. Using x_{ij} to denote the j -th variable in the i -th conjunct we obtain the condition

$$\exists \text{ finite set } u. \exists a_1, \dots, a_p \subseteq u. \bigwedge_{i=1}^m C_i(u, a_1, \dots, a_p) \quad (3)$$

where $C_i(u, a_1, \dots, a_p)$ is

$$\exists v_{i1} \dots \exists v_{iw_i}. \{\text{obj} \rightarrow u, A_1 \mapsto a_1, \dots, A_p \mapsto a_p, x_{i1} \mapsto v_{i1}, \dots, x_{iw_i} \mapsto v_{iw_i}\}(\{F_i\} \cup \mathcal{T}_i)$$

The idea of our combination method is to simplify each condition $C_i(u, a_1, \dots, a_p)$ into the truth value of a BAPA formula. If this is possible, we say that there exists a BAPA reduction.

Definition 2 (BAPA Reduction). *If \mathcal{F}_i is a set of formulas and $\mathcal{T}_i \subseteq \mathcal{F}_i$ a theory, we call a function $\rho : \mathcal{F}_i \rightarrow \mathcal{F}_{\text{BAPA}}$ a BAPA reduction for $(\mathcal{F}_i, \mathcal{T}_i)$ iff for every formula $F_i \in \mathcal{F}_i$ and for all finite u and $a_1, \dots, a_p \subseteq u$, the condition*

$$\exists v_{i1} \dots \exists v_{iw_i}. \{\text{obj} \rightarrow u, A_1 \mapsto a_1, \dots, A_p \mapsto a_p, x_{i1} \mapsto v_{i1}, \dots, x_{iw_i} \mapsto v_{iw_i}\}(\{F_i\} \cup \mathcal{T}_i)$$

is equivalent to the condition $\{\text{obj} \rightarrow u, A_1 \mapsto a_1, \dots, A_p \mapsto a_p\}(\rho(F_i))$.

A computable BAPA reduction is a BAPA reduction which is computable as a function on formula syntax trees.

Theorem 3. *Suppose that for every $1 \leq i \leq n$ for $(\mathcal{F}_i, \mathcal{T}_i)$ there exists a computable BAPA reduction ρ_i . Then the problem (1) in Section 4 is decidable.*

Specifically, to check satisfiability of $B(F_1, \dots, F_n)$, compute $B(\rho_1(F_1), \dots, \rho_n(F_n))$ and then check its satisfiability using a BAPA decision procedure [20, 21].

5 BAPA Reductions

5.1 Monadic Second-Order Logic of Finite Trees

Figure 4 shows the syntax of (our presentation of) monadic second-order logic of finite trees (FT), a variant of weak monadic second-order logic of two successors

$$\begin{aligned}
F &::= P \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid \neg F \mid \forall x:s.F \mid \exists x:s.F \\
s &::= \mathbf{obj} \mid \mathbf{set} \\
P &::= B_1 = B_2 \mid B_1 \subseteq B_2 \mid r(x, y) \\
r &::= \mathbf{succ}_L \mid \mathbf{succ}_R \\
B &::= x \mid \epsilon \mid \emptyset \mid \mathbf{Univ} \mid \{x\} \mid B_1 \cup B_2 \mid B_1 \cap B_2 \mid B^c
\end{aligned}$$

Fig. 4. Monadic Second-Order Logic of Finite Trees (FT)

(WS2S) [18,31]. The following are the sorts of variables specific to FT formulas: $\mathbf{succ}_L, \mathbf{succ}_R : \mathbf{obj}^2 \rightarrow \mathbf{bool}$.

We interpret the sort \mathbf{obj} over finite, prefix-closed sets of binary strings. More precisely, we use $\{1, 2\}$ as the binary alphabet, and we let $\llbracket \mathbf{obj} \rrbracket \subset \{1, 2\}^*$ such that

$$\forall w \in \{1, 2\}^*. (w1 \in \llbracket \mathbf{obj} \rrbracket \vee w2 \in \llbracket \mathbf{obj} \rrbracket) \rightarrow w \in \llbracket \mathbf{obj} \rrbracket$$

In each model, $\llbracket \mathbf{set} \rrbracket$ is the set of all subsets of $\llbracket \mathbf{obj} \rrbracket$. We let $\llbracket \epsilon \rrbracket$ be the empty string which we also denote by ϵ . We define

$$\llbracket \mathbf{succ}_L \rrbracket = \{(w, w1) \mid w1 \in \llbracket \mathbf{obj} \rrbracket\} \quad \text{and} \quad \llbracket \mathbf{succ}_R \rrbracket = \{(w, w2) \mid w2 \in \llbracket \mathbf{obj} \rrbracket\}$$

The remaining constants and operations on sets are interpreted as in BAPA.

Let \mathcal{F}_{FT} be the set of all formulas in Figure 4. Let \mathcal{M}_{FT} be the set of all (finite) structures described above. We define \mathcal{T}_{FT} as the set of all formulas $F \in \mathcal{F}_{\text{FT}}$ such that F is true in all structures from \mathcal{M}_{FT} .

Note that any FT formula $F(x)$ with a free variable x of sort \mathbf{obj} can be transformed into the equisatisfiable formula $\exists x : \mathbf{obj}. y = \{x\} \wedge F(x)$ where y is a fresh variable of sort \mathbf{set} . For conciseness of presentation, in the rest of this section we only consider FT formulas F with $\mathbf{FV}_{\mathbf{obj}}(F) = \emptyset$.

Finite tree automata. In the following, we recall the connection between FT formulas and finite tree automata. Let Σ be a finite ranked alphabet. We call symbols of rank 0 constant symbols and a symbol of rank $k > 0$ a k -ary function symbol. We denote by $\mathbf{Terms}(\Sigma)$ the set of all terms over Σ . We associate a position $p \in \{1, \dots, r_{\max}\}^*$ with each subterm in a term t where r_{\max} is the maximal rank of all symbols in Σ . We denote by $t[p]$ the topmost symbol of the subterm at position p . For instance, consider the term $t = f(g(a, b, c), a)$ then we have $t[\epsilon] = f$ and $t[13] = c$.

A finite (deterministic bottom-up) tree automaton A for alphabet Σ is a tuple (Q, Q_f, ι) where Q is a finite set of states, $Q_f \subseteq Q$ is a set of final states, and ι is a function that associates with each constant symbol $c \in \Sigma$ a state $\iota(c) \in Q$ and with each k -ary function symbol $f \in \Sigma$ a function $\iota(f) : Q^k \rightarrow Q$. We homomorphically extend ι from symbols in Σ to Σ -terms. We say that A accepts a term $t \in \mathbf{Terms}(\Sigma)$ if $\iota(t) \in Q_f$. The language $\mathcal{L}(A)$ accepted by A is the set of all Σ -terms accepted by A .

Let F be an FT formula and let $\mathbf{SV}(F)$ be the set $\mathbf{SV}(F) = \mathbf{FV}(F) \cup \{\mathbf{Univ}\}$. We denote by Σ_F the alphabet consisting of the constant symbol \perp and all binary function symbols f_ν where ν is a function $\nu : \mathbf{SV}(F) \rightarrow \{0, 1\}$. We inductively

associate a Σ_F -term $t_{\alpha,w}$ with every structure $\alpha \in \mathcal{M}_{\text{FT}}$ and string $w \in \{1, 2\}^*$ as follows:

$$t_{\alpha,w} = \begin{cases} f_{\nu_{\alpha,w}}(t_{\alpha,w1}, t_{\alpha,w2}) & \text{if } w \in \alpha(\text{obj}) \\ \perp & \text{otherwise} \end{cases}$$

such that for all $x \in \text{SV}(F)$, $\nu_{\alpha,w}(x) = 1$ iff $w \in \alpha(x)$. The language $\mathcal{L}(F) \subseteq \text{Terms}(\Sigma_F)$ of F is then defined by $\mathcal{L}(F) = \{t_{\alpha,\epsilon} \mid \alpha \in \mathcal{M}_{\text{FT}} \wedge \alpha(F)\}$.

Parikh image. We recall Parikh's commutative image [26]. The Parikh image for an alphabet Σ is the function $\text{Parikh} : \Sigma^* \rightarrow \Sigma \rightarrow \mathbb{N}$ such that for any word $w \in \Sigma^*$ and symbol $\sigma \in \Sigma$, $\text{Parikh}(w)(\sigma)$ is the number of occurrences of σ in w . The Parikh image is extended pointwise from words to sets of words: $\text{Parikh}(W) = \{\text{Parikh}(w) \mid w \in W\}$. In the following, we implicitly identify $\text{Parikh}(W)$ with the set of integer vectors $\{(\chi(\sigma_1), \dots, \chi(\sigma_n)) \mid \chi \in \text{Parikh}(W)\}$ where we assume some fixed order on the symbols $\sigma_1, \dots, \sigma_n$ in Σ .

We generalize the Parikh image from words to terms as expected: the Parikh image for a ranked alphabet Σ is the function $\text{Parikh} : \text{Terms}(\Sigma) \rightarrow \Sigma \rightarrow \mathbb{N}$ such that for all $t \in \text{Terms}(\Sigma)$ and $\sigma \in \Sigma$, $\text{Parikh}(t)(\sigma)$ is the number of positions p in t such that $t[p] = \sigma$. Again we extend this function pointwise from terms to sets of terms.

BAPA Reduction. In the following, we prove the existence of a computable BAPA reduction for the theory of monadic second-order logic of finite trees.

Let F be an FT formula and let Σ_F^2 be the set of all binary function symbols in Σ_F , i.e., $\Sigma_F^2 \stackrel{\text{def}}{=} \Sigma_F \setminus \{\perp\}$. We associate with each $\sigma_\nu \in \Sigma_F^2$ the *Venn region* $\text{vr}(\sigma_\nu)$, which is given by a set-algebraic expression over $\text{SV}(F)$: let $\text{SV}(F) = \{x_1, \dots, x_n\}$ then

$$\text{vr}(\sigma_\nu) \stackrel{\text{def}}{=} x_1^{\nu(x_1)} \cap \dots \cap x_n^{\nu(x_n)} .$$

Hereby x_i^0 denotes x_i^c and x_i^1 denotes x_i . Let $\alpha \in \mathcal{M}_{\text{FT}}$ be a model of F . Then the term $t_{\alpha,\epsilon}$ encodes for each $w \in \alpha(\text{obj})$ the Venn region to which w belongs in α , namely $\text{vr}(t_{\alpha,\epsilon}[w])$. Thus, the Parikh image $\text{Parikh}(t_{\alpha,\epsilon})$ encodes the cardinality of each Venn region over $\text{SV}(F)$ in α .

Lemma 4. *Let F be an FT formula then*

$$\text{Parikh}(\mathcal{L}(F))|_{\Sigma_F^2} = \left\{ \left\{ \sigma \mapsto |\alpha(\text{vr}(\sigma))| \mid \sigma \in \Sigma_F^2 \right\} \mid \alpha \in \mathcal{M}_{\text{FT}} \wedge \alpha(F) \right\} .$$

Following [31, Theorem 17] we can construct a finite tree automaton A_F over Σ_F such that $\mathcal{L}(F) = \mathcal{L}(A_F)$. From [26, Theorem 2] follows that $\text{Parikh}(\mathcal{L}(F))$ is a semilinear set whose finite representation in terms of base and step vectors is effectively computable from A_F . From this finite representation we can construct a Presburger arithmetic formula ϕ_F over free integer variables $\{x_\sigma \mid \sigma \in \Sigma_F\}$ whose set of solutions is the Parikh image of $\mathcal{L}(F)$, i.e.

$$\text{Parikh}(\mathcal{L}(F)) = \left\{ \left\{ \sigma \mapsto k_\sigma \mid \sigma \in \Sigma_F \right\} \mid \left\{ x_\sigma \mapsto k_\sigma \mid \sigma \in \Sigma \right\} (\phi_F) \right\} \quad (4)$$

Using the above construction of the Presburger arithmetic formula ϕ_F for a given FT formula F , we define the function $\rho_{\text{FT}} : \mathcal{F}_{\text{FT}} \rightarrow \mathcal{F}_{\text{BAPA}}$ as follows:

$$\rho_{\text{FT}}(F) \stackrel{\text{def}}{=} \exists \mathbf{x}_\sigma. \phi_F \wedge \bigwedge_{\sigma \in \Sigma_F^2} \text{card}(\text{vr}(\sigma)) = x_\sigma$$

where x_σ are the free integer variables of ϕ_F .

Theorem 5. *The function $\rho_{\mathcal{F}_T}$ is a BAPA reduction for $(\mathcal{F}_T, \mathcal{T}_T)$.*

5.2 Two-Variable Logic with Counting

Figure 5 shows the syntax of (our presentation of) two-variable logic with counting (denoted C^2) [29]. As usual in C^2 , we require that every sub-formula of a formula has at most two free variables. In the atomic formula $r(x_1, x_2)$, variables x_1, x_2 are of sort \mathbf{obj} and r is a relation variable of sort $\mathbf{obj}^2 \rightarrow \mathbf{bool}$. The formula $\{x\} \subseteq A$ replaces $A(x)$ in predicate-logic notation, and has the expected meaning, with the variable x is of sort \mathbf{obj} and A of sort \mathbf{set} . The interpretation of the counting quantifier $\exists^K x:\mathbf{obj}.F$ for a positive constant K is that there exist at least K distinct elements x for which the formula F holds.

$$\begin{aligned} F &::= P \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid \neg F \mid \exists^K x:\mathbf{obj}.F \\ P &::= x_1 = x_2 \mid \{x\} \subseteq A \mid r(x_1, x_2) \end{aligned}$$

Fig. 5. Two-Variable Logic with Counting (C^2)

Let \mathcal{F}_{C^2} be the set of all formulas in Figure 5. Let \mathcal{M}_{C^2} be the set of structures that interpret formulas in \mathcal{F}_{C^2} . We define \mathcal{T}_{C^2} as the set of all formulas $F \in \mathcal{F}_{C^2}$ such that F is true in all structures from \mathcal{M}_{C^2} . Modulo our minor variation in syntax and terminology (using relation and set variables instead of predicate symbols), \mathcal{T}_{C^2} corresponds to the standard set of valid C^2 formulas over finite structures.

BAPA Reduction. We next build on the results in [29] to define a BAPA reduction for C^2 . We fix set variables A_1, \dots, A_p and relation variables r_1, \dots, r_q . Throughout this section, let $\Sigma_A = \{A_1, \dots, A_p\}$, $\Sigma_R = \{r_1, \dots, r_q\}$, and $\Sigma_0 = \Sigma_A \cup \Sigma_R$. We call $\Sigma_A, \Sigma_R, \Sigma_0$ signatures because they correspond to the notion of signature in the traditional first-order logic formulation of C^2 .

Model theoretic types. Define the model-theoretic notion of n -type $\pi_\Sigma(x_1, \dots, x_n)$ in the signature Σ as the maximal consistent set of non-equality literals in Σ whose \mathbf{obj} -sort variables are included in $\{x_1, \dots, x_n\}$. Given a structure α such that $\alpha(x_1), \dots, \alpha(x_n)$ are all distinct, α induces an n -type. We also define the set of n -tuples for which a type π holds in a structure α :

$$S^\alpha(\pi(x_1, \dots, x_n)) = \{(e_1, \dots, e_n) \in \alpha(\mathbf{obj})^n \mid \alpha(x_1 := e_1, \dots, x_n := e_n)(\pi)\}$$

If $\Sigma \subseteq \Sigma'$ and π' is an n -type in signature Σ' , by $\pi'|_\Sigma$ we denote the subset of π containing precisely those literals from π whose sets and relations belong to Σ . The family of sets $\{S^\alpha(\pi') \mid \pi'|_\Sigma = \pi\}$ is a partition of $S^\alpha(\pi')$. We will be particularly interested in 1-types. We identify a 1-type $\pi(x)$ in the signature Σ_A with the corresponding Venn region

$$\bigcap \{A_i \mid (\{x\} \subseteq A_i) \in \pi(x)\} \cap \bigcap \{A_i^c \mid (\neg(\{x\} \subseteq A_i)) \in \pi(x)\}.$$

If π_1, \dots, π_m is the sequence of all 1-types in the signature Σ and α is a structure, let $I^\alpha(\Sigma) = (|S^\alpha(\pi_1)|, \dots, |S^\alpha(\pi_m)|)$. If \mathcal{M} is a set of structures let $I^\mathcal{M}(\Sigma) = \{I^\alpha(\Sigma) \mid \alpha \in \mathcal{M}\}$.

Observation 6 *If π is a 1-type in Σ and π' a 1-type in Σ' for $\Sigma \subseteq \Sigma'$, then*

$$|I^\alpha(\pi)| = \sum_{\pi'|_{\Sigma}=\pi} |I^\alpha(\pi')|$$

Making structures differentiated, chromatic, sparse preserves 1-types.

Let ϕ be a C^2 formula with signature Σ_0 of relation symbols. By Scott normal form transformation [29, Lemma 1] it is possible to introduce fresh set variables and compute another C^2 formula ϕ^* in an extended signature $\Sigma^* \supseteq \Sigma_0$, and compute a constant C_ϕ such that, for all sets u with $|u| \geq C_\phi$: 1) if α_0 is a Σ_0 interpretation with domain u such that $\alpha_0(\phi)$, then there exists its Σ^* extension $\alpha^* \supseteq \alpha_0$ such that $\alpha^*(\phi^*)$, and 2) if α^* is a Σ^* interpretation with domain u such that $\alpha^*(\phi^*)$, then for its restriction $\alpha_0 = \alpha^*|_{\Sigma}$ we have $\alpha_0(\phi)$. By introducing further fresh set- and relation- symbols, [29, lemmas 2 and 3] shows that we can extend the signature from Σ^* to Σ such that each model α^* in Σ^* extends to a model α in Σ , where α satisfies some further conditions of interest: α is *chromatic* and *differentiated*. [29, Lemma 10] then shows that it is possible to transform a model of a formula into a so-called *X-sparse* model for an appropriately computed integer constant X . What is important for us is the following.

Observation 7 *The transformations that start from α_0 with $\alpha_0(\phi)$, and that produce a chromatic, differentiated, X-sparse structure α with $\alpha(\phi)$, have the property that, for structures of size C_ϕ or more,*

1. *the domain remains the same: $\alpha_0(\text{obj}) = \alpha(\text{obj})$,*
2. *the induced 1-types in the signature Σ_0 remain the same: for each 1-type π in signature Σ_0 , $S^{\alpha_0}(\pi) = S^\alpha(\pi)$.*

Star types. [29, Definition 9] introduces a star-type (π, \mathbf{v}) (denoted by letter σ) as a description of a local neighborhood of a domain element, containing its induced 1-type π as well as an integer vector $\mathbf{v} \subseteq \mathbb{Z}^N$ that counts 2-types in which the element participates, where N is a function of the signature Σ . A star type thus gives a more precise description of the properties of a domain element than a 1-type. Without repeating the definition of star type [29, Definition 9], we note that we can similarly define the set $S^\alpha((\pi, \mathbf{v}))$ of elements that realize a given star type (π, \mathbf{v}) . Moreover, for a given 1-type π , the family of the non-empty among the sets $S^\alpha((\pi, \mathbf{v}))$ partitions the set $S^\alpha(\pi)$.

Frames. The notion of *Y-bounded chromatic frame* [29, Definition 11] can be thought of as a representation of a disjunct in a normal form for the formula ϕ^* . It summarizes the properties of elements in the structure and specifies (among others), the list of possible star types $\sigma_1, \dots, \sigma_N$ whose integer vectors \mathbf{v} are bounded by Y . For a given ϕ^* , it is possible to effectively compute the set of

C_ϕ -bounded frames \mathcal{F} such that $\mathcal{F} \models \phi^*$ holds. The ‘ \models ’ in $\mathcal{F} \models \phi^*$ is a certain syntactic relation defined in [29, Definition 13].

For each frame \mathcal{F} with star-types $\sigma_1, \dots, \sigma_N$, [29, Definition 14] introduces an effectively computable Presburger arithmetic formula $P_{\mathcal{F}}$ with N free variables. We write $P_{\mathcal{F}}(w_1, \dots, w_N)$ if $P_{\mathcal{F}}$ is true when these variables take the values w_1, \dots, w_N . The following statement is similar to the main [29, Theorem 1], and can be directly recovered from its proof and the proofs of the underlying [29, lemmas 12,13,14].

Theorem 8. *Given a formula ϕ^* , and the corresponding integer constant C_ϕ , there exists a computable constant X such that if $N \leq X$, if $\sigma_1, \dots, \sigma_N$ is a sequence of star types in Σ whose integer vectors are bounded by C_ϕ , and w_1, \dots, w_N are integers, then the following are equivalent:*

1. *There exists a chromatic differentiated structure α such that $\alpha(\phi^*)$, $w_i = |S^\alpha(\sigma_i)|$ for $1 \leq i \leq N$, and $\alpha(\text{obj}) = \bigcup_{i=1}^N S^\alpha(\sigma_i)$.*
2. *There exists a chromatic frame \mathcal{F} with star types $\sigma_1, \dots, \sigma_N$, such that $\mathcal{F} \models \phi^*$ and $P_{\mathcal{F}}(w_1, \dots, w_N)$.*

We are now ready to describe our BAPA reduction. Fix V_1, \dots, V_M to be the list of all 1-types in signature Σ_A ; let s_1, \dots, s_M be variables corresponding to their counts. By the transformation of models into chromatic, differentiated, X -sparse ones, the observations 7, 6, and Theorem 8, we obtain

Corollary 9. *If $\mathcal{M} = \{\alpha \mid \alpha(\phi^*)\}$, then there is a computable constant X such that $I^{\mathcal{M}}(\Sigma_A) = \{(s_1, \dots, s_M) \mid F_{\phi^*}(s_1, \dots, s_M)\}$ where $F_{\phi^*}(s_1, \dots, s_M)$ is the following Presburger arithmetic formula*

$$\bigvee_{N, \sigma_1, \dots, \sigma_N, \mathcal{F}} \exists w_1, \dots, w_N. P_{\mathcal{F}}(w_1, \dots, w_N) \wedge \bigwedge_{j=1}^M s_j = \sum \{w_i \mid V_j = (\pi_i \upharpoonright_{\Sigma_A})\}$$

where N ranges over $\{0, 1, \dots, X\}$, $\sigma_1, \dots, \sigma_N$ range over sequences of C_ϕ -bounded star types, and where \mathcal{F} ranges over the C_ϕ -bounded frames with star types $\sigma_1, \dots, \sigma_N$ such that $\mathcal{F} \models \phi^*$.

By adjusting for the small structures to take into account Scott normal form transformation, we further obtain

Corollary 10. *If $\mathcal{M} = \{\alpha \mid \alpha(\phi)\}$, then $I^{\mathcal{M}}(\Sigma_A) = \{(s_1, \dots, s_M) \mid G_\phi(s_1, \dots, s_M)\}$ where $G_\phi(s_1, \dots, s_M)$ is the Presburger arithmetic formula*

$$\begin{aligned} & \sum_{i=1}^M s_i \geq C_\phi \wedge F_{\phi^*}(s_1, \dots, s_M) \quad \bigvee \\ & \bigvee \{ \bigwedge_{i=1}^M s_i = d_i \mid \exists \alpha. |\alpha(\text{obj})| < C_\phi \wedge (d_1, \dots, d_M) \in I^\alpha(\Sigma_A) \} \end{aligned}$$

Theorem 11. *The following is a BAPA reduction for C^2 over finite models to variables Σ_A : given a two-variable logic formula ϕ , compute the BAPA formula $\exists s_1, \dots, s_M. G_\phi(s_1, \dots, s_M) \wedge \bigwedge_{i=1}^M \text{card}(V_i) = s_i$.*

5.3 Bernays-Schönfinkel-Ramsey Fragment of First-Order Logic

Figure 6 shows the syntax of (our presentation of) the Bernays-Schönfinkel-Ramsey fragment of first-order logic with equality [6], often called effectively propositional logic (EPR). The interpretation of atomic formulas is analogous as for C^2 in previous section. Quantification is restricted to variables of sort `obj` and must obey the usual restriction of $\exists^*\forall^*$ -prenex form that characterizes the Bernays-Schönfinkel-Ramsey class.

$$\begin{aligned}
 F &::= \exists z_1:\text{obj} \dots \exists z_n:\text{obj} \cdot \forall y_1:\text{obj} \dots \forall y_m:\text{obj} \cdot B \\
 B &::= P \mid B_1 \wedge B_2 \mid B_1 \vee B_2 \mid \neg B \\
 P &::= x_1 = x_2 \mid \{x\} \subseteq A \mid r(x_1, \dots, x_k)
 \end{aligned}$$

Fig. 6. Bernays-Schönfinkel-Ramsey Fragment of First-Order Logic

BAPA Reduction. Our BAPA reduction for the Bernays-Schönfinkel-Ramsey fragment (EPR) is in fact a reduction from EPR formulas to *unary* EPR formulas, in which all free variables have the sort `set`. To convert a unary EPR formula into BAPA, treat first-order variables as singleton sets and apply quantifier elimination for BAPA [20].

Theorem 12 (BAPA Reduction for EPR). *Let ϕ be a quantifier-free formula whose free variables are: 1) A_1, \dots, A_p , of sort `set`, 2) r_1, \dots, r_q , each r_i of sorts `obj` ^{$K(i)$} \rightarrow `bool` for some $K(i) \geq 2$, 3) $z_1, \dots, z_n, y_1, \dots, y_m$, of sort `obj`. Then*

$$\exists r_1, \dots, r_q \cdot \exists z_1, \dots, z_n \cdot \forall y_1, \dots, y_m \cdot \phi$$

is equivalent to an effectively computable BAPA formula.

The proof of Theorem 12 builds on and generalizes, for finite models, the results on the spectra of EPR formulas [10, 11, 30]. We here provide some intuition. The key insight [30] is that, when a domain of a model of an EPR formula has sufficiently many elements, then the model contains an induced submodel S of m nodes such that for every $0 \leq k < m$ elements e_1, \dots, e_k outside S the m -type induced by e_1, \dots, e_k and any $m - k$ elements in S is the same. Then an element of S can be replicated to create a model with more elements, without changing the set of all m -types in the model and thus without changing the truth value of the formula. Moreover, every sufficiently large model of the EPR formula that has a submodel S with more than m such symmetric elements can be shrunk to a model by whose expansion it can be generated. This allows us to enumerate a finite (even if very large) number of characteristic models whose expansion generates all models. The expansion of a characteristic model increases by one the number of elements of some existing 1-type, so the cardinalities of Venn regions of models are a semilinear set whose base vectors are given by characteristic models and whose step vectors are given by the 1-types being replicated.

5.4 Quantifier-free Multisets with Cardinality Constraints

The satisfiability of the quantifier-free fragment of multisets with cardinality operators is decidable [27]. There is, in fact, also a BAPA reduction from a quantifier-free multiset formula over multiset and set variables to a BAPA formula ranging only over the set variables. Multiset formula are built from multiset variables, multiset operations, and cardinality constraints on multisets. Multiset variables have sort $\text{obj} \rightarrow \text{int}$ and are interpreted as functions from $\llbracket \text{obj} \rrbracket$ to the nonnegative integers (for details see [35]).

Let F be a multiset constraint containing sets A_1, \dots, A_p and multisets M_1, \dots, M_q . To obtain a BAPA reduction, we apply the decision procedure in [27] to the formula $F \wedge \bigwedge_{i=1}^w \text{card}(V_i) = k_i$ with fresh integer variables k_1, \dots, k_w , where V_1, \dots, V_w are the Venn regions over sets A_1, \dots, A_p . The result is a Presburger arithmetic formula P with $\{k_1, \dots, k_w\} \subseteq \text{FV}(P)$. If x_1, \dots, x_n are the variables in P other than k_1, \dots, k_w , the result of the BAPA reduction is then the formula

$$P_F \stackrel{\text{def}}{=} \exists k_1 \dots \exists k_w. \left(\bigwedge_{i=1}^w \text{card}(V_i) = k_i \right) \wedge (\exists x_1 : \text{int} \dots \exists x_n : \text{int}. P)$$

Theorem 13. *The function mapping a multiset formula F to the BAPA formula P_F is a BAPA reduction for the theory of multisets with cardinalities.*

6 Further Related Work

There are combination results for the disjoint combinations of non-stably infinite theories [10, 11, 19, 33]. These results are based on the observation that such combinations are possible whenever one can decide for each component theory whether a model of a specific cardinality exists. Our combination result takes into account not only the cardinality of the models, i.e. the interpretation of the universal set, but cardinalities of Venn regions over the interpretations of arbitrary shared set variables. It is a natural generalization of the disjoint case restricted to theories that share the theory of finite sets, thus, leading to a non-disjoint combination of non-stably infinite theories.

Ghilardi [14] proposes a model-theoretic condition for decidability of the non-disjoint combination of theories based on quantifier elimination and local finiteness of the shared theory. Note that BAPA is not locally finite and that, in general, we need the full expressive power of BAPA to compute the projections on the shared set variables. For instance, consider the C^2 formula

$$(\forall x. \exists^{=1} y. r(x, y)) \wedge (\forall x. \exists^{=1} y. r(y, x)) \wedge (\forall y. y \in B \leftrightarrow (\exists x. x \in A \wedge r(x, y)))$$

where r is a binary relation variable establishing the bijection between A and B . This constraint expresses $|A| = |B|$ without imposing any additional constraint on A and B . Similar examples can be given for weak monadic second-order logic of finite trees.

The reduction approach to combination of decision procedures has previously been applied in the simpler scenario of reduction to propositional logic [22].

Like propositional logic, quantifier-free BAPA is NP-complete, so it presents an appealing alternative for combination of theories that share sets.

Gabbay and Ohlbach [12] present a procedure, called SCAN, for second-order quantifier elimination. However, [12] gives no characterization of when SCAN terminates. We were therefore unable to use SCAN to derive any BAPA reductions.

The general combination of weak monadic second-order logics with linear cardinality constraints has been proven undecidable by Klaedtke and Rueß [16, 17]. They introduce the notion of Parikh automata to identify decidable fragments of this logic which inspired our BAPA reduction of MSOL of finite trees. Our combined logic is incomparable to the decidable fragments identified by Klaedtke and Rueß because it supports non-tree structures as well. However, by applying projection to C^2 and the Bernays-Schönfinkel-Ramsey class, we can combine our logic with [16, 17], obtaining an even more expressive decidable logic.

7 Conclusion

We have presented a combination result for logics that share operations on sets. This result yields an expressive decidable logic that is useful for software verification. We therefore believe that we made an important step in increasing the class of properties that are amenable to automated verification.

References

1. P. B. Andrews. *An Introduction to Mathematical Logic and Type Theory: To Truth Through Proof*. Springer (Kluwer), 2nd edition, 2002.
2. T. Ball, A. Podelski, and S. K. Rajamani. Relative completeness of abstraction refinement for software model checking. In *TACAS'02*, 2002.
3. M. Barnett, R. DeLine, M. Fähndrich, K. R. M. Leino, and W. Schulte. Verification of object-oriented programs with invariants. *Journal of Object Technology*, 3(6):27–56, 2004.
4. C. Barrett and C. Tinelli. CVC3. In *CAV*, volume 4590 of *LNCS*, 2007.
5. D. Basin and S. Friedrich. Combining WS1S and HOL. In *FroCoS*, 1998.
6. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer-Verlag, 1997.
7. R. S. Boyer and J. S. Moore. Integrating decision procedures into heuristic theorem provers: A case study of linear arithmetic. In *Machine Intelligence*, volume 11. Oxford University Press, 1988.
8. L. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
9. S. Feferman and R. L. Vaught. The first order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.
10. P. Fontaine. Combinations of theories and the bernays-schönfinkel-ramsey class. In *VERIFY*, 2007.
11. P. Fontaine. Combinations of decidable fragments of first-order logic. In *FroCoS*, 2009.

12. D. M. Gabbay and H. J. Ohlbach. Quantifier elimination in second-order predicate logic. In B. Nebel, C. Rich, and W. Swartout, editors, *Principles of Knowledge Representation and Reasoning*. Morgan-Kaufmann, 1992.
13. Y. Ge, C. Barrett, and C. Tinelli. Solving quantified verification conditions using satisfiability modulo theories. In *CADE*, 2007.
14. S. Ghilardi. Model theoretic methods in combined constraint satisfiability. *Journal of Automated Reasoning*, 33(3-4):221–249, 2005.
15. S. Ginsburg and E. Spanier. Semigroups, Pressburger formulas and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.
16. F. Klaedtke and H. Rueß. Parikh automata and monadic second-order logics with linear cardinality constraints. Technical Report 177, Institute of Computer Science at Freiburg University, 2002.
17. F. Klaedtke and H. Rueß. Monadic second-order logics with cardinalities. In *ICALP*, volume 2719 of *LNCS*, 2003.
18. N. Klarlund and A. Møller. *MONA Version 1.4 User Manual*. BRICS Notes Series NS-01-1, Department of Computer Science, University of Aarhus, January 2001.
19. S. Krstic, A. Goel, J. Grundy, and C. Tinelli. Combined satisfiability modulo parametric theories. In *TACAS*, volume 4424 of *LNCS*, pages 602–617, 2007.
20. V. Kuncak, H. H. Nguyen, and M. Rinard. Deciding Boolean Algebra with Presburger Arithmetic. *J. of Automated Reasoning*, 2006.
21. V. Kuncak and M. Rinard. Towards efficient satisfiability checking for Boolean Algebra with Presburger Arithmetic. In *CADE-21*, 2007.
22. S. K. Lahiri and S. A. Seshia. The UCLID decision procedure. In *CAV'04*, 2004.
23. S. McLaughlin, C. Barrett, and Y. Ge. Cooperating theorem provers: A case study combining HOL-Light and CVC Lite. In *PDPAR*, volume 144(2) of *ENTCS*, 2006.
24. G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM TOPLAS*, 1(2):245–257, 1979.
25. S. Owre, J. M. Rushby, and N. Shankar. PVS: A prototype verification system. In D. Kapur, editor, *11th CADE*, volume 607 of *LNAI*, pages 748–752, jun 1992.
26. R. J. Parikh. On context-free languages. *J. ACM*, 13(4):570–581, 1966.
27. R. Piskac and V. Kuncak. Decision procedures for multisets with cardinality constraints. In *VMCAI*, number 4905 in *LNCS*, 2008.
28. R. Piskac and V. Kuncak. Linear arithmetic with stars. In *CAV*, 2008.
29. I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395, 2005.
30. F. P. Ramsey. On a problem of formal logic. *Proc. London Math. Soc.*, s2-30:264–286, 1930. doi:10.1112/plms/s2-30.1.264.
31. J. W. Thatcher and J. B. Wright. Generalized finite automata theory with an application to a decision problem of second-order logic. *Mathematical Systems Theory*, 2(1):57–81, 1968.
32. C. Tinelli and C. Ringeissen. Unions of non-disjoint theories and combinations of satisfiability procedures. *Th. Comp. Sc.*, 290(1):291–353, Jan. 2003.
33. C. Tinelli and C. Zarba. Combining nonstably infinite theories. *Journal of Automated Reasoning*, 34(3), 2005.
34. T. Wies. *Symbolic Shape Analysis*. PhD thesis, University of Freiburg, 2009.
35. T. Wies, R. Piskac, and V. Kuncak. On Combining Theories with Shared Set Operations. Technical Report LARA-REPORT-2009-002, EPFL, May 2009.
36. K. Zee, V. Kuncak, and M. Rinard. Full functional verification of linked data structures. In *ACM Conf. Programming Language Design and Implementation (PLDI)*, 2008.