


Counting Constraints in Flat Array Fragments

Francesco Alberti¹, Silvio Ghilardi², and Elena Pagani²

¹ Fondazione Centro San Raffaele, Milano, Italy

² Università Degli Studi di Milano, Milano, Italy
silvio.ghilardi@unimi.it

Abstract. We identify a fragment of Presburger arithmetic enriched with free function symbols and cardinality constraints for interpreted sets, which is amenable to automated analysis. We establish decidability and complexity results for such a fragment and we implement our algorithms. The experiments run in discharging proof obligations coming from invariant checking and bounded model-checking benchmarks show the practical feasibility of our decision procedure.

1 Introduction

Enriching logic formalisms with counting capabilities is an important task in view of the needs of many application areas, ranging from database theory to formal verification. Such enrichments have been designed both in the description logics area and in the area of Satisfiability Modulo Theories (SMT), where some of the most important recent achievements were decidability and complexity bounds for BAPA [14] - the enrichment of Presburger arithmetic with the ability of talking about finite sets and their cardinalities. As pointed out in [15], BAPA constraints can be used for program analysis and verification by expressing data structure invariants, simulations between program fragments or termination conditions. The analysis of BAPA constraints was successfully extended also to formalisms encompassing multisets [18] as well as direct/inverse images along relations and functions [23].

A limitation of BAPA and its extensions lies in the fact that only uninterpreted symbols (for sets, relations, functions, etc.) are allowed. On the other hand, it is well-known that a different logical formalism, namely unary counting quantifiers, can be used in order to reason about the cardinality of definable (i.e. of interpreted) sets. Unary counting quantifiers can be added to Presburger arithmetic without compromising decidability, see [19], however they might be quite problematic if combined in an unlimited way with free function symbols. In this paper, we investigate the extension of Presburger arithmetic including both counting quantifiers and uninterpreted function symbols, and we isolate fragments where we can achieve decidability and in some cases also relatively good complexity bounds. The key ingredient to isolate such fragments is the notion of flatness: roughly, in a flat formula, subterms of the kind $a(t)$ (where a is a free function symbol) can occur only if t is a variable. By itself, this naive

flatness requirement is useless (any formula can match it to the price of introducing extra quantified variables); in order to make it effective, further syntactic restrictions should be incorporated in it, as witnessed in [2]. This is what we are going to do in this paper, where suitable notions of ‘flat’ and ‘simple flat’ formulæ are introduced in the rich context of Presburger arithmetic enriched with free function symbols and with unary counting quantifiers (we use free function symbols to model arrays, see below).

The fragments we design are all obviously more expressive than BAPA, but they do not come from pure logic motivations, on the contrary they are suggested by an emerging application area, namely the area of verification of fault-tolerant distributed systems. Such systems (see [8] for a good account) are modeled as partially synchronous systems, where a finite number of identical processes operate in lock-step (in each *round* they send messages, receive messages, and update their local state depending on the local state at the beginning of the round and the received messages). Messages can be lost, processes may omit to perform some tasks or also behave in a malicious way; for these reasons, the fact that some actions are enabled or not, and the correctness of the algorithms themselves, are subject to threshold conditions saying for instance that some qualified majority of processes are in a certain status or behave in a non-faulty way. Verifications tasks thus have to handle cardinality constraints of the kind studied in this paper (the reader interested in full formalization examples can directly go to Sect. 5).

The paper is organized as follows: we first present basic syntax (Sect. 2), then decidability (Sect. 3) and complexity (Sect. 4) results; experiments with our prototypical implementation are supplied in Sect. 5, and Sect. 6 concludes the work.

2 Preliminaries

We work within Presburger arithmetic enriched with free function symbols and cardinality constraints. This is a rather expressive logic, whose syntax is summarized in Fig. 1. Terms and formulæ are interpreted in the natural way over the domain of integers \mathbb{Z} ; as a consequence, satisfiability of a formula ϕ means that it is possible to assign values to parameters, free variables and array-ids so as to make ϕ true in \mathbb{Z} (validity of ϕ means that $\neg\phi$ is not satisfiable, equivalence of ϕ and ψ means that $\phi \leftrightarrow \psi$ is valid, etc.). We nevertheless implicitly assume few constraints (to be explained below) about our intended semantics.

To denote integer numbers, we have (besides variables and numerals) also parameters: the latter denote unspecified integers. Among parameters, we always include a specific parameter (named N) identifying the dimension of the system - alias the length of our arrays: in other words, it is assumed that for all array identifiers $a \in Arr$, the value $a(x)$ is conventional (say, zero) outside the interval $[0, N) = \{n \in \mathbb{Z} \mid 0 \leq n < N\}$. Although binary free function symbols are quite useful in some applications, in this paper we prefer not to deal with them. The operator $\#\{x \mid \phi\}$ indicates the cardinality of the finite set formed by the $x \in [0, N)$ such that $\phi(x)$ holds.

$0, 1, \dots$	$\in \mathbb{Z}$	numerals (numeric constants)
x, y, z, \dots	$\in Var$	individual variables
M, N, \dots	$\in Par$	parameters (free constants)
a, b, \dots	$\in Arr$	array ids (free unary function symbols)
t, u, \dots	$::= n \mid M \mid x \mid t + t \mid -t \mid a(t) \mid \#\{x \mid \phi\}$	terms
A, B, \dots	$::= t < t \mid t = t \mid t \equiv_n t$	atoms
ϕ, ψ, \dots	$::= A \mid \phi \wedge \phi \mid \neg\phi \mid \exists x \phi$	formulae

Fig. 1. Syntax

Notice that the cardinality constraint operator $\#\{x \mid -\}$, as well the quantifier $\exists x$, bind the variable x ; below, we indicate with $\psi(\underline{x})$ (resp. $t(\underline{x})$) the fact that the formula ψ (the term t) has free individual variables included in the list \underline{x} . When we speak of a substitution, we always mean ‘substitution without capture’, meaning that, when we replace the free occurrences of a variable x with a term u in a formula ϕ or in a term t , the term u should not contain free variables that might be located inside the scope of a binder for them once the substitution is performed; the result of the substitution is denoted with $\phi(u/x)$ and $t(u/x)$.

The logic of Fig. 1 is far from being tractable, because even the combination of free function symbols and Presburger arithmetic lands in a highly undecidable class [10]. We are looking for a mild fragment, nevertheless sufficiently expressive for our intended applications. These applications mostly come from verification tasks, like bounded model checking or invariant checking. Our aim is to design a decidable fragment (so as to be able not only to produce certifications, but also to find bugs) with some minimal closure properties; from this point of view, notice that for bounded model checking closure under conjunctions is sufficient, but for invariant checking we need also closure under negations in order to discharge entailments.

2.1 Flat Formulae

We now introduce some useful subclasses of the formulae built up according to the grammar of Fig. 1 (all subclasses are closed under Boolean operations):

- *Arithmetic formulae* : these are built up from the grammar of Fig. 1 without using neither array-ids nor cardinality constraint operators; we use letter α, β, \dots for arithmetic formulae. Recall that, according to the well-known quantifier elimination result, arithmetic formulae are equivalent to *quantifier-free* arithmetic formulae.
- *Constraint formulae* : these are built up from the grammar of Fig. 1 without using array-ids.
- *Basic formulae* : these are obtained from an arithmetic formula by simultaneously replacing some free variables by terms of the kind $a(y)$, where y is a variable and a an array-id. When we need to display full information,

we may use the notation $\alpha(\underline{y}, \mathbf{a}(\underline{y}))$ to indicate basic formulæ. By this notation, we mean that $\underline{y} = y_1, \dots, y_n$ are variables, $\mathbf{a} = a_1, \dots, a_s$ are array-ids and that $\alpha(\underline{y}, \mathbf{a}(\underline{y}))$ is obtained from an arithmetic formula $\alpha(\underline{x}, \underline{z})$ (where $\underline{z} = z_{11}, \dots, z_{sn}$) by replacing z_{ij} with $a_i(y_j)$ ($i = 1, \dots, s$ and $j = 1, \dots, n$).

– *Flat formulæ* : these are recursively defined as follows (i) basic formulæ are flat formulæ; (ii) if ϕ is a flat formula, β is a basic formula, z and x are variables, then $\phi(\#\{x \mid \beta\} / z)$ is a flat formula. Thus in flat formulæ all dereferenced indexes are either free or the ones defining the comprehension.¹

Example 1. The formulæ $a(y) + a(z) \leq z$ and $z = \#\{x \mid \#\{x' \mid a(x') < 1\} = a(x)\}$ are flat (the former also basic) whereas $z = \#\{x \mid \#\{x' \mid a(x') < x\} = a(x)\}$ is not such (the binder $\#\{x \mid \dots$ captures a free occurrence of x in $\#\{x' \mid a(x') < x\}$).

The following result is proved in [19] (see also the Appendix of [1]):

Theorem 1. *For every constraint formula one can compute an arithmetic formula equivalent to it.*

3 Satisfiability for Flat formulæ

We shall show that flat formulæ are decidable for satisfiability. In fact, we shall show decidability of the slightly larger class covered by the following

Definition 1. *Extended flat formulæ (briefly, E-flat formulæ) are formulæ of the kind*

$$\exists \underline{z}. \alpha \wedge \#\{x \mid \beta_1\} = z_1 \wedge \dots \wedge \#\{x \mid \beta_K\} = z_K \quad (1)$$

where $\underline{z} = z_1, \dots, z_K$ and $\alpha, \beta_1, \dots, \beta_K$ are basic formulæ and x does not occur in α .

Notice that α and the β_j in (1) above may contain further free variables \underline{y} (besides \underline{z}) as well as the terms $\mathbf{a}(\underline{y})$ and $\mathbf{a}(\underline{z})$; the β_j may contain occurrences of x and of $\mathbf{a}(x)$.

That flat formulæ are also E-flat can be seen as follows: due to the fact that our substitutions avoid captures, we can use equivalences like $\phi(t/z) \leftrightarrow \exists z (t = z \wedge \phi)$ in order to abstract out the terms $t := \#\{x \mid \alpha\}$ occurring in the recursive construction of a flat formula ϕ . By repeating this linear time transformation, we end up in a formula of the kind (1). However, not all E-flat formulæ are flat because the dependency graph associated to (1) might not be acyclic (the graph we are talking about has the z_j as nodes and has an arc $z_j \rightarrow z_i$ when z_i occurs in β_j). The above conversion of a flat formula into a formula of the form (1) on the other hand produces an E-flat formula whose associated graph is acyclic.

We prove a technical lemma showing how we can manipulate E-flat formulæ without loss of generality. Formulae $\varphi_1, \dots, \varphi_K$ are said to be a partition iff the

¹ If we want to emphasize the way the basic formula β is built up, following the above conventions, we may write it as $\beta(x, \underline{y}, \mathbf{a}(x), \mathbf{a}(\underline{y}))$; here, supposing that \mathbf{a} is a_1, \dots, a_s , since x is a singleton, the tuple $\mathbf{a}(x)$ is $a_1(x), \dots, a_s(x)$.

formulae $\bigvee_{l=1}^K \varphi_l$ and $\neg(\varphi_l \wedge \varphi_h)$ (for $h \neq l$) are valid. Recall that the existential closure of a formula is the sentence obtained by prefixing it with a string of existential quantifiers binding all variables having a free occurrence in it.

Lemma 1. *The existential closure of an E-flat formula is equivalent to a sentence of the kind*

$$\exists \underline{z} \exists \underline{y}. \alpha(\underline{y}, \underline{z}) \wedge \sharp\{x \mid \beta_1(x, \mathbf{a}(x), \underline{y}, \underline{z})\} = z_1 \wedge \dots \wedge \sharp\{x \mid \beta_K(x, \mathbf{a}(x), \underline{y}, \underline{z})\} = z_K \quad (2)$$

where \underline{y} and $\underline{z} := z_1, \dots, z_K$ are variables, α is arithmetical, and the formulae β_1, \dots, β_K are basic and form a partition.

Proof. The differences between (the matrices of) (2) and (1) are twofold: first in (2), the β_l form a partition and, second, in (1) the terms $a_s(y_i)$ and $a_s(z_h)$ (for $a_s \in \mathbf{a}$ and $y_i \in \underline{y}$, $z_h \in \underline{z}$) may occur in α and in the β_l .

We may disregard the $a_s(z_h)$ without loss of generality, because we can include them in the $a_s(y_i)$: to this aim, it is sufficient to take a fresh y , to add the conjunct $y = z_h$ to α and to replace everywhere $a_s(z_h)$ by $a_s(y)$. In order to eliminate also a term like $a_s(y_i)$, we make a guess and distinguish the case where $y_i \geq N$ and the case where $y_i < N$ (formally, ‘making a guess’ means to replace (1) with a disjunction - the two disjuncts being obtained by adding to α the case description). According to the semantics conventions we made in Sect. 2, the first case is trivial because we can just replace $a_s(y_i)$ by 0. In the other case, we first take a fresh variable u and apply the equivalence $\gamma(\dots a_s(y_j) \dots) \leftrightarrow \exists u (a_s(y_j) = u \wedge \gamma(\dots u \dots))$ (here γ is the whole (1)); then we replace $a_s(y_j) = u$ by the equivalent formula $\sharp\{x \mid x = y_j \wedge a_s[x] = u\} = 1$ and finally the latter by $\exists u' (u' = 1 \wedge \sharp\{x \mid x = y_j \wedge a_s[x] = u\} = u')$ (the result has the desired shape once we move the new existential quantifiers in front).

After this, we still need to modify the β_l so that they form a partition (this further step produces an exponential blow-up). Let $\psi(\underline{y})$ be the matrix of a formula of the kind (2), where the β_l are not a partition. Let us set $\underline{K} := \{1, \dots, K\}$ and let us consider further variables $\underline{u} = \langle u_\sigma \rangle_\sigma$, for $\sigma \in 2^{\underline{K}}$. Then it is clear that the existential closure of ψ is equivalent to the formula obtained by prefixing the existential quantifiers $\exists \underline{u} \exists \underline{z}$ to the formula

$$\left(\alpha \wedge \bigwedge_{l=1}^K z_l = \sum_{\sigma \in 2^{\underline{K}}, \sigma(l)=1} u_\sigma \right) \wedge \bigwedge_{\sigma \in 2^{\underline{K}}} \sharp\{x \mid \beta_\sigma\} = u_\sigma \quad (3)$$

where $\beta_\sigma := \bigwedge_{l=1}^K \epsilon_{\sigma(l)} \beta_l$ (here $\epsilon_{\sigma(l)}$ is ‘ \neg ’ if $\sigma(l) = 0$, it is a blank space otherwise). ⊣

Theorem 2. *Satisfiability of E-flat formulae is decidable.*

Proof. We reduce satisfiability of (2) to satisfiability of a constraint formula (4) which is decidable by Theorem 1. The matrix of (2) has free variables $\underline{z}, \underline{y}$ and these are inherited by the equi-satisfiable formula (4), but the latter contains

extra free variables $z_S, z_{l,S}$: variables z_S count new Venn regions, whereas variables $z_{l,S}$ counts how many elements are taken from z_S to contribute to the old Venn region counted by z_l . In detail, we show that (2) is equisatisfiable with

$$\begin{aligned} & \alpha \wedge \bigwedge_{S \in \wp(\underline{K})} \left(z_S = \#\{x \mid \bigwedge_{l \in S} \exists \underline{u} \beta_l(x, \underline{u}, \underline{y}, \underline{z}) \wedge \bigwedge_{l \notin S} \forall \underline{u} \neg \beta_l(x, \underline{u}, \underline{y}, \underline{z})\} \right) \wedge \\ & \wedge \bigwedge_{S \in \wp(\underline{K})} \left(z_S = \sum_{l \in S} z_{l,S} \right) \wedge \bigwedge_{l=1}^K \left(z_l = \sum_{S \in \wp(\underline{K}), l \in S} z_{l,S} \right) \wedge \bigwedge_{l \in S \in \wp(\underline{K})} z_{l,S} \geq 0 \end{aligned} \quad (4)$$

(according to our notations, the basic formulæ $\beta_l(x, \mathbf{a}(x), \underline{y}, \underline{z})$ from (2) were supposed to be built up from the arithmetic formulæ $\beta_l(x, \underline{u}, \underline{y}, \underline{z})$ by replacing the variables $\underline{u} = u_1, \dots, u_s$ with the terms $\mathbf{a}(x) = a_1(x), \dots, a_s(x)$).

Suppose that (4) is satisfiable. Then there is an assignment V to the free variables occurring in it so that (4) is true in the standard structure of the integers (for simplicity, we use the same name for a free variable and for the integer assigned to it by V). If $\mathbf{a} = a_1, \dots, a_s$, we need to define $a_s(i)$ for all s and for all $i \in [0, N)$. For every $l = 1, \dots, K$ this must be done in such a way that there are exactly z_l integer numbers taken from $[0, N)$ satisfying $\beta_l(x, \mathbf{a}(x), \underline{y}, \underline{z})$. The interval $[0, N)$ can be partitioned by associating with each $i \in [0, N)$ the set $i_S = \{l \in \underline{K} \mid \exists \underline{u} \beta_l(i, \underline{u}, \underline{y}, \underline{z}) \text{ holds under } V\}$. For every $S \in \wp(\underline{K})$ the number of the i such that $i_S = S$ is z_S ; for every $l \in S$, pick $z_{l,S}$ among them and, for these selected i , let the s -tuple $\mathbf{a}(i)$ be equal to an s -tuple \underline{y} such that $\beta_l(i, \underline{u}, \underline{y}, \underline{z})$ holds (for this tuple \underline{y} , since the β_l are a partition, $\beta_h(i, \underline{u}, \underline{y}, \underline{z})$ does not hold, if $h \neq l$). Since $z_S = \sum_{l \in S} z_{l,S}$ and since $\sum_S z_S$ is equal to the length of the interval $[0, N)$ (because the formulæ $\bigwedge_{l \in S} \exists \underline{u} \beta_l \wedge \bigwedge_{l \notin S} \forall \underline{u} \neg \beta_l$ are a partition), the definition of the \mathbf{a} is complete. The formula (2) is true by construction.

On the other hand *suppose that (2) is satisfiable* under an assignment V ; we need to find $V(z_S), V(z_{l,S})$ (we again indicate them simply as $z_S, z_{l,S}$) so that (4) is true. For z_S there is no choice, since $z_S = \#\{i \mid \bigwedge_{l \in S} \exists \underline{u} \beta_l(i, \underline{u}, \underline{y}, \underline{z}) \wedge \bigwedge_{l \notin S} \forall \underline{u} \neg \beta_l(i, \underline{u}, \underline{y}, \underline{z})\}$ must be true; for $z_{l,S}$, we take it to be the cardinality of the set of the i such that $\beta_l(i, \mathbf{a}(i), \underline{y}, \underline{z})$ holds under V and $S = \{h \in \underline{K} \mid \exists \underline{u} \beta_h(i, \underline{u}, \underline{y}, \underline{z}) \text{ holds under } V\}$. In this way, for every S , the equality $z_S = \sum_{l \in S} z_{l,S}$ holds and for every l , the equality $z_l = \sum_{S \in \wp(\underline{K}), l \in S} z_{l,S}$ holds too. Thus the formula (2) becomes true under our extended V . \dashv

Example 2. Let us test the satisfiability of

$$N > 3 \wedge z_2 \equiv_5 1 \wedge z_1 = \#\{x \mid x + |a(x)| < 3\} \wedge z_2 = \#\{x \mid x + |a(x)| \geq 3\} \quad (5)$$

We have $K = 2$ and let us put $S_1 := \{1, 2\}, S_2 := \emptyset, S_3 := \{1\}, S_4 := \{2\}$. Since the absolute value is a positive number, when writing down (4), we easily realize that we must have $z_{S_1} = \#\{0, 1, 2\} = 3, z_{S_2} = z_{S_3} = 0, z_{S_4} = \#\{3, \dots, N-1\} = N-3$. Thus (5) is satisfiable iff there are $z_{1S_1}, z_{2S_1}, z_{2S_4} \geq 0$ such that

$$N > 3 \wedge z_2 \equiv_5 1 \wedge z_{1S_1} + z_{2S_1} = 3 \wedge z_{2S_4} = N-3 \wedge z_1 = z_{1S_1} \wedge z_2 = z_{2S_1} + z_{2S_4}$$

which is in fact the case (but notice that an additional conjunct like $N \equiv_5 0$ would make (5) unsatisfiable).

4 A More Tractable Subcase

We saw that satisfiability of flat formulæ is decidable, but the complexity of the decision procedure is very high: Lemma 1 introduces an exponential blow-up and other exponential blow-ups are introduced by Theorem 2 and by the decision procedure (via quantifier elimination) from [19]. Of course, all this might be subject to dramatic optimizations (to be investigated by future research); in this paper we show that there is a much milder (and still practically useful) fragment.

Definition 2. *Simple flat formulæ are recursively defined as follows: (i) basic formulæ are simple flat formulæ; (ii) if ϕ is a simple flat formula, $\beta(\mathbf{a}(x), \mathbf{a}(\underline{y}), \underline{y})$ is a basic formula and x, z are variables, then $\phi(\#\{x \mid \beta\} / z)$ is a simple flat formula.*

As an example of a simple flat formula consider the following one

$$a'(y) = z \wedge \#\{x \mid a'(x) = a(x)\} \geq N-1 \wedge (\#\{x \mid a'(x) = a(x)\} < N \rightarrow a(y) \neq z)$$

expressing that $a' = \text{write}(a, y, z)$ (i.e. that the array a' is obtained from a by over-writing z in the entry y).

Definition 3. *Simple E-flat formulæ are formulæ of the kind*

$$\begin{aligned} \exists \underline{z}. \alpha(\mathbf{a}(\underline{y}), \mathbf{a}(\underline{z}), \underline{y}, \underline{z}) \wedge \#\{x \mid \beta_1(\mathbf{a}(x), \mathbf{a}(\underline{y}), \mathbf{a}(\underline{z}), \underline{y}, \underline{z})\} = z_1 \wedge \cdots \\ \cdots \wedge \#\{x \mid \beta_K(\mathbf{a}(x), \mathbf{a}(\underline{y}), \mathbf{a}(\underline{z}), \underline{y}, \underline{z})\} = z_K \end{aligned} \quad (6)$$

where α and the β_i are basic.

It is easily seen that (once again) simple flat formulæ are closed under Boolean combinations and that simple flat formulæ are simple E-flat formulæ (the converse is not true, for cyclicity of the dependence graph of the z_i 's in (6)).

The difference between simple and non-simple flat/E-flat formulæ is that in simple formulæ *the abstraction variable cannot occur outside the read of an array symbol* (in other words, the β, β_i from the above definition are of the kind $\beta_i(\mathbf{a}(x), \mathbf{a}(\underline{y}), \mathbf{a}(\underline{z}), \underline{y}, \underline{z})$ and not of the kind $\beta_i(\mathbf{a}(x), \mathbf{a}(\underline{y}), \mathbf{a}(\underline{z}), x, \underline{y}, \underline{z})$). This restriction has an important semantic effect, namely that formulæ (6) are equi-satisfiable to formulæ which are *permutation-invariant*, in the following sense. The truth value of an arithmetical formula or of a formula like $z = \#\{x \mid \alpha(\mathbf{a}(x), \underline{y})\}$ is not affected by a permutation of the values of the $\mathbf{a}(x)$ for $x \in [0, N)$, because x does not occur free in α (permuting the values of the $\mathbf{a}(x)$ may on the contrary change the value of a flat non-simple sentence like $z = \#\{x \mid a(x) \leq x\}$). This ‘permutation invariance’ will be exploited in the argument proving the correctness of decision procedure of Theorem 3 below. Formulae (6) themselves are not permutation-invariant because of subterms $\mathbf{a}(\underline{z}), \mathbf{a}(\underline{y})$, so we first show how to eliminate them up to satisfiability:

Lemma 2. *Simple E-flat formulæ are equi-satisfiable to disjunctions of permutation-invariant formulæ of the kind*

$$\exists \underline{z}. \alpha(\underline{y}, \underline{z}) \wedge \#\{x \mid \beta_1(\mathbf{a}(x), \underline{y}, \underline{z})\} = z_1 \wedge \cdots \wedge \#\{x \mid \beta_K(\mathbf{a}(x), \underline{y}, \underline{z})\} = z_K \quad (7)$$

Proof. Let us take a formula like (6): we convert it to an equi-satisfiable disjunction of formulæ of the kind (7). The task is to eliminate terms $\mathbf{a}(\underline{z})$, $\mathbf{a}(\underline{y})$ by a series of guessings (each guessing will form the content of a disjunct). Notice that we can apply the procedure of Lemma 1 to eliminate the $\mathbf{a}(\underline{z})$, but for the $\mathbf{a}(\underline{y})$ we must operate differently (the method used in Lemma 1 introduced non-simple abstraction terms).

Let us suppose that $\underline{y} := y_1, \dots, y_m$ and that, after a first guess, α contains the conjunct $y_j < N$ for each $j = 1, \dots, m$ (if it contains $y_j \geq N$, we replace $a_s(y_j)$ by 0); after a second series of guesses, we can suppose also that α contains the conjuncts $y_{j_1} \neq y_{j_2}$ for $j_1 \neq j_2$ (if it contains $y_{j_1} = y_{j_2}$, we replace y_{j_1} by y_{j_2} everywhere, making y_{j_1} to disappear from the whole formula). In the next step, (i) we introduce for every $a \in \mathbf{a}$ and for every $j = 1, \dots, m$ a fresh variable u_{aj} , (ii) we replace everywhere $a(y_j)$ by u_{aj} and (iii) we conjoin to α the equalities $a(y_j) = u_{aj}$. In this way we get a formula of the following kind

$$\exists \underline{z}. \bigwedge_{a \in \mathbf{a}, y_j \in \underline{y}} a(y_j) = u_{aj} \wedge \alpha(\underline{y}, \underline{u}, \underline{z}) \wedge \bigwedge_{l=1}^K \#\{x \mid \beta_l(\mathbf{a}(x), \underline{y}, \underline{u}, \underline{z})\} = z_l \quad (8)$$

where \underline{u} is the tuple formed by the u_{aj} (varying a and j). We now make another series of guesses and conjoin to α either $u_{aj} = u_{a'j'}$ or $u_{aj} \neq u_{a'j'}$ for $(a, j) \neq (a', j')$. Whenever $u_{aj} = u_{a'j'}$ is conjoined, u_{aj} is replaced by $u_{a'j'}$ everywhere, so that u_{aj} disappears completely. The resulting formula still has the form (8), but now the map $(a, j) \mapsto u_{aj}$ is not injective anymore (otherwise said, u_{aj} now indicates the element from the tuple \underline{u} associated with the pair (a, j) and we might have that the same u_{aj} is associated with different pairs (a, j)).

Starting from (8) modified in this way, let us define now the equivalence relation among the y_j that holds between y_j and $y_{j'}$ whenever for all $a \in \mathbf{a}$ there is $u_a \in \underline{u}$ such that α contains the equalities $a(y_j) = u_a$ and $a(y_{j'}) = u_a$. Each equivalence class E is uniquely identified by the corresponding function f_E from \mathbf{a} into \underline{u} (it is the function that for each $y_j \in E$ maps $a \in \mathbf{a}$ to the $u_a \in \underline{u}$ such that α contains the equality $a(y_j) = u_a$ as a conjunct). Let E_1, \dots, E_r be the equivalence classes and let n_1, \dots, n_r be their cardinalities. We claim that (8) is equisatisfiable to

$$\begin{aligned} \exists \underline{z}. \alpha(\underline{y}, \underline{u}, \underline{z}) \wedge \bigwedge_{q=1}^r \#\{x \mid \bigwedge_{a \in \mathbf{a}} a(x) = f_{E_q}(a)\} \geq n_q \wedge \\ \bigwedge_{l=1}^K \#\{x \mid \beta_l(\mathbf{a}(x), \underline{y}, \underline{u}, \underline{z})\} = z_l \end{aligned} \quad (9)$$

In fact, satisfiability of (8) trivially implies the satisfiability of the formula (9); vice versa, since (9) is permutation-invariant, if it is satisfiable we can modify

any assignment satisfying it via a simultaneous permutation of the values of the $a \in \mathbf{a}$ so as to produce an assignment satisfying (8).

We now need just the trivial observation that the inequalities $\#\{x \mid \bigwedge_{a \in \mathbf{a}} a(x) = f_{E_q}(a)\} \geq n_q$ can be replaced by the formulæ $\#\{x \mid \bigwedge_{a \in \mathbf{a}} a(x) = f_{E_q}(a)\} = z'_q \wedge z'_q \geq n_q$ (for fresh z'_q) in order to match the syntactic shape of (7). \dashv

We can freely assume that quantifiers do not occur in simple flat formulæ: this is without loss of generality because such formulæ are built up from arithmetic and basic formulæ.²

Theorem 3. *Satisfiability of simple flat formulæ can be decided in NP (and thus it is an NP-complete problem).*

Proof. First, by applying the procedure of the previous Lemma we can reduce to the problem of checking the satisfiability of formulæ of the kind

$$\alpha(\underline{y}, \underline{z}) \wedge \#\{x \mid \beta_1(\mathbf{a}(x), \underline{y}, \underline{z})\} = z_1 \wedge \cdots \wedge \#\{x \mid \beta_K(\mathbf{a}(x), \underline{y}, \underline{z})\} = z_K \quad (10)$$

where $\alpha, \beta_1, \dots, \beta_K$ are basic (notice also that each formula in the output of the procedure of the previous Lemma comes from a polynomial guess).

Suppose that $A_1(\mathbf{a}(x), \underline{y}, \underline{z}), \dots, A_L(\mathbf{a}(x), \underline{y}, \underline{z})$ are the atoms occurring in β_1, \dots, β_K . For a Boolean assignment σ to these atoms, we indicate with $\llbracket \beta_j \rrbracket^\sigma$ the Boolean value (0 or 1) the formula β_j has under such assignment. We first claim that (10) is satisfiable iff there exists a set of assignments Σ such that the formula

$$\alpha(\underline{y}, \underline{z}) \wedge \bigwedge_{\sigma \in \Sigma} \exists \underline{u} \left(\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\underline{u}, \underline{y}, \underline{z}) \right) \wedge \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_K \end{bmatrix} = \sum_{\sigma \in \Sigma} v_\sigma \begin{bmatrix} \llbracket \beta_1 \rrbracket^\sigma \\ \llbracket \beta_2 \rrbracket^\sigma \\ \vdots \\ \llbracket \beta_K \rrbracket^\sigma \end{bmatrix} \wedge \quad (11)$$

$$\wedge \sum_{\sigma \in \Sigma} v_\sigma = N \wedge \bigwedge_{\sigma \in \Sigma} v_\sigma > 0$$

is satisfiable (we introduced extra fresh variables v_σ , for $\sigma \in \Sigma$; notation $\epsilon_{\sigma(A_j)}$ is the same as in the proof of Lemma 1). In fact, on one side, if (10) is satisfiable under V , we can take as Σ the set of assignments for which $\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\mathbf{a}(i), \underline{y}, \underline{z})$ is true under V for some $i \in [0, N)$ and for v_σ the cardinality of the set of the $i \in [0, N)$ for which $\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\mathbf{a}(i), \underline{y}, \underline{z})$ holds. This choice makes (11) true. Vice versa, if (11) is true under V , in order to define the value of the tuple $\mathbf{a}(i)$ (for $i \in [0, N)$), pick for every $\sigma \in \Sigma$ some \underline{u}_σ such that $\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\underline{u}_\sigma, \underline{y}, \underline{z})$ holds; then, supposing $\Sigma = \{\sigma_1, \dots, \sigma_h\}$, let

² By the quantifier-elimination result for Presburger arithmetic, it is well-known that arithmetic formulæ are equivalent to quantifier-free ones. The same is true for basic formulæ because they are obtained from arithmetic formulae by substitutions without capture.

$\mathbf{a}(i)$ be equal to \underline{u}_{σ_1} for $i \in [0, v_{\sigma_1})$, to \underline{u}_{σ_2} for $i \in [v_{\sigma_1}, v_{\sigma_2})$, etc. Since we have that $\sum_{\sigma \in \Sigma} v_\sigma = N$, the definition of the interpretation of the \mathbf{a} is complete (any other permutation of the values $\mathbf{a}(x)$ inside $[0, N)$ would fit as well). In this way, formula (10) turns out to be true.

We so established that our original formula is satisfiable iff there is some Σ such that (11) is satisfiable; the only problem we still have to face is that Σ might be exponentially large. To reduce to a polynomial Σ , we use the same technique as in [16]. In fact, if (11) is satisfiable, then the column vector $(z_1, \dots, z_K, N)^T$ is a linear combination with positive integer coefficients of the 0/1-vectors $(\llbracket \beta_1 \rrbracket^\sigma, \dots, \llbracket \beta_K \rrbracket^\sigma, 1)^T$ and it is known from [9] that, if this is the case, the same result can be achieved by assuming that at most $2K' \log_2(4K')$ of the v_σ are nonzero (we put $K' := K + 1$). Thus polynomially many Σ are sufficient and for such Σ , a satisfying assignment for the existential Presburger formula (11) is a polynomial certificate. \dashv

4.1 Some Heuristics

We discuss here some useful heuristics for the satisfiability algorithm for simple flat formulæ (most of these heuristics have been implemented in our prototype).

1.- The satisfiability test involves all formulæ (11) for each set of assignments Σ having cardinality *at most* $M = \lceil 2K' \log_2(4K') \rceil$ (actually, one can improve this bound, see [16]). If we replace in (11), for every σ , the conjunct $v_\sigma > 0$ by $v_\sigma \geq 0$ and the conjunct³ $\exists \underline{u} (\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\underline{u}, \underline{y}, \underline{z}))$ by $v_\sigma > 0 \rightarrow \exists \underline{u} (\bigwedge_{j=1}^L \epsilon_{\sigma(A_j)} A_j(\underline{u}, \underline{y}, \underline{z}))$, we can limit ourselves to the Σ having cardinality *equal to* M . This trick is useful if, for some reason, we prefer to go through any sufficient set of assignments (like the set of all assignments supplied by some Boolean propagation, see below).

2.- There is no need to consider assignments σ over the set of the atoms A_j occurring in the β_1, \dots, β_K : any set of formulæ generating the β_1, \dots, β_K by Boolean combinations fits our purposes. As a consequence, the choice of these ‘atoms’ is subject to case-by-case evaluations.

3.- Universally quantified formulæ of the kind $\forall x (0 \leq x \wedge x < N \rightarrow \beta)$ can be turned into flat formulæ by rewriting them as $N = \#\{x \mid \beta\}$ (and in fact such universally quantified formulæ often occur in our benchmarks suite). These formulæ contribute to (10) via the conjuncts of the kind $z_i = N \wedge \#\{x \mid \beta_i(\mathbf{a}(x), \underline{y}, \underline{z})\} = z_i$. It is quite useful to consider the $\{\beta_{i_1}, \dots, \beta_{i_L}\}$ arising in this way as atoms (in the sense of point 2 above) and restrict to the assignments σ such that $\sigma(\beta_{i_1}) = \dots = \sigma(\beta_{i_L}) = 1$.

4.- Boolean propagation is a quite effective strategy to prune useless assignments. In our context, as soon as a partial assignment σ is produced inside the assignments enumeration subroutine, an SMT solver is invoked to test the satisfiability of $\alpha(\underline{y}, \underline{z}) \wedge \bigwedge_{j \in \text{dom}(\sigma)} \epsilon_{\sigma(A_j)} A_j(\underline{u}, \underline{y}, \underline{z})$. Since this is implied by a

³ These conjuncts (varying $\sigma \in \Sigma$) are needed in (11) to ensure that the assignments we are using can coexist in a model.

(skolemized) conjunct of (11), if the test is negative the current partial assignment is discarded and next partial assignment (obtained by complementing the value of the last assigned literal) is taken instead.

5 Examples and Experiments

We implemented a prototype ARCA-SAT⁴ producing out of simple E-flat formulæ (10) the proof obligations (11) (written in SMT-LIB2 format), exploiting the heuristics explained in Sect. 4.1. To experiment the feasibility of our approach for concrete verification problems, we also implemented a (beta) version of a tool called ARCA producing out of the specification of a parametric distributed system and of a safety-like problem, some E-flat simple formulæ whose unsatisfiability formalizes invariant-checking and bounded-model checking problems. A script executing in sequence ARCA, ARCA-SAT and z3 can then solve such problems by reporting a ‘sat/unsat’ answer.

A system is specified via a pair of flat (simple) formulæ $\iota(\underline{p})$ and $\tau(\underline{p}, \underline{p}')$ and a safety problem via a further formula $\nu(\underline{p})$ (here the \underline{p} are parameters and array-ids, the \underline{p}' are renamed copies of the \underline{p}). A bounded model checking problem is the problem of checking whether the formula

$$\iota(\underline{p}_0) \wedge \tau(\underline{p}_0, \underline{p}_1) \wedge \cdots \wedge \tau(\underline{p}_n, \underline{p}_{n+1}) \wedge \nu(\underline{p}_{n+1})$$

is satisfiable for a fixed n . An invariant-checking problem, given also a formula $\phi(\underline{p})$, is the problem of checking whether the three formulæ

$$\iota(\underline{p}) \wedge \neg\phi(\underline{p}), \quad \phi(\underline{p}) \wedge \tau(\underline{p}, \underline{p}') \wedge \neg\phi(\underline{p}'), \quad \phi(\underline{p}) \wedge \nu(\underline{p})$$

are unsatisfiable. Notice that since all our algorithms terminate and are sound and complete, the above problems are always solved by the above tool combination (if enough computation resources are available). Thus, our technique is able *both to make safety certifications and to find bugs*.

To validate our technique, in the following we describe in detail the formalization of the send-receive broadcast primitive (SRBP) in [21]. SRBP is used as a basis to synchronize clocks in systems where processes may fail in sending and/or receiving messages. Periodically, processes broadcast the virtual time to be adopted by all, as a (*session* s) message. Processes that accept this message set s as their current time. SRBP aims at guaranteeing the following properties:

Correctness: if at least $f + 1$ correct processes broadcast the message (*session* s), all correct processes accept the message.

Unforgeability: if no correct process broadcasts (*session* s), no correct process accepts the message.

Relay: if a correct process accepts (*session* s), all correct processes accept it.

⁴ ARCA stands for *Array with Cardinalities*.

where $f < N/2$ is the number of processes failing during an algorithm run, with N the number of processes in the system. Algorithm 1 shows the pseudo-code.

We model SRBP as follows: $IT(x)$ is the initial state of a process x ; it is s when x broadcasts a (*init, session* s) message, and 0 otherwise. $SE(x) = s$ indicates that x has broadcast its own echo. $AC(x) = s$ indicates that x has accepted (*session* s). Let pc be the program counter, r the round number, and G a flag indicating whether one round has been executed. We indicate with $F(x) = 1$ the fact that x is faulty, and $F(x) = 0$ otherwise. Finally, $CI(x)$ and $CE(x)$ are the number of respectively inits and echoes received. In the following, $\forall x$ means $\forall x \in [0, N)$. Some sentences are conjoined to all our proof obligations, namely: $\#\{x|F(x) = 0\} + \#\{x|F(x) = 1\} = N \wedge \#\{x|F(x) = 1\} < N/2$. For the Correctness property, we write ι_c as follows:

$$\iota_c := pc = 1 \wedge r = 0 \wedge G = 0 \wedge s \neq 0 \wedge \quad (12)$$

$$\#\{x|IT(x) = 0\} + \#\{x|IT(x) = s\} = N \wedge \quad (13)$$

$$\#\{x|F(x) = 0 \wedge IT(x) = s\} \geq (\#\{x|F(x) = 1\} + 1) \wedge \quad (14)$$

$$\forall x.SE(x) = 0 \wedge AC(x) = 0 \wedge CI(x) = 0 \wedge CE(x) = 0 \quad (15)$$

where we impose that the number of correct processes broadcasting the init message is at least the number of faulty processes, f , plus 1. It is worth to notice that – from the above definition – our tool produces a specification that is checked for any $N \in \mathbb{N}$ number of processes. The constraints on IT allow to verify all admissible assignments of 0 or s to the variables. Similarly for $F(x)$.

The algorithm safety is verified by checking that the bad properties cannot be reached from the initial state. For Correctness, we set $\nu_c := pc = 1 \wedge G = 1 \wedge \#\{x|F(x) = 0 \wedge AC(x) = 0\} > 0$, that is, Correctness is not satisfied if – after one round – some correct process exists that has yet to accept. The algorithm evolution is described by two transitions: τ_1 and τ_2 . The former allows to choose the number of both inits and echoes received by each process. The latter describes the actions in Algorithm 1.

$$\tau_1 := pc = 1 \wedge pc' = 2 \wedge r' = r \wedge G' = G \wedge s' = s \wedge \exists K1, K2, K3, K4.$$

$$K1 = \#\{x|F(x) = 0 \wedge IT(x) = s\} \wedge K2 = \#\{x|F(x) = 0 \wedge SE(x) = s\} \wedge$$

$$K3 = \#\{x|F(x) = 1 \wedge IT(x) = s\} \wedge K4 = \#\{x|F(x) = 1 \wedge SE(x) = s\} \wedge$$

$$\forall x.F(x) = 0 \Rightarrow (CI'(x) \geq K1 \wedge CI'(x) \leq (K1 + K3) \wedge CE'(x) \geq K2 \wedge$$

$$CE'(x) \leq (K2 + K4)) \wedge$$

$$\forall x.F(x) = 1 \Rightarrow (CI'(x) \geq 0 \wedge CI'(x) \leq (K1 + K3) \wedge CE'(x) \geq 0 \wedge$$

$$CE'(x) \leq (K2 + K4)) \wedge$$

$$\forall x.IT'(x) = IT(x) \wedge SE'(x) = SE(x) \wedge AC'(x) = AC(x)$$

$$\tau_2 := pc = 2 \wedge pc' = 1 \wedge r' = (r + 1) \wedge s' = s \wedge G' = 1 \wedge$$

$$\forall x.(CI(x) \geq \#\{x|F(x) = 1\} + 1 \Rightarrow SE'(x) = s \wedge AC'(x) = s) \wedge$$

$$\forall x.(CI(x) < \#\{x|F(x) = 1\} + 1 \wedge CE(x) \geq 1 \Rightarrow SE'(x) = s \wedge AC'(x) = s) \wedge$$

$$\forall x.(CI(x) < \#\{x|F(x) = 1\} + 1 \wedge CE(x) < 1 \Rightarrow SE'(x) = 0 \wedge AC'(x) = 0) \wedge$$

$$\forall x.IT'(x) = IT(x) \wedge CI'(x) = CI(x) \wedge CE'(x) = CE(x)$$

Algorithm 1. Pseudo-code for the send-receive broadcast primitive.

Initialization:

To broadcast a (*session s*) message, a correct process sends (*init, session s*) to all.

End Initialization

for each correct process:

1. if received (*init, session s*) from at least $f + 1$ distinct processes or
2. received (*echo, session s*) from any process then
3. accept (*session s*);
4. send (*echo, session s*) to all;
5. endif

end for

The same two transitions are used to verify both the Unforgeability and the Relay properties, for which however we have to change the initial and final formula. For Unforgeability, (13) in ι changes as $\dots \wedge \#\{x|F(x) = 0 \wedge IT(x) = 0\} = \#\{x|F(x) = 0\} \wedge \dots$; while $v_u := pc = 1 \wedge G = 1 \wedge \#\{x|F(x) = 0 \wedge AC(x) = s\} > 0$. In ι_u we say that all non-faulty processes have $IT(x) = 0$. Unforgeability is not satisfied if some correct process accepts. For Relay, we use:

$$\begin{aligned} \iota_r := & pc = 1 \wedge r = 0 \wedge s \neq 0 \wedge G = 0 \wedge \\ & \#\{x|F(x) = 0 \wedge AC(x) = s \wedge SE(x) = s\} = 1 \wedge \\ & \#\{x|AC(x) = 0 \wedge SE(x) = 0\} = (N - 1) \wedge \#\{x|AC(x) = s \wedge SE(x) = s\} = 1 \wedge \\ & \forall x.IT(x) = 0 \wedge CI(x) = 0 \wedge CE(x) = 0 \end{aligned}$$

while $v_r = v_c$. In this case, we start the system in the worst condition: by the hypothesis, we just know that one correct process has accepted. Upon acceptance, by the pseudo-code, it must have sent an echo. All the other processes are initialized in an idle state. We also produce an unsafe model of Correctness: we modify ι_c by imposing that just f correct processes broadcast the init message.

In Table 1, we report the results of validating these and other models with our tool. In the first column, the considered algorithm is indicated. The second column indicates the property to be verified; the third column reports the conditions of verification. In the fourth column, we indicate whether we consider either a bounded model checking (BMC) or an invariant-checking (IC) problem. The fifth column supplies the obtained results (for BMC problems, ‘safe’ means of course ‘safe up to the analyzed bound’). The sixth column shows the time jointly spent by ARCA, ARCA-SAT and z3 for the verification, considering for BMC the sum of the times spent for every traces of length up to 10. We used a PC equipped with Intel Core i7 processor and operating system Linux Ubuntu 14.04 64 bits. We focused on BMC problems as they produce longer formulas thus stressing more the tools. Specifically, following the example above, we modeled:

Table 1. Evaluated algorithms and experimental results.

Algorithm	Property	Condition	Problem	Outcome	Time (s.)
SRBP [21]	Correctness	$\geq (f + 1)$ init's	BMC	safe	0.82
SRBP [21]	Correctness	$\leq f$ init's	BMC	unsafe	2.21
SRBP [21]	Unforgeability	$\geq (f + 1)$ init's	BMC	safe	0.85
SRBP [21]	Relay	$\geq (f + 1)$ init's	BMC	safe	1.93
BBP [22]	Correctness	$N > 3f$	BMC	safe	6.17
BBP [22]	Unforgeability	$N > 3f$	BMC	safe	0.25
BBP [22]	Unforgeability	$N \geq 3f$	BMC	unsafe	0.25
BBP [22]	Relay	$N > 3f$	BMC	safe	1.01
OT [4]	Agreement	threshold $> 2N/3$	IC	safe	4.20
OT [4]	Agreement	threshold $> 2N/3$	BMC	safe	278.95
OT [4]	Agreement	threshold $\leq 2N/3$	BMC	unsafe	17.75
OT [4]	Irrevocability	threshold $> 2N/3$	BMC	safe	8.72
OT [4]	Irrevocability	threshold $\leq 2N/3$	BMC	unsafe	9.51
OT [4]	Weak Validity	threshold $> 2N/3$	BMC	safe	0.45
OT [4]	Weak Validity	threshold $\leq 2N/3$	BMC	unsafe	0.59
UV [5]	Agreement	$\mathcal{P}_{nosplit}$ violated	BMC	unsafe	4.18
UV [5]	Irrevocability	$\mathcal{P}_{nosplit}$ violated	BMC	unsafe	2.04
UV [5]	Integrity	-	BMC	safe	1.02
$U_{T,E,\alpha}$ [3]	Integrity	$\alpha = 0 \wedge \mathcal{P}_{safe}$	BMC	safe	1.16
$U_{T,E,\alpha}$ [3]	Integrity	$\alpha = 0 \wedge \neg \mathcal{P}_{safe}$	BMC	unsafe	0.83
$U_{T,E,\alpha}$ [3]	Integrity	$\alpha = 1 \wedge \mathcal{P}_{safe}$	BMC	safe	5.20
$U_{T,E,\alpha}$ [3]	Integrity	$\alpha = 1 \wedge \neg \mathcal{P}_{safe}$	BMC	unsafe	4.93
$U_{T,E,\alpha}$ [3]	Agreement	$\alpha = 0 \wedge \mathcal{P}_{safe}$	BMC	safe	59.80
$U_{T,E,\alpha}$ [3]	Agreement	$\alpha = 0 \wedge \neg \mathcal{P}_{safe}$	BMC	unsafe	7.78
$U_{T,E,\alpha}$ [3]	Agreement	$\alpha = 1 \wedge \mathcal{P}_{safe}$	BMC	safe	179.67
$U_{T,E,\alpha}$ [3]	Agreement	$\alpha = 1 \wedge \neg \mathcal{P}_{safe}$	BMC	unsafe	31.94
MESI [17]	cache coherence	-	IC	safe	0.11
MOESI [20]	cache coherence	-	IC	safe	0.08
Dekker [6]	mutual exclusion	-	IC	safe	2.05

- the byzantine broadcast primitive (BBP) [22] used to simulate authenticated broadcast in the presence of malicious failures of the processes,
- the one-third algorithm (OT) [4] for consensus in the presence of benign transmission failures,
- the Uniform Voting (UV) algorithm [5] for consensus in the presence of benign transmission failures,

- the $U_{T,E,\alpha}$ algorithm [3] for consensus in the presence of malicious transmission failures,
- the MESI [17] and MOESI [20] algorithms for cache coherence,
- the Dekker’s algorithm [6] for mutual exclusion.

All the models, together with our tools to verify them, are available at <http://users.mat.unimi.it/users/ghilardi/arca> (for the z3 solver see <http://rise4fun.com/z3>).

As far as the processing times are concerned, we observed that on average z3 accounts for around 68% of the processing time, while ARCA and ARCA-SAT together account for the remaining 32%. Indeed, the SMT tests performed by ARCA-SAT are lightweight – as they only prune assignments – yet effective, as they succeed in reducing the number of assignments of at least one order of magnitude.

6 Conclusions, Related and Further Work

We identified two fragments of the rich syntax of Fig. 1 and we showed their decidability (for the second fragment we showed also a tight complexity bound). Since our fragments are closed under Boolean connectives, it is possible to use them not only in bounded model checking (where they can both give certifications and find bugs), but also in order to decide whether an invariant holds or not. We implemented our algorithm for the weaker fragment and used it in some experiments. As far as we know, this is the first implementation of a *complete* algorithm for a fragment of arithmetic with arrays and counting capabilities for interpreted sets. In future, we plan to extend both our tool ARCA and our results in order to deal with more complex verification problems.

Since one of the major intended applications concerns fault-tolerant distributed systems, we briefly review and compare here some recent work in the area. Papers [11–13] represent a very interesting and effective research line, where cardinality constraints are not directly handled but abstracted away using interval abstract domains and counters. As a result, a remarkable amount of algorithms are certified, although the method might suffer of some lack of expressiveness for more complex examples. On the contrary, paper [4] directly handles cardinality constraints for interpreted sets; nontrivial invariant properties are synthesized and checked, based on Horn constraint solving technology. At the level of decision procedures, some incomplete inference schemata are employed (completeness is nevertheless showed for array updates against difference bounds constraints). Paper [7] introduces a very expressive logic, specifically tailored to handle consensus problems (whence the name ‘consensus logic’ *CL*). Such logic employs arrays with values into power set types, hence it is situated in a higher order logic context. Despite this, our flat fragment is not fully included into *CL*, because we allow arithmetic constraints on the sort of indexes and also mixed constraints between indexes and data: in fact, we have a unique sort for indexes and data, leading to the possibility of writing typically non permutation-invariant formulæ

like $\sharp \{x \mid a(x) + x = N\} = z$. As pointed out in [2], this mono-sorted approach is useful in the analysis of programs, when pointers to the memory (modeled as an array) are stored into array variables. From the point of view of deduction, the paper [7] uses an incomplete algorithm in order to certify invariants. A smaller decidable fragment (identified via several syntactic restrictions) is introduced in the final part of the paper; the sketch of the decidability proof supplied for this smaller fragment uses bounds for minimal solutions of Presburger formulæ as well as Venn regions decompositions in order to build models where all nodes in the same Venn region share the same value for their function symbols.

References

1. Alberti, F., Ghilardi, S., Pagani, E.: Counting constraints in flat array fragments (2016). CoRR, abs/1602.00458
2. Alberti, F., Ghilardi, S., Sharygina, N.: Decision procedures for flat array properties. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 15–30. Springer, Heidelberg (2014)
3. Biely, M., Charron-Bost, B., Gaillard, A., Hutle, M., Schiper, A., Widder, J.: Tolerating corrupted communication. In: Proceedings of PODC, pp. 244–253 (2007)
4. Bjørner, N., von Gleissenthall, K., Rybalchenko, A.: Synthesizing cardinality invariants for parameterized systems (2015). <https://www7.in.tum.de/~gleissen/papers/sharpie.pdf>
5. Charron-Bost, B., Schiper, A.: The heard-of model: computing in distributed systems with benign faults. *Distrib. Comput.* **22**, 49–71 (2009)
6. Dijkstra, E.W.: Cooperating sequential processes. In: Genuys, F. (ed.) *Programming Languages*, pp. 43–112. Academic Press, New York (1968)
7. Drăgoi, C., Henzinger, T.A., Veith, H., Widder, J., Zufferey, D.: A logic-based framework for verifying consensus algorithms. In: McMillan, K.L., Rival, X. (eds.) *VMCAI 2014*. LNCS, vol. 8318, pp. 161–181. Springer, Heidelberg (2014)
8. Dragoi, C., Henzinger, T.A., Zufferey, D.: The need for language support for fault-tolerant distributed systems. In: *Proceedings of SNAPL (2015)*
9. Eisenbrand, F., Shmonin, G.: Carathéodory bounds for integer cones. *Oper. Res. Lett.* **34**(5), 564–568 (2006)
10. Halpern, J.Y.: Presburger arithmetic with unary predicates is Π_1^1 complete. *J. Symbol. Logic* **56**(2), 637–642 (1991)
11. John, A., Konnov, I., Schmid, U., Veith, H., Widder, J.: Parameterized model checking of fault-tolerant distributed algorithms by abstraction. In: *Proceedings of FMCAD*, pp. 201–209, August 2013
12. Konnov, I., Veith, H., Widder, J.: On the completeness of bounded model checking for threshold-based distributed algorithms: reachability. In: Baldan, P., Gorla, D. (eds.) *CONCUR 2014*. LNCS, vol. 8704, pp. 125–140. Springer, Heidelberg (2014)
13. Konnov, I., Veith, H., Widder, J.: SMT and POR beat counter abstraction: parameterized model checking of threshold-based distributed algorithms. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015*. LNCS, vol. 9206, pp. 85–102. Springer, Heidelberg (2015)
14. Kuncak, V., Nguyen, H.H., Rinard, M.: An algorithm for deciding BAPA: boolean algebra with Presburger arithmetic. In: Nieuwenhuis, R. (ed.) *CADE 2005*. LNCS (LNAI), vol. 3632, pp. 260–277. Springer, Heidelberg (2005)

15. Kuncak, V., Nguyen, H.H., Rinard, M.: Deciding boolean algebra with Presburger arithmetic. *J. Autom. Reasoning* **36**(3), 213–239 (2006)
16. Kuncak, V., Rinard, M.: Towards efficient satisfiability checking for boolean algebra with Presburger arithmetic. In: Pfenning, F. (ed.) *CADE 2007*. LNCS (LNAI), vol. 4603, pp. 215–230. Springer, Heidelberg (2007)
17. Papamarcos, M.S., Patel, J.H.: A low-overhead coherence solution for multiprocessors with private cache memories. In: *Proceedings of ISCA*, p. 348 (1984)
18. Piskac, R., Kuncak, V.: Decision procedures for multisets with cardinality constraints. In: Logozzo, F., Peled, D.A., Zuck, L.D. (eds.) *VMCAI 2008*. LNCS, vol. 4905, pp. 218–232. Springer, Heidelberg (2008)
19. Schweikhart, N.: Arithmetic, first-order logic, and counting quantifiers. *ACM TOCL* **6**, 1–35 (2004)
20. Solihin, Y.: *Fundamentals of Parallel Computer Architecture Multichip and Multicore Systems*. Solihin Publishing & Consulting LLC (2008)
21. Srikanth, T.K., Toueg, S.: Optimal clock synchronization. *J. ACM* **34**(3), 626–645 (1987)
22. Srikanth, T.K., Toueg, S.: Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distrib. Comput.* **2**(2), 80–94 (1987)
23. Yessenov, K., Piskac, R., Kuncak, V.: Collections, cardinalities, and relations. In: Barthe, G., Hermenegildo, M. (eds.) *VMCAI 2010*. LNCS, vol. 5944, pp. 380–395. Springer, Heidelberg (2010)