Università
della
Svizzera
italiana

Faculty
of Informatics

USI Technical Report Series in Informatics

Abstraction and Acceleration in SMT-based Model-Checking for Array Programs

Francesco Alberti[1], Silvio Ghilardi[2], Natasha Sharygina[1]

[1] Faculty of Informatics, University of Lugano, Switzerland
[2] Università degli Studi di Milano, Milan, Italy

arXiv:1304.4499v2 [cs.LO] 3 Oct 2013

Abstract

Abstraction (in its various forms) is a powerful established technique in model-checking; still, when unbounded data-structures are concerned, it cannot always cope with divergence phenomena in a satisfactory way. Acceleration is an approach which is widely used to avoid divergence, but it has been applied mostly to integer programs. This paper addresses the problem of accelerating transition relations for unbounded arrays with the ultimate goal of avoiding divergence during reachability analysis of abstract programs. For this, we first design a format to compute accelerations in this domain; then we show how to adapt the so-called 'monotonic abstraction' technique to efficiently handle complex formulæ with nested quantifiers generated by the acceleration preprocessing. Notably, our technique can be easily plugged-in into abstraction/refinement loops, and strongly contributes to avoid divergence: experiments conducted with the MCMT model checker attest the effectiveness of our approach on programs with unbounded arrays, where acceleration and abstraction/refinement technologies fail if applied alone.

Report Info

Published
October 2012

Revised
April 2013

Number
USI-INF-TR-2012-1

Institution
Faculty of Informatics
University of Lugano
Lugano, Switzerland

Online Access
www.inf.usi.ch/techreports

1 Introduction

Transitive closure is a logical construct that is far beyond first order logic: either infinite disjunctions or higher order quantifiers or, at least, fixpoints operators are required to express it. Indeed, due to the compactness of first order logic, transitive closure (even modulo the axioms of a first order theory) is first-order definable only in trivial cases. These general results do not hold if we define a theory as a class of structures $\mathscr{C}$ over a given signature[1]. Such definition is different from the "classical" one where a theory is identified as a set of axioms. By taking a theory as a class of structures the property of compactness breaks, and it might well happen that transitive closure becomes first-order definable (the first order definition being valid just inside the class $\mathscr{C}$ - which is often reduced to a single structure).

In this paper we consider the extension of Presburger arithmetic with free unary function symbols. Inside Presburger arithmetic, various classes of relations are known to have definable *acceleration*[2] (see related work section below). In our combined setting, the presence of free function symbols introduces a novel feature that, for instance, limits decidability to controlled extensions of the quantifier-free fragment [15, 22]. In this paper we show that in such theory some classes of relations admit a definable acceleration.

The theoretical problem of studying the definability of accelerated relations has an important application in program verification. The theory we focus on is widely adopted to represent programs handling arrays, where free functions model arrays of integers. In this application domain, the accelerated counterpart of

[1] Such definition is widely adopted in the SMT literature [7].
[2] 'acceleration' is the name usually adopted in the formal methods literature to indicate transitive closure.

relations encoding systems evolution (e.g., loops in programs) allows to compute 'in one shot' the reachable set of states after an arbitrary but finite number of execution steps. This has the great advantage of keeping under control sources of (possible) divergence arising in the reachability analysis.

The contributions of the paper are many-fold. First, we show that inside the combined theory of Presburger arithmetic augmented with free function symbols, the acceleration of some classes of relations – corresponding, in our application domain, to relations involving arrays and counters – can be expressed in first order language. This result comes at a price of allowing nested quantifiers. Such nested quantification can be problematic in practical applications. To address this complication, as a second contribution of the paper, we show how to take care of the quantifiers added by the accelerating procedure: the idea is to import in this setting the so-called *monotonic abstraction* technique [1, 2]. Such technique has been reinterpreted and analyzed in a declarative context in [5]: from a logical point of view, it amounts to a restricted form of *instantiation for universal quantifiers*. Third, we show that the ability to compute accelerated relations is greatly beneficial in program verification. In particular, one of the biggest problems in verifying safety properties of array programs is designing procedures for the synthesis of relevant quantified predicates. In typical sequential programs (like those illustrated in Figure1), the guarded assignments used to model the program instructions are ground and, as a consequence, the formulae representing backward reachable states are ground too. However, the invariants required to certify the safety of such programs contain quantifiers. Our acceleration procedure is able to supply the required quantified predicates. Our experimentation attests that abstraction/refinement-based strategies widely used in verification benefit from accelerated transitions. In programs with nested loops, as the `allDiff` procedure of Figure1 for example, the ability to accelerate the inner loop simplifies the structure of the problem, allowing abstraction to converge during verification of the entire program. For such programs, abstraction/refinement or acceleration approaches taken in isolation are not sufficient, reachability analysis converges only if they are combined together.

Related Work. To the best of our knowledge, the only work addressing the problem of accelerating relations involving arrays is [12]. Such approach seems to be unable to handle properties of common interest with more than one quantified variable (e.g., "sortedness") and is limited to programs without nested loops. Our technique is not affected by such limitations and can successfully handle examples outside the scope of [12].

Inside Presburger arithmetic, various classes of relations are known to have definable acceleration: these include relations that can be formalized as difference bounds constraints [14, 19], octagons [11] and finite monoid affine transformations [20] (paper [13] presents a general approach covering all these domains). Acceleration for relations over Presburger arithmetic has been also plugged into abstraction/refinement loop for verifying integer programs [16, 26].

We recall that acceleration has also been applied proficiently in the analysis of real time systems (e.g., [8, 25]), to compactly represent the iterated execution of cyclic actions (e.g., polling-based systems) and address fragmentation problems.

Our work can be proficiently combined with SMT-based techniques for the verification of programs, as it helps helps avoiding the reachability analysis divergence when it comes to abstraction of programs with arrays of unknown length. Since the technique mostly operates at pre-processing level (we add to the system accelerated transitions by collapsing branches of loops handling arrays), *we believe that our technique is compatible with most approaches* proposed in array-based software model checking. We summarize some of these approaches below, without pretending of being exhaustive.

The vast majority of software model-checkers implement abstraction-refinement algorithms (e.g., [6, 18, 24]). *Lazy Abstraction with Interpolants* [30] is one of the most effective frameworks for unbounded reachability analysis of programs. It relies on the availability of interpolation procedures (nowadays efficiently embedded in SMT-Solvers [17]) to generate new predicates as (quantifier-free) interpolants for refining infeasible counterexamples.

For programs with arrays of unknown length the classical interpolation-based lazy abstraction works only if there is a support to handle quantified predicates [3] (the approach of [3] is the basis of our experiments below). Effectiveness and performances of abstraction/refinement approaches strongly depend on their ability in generating the "right" predicates to stop divergence of verification procedures. In case of programs with arrays, this quest can rely on *ghost variables* [21] retrieved from the post-conditions, on the backward propagation of post-conditions along spurious counterexamples [33] or can be constraint-based [9, 34]. Recently, constraint-based techniques have been significantly extended to the generation of loop invariants outside the array property fragment [29]. This solution exploits recent advantages in SMT-Solving, namely those devoted to finding solutions of constraints over non-linear integer arithmetic [10]. Other ways to generate predicates are by means of *saturation-based* theorem provers [28, 31] or interpolation procedures [3, 27].

```
function allDiff ( int a[N] ):
1 r = true;
2 for (i = 1; i < N ∧ r; i++)
3   for (j = i-1; j ≥ 0 ∧ r; j--)
4     if (a[i] = a[j]) r = false;
5 assert ( r → ( ∀x, y(0 ≤ x < y < N)
                  → (a[x] ≠ a[y]) ) )
```

(a)

```
function Reverse ( int I[N+1]; int O[N+1]; int c ):
1 c = 0;
2 while (c ≠ N+1) {O[c] = I[N−c]; c++;}
3 assert ( ∀x ≥ 0, y ≥ 0
           (x + y = N → I[x] = O[y] ) )
```

(b)

Figure 1: Motivating examples.

All the aforementioned techniques suffer from a certain degree of randomness due to the fact that detecting the "right" predicate is an undecidable problem. For example, predicate abstraction approaches (i.e., [3, 4, 33]) fail verifying the procedures in Figure1, which are commonly considered to be challenging for verifiers because they cause divergence[3]. Acceleration, on the other side, provides a precise and systematic way for addressing the verification of programs. Its combination, as a preprocessing procedure, with standard abstraction-refinement techniques allows to successfully solve challenging problems like the ones in Figure1.

The paper is structured as follows: Section 2 recalls the background notions about Presburger arithmetic and extensions. In order to identify the classes of relations whose acceleration we want to study, we are guided by software model checking applications. To this end, we provide in Section 3 classification of the guarded assignments we are interested in. Section 4 demonstrates the practical application of the theoretical results. In particular, it presents a backward reachability procedure and shows how to plug acceleration with monotonic abstraction in it. The details of the theoretical results are presented later. The main definability result for accelerations is in Section 6, while Section 5 introduces the abstract notion of an iterator. Section 7 discusses our experiments and Section 8 concludes the paper.

2 Preliminaries

We work in Presburger arithmetic enriched with free function symbols and with definable function symbols (see below); when we speak about validity or satisfiability of a formula, we mean *satisfiability and validity in all structures having the standard structure of natural numbers as reduct*. Thus, satisfiability and validity are decidable if we limit to quantifier-free formulæ (by adapting Nelson-Oppen combination results [32,35]), but may become undecidable otherwise (because of the presence of free function symbols).

We use $x, y, z, \dots$ or $i, j, k, \dots$ for variables; $t, u, \dots$ for terms, $c, d, \dots$ for free constants, $a, b, \dots$ for free function symbols, $\phi, \psi, \dots$ for *quantifier-free* formulæ. Bold letters are used for tuples and $|-|$ indicates tuples length; hence for instance $\mathbf{u}$ indicates a tuple of terms like $u_1, \dots, u_m$, where $m = |\mathbf{u}|$ (these tuples may contain repetitions). For variables, we use underline letters $\underline{x}, \underline{y}, \dots, \underline{i}, \underline{j}, \dots$ to indicates tuples without repetitions. Vector notation can also be used for equalities: if $\mathbf{u} = u_1, \dots, u_n$ and $\mathbf{v} = v_1, \dots, v_n$, we may use $\mathbf{u} = \mathbf{v}$ to mean the formula $\bigwedge_{i=1}^{n} u_i = v_i$.

If we write $t(x_1, \dots, x_n), \mathbf{u}(x_1, \dots, x_n), \phi(x_1, \dots, x_n)$ (or $t(\underline{x}), \mathbf{u}(\underline{x}), \phi(\underline{x}), \dots$, in case $\underline{x} = x_1, \dots, x_n$), we mean that the term t, the tuple of terms $\mathbf{u}$, the quantifier-free formula ϕ contain variables only from the tuple $x_1, \dots, x_n$. Similarly, we may use $t(\mathbf{a}, \mathbf{c}, \underline{x}), \phi(\mathbf{a}, \mathbf{c}, \underline{x}), \dots$ to mean *both* that the term t or the quantifier-free formula ϕ have free variables included in $\underline{x}$ *and* that the free function, free constants symbols occurring in them are among $\mathbf{a}, \mathbf{c}$. Notations like $t(\mathbf{u}/\underline{x}), \phi(\mathbf{u}/\underline{x}), \dots$ or $t(u_1/x_1, \dots, u_n/x_n), \phi(u_1/x_1, \dots, u_n/x_n), \dots$ - or occasionally just $t(\mathbf{u}), \phi(\mathbf{u}), \dots$ if confusion does not arise - are used for simultaneous substitutions within terms and formulæ. For a given natural number n, we use the standard abbreviations $\bar{n}$ and $n * y$ to denote the numeral of n (i.e. the term $s^n(0)$, where s is the successor function) and the sum of n addends all equal to y, respectively. If confusion does not arise, we may write just n for $\bar{n}$.

By a *definable function symbol*, we mean the following. Take a quantifier-free formula $\phi(\underline{j}, y)$ such that $\forall \underline{j} \exists! y\, \phi(\underline{j}, y)$ is valid ($\exists! y$ stands for 'there is a unique y such that ...'). Then a definable function symbol

[3] The procedure `Reverse` outputs to the array `O` the reverse of the array `I`; the procedure `allDiff` checks whether the entries of the array `a` are all different. Many thanks to Madhusudan Parthasarath and his group for pointing us to challenging problems with arrays of unknown length, including the `allDiff` example.

F (defined by ϕ) is a fresh function symbol, matching the length of $\underline{j}$ as arity, which is constrained to be interpreted in such a way that the formula $\forall y. F(\underline{j}) = y \longleftrightarrow \phi(\underline{j}, y)$ is true. The addition of definable function symbols does not affect decidability of quantifier-free formulæ and can be used for various purposes, for instance in order to express directly case-defined functions, array updates, etc. For instance, if a is a unary free function symbol, the term $wr(a, i, x)$ (expressing the update of the array a at position i by over-writing x) is a definable function; formally, we have $\underline{j} := i, x, j$ and $\phi(\underline{j}, y)$ is given by $(j = i \wedge y = x) \vee (j \neq i \wedge y = a(j))$. This formula $\phi(\underline{j}, y)$ (and similar ones) can be abbreviated like

$$y \ = (\texttt{if } j = i \texttt{ then } x \texttt{ else } a(j))$$

to improve readability. Another useful definable function is integer division by a fixed natural number n: to show that integer division by n is definable, recall that in Presburger arithmetic we have that $\forall x \ \exists! y \ \bigvee_{r=0}^{n-1}(x = n * y + r)$ is valid.

3 Programs representation

As a first step towards our main definability result, we provide a classification of the relations we are interested in. Such relations are guarded assignments required to model programs handling arrays of unknown length.

In our framework a *program* $\mathscr{P}$ is represented by a tuple $(\mathbf{v}, l_I, l_E, T)$; the tuple $\mathbf{v} := \mathbf{a}, \mathbf{c}, pc$ models system variables; formally, we have that

- the tuple $\mathbf{a} = a_1, \dots, a_s$ contains free unary function symbols, i.e., the arrays manipulated by the program;
- the tuple $\mathbf{c} = c_1, \dots, c_t$ contains free constants, i.e., the integer data manipulated by the program;
- the additional free constant pc (called *program counter*) is constrained to range over a finite set $L = \{l_1, ..., l_n\}$ of *program locations* over which we distinguish the *initial* and *error* locations denoted by l_I and l_E, respectively.

T is a set of finitely many formulæ $\{\tau_1(\mathbf{v}, \mathbf{v}'), \dots, \tau_r(\mathbf{v}, \mathbf{v}')\}$ called *transition formulæ* representing the program's body (here $\mathbf{v}'$ are renamed copies of the $\mathbf{v}$ representing the next-state variables). $\mathscr{P} = (\mathbf{v}, l_I, l_E, T)$ is *safe* iff there is no satisfiable formula like

$$(pc^0 = l_I) \wedge \tau_{i_1}(\mathbf{v}^0, \mathbf{v}^1) \wedge \dots \wedge \tau_{i_N}(\mathbf{v}^{N-1}, \mathbf{v}^N) \wedge (pc^N = l_E)$$

where $\mathbf{v}^0, \dots, \mathbf{v}^N$ are renamed copies of the $\mathbf{v}$ and each τ_{i_h} belongs to T.

Sentences denoting sets of states reachable by $\mathscr{P}$ can be:

- *ground* sentences, i.e., sentences of the kind $\phi(\mathbf{c}, \mathbf{a}, pc)$;
- Σ_1^0*-sentences*, i.e., sentences of the form $\exists \underline{i}.\ \phi(\underline{i}, \mathbf{a}, \mathbf{c}, pc)$;
- Σ_2^0*-sentences*, i.e., sentences of the form $\exists \underline{i} \forall \underline{j}.\ \phi(\underline{i}, \underline{j}, \mathbf{a}, \mathbf{c}, pc)$.

We remark that in our context satisfiability can be fully decided only for ground sentences and Σ_1^0-sentences (by Skolemization, as a consequence of the general combination results [32, 35]), while only subclasses of Σ_2^0-sentences enjoy a decision procedure [15, 22]. Transition formulæ can also be classified in three groups:

- *ground assignments*, i.e., transitions of the form

$$pc = l \ \wedge \ \phi_L(\mathbf{c}, \mathbf{a}) \ \wedge \ pc' = l' \ \wedge \ \mathbf{a}' = \lambda j.\ G(\mathbf{c}, \mathbf{a}, j) \ \wedge \ \mathbf{c}' = H(\mathbf{c}, \mathbf{a}) \tag{1}$$

- Σ_1^0*-assignments*, i.e., transitions of the form

$$\exists \underline{k} \begin{pmatrix} pc = l \ \wedge \ \phi_L(\mathbf{c}, \mathbf{a}, \underline{k}) \ \wedge \ pc' = l' \ \wedge \\ \mathbf{a}' = \lambda j.\ G(\mathbf{c}, \mathbf{a}, \underline{k}, j) \ \wedge \ \mathbf{c}' = H(\mathbf{c}, \mathbf{a}, \underline{k}) \end{pmatrix} \tag{2}$$

- Σ_2^0*-assignments*, i.e., transitions of the form

$$\exists \underline{k} \begin{pmatrix} pc = l \ \wedge \ \phi_L(\mathbf{c}, \mathbf{a}, \underline{k}) \ \wedge \ \forall \underline{j} \ \psi_U(\mathbf{c}, \mathbf{a}, \underline{k}, \underline{j}) \wedge \\ pc' = l' \ \wedge \ \mathbf{a}' = \lambda j.\ G(\mathbf{c}, \mathbf{a}, \underline{k}, j) \wedge \ \mathbf{c}' = H(\mathbf{c}, \mathbf{a}, \underline{k}) \end{pmatrix} \tag{3}$$

where $G = G_1, \ldots, G_s$, $H = H_1, \ldots, H_t$ are tuples of definable functions (vectors of equations like $\mathbf{a}' = \lambda j.\ G(\mathbf{c}, \mathbf{a}, \underline{k} j)$ can be replaced by the corresponding first order sentences $\forall j.\ \bigwedge_{h=1}^{s} a'_h(j) = G_h(\mathbf{c}, \mathbf{a}, \underline{k}, j)$).

The *composition* $\tau_1 \circ \tau_2$ of two transitions $\tau_1(\mathbf{v}, \mathbf{v}')$ and $\tau_2(\mathbf{v}, \mathbf{v}')$ is expressed by the formula $\exists \mathbf{v}_1(\tau_1(\mathbf{v}, \mathbf{v}_1) \wedge \tau_2(\mathbf{v}_1, \mathbf{v}'))$ (notice that composition may result in an inconsistent formula, e.g., in case of location mismatch). The *preimage* $Pre(\tau, K)$ of the set of states satisfying the formula $K(\mathbf{v})$ along the transition $\tau(\mathbf{v}, \mathbf{v}')$ is the set of states satisfying the formula $\exists \mathbf{v}'(\tau(\mathbf{v}, \mathbf{v}') \wedge K(\mathbf{v}'))$. The following proposition is immediate by straightforward syntactic manipulations:

Proposition 3.1. *Let τ, τ_1, τ_2 be transition formulæ and let $K(\mathbf{v})$ be a formula. We have that:* (i) *if τ_1, τ_2, τ, K are ground, then $\tau_1 \circ \tau_2$ is a ground assignment and $Pre(\tau, K)$ is a ground formula;* (ii) *if τ_1, τ_2, τ, K are Σ_1^0, then $\tau_1 \circ \tau_2$ is a Σ_1^0-assignment and $Pre(\tau, K)$ is a Σ_1^0-sentence;* (iii) *if τ_1, τ_2, τ, K are Σ_2^0, then $\tau_1 \circ \tau_2$ is a Σ_2^0-assignment and $Pre(\tau, K)$ is a Σ_2^0-sentence.*

4 Backward search and acceleration

This section demonstrates the practical applicability of the theoretical results of the paper in program verification. In particular, it presents the application of the accelerated transitions during reachability analysis for guarded-assignments representing programs handling arrays. For readability, we first present a basic reachability procedure. We subsequently analyze the divergence problems and show how acceleration can be applied to solve them. Acceleration application is not straightforward, though. The presence of accelerated transitions might generate undesirable Σ_2^0-sentences. The solution we propose is to over-approximate such sentences by adopting a selective instantiation schema, known in literature as *monotonic abstraction*. An enhanced reachability procedure integrating acceleration and monotonic abstraction concludes the Section.

The methodology we exploit to check safety of a program $\mathscr{P} = (\mathbf{v}, l_I, l_E, T)$ is *backward search*: we successively explore, through symbolic representation, all states leading to the error location l_E in one step, then in two steps, in three steps, etc. until either we find a fixpoint or until we reach l_I. To do this properly, it is convenient to build a tree: the tree has arcs labeled by transitions and nodes labeled by formulæ over $\mathbf{v}$. Leaves of the tree might be marked 'checked', 'unchecked' or 'covered'. The tree is built according to the following non-deterministic rules.

BACKWARD SEARCH

INITIALIZATION: a single node tree labeled by $pc = l_E$ and is marked 'unchecked'.

CHECK: pick an unchecked leaf L labeled with K. If $K \wedge pc = l_I$ is satisfiable ('safety test'), exit and return unsafe. If it is not satisfiable, check whether there is a set S of uncovered nodes such that (i) $L \notin S$ and (ii) K is inconsistent with the conjunction of the negations of the formulæ labeling the nodes in S ('fixpoint check'). If it is so, mark L as 'covered' (by S). Otherwise, mark L as 'checked'.

EXPANSION: pick a checked leaf L labeled with K. For each transition $\tau_i \in T$, add a new leaf below L labeled with $Pre(\tau_i, L)$ and marked as 'unchecked'. The arc between L and the new leaf is labeled with τ_i.

SAFETY EXIT: if all leaves are covered, exit and return safe.

The algorithm may not terminate (this is unavoidable by well-known undecidability results). Its correctness depends on the possibility of discharging safety tests with complete algorithms. By Proposition 3.1, if transitions are ground- or Σ_1^0-assignments, completeness of safety tests arising during the backward reachability procedure is guaranteed by the fact that satisfiability of Σ_1^0-formulæ is decidable. For fixpoint tests, sound but incomplete algorithms may compromise termination, but not correctness of the answer; hence for fixpoint tests, we can adopt incomplete pragmatic algorithms (e.g. if in fixpoint tests we need to test satisfiability of Σ_2^0-sentences, the obvious strategy is to Skolemize existentially quantified variables and to instantiate the universally quantified ones over sets of terms chosen according to suitable heuristics). To sum up, we have:

Proposition 4.1. *The above* BACKWARD SEARCH *procedure is partially correct for programs whose transitions are Σ_1^0-assignments, i.e., when the procedure terminates it gives a correct information about the safety of the input program.*

Divergence phenomena are usually not due to incomplete algorithms for fixpoint tests (in fact, divergence persists even in cases where fixpoint tests are precise).

Example 4.1. Consider a running example in Figure 1(b): it reverses the content of the array I into O. In our formalism, it is represented by the following transitions[4]:

$$\begin{aligned}
\tau_1 &\equiv \texttt{pc} = 1 \land \texttt{pc}' = 2 \land \texttt{c}' = 0 \\
\tau_2 &\equiv \texttt{pc} = 2 \land \texttt{c} \neq N+1 \land \texttt{c}' = \texttt{c}+1 \land O' = wr(O, \texttt{c}, I(N-\texttt{c})) \\
\tau_3 &\equiv \texttt{pc} = 2 \land \texttt{c} = N+1 \land \texttt{pc}' = 3 \\
\tau_4 &\equiv \texttt{pc} = 3 \land \exists z_1 \geq 0, z_2 \geq 0\, (z_1 + z_2 = N \land I(z_1) \neq O(z_2)\,) \land \texttt{pc}' = 4.
\end{aligned}$$

Notice that $\tau_1 - \tau_3$ all are ground assignments; only τ_4 (that translates the error condition) is a Σ_1^0-assignment. If we apply our tree generation procedure, we get an infinite branch, whose nodes - after routine simplifications - are labeled as follows

$$\begin{aligned}
&\ldots \\
&(K_i)\ \texttt{pc} = 2 \land \exists z_1, z_2\ \psi(z_1, z_2) \land \texttt{c} = N - i \land z_2 \neq N \land \cdots \land z_2 \neq N - i \\
&\ldots
\end{aligned}$$

where $\psi(z_1, z_2)$ stands for $z_1 \geq 0 \land z_2 \geq 0 \land z_1 + z_2 = N \land I(z_1) \neq O(z_2)$. □

As demonstrated by the above example, a divergence source comes from the fact that we are unable to represent *in one shot* the effect of executing finitely many times a given sequence of transitions. Acceleration can solve this problem.

Definition 4.1. The n-th composition of a transition $\tau(\mathbf{v}, \mathbf{v}')$ with itself is recursively defined by $\tau^1 := \tau$ and $\tau^{n+1} := \tau \circ \tau^n$. The *acceleration* τ^+ of τ is $\bigvee_{n \geq 1} \tau^n$.

In general, acceleration requires a logic supporting infinite disjunctions. Notable exceptions are witnessed by Theorem 6.1. For now we focus on examples where accelerations yield Σ_2^0-assignments starting from ground assignments.

Example 4.2. Recall transition τ_2 from the running example.

$$\tau_2 \equiv \texttt{pc} = 2 \land \texttt{c} \neq N+1 \land pc' = 2 \land \texttt{c}' = \texttt{c}+1 \land I' = I \land O' = wr(O, \texttt{c}, I(N-\texttt{c}))$$

(here we displayed identical updates for completeness). Notice that the variable pc is left unchanged in this transition (this is essential, otherwise the acceleration gives an inconsistent transition that can never fire). If we accelerate it, we get the Σ_2^0-assignment[5]

$$\exists n > 0 \left(\begin{array}{c} \texttt{pc} = 2\ \land\ \forall j\, (\texttt{c} \leq j < \texttt{c}+n \to j \neq N+1)\ \land\ \texttt{c}' = \texttt{c}+n\ \land \\ \land\ \texttt{pc}' = 2\ \land\ O' = \lambda j\, (\texttt{if}\ \texttt{c} \leq j < \texttt{c}+n\ \texttt{then}\ I(N-j)\ \texttt{else}\ O(j)) \end{array} \right) \tag{4}$$

□

In presence of these accelerated Σ_2^0-assignments, Backward Search can produce problematic Σ_2^0-sentences (see Proposition 3.1 above) which cannot be handled precisely by existing solvers. As a solution to this problem we propose applying to such sentences a suitable abstraction, namely *monotonic abstraction.*

Definition 4.2. Let $\psi :\equiv \exists \underline{i}\, \forall \underline{j}.\ \phi(\underline{i}, \underline{j}, \mathbf{a}, \mathbf{c}, pc)$ be a Σ_2^0-sentences and let $\mathscr{S}$ be a finite set of terms of the kind $t(\underline{i}, \mathbf{v})$. The *monotonic $\mathscr{S}$-approximation of* ψ is the Σ_1^0-sentence

$$\exists \underline{i} \bigwedge_{\sigma : \underline{j} \to \mathscr{S}} \phi(\underline{i}, \underline{j}\sigma / \underline{j}, \mathbf{a}, \mathbf{c}, pc) \tag{5}$$

(here $\underline{j}\sigma$, if $\underline{j} = j_1, \ldots, j_n$, is the tuple of terms $\sigma(j_1), \ldots, \sigma(j_n)$).

By Definition 4.2, universally quantified variables are *eliminated through instantiation*; the larger the set $\mathscr{S}$ is, the better approximation you get. In practice, the natural choices for $\mathscr{S}$ are $\underline{i}$ or the set of terms of the kind $t(\underline{i}, \mathbf{v})$ occurring in ψ (we adopted the former choice in our implementation). As a result of replacing Σ_2^0-sentences by their monotonic approximation, spurious unsafe traces might occur. However, **those can be disregarded if accelerated transitions contribute to their generation**. This is because if $\mathscr{P}$ is unsafe, then unsafety can be discovered without appealing to accelerated transitions.

To integrate monotonic abstraction, the above Backward Search procedure is modified as follows. In a Preprocessing step, we add some accelerated transitions of the kind $(\tau_1 \circ \cdots \circ \tau_n)^+$ to T. These transitions

[4] For readability, we omit identical updates like $I' = I$, etc. Notice that we have $l_I = 1$ and $l_E = 4$.
[5] This Σ_2^0-assignment can be automatically computed using procedures outlined in the proof of Theorem 6.1.

can be found by inspecting cycles in the control flow graph of the program and accelerating them following the procedure described in Sections 5, 6. The natural cycles to inspect are those corresponding to loop branches in the source code. It should be noticed, however, that identifying the good cycles to accelerate is subject to specific heuristics that deserve separate investigation in case the program has infinitely many cycles. (choosing cycles from branches of innermost loops is the simplest example of such heuristics and the one we implemented).

After this extra preprocessing step, the remaining instructions are left unchanged, with the exception of CHECK that is modified as follows:

CHECK': pick an unchecked leaf L labeled by a formula K. If K is a Σ^0_2-sentence, choose a suitable $\mathscr{S}$ and replace K by its monotonic $\mathscr{S}$-abstraction K'. If $K' \wedge pc = l_I$ is inconsistent, mark L as 'covered' or 'checked' according to the outcome of the fixpoint check, as was done in the original CHECK. If $K' \wedge pc = l_I$ is satisfiable, analyze the path from the root to L. If no accelerated transition τ^+ is found in it return $\mathtt{unsafe}$, otherwise remove the sub-tree D from the target of τ^+ to the leaves. Each node N covered by a node in D will be flagged as 'unchecked' (to make it eligible in future for the EXPANSION instruction).

The new procedure will be referred as BACKWARD SEARCH'. It is quite straightforward to see that Proposition 4.1 still applies to the modified algorithm. Notice that, although termination cannot be ensured (given well-known undecidability results), spurious traces containing approximated accelerated transitions cannot be produced again and again: when the sub-tree D from the target node v of τ^+ is removed by CHECK', the node v is not a leaf (the arcs labeled by the transitions τ are still there), hence it cannot be expanded anymore according to the EXPANSION instruction.

Example 4.3. Let again consider our running example and demonstrate how acceleration and monotonic abstraction work. In the preprocessing step, we add the accelerated transition τ_2^+ given by (4) to the transitions we already have. After having computed $(K') \equiv Pre(\tau_4, K), (K'') \equiv Pre(\tau_3, K')$, we compute $(\tilde{K}) \equiv Pre(\tau_2^+, K'')$ and get

$$\exists n > 0 \exists z_1, z_2 \begin{pmatrix} \mathtt{pc} = 2 \wedge \forall j\, (\mathtt{c} \leq j < \mathtt{c}+n \rightarrow j \neq N+1) \wedge \\ \wedge\, \mathtt{c}+n = N+1 \wedge z_1 \geq 0 \wedge z_2 \geq 0 \wedge z_1 + z_2 = N \wedge \\ \wedge\, I(z_1) \neq \lambda j\, (\mathtt{if}\ \mathtt{c} \leq j < \mathtt{c}+n\ \mathtt{then}\ I(N-j)\ \mathtt{else}\ O(j))(z_2) \end{pmatrix}$$

We approximate using the set of terms $\mathscr{S} = \{z_1, z_2, n\}$. After simplifications we get

$$\exists z_1, z_2 (\mathtt{pc} = 2 \wedge \mathtt{c} \leq N \wedge z_1 \geq 0 \wedge z_2 \geq 0 \wedge z_1 + z_2 = N \wedge O(z_2) \neq I(z_1) \wedge \mathtt{c} > z_2)$$

Generating this formula is enough to stop divergence. □

Notice that in the computations of the above example we eventually succeeded in eliminating the extra quantifier $\exists n$ introduced by the accelerated transition. This is not always possible: sometimes in fact, to get the good invariant one needs more quantified variables than those occurring in the annotated program and accelerated transitions might be the way of getting such additional quantified variables. As an example of this phenomenon, consider the `init+test` program included in our benchmark suite of Section 7 below.

5 Iterators

This Section introduces *iterators* and *selectors*, two main ingredients used to supply a useful format to compute accelerated transitions. Iterators are meant to formalize the notion of a counter scanning the indexes of an array: the most simple iterators are increments and decrements, but one may also build more complex ones for different scans, like in binary search. We give their formal definition and then we supply some examples. We need to handle tuples of terms because we want to consider the case in which we deal with different arrays with possibly different scanning variables. Given a m-tuple of terms

$$\mathbf{u}(\underline{x}) := u_1(x_1, \ldots, x_m), \ldots, u_m(x_1, \ldots, x_m) \tag{6}$$

containing the m variables $\underline{x} = x_1, \ldots, x_m$, we indicate with $\mathbf{u}^n$ the term expressing the n-times composition of (the function denoted by) $\mathbf{u}$ with itself. Formally, we have $\mathbf{u}^0(\underline{x}) := \underline{x}$ and

$$\mathbf{u}^{n+1}(\underline{x}) := u_1(\mathbf{u}^n(\underline{x})), \ldots, u_m(\mathbf{u}^n(\underline{x}))\,.$$

Definition 5.1. A tuple of terms $\mathbf{u}$ like (6) is said to be an *iterator* iff there exists an m-tuple of $m+1$-ary terms $\mathbf{u}^*(\underline{x}, y) := u_1^*(x_1, \ldots, x_m, y), \ldots, u_m^*(x_1, \ldots, x_m, y)$ such that for any natural number n it happens that the formula

$$\mathbf{u}^n(\underline{x}) = \mathbf{u}^*(\underline{x}, \bar{n}) \tag{7}$$

is valid.[6] Given an iterator $\mathbf{u}$ as above, we say that an m-ary term $\kappa(x_1, \ldots, x_m)$ is a *selector* for $\mathbf{u}$ iff there is an $m+1$-ary term $\iota(x_1, \ldots, x_m, y)$ yielding the validity of the formula

$$z = \kappa(\mathbf{u}^*(\underline{x}, y)) \to y = \iota(\underline{x}, z) \ . \tag{8}$$

The meaning of condition (8) is that, once the input $\underline{x}$ and the selected output z are known, it is possible to identify uniquely (through ι) the number of iterations y that are needed to get z by applying κ to $\mathbf{u}^*(\underline{x}, y)$. The term κ is a selector function that selects (and possibly modifies) one of the $\mathbf{u}$; in most applications (though not always) κ is a projection, represented as a variable x_i (for $1 \le i \le m$), so that $\kappa(\mathbf{u}^*(\underline{x}, y))$ is just the i-th component $u_i^*(\underline{x}, y)$ of the tuple of terms $\mathbf{u}^*(\underline{x}, y)$. In these cases, the formula (8) reads as

$$z = u_i^*(\underline{x}, y) \to y = \iota(\underline{x}, z) \, . \tag{9}$$

Example 5.1. The canonical example is when we have $m = 1$ and $\mathbf{u} := u_1(x_1) := x_1 + 1$; this is an iterator with $u_1^*(x_1, y) := x_1 + y$; as a selector, we can take $\kappa(x_1) := x_1$ and $\iota(x_1, z) := z - x_1$. □

Example 5.2. The previous example can be modified, by choosing $\mathbf{u}$ to be $x_1 + \bar{n}$, for some integer $n \neq 0$: then we have $u^*(x_1, y) := x_1 + n * y$, $\kappa(x_1) := x_1$, and $\iota(x_1, z) = (z - x_1)//n$ where $//$ is integer division (recall that integer division by a given n is definable in Presburger arithmetic). □

Example 5.3. If we move to more expressive arithmetic theories, like Primitive Recursive Arithmetic (where we have a symbol for every primitive recursive function), we can get much more examples. As an example with $m > 1$, we can take $\mathbf{u} := x_1 + x_2, x_2$ and get $u_1^*(x_1, x_2, y) = x_1 + y * x_2$, $u_2^*(x_1, x_2, y) = x_2$. Here a selector is for instance $\kappa_1(x_1, x_2) := \bar{7} + x_1$, $\iota(x_1, x_2, z) := (z - x_1 - \bar{7})//x_2$. □

6 Accelerating local ground assignments

Back to our program $\mathscr{P} = (\mathbf{v}, l_I, l_E, T)$, we look for conditions on transitions from T allowing to accelerate them via a Σ_2^0-assignment. Given an iterator $\mathbf{u}(\underline{x})$, a *selector assignment* for $\mathbf{a} := a_1, \ldots, a_s$ (relative to $\mathbf{u}$) is a tuple of selectors $\kappa := \kappa_1, \ldots, \kappa_s$ for $\mathbf{u}$. Intuitively, the components of the tuple are meant to indicate the scanners of the arrays $\mathbf{a}$ and as such might not be distinct (although, of course, just *one* selector is assigned to each array). A formula ψ (resp. a term t) is said to be *purely arithmetical* over a finite set of terms V iff it is obtained from a formula (resp. a term) *not containing the extra free function symbols* $\mathbf{a}, \mathbf{c}$ by replacing some free variables in it by terms from V. Let $\mathbf{v} = v_1, \ldots, v_s$ and $\mathbf{w} = w_1, \ldots, w_s$ be s-tuples of terms; below $wr(\mathbf{a}, \mathbf{v}, \mathbf{w})$ and $\mathbf{a}(\mathbf{v})$ indicate the tuples $wr(a_1, v_1, w_1), \ldots, wr(a_s, v_s, w_s)$ and $a_1(v_1), \ldots, a_s(v_s)$, respectively (recall from Section 3 that $s = |\mathbf{a}|$).

Definition 6.1. A *local ground assignment* is a ground assignment of the form

$$pc = l \ \wedge \ \phi_L(\mathbf{c}, \mathbf{a}) \ \wedge \ pc' = l \ \wedge \mathbf{a}' = wr(\mathbf{a}, \kappa(\tilde{\mathbf{c}}), \mathbf{t}(\mathbf{c}, \mathbf{a})) \ \wedge \ \tilde{\mathbf{c}}' = \mathbf{u}(\tilde{\mathbf{c}}) \ \wedge \mathbf{d}' = \mathbf{d} \tag{10}$$

where (i) $\mathbf{c} = \tilde{\mathbf{c}}, \mathbf{d}$; (ii) $\mathbf{u} = u_1, \ldots, u_{|\tilde{\mathbf{c}}|}$ is an iterator; (iii) the terms κ are a selector assignment for $\mathbf{a}$ relative to $\mathbf{u}$; (iv) the formula $\phi_L(\mathbf{c}, \mathbf{a})$ and the terms $\mathbf{t}(\mathbf{c}, \mathbf{a})$ are purely arithmetical over the set of terms $\{\mathbf{c}, \mathbf{a}(\kappa(\tilde{\mathbf{c}}))\} \cup \{a_i(d_j)\}_{1 \le i \le s, 1 \le j \le |\mathbf{d}|}$; (v) the guard ϕ_L contains the conjuncts $\kappa_i(\tilde{\mathbf{c}}) \neq d_j$, for $1 \le i \le s$ and $1 \le j \le |\mathbf{d}|$.

Thus in a local ground assignment, there are various restrictions: (a) the numerical variables are split into 'idle' variables $\mathbf{d}$ and variables $\tilde{\mathbf{c}}$ subject to update via an iterator $\mathbf{u}$; (b) the program counter is not modified; (c) the guard does not depend on the values of the a_i at cells different from $\kappa_i(\tilde{\mathbf{c}}), \mathbf{d}$; (d) the update of the $\mathbf{a}$ are simultaneous writing operations modifying only the entries $\kappa(\tilde{\mathbf{c}})$. Thus, the assignment is local and the relevant modifications it makes are determined by the selectors locations. The 'idle' variables $\mathbf{d}$ are useful to accelerate branches of nested loops; the inequalities mentioned in (v) are automatically generated by making case distinctions in assignment guards.

[6] Recall that $\bar{n}$ is the numeral of n, i.e. it is $s^n(0)$.

Example 6.1. For our running example, we show that transition τ_2 (the one we want to accelerate) is a local ground assignment. We have $\mathbf{d} = \emptyset$ and $\tilde{\mathbf{c}} = \mathtt{c}$ and $\mathbf{a} = I, O$. The counter $\mathtt{c}$ is incremented by 1 at each application of τ_2. Thus, our iterator is $\mathbf{u} := x_1 + 1$ and the selector assignment assigns $\kappa_1 := N - x_1$ to I and $\kappa_2 := x_1$ to O. In this way, I is modified (identically) at $N - \mathtt{c}$ via $I' = wr(I, N - \mathtt{c}, I(N - \mathtt{c}))$ and O is modified at $\mathtt{c}$ via $O' = wr(O, \mathtt{c}, I(N - \mathtt{c}))$. The guard τ_2 is $\mathtt{c} \neq N + 1$. Since the formula $\mathtt{c} \neq N + 1$ and the term $I(N - \mathtt{c})$ are purely arithmetical over $\{\mathtt{c}, I(N - \mathtt{c}), O(\mathtt{c})\}$, we conclude that τ_2 is local. □

Theorem 6.1. *If τ is a local ground assignment, then τ^+ is a Σ_2^0-assignment.*

Proof. (Sketch, see Appendix A for full details). Let us fix the local ground assignment (10); let $\mathbf{a}[\mathbf{d}]$ indicate the $s * |\mathbf{d}|$-tuple of terms $\{a_i(d_j)\}_{1 \le i \le s, 1 \le j \le |\mathbf{d}|}$; since ϕ_L and $\mathbf{t} := t_1, \ldots, t_s$ are purely arithmetical over $\{\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}]\}$, we have that they can be written as $\tilde{\phi}_L(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}])$, $\tilde{\mathbf{t}}(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}])$, respectively, where $\tilde{\phi}_L, \tilde{\mathbf{t}}$ do not contain occurrences of the free function and constant symbols $\mathbf{a}, \mathbf{c}$. The transition τ^+ can be expressed as a Σ_2^0-assignment by

$$\exists y > 0 \begin{pmatrix} \forall z\,(0 \le z < y \to \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, z), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\mathbf{c}, z))), \mathbf{a}[\mathbf{d}]) \wedge \mathbf{d}' = \mathbf{d} \wedge \\ \wedge\, pc = l \,\wedge\, pc' = l \,\wedge\, \tilde{\mathbf{c}}' = \mathbf{u}^*(\tilde{\mathbf{c}}, y) \,\wedge\, \mathbf{a}' = \lambda j.\, F(\mathbf{c}, \mathbf{a}, y, j) \end{pmatrix}$$

where the tuple $F = F_1, \ldots, F_s$ of definable functions is given by

$$\begin{aligned} F_h(\mathbf{c}, \mathbf{a}, y, j) \;=\; \lambda j.\; & \mathtt{if}\ 0 \le \iota_h(\tilde{\mathbf{c}}, j) < y\ \wedge j = \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j)))\ \mathtt{then} \\ & \tilde{t}_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \iota_h(\tilde{\mathbf{c}}, j)), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \iota_h(\tilde{\mathbf{c}}, j)))), \mathbf{a}[\mathbf{d}])\ \mathtt{else}\ a_h[j] \end{aligned}$$

for $h = 1, \ldots, s$ (here $\iota_1, \ldots, \iota_s$ are the terms corresponding to $\kappa_1, \ldots, \kappa_s$ according to the definition of a selector for the iterator $\mathbf{u}$). ⊣ □

We point out that the effective use of Theorem 6.1 relies on the implementation of a repository of iterators and selectors and of algorithms recognizing them. The larger the repository is, the more possibilities the model checker has to exploit the full power of acceleration.

In most applications it is sufficient to consider accelerated transitions of the canonical form of Example 5.1. Let us examine in details this special case; here $\mathbf{c}$ is a single counter $\mathtt{c}$ that is incremented by one (otherwise said, the iterator is $x_1 + 1$) and the selector assignment is trivial, namely it is just x_1. We call these local ground assignments *simple*. Thus, a simple local ground assignment has the form

$$pc = l \,\wedge\, \phi_L(\mathtt{c}, \mathbf{a}) \,\wedge\, pc' = l \wedge \mathtt{c}' = \mathtt{c} + 1 \,\wedge\, \mathbf{a}' = wr(\mathbf{a}, \mathtt{c}, \mathbf{t}(\mathtt{c}, \mathbf{a})) \tag{11}$$

where the first occurrence of $\mathtt{c}$ in $wr(\mathbf{a}, \mathtt{c}, \mathbf{t}(\mathtt{c}, \mathbf{a}))$ stands in fact for an s-tuple of terms all identical to $\mathtt{c}$, and where $\phi_L, \mathbf{t}$ are purely arithmetical over the terms $\mathtt{c}, a_1[\mathtt{c}], \ldots, a_s[\mathtt{c}]$. The accelerated transition computed in the proof of Theorem 6.1 for (11) can be rewritten as follows:

$$\exists k \begin{pmatrix} k > 0 \wedge pc = l \,\wedge\, \forall j\,(\mathtt{c} \le j < \mathtt{c} + k \to \phi_L(j, \mathbf{a})) \,\wedge\, pc' = l \,\wedge \\ \wedge\, \mathtt{c}' = \mathtt{c} + k \,\wedge\, \mathbf{a}' = \lambda j.(\mathtt{if}\ \mathtt{c} \le j < \mathtt{c} + k\ \mathtt{then}\ \mathbf{t}(j, \mathbf{a})\ \mathtt{else}\ \mathbf{a}[j]) \end{pmatrix} \tag{12}$$

A slight extension of the notion of a simple assignment leads to a further subclass of local ground assignments useful to accelerated branches of nested loops (see Appendix B for more details).

7 Experimental evaluation

We implemented the algorithm described in Section 4 - 6 as a preprocessing module inside the MCMT model checker [23]. To perform a feasibility study, we intentionally focused our implementation on simple and simple+ local ground assignments. For a thorough and unbiased evaluation we compared/combined the new technique with an abstraction algorithm suited for array programs [3] implemented in the same tool. This section describes benchmarks and discusses experimental results. A clear outcome from our experiments is that abstraction/refinement and acceleration techniques can be gainfully combined.

Benchmarks. We evaluated the new algorithm on 55 programs with arrays, each annotated with an assertion. We considered only quantifier-free or ∀-assertions. Our set of benchmarks comprises programs used to evaluate the Lazy Abstraction with Interpolation for Arrays framework [4] and other focused benchmarks where

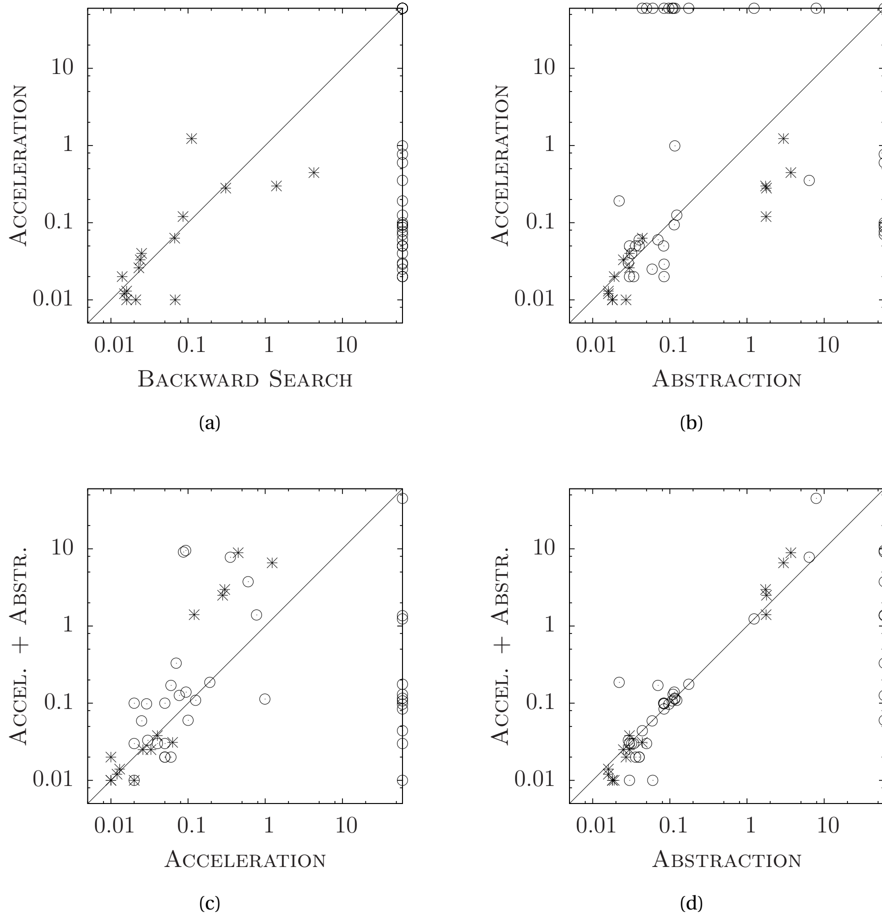

Figure 2: Comparison of time for different options of BACKWARD SEARCH. Stars and circles represent buggy and correct programs respectively.

abstraction diverges. These are problems involving array manipulations as copying, comparing, searching, sorting, initializing, testing, etc. About one third of the programs contain bugs.[7]

Evaluation. Experiments have been run on a machine equipped with a i7@2.66 GHz CPU and 4GB of RAM running OS X. Time limit for each experiment has been set to 60 seconds. We run MCMT with four different configurations:

- BACKWARD SEARCH - MCMT executes the procedure described at the beginning of Section 4.
- ABSTRACTION - MCMT integrates the backward reachability algorithm with the abstraction/refinement loop [3].
- ACCELERATION - The transition system is pre-processed in order to compute accelerated transitions (when it is possible) and then the BACKWARD SEARCH' procedure is executed.
- ACCEL. + ABSTR. - This configuration enables both the preprocessing step in charge of computing accelerated transitions and the abstraction/refinement engine on the top of the BACKWARD SEARCH' procedure.

The complete statistics can be found in Appendix C. In summary, the comparative analysis of timings presented in Figure2 confirms that acceleration indeed helps to avoid divergence for problematic programs where abstraction fails. The first comparison (Figure2(a)) highlights the benefits of using acceleration: BACKWARD SEARCH diverges on all 39 safe instances. Acceleration stops divergence in 23 cases, and moreover the

[7] The set of benchmarks can be downloaded from http://www.inf.usi.ch/phd/alberti/prj/acc; the tool set MCMT is available at http://users.mat.unimi.it/users/ghilardi/mcmt/.

overhead introduced by the preprocessing step does not affect unsafe instances. Figure2(b) shows that acceleration and abstraction are two complementary techniques, since MCMT times out in both cases but for two different sets of programs. Figure2(c) and Figure2(d) attest that acceleration and abstraction/refinement techniques mutually benefit from each other: with both techniques MCMT solves all the 55 benchmarks.

8 Conclusion and Future Work

We identified a class of transition relations involving array updates that can be accelerated, showed how it is possible to compute the accelerated transition and describe a solution for dealing with universal quantifiers arising from the acceleration process. Our paper lays theoretical foundations for this interesting research topic and confirms by our prototype experiments on challenging benchmarks its advantages over standalone verification approaches since it's able to solve problems on which other techniques fail to converge.

As future directions, a challenging task is to enlarge the definability result of Theorem 6.1 so as to cover classes of transitions modeling more and more loop branches arising from concrete programs. In addition, one may want to consider more sophisticated strategies for instantiation in order to support acceleration. Considering increasing larger $\mathscr{S}$ or handling Σ^0_2-sentences when they belong to decidable fragments [15, 22] may lead to further improvements.

References

[1] P.A. Abdulla, G. Delzanno, N.B. Henda, and A. Rezine. Regular model checking without transducers. In *TACAS*, volume 4424 of *LNCS*, pages 721–736, 2007.

[2] P.A. Abdulla, G. Delzanno, and A. Rezine. Parameterized verification of infinite-state processes with global conditions. In *CAV*, LNCS, pages 145–157, 2007.

[3] F. Alberti, R. Bruttomesso, S. Ghilardi, S. Ranise, and N. Sharygina. Lazy Abstraction with Interpolants for Arrays. In *LPAR*, pages 46–61, 2012.

[4] F. Alberti, R. Bruttomesso, S. Ghilardi, S. Ranise, and N. Sharygina. SAFARI: SMT-Based Abstraction for Arrays with Interpolants. In *CAV*, 2012.

[5] F. Alberti, S. Ghilardi, E. Pagani, S. Ranise, and G.P. Rossi. Universal Guards, Relativization of Quantifiers, and Failure Models in Model Checking Modulo Theories. *JSAT*, pages 29–61, 2012.

[6] Thomas Ball and Sriram K. Rajamani. The slam toolkit. In *CAV*, pages 260–264, 2001.

[7] Clark Barrett, Aaron Stump, and Cesare Tinelli. The SMT-LIB Standard: Version 2.0. `www.SMT-LIB.org`, 2010.

[8] G. Behrmann, J. Bengtsson, A. David, K.G. Larsen, P. Pettersson, and W. Yi. Uppaal implementation secrets. In *FTRTFT*, pages 3–22, 2002.

[9] Dirk Beyer, Thomas A. Henzinger, Rupak Majumdar, and Andrey Rybalchenko. Path invariants. In *PLDI*, pages 300–309, 2007.

[10] Cristina Borralleras, Salvador Lucas, Albert Oliveras, Enric Rodríguez-Carbonell, and Albert Rubio. Sat modulo linear arithmetic for solving polynomial constraints. *J. Autom. Reasoning*, 48(1):107–131, 2012.

[11] M. Bozga, C. Girlea, and R. Iosif. Iterating octagons. In *TACAS*, LNCS, pages 337–351, 2009.

[12] M. Bozga, P. Habermehl, R. Iosif, F. Konecný, and T. Vojnar. Automatic verification of integer array programs. In *CAV*, pages 157–172, 2009.

[13] M. Bozga, R. Iosif, and F. Konecny. Fast acceleration of ultimately periodic relations. In *CAV*, LNCS, 2010.

[14] M. Bozga, R. Iosif, and Y. Lakhnech. Flat parametric counter automata. *Fundamenta Informaticae*, (91):275–303, 2009.

[15] A.R. Bradley, Z. Manna, and H.B. Sipma. What's decidable about arrays? In *VMCAI*, pages 427–442, 2006.

[16] Nicolas Caniart, Emmanuel Fleury, Jérôme Leroux, and Marc Zeitoun. Accelerating interpolation-based model-checking. In *TACAS*, pages 428–442, 2008.

[17] Alessandro Cimatti, Alberto Griggio, and Roberto Sebastiani. Efficient generation of craig interpolants in satisfiability modulo theories. *ACM Trans. Comput. Log.*, 12(1):7, 2010.

[18] E.M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-Guided Abstraction Refinement. In *CAV*, pages 154–169, 2000.

[19] H. Comon and Y. Jurski. Multiple counters automata, safety analysis and presburger arithmetic. In *CAV*, volume 1427 of *LNCS*, pages 268–279. Springer, 1998.

[20] A. Finkel and J. Leroux. How to compose presburger-accelerations: Applications to broadcast protocols. In *FST TCS '02*, pages 145–156. Springer, 2002.

[21] C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *POPL*, pages 191–202, 2002.

[22] Y. Ge and L. de Moura. Complete instantiation for quantified formulas in satisfiabiliby modulo theories. In *CAV*, pages 306–320, 2009.

[23] S. Ghilardi and S. Ranise. MCMT: A Model Checker Modulo Theories. In *IJCAR*, pages 22–29, 2010.

[24] S. Graf and H. Saïdi. Construction of Abstract State Graphs with PVS. In *CAV*, pages 72–83, 1997.

[25] M. Hendriks and K.G. Larsen. Exact acceleration of real-time model checking. *Electr. Notes Theor. Comput. Sci.*, 65(6):120–139, 2002.

[26] H. Hojjat, R. Josif, F. Konecny, V. Kuncak, and P. Rümmer. On accelerating interpolants. In *ATVA*, 2012.

[27] R. Jhala and K.L. McMillan. Array Abstractions from Proofs. In *CAV*, 2007.

[28] L. Kovács and A. Voronkov. Interpolation and Symbol Elimination. In *CADE*, 2009.

[29] Daniel Larraz, Enric Rodríguez-Carbonell, and Albert Rubio. Smt-based array invariant generation. In *VMCAI*, pages 169–188, 2013.

[30] K.L. McMillan. Lazy Abstraction with Interpolants. In *CAV*, 2006.

[31] K.L. McMillan. Quantified Invariant Generation Using an Interpolating Saturation Prover. In *TACAS*, 2008.

[32] G. Nelson and D.C. Oppen. Simplification by cooperating decision procedures. *ACM Transaction on Programming Languages and Systems*, 1(2):245–257, 1979.

[33] M. N. Seghir, A. Podelski, and T. Wies. Abstraction Refinement for Quantified Array Assertions. In *SAS*, pages 3–18, 2009.

[34] S. Srivastava and S. Gulwani. Program Verification using Templates over Predicate Abstraction. In *PLDI*, 2009.

[35] C. Tinelli and M. T. Harandi. A new correctness proof of the Nelson-Oppen combination procedure. In *Proc. of FroCoS 1996*, pages 103–119. Kluwer, 1996.

A Proof of Theorem 6.1

In this technical Appendix, we supply the proof of Theorem 6.1.

Proof. As a preliminary observation, we notice that the bi-implications of the kind

$$(\bigvee_{n\geq 0} \psi(\underline{x}, \bar{n})) \longleftrightarrow \exists y\,(y \geq 0 \wedge \psi(\underline{x}, y))\,. \tag{13}$$

are valid because we interpret our formulæ in the *standard* structure of natural numbers (enriched with extra free symbols).

As a second preliminary observation, we notice that (8) can be equivalently re-writtem in the form of a bi-implication as:

$$z = \kappa(\mathbf{u}^*(\underline{x}, y)) \quad \longleftrightarrow \quad [\, y = \iota(\underline{x}, z) \wedge z = \kappa(\mathbf{u}^*(\underline{x}, \iota(\underline{x}, z)))\,] \tag{14}$$

(to see why (14) is equivalent to (8) it is sufficient to apply the logical laws of pure identity).

Let us fix a local ground assignment of the form (10); let $\mathbf{a}[\mathbf{d}]$ indicate the $s*|\mathbf{d}|$-tuple of terms $\{a_i(d_j)\}_{1\leq i\leq s, 1\leq j\leq |\mathbf{d}|}$; since ϕ_L and $\mathbf{t}$ are purely arithmetical over $\{\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}]\}$, we have that they can be written as $\tilde{\phi}_L(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}])$, $\tilde{\mathbf{t}}(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}])$, respectively, where $\tilde{\phi}_L, \tilde{\mathbf{t}}$ do not contain occurrences of the free function and constant symbols $\mathbf{a}, \mathbf{c}$.

Claim. As a first step, we show by induction on n that τ^n can be expressed as follows (we omit here and below the conjuncts $pc = l \wedge pc' = l \wedge \mathbf{d}' = \mathbf{d}$ that do not play any role)

$$\bigwedge_{0\leq k<n} \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k}), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k}))), \mathbf{a}[\mathbf{d}]) \wedge \tilde{\mathbf{c}}' = \mathbf{u}^*(\tilde{\mathbf{c}}, \bar{n}) \wedge \mathbf{a}' = \lambda j.\, F(\mathbf{c}, \mathbf{a}, \bar{n}, j) \tag{15}$$

where the tuple $F = F_1, \ldots, F_s$ of definable functions is given by[8]

$$\begin{aligned} F_h(\mathbf{c}, \mathbf{a}, y, j) \;=\; \lambda j.\; & \texttt{if } 0 \leq \iota_h(\tilde{\mathbf{c}}, j) < y \wedge j = \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j))) \texttt{ then} \\ & \tilde{t}_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \iota_h(\tilde{\mathbf{c}}, j)), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \iota_h(\tilde{\mathbf{c}}, j)))), \mathbf{a}[\mathbf{d}]) \texttt{ else } a_h[j] \end{aligned} \tag{16}$$

for $h = 1, \ldots, s$ (here $\iota_1, \ldots, \iota_s$ are the terms corresponding to $\kappa_1, \ldots, \kappa_s$ according to the definition of a selector for the iterator $\mathbf{u}$).

Proof of the Claim. For $n = 1$, notice that $\tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, 0), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, 0))), \mathbf{a}[\mathbf{d}])$ is equivalent to $\tilde{\phi}_L(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}])$, that $\tilde{\mathbf{c}}' = \mathbf{u}^*(\tilde{\mathbf{c}}, \bar{1})$ is equivalent to $\tilde{\mathbf{c}}' = \mathbf{u}(\tilde{\mathbf{c}})$ and that $\lambda j.\, F(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}, \bar{1}, j) = wr(\mathbf{a}, \kappa(\tilde{\mathbf{c}}), \mathbf{t}(\tilde{\mathbf{c}}, \mathbf{d}, \mathbf{a}(\kappa(\tilde{\mathbf{c}})), \mathbf{a}[\mathbf{d}]))$ holds (the latter because for every h, $\iota_h(\tilde{\mathbf{c}}, j) = 0 \wedge j = \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j)))$ is equivalent to $j = \kappa_h(\mathbf{u}^*(\tilde{\mathbf{c}}, 0)) = \kappa_h(\tilde{\mathbf{c}})$ by (14)).

For the induction step, we suppose the Claim holds for n and show it for $n+1$. As a preliminary remark, notice that from (10), we get not only $\mathbf{d}' = \mathbf{d}$, but also $\mathbf{a}'[\mathbf{d}'] = \mathbf{a}[\mathbf{d}]$, because of (v) of Definition 6.1. As a consequence, after n iterations of τ, the values $\mathbf{d}, \mathbf{a}[\mathbf{d}]$ are left unchanged; thus, for notation simplicity, we will not display anymore below the dependence of $\phi_L, \tilde{\mathbf{t}}$ on $\mathbf{d}, \mathbf{a}[\mathbf{d}]$. We need to show that $\tau \circ \tau^n$ matches the required shape (15)-(16) with $n+1$ instead of n. After unraveling the definitions, this splits into three sub-claims, concerning the update of the $\mathbf{c}$, the guard and the update of the $\mathbf{a}$, respectively:

(i) the equality $\mathbf{u}(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{n})) = \mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n+1})$ is valid;

(ii)

$$\bigwedge_{0\leq k<n} \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k}), \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k})))) \wedge \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}), \lambda j.\, F(\mathbf{c}, \mathbf{a}, \bar{n}, j)(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}))))$$

is equivalent to

$$\bigwedge_{0\leq k<n+1} \tilde{\phi}_L(\mathbf{u}^*(\mathbf{c}, \bar{k}), \mathbf{a}(\kappa(\mathbf{u}^*(\mathbf{c}, \bar{k}))));$$

(iii) $wr(\lambda j.\, F(\mathbf{c}, \mathbf{a}, \bar{n}, j), \kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})), \tilde{\mathbf{t}}(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}), \lambda j.\, F(\mathbf{c}, \mathbf{a}, \bar{n}, j)(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{n})))))$ is the same function as $\lambda j.\, F(\mathbf{c}, \mathbf{a}, \overline{n+1}, j)$.

[8] The following is an informal explanation of the formula (16) expressing iterated updates. The point is to recognize whether a given cell j has been over-written or not within the first y iterations. The number $\iota_h(\tilde{\mathbf{c}}, j)$ gives the candidate number of iterations needed to get j and the further condition $j = \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j)))$ checks whether this number is correct or not. Take for instance Example 5.2 with $n = 2$. Then if we have a single counter initialized to say 4, our iterations give values $4+2, 4+2+2, \ldots$ for the updated counter. If we want to know whether j can be reached within less than 5 iterations, we just compute $\iota(4, j)$ which is the quotient of the integer division of $j-4$ by 2. The we need to check that $\iota(4, j)$ is among $0, \ldots, 4 = 5-1$ and *also* that j can be really reached from $\tilde{\mathbf{c}} = 4$ by adding 2 to it $\iota(4, j)$-times (the latter won't be true if j is odd).

Indeed statement (i) is trivial, because $\mathbf{u}(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{n})) = \mathbf{u}(\mathbf{u}^n(\tilde{\mathbf{c}})) = \mathbf{u}^{n+1}(\tilde{\mathbf{c}}) = \mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n+1})$ holds by (7).
To show (ii), it is sufficient to check that

$$\mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}))) \;=\; \lambda j.\; F(\mathbf{c}, \mathbf{a}, \bar{n}, j)(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}))) \tag{17}$$

is true. In turn, this follows from (16) and the validity of the following implications (varying $h = 1, \ldots, s$)

$$\iota_h(\tilde{\mathbf{c}}, j) \neq \bar{n} \;\to\; j \neq \kappa_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})) \tag{18}$$

(in fact, a_h and F_h can possibly differ only for the j satisfying $0 \leq \iota_h(\tilde{\mathbf{c}}, j) < \bar{n}$, i.e. in particular for the j such that $\iota_h(\tilde{\mathbf{c}}, j) \neq n$). To see why (18) is valid, notice that in view of (8), what (18) says is that we cannot have simultaneously both $\iota_h(\tilde{\mathbf{c}}, j) = \overline{n}$ and $\iota_h(\tilde{\mathbf{c}}, j) = \bar{m}$, for some $m \neq n$: indeed it is so by the definition of a function.

It remains to prove (iii); in view of (17) just shown, we need to check that

$$wr(\lambda j.\, F(\mathbf{c}, \mathbf{a}, \bar{n}, j), \kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})), \tilde{\mathbf{t}}(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}), \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})))))$$

is the same as $\lambda j.\; F(\mathbf{c}, \mathbf{a}, \overline{n+1}, j)$. For every $h = 1, \ldots, s$, this is split into three cases, corresponding to the validity check for the three implications:

$$\begin{aligned}
i_h(\tilde{\mathbf{c}}, j) < \bar{n} &\to wr(\lambda j.\, F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j), \kappa_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})), \tilde{t}_h)(j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j) \\
i_h(\tilde{\mathbf{c}}, j) = \overline{n} &\to wr(\lambda j.\, F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j), \kappa_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})), \tilde{t}_h)(j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j) \\
i_h(\tilde{\mathbf{c}}, j) > \overline{n} &\to wr(\lambda j.\, F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j), \kappa_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n})), \tilde{t}_h)(j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j)
\end{aligned}$$

where we wrote simply $\tilde{t}_h$ instead of $\tilde{t}_h(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}), \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \overline{n}))))$. However, keeping in mind (18) and (14), the three implications can be rewritten as follows (the second one is split into two subcases)

$$\begin{aligned}
i_h(\tilde{\mathbf{c}}, j) < \bar{n} &\to F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j) \\
i_h(\tilde{\mathbf{c}}, j) = \overline{n} \wedge j = \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j))) &\to \tilde{t}_h = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j) \\
i_h(\tilde{\mathbf{c}}, j) = \overline{n} \wedge j \neq \kappa_h(\mathbf{u}^*(\mathbf{c}, \iota_h(\tilde{\mathbf{c}}, j))) &\to F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j) \\
i_h(\tilde{\mathbf{c}}, j) > \overline{n} &\to F_h(\mathbf{c}, \mathbf{a}, \bar{n}, j) = F_h(\mathbf{c}, \mathbf{a}, \overline{n+1}, j)
\end{aligned}$$

The above four implications all hold by the definitions (16) of the F_h.

Proof of Theorem 6.1 (continued). As a consequence of the Claim, since the formula

$$\bigwedge_{0 \leq k < n} \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k}), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\tilde{\mathbf{c}}, \bar{k}))), \mathbf{a}[\mathbf{d}])$$

is equivalent to $\forall z\, (0 \leq z < \bar{n} \to \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, z), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\mathbf{c}, z))), \mathbf{a}[\mathbf{d}])$, we can use (13) to express τ^+ as

$$\exists y > 0 \begin{pmatrix} \forall z\, (0 \leq z < y \to \tilde{\phi}_L(\mathbf{u}^*(\tilde{\mathbf{c}}, z), \mathbf{d}, \mathbf{a}(\kappa(\mathbf{u}^*(\mathbf{c}, z))), \mathbf{a}[\mathbf{d}]) \wedge \mathbf{d}' = \mathbf{d} \wedge \\ \wedge\, pc = l \;\wedge\; pc' = l \;\wedge\; \tilde{\mathbf{c}}' = \mathbf{u}^*(\tilde{\mathbf{c}}, y) \;\wedge\; \mathbf{a}' = \lambda j.\; F(\mathbf{c}, \mathbf{a}, y, j) \end{pmatrix} \tag{19}$$

The latter shows that τ^+ is a Σ^0_2-assignment, as desired. ⊣

□

B A worked out example

Simple assignements might not be sufficient for nested loops where an array is scanned by a couple of counters, one of which is kept fixed (think for instance of inner loops of sorting algorithms). To cope with these more complicated cases, we introduce a larger class of assignments (these assignments are still local, hence covered by Theorem 6.1). We call *simple+* the ground assignments of the form

$$pc = l \;\wedge\; \phi_L(\mathtt{c}, \mathbf{d}, \mathbf{a}) \;\wedge\; pc' = l \wedge \mathtt{c}' = \mathtt{c} \pm 1 \wedge \mathbf{d}' = \mathbf{d} \;\wedge\; \mathbf{a}' = wr(\mathbf{a}, \mathtt{c}, \mathbf{t}(\mathtt{c}, \mathbf{d}, \mathbf{a})) \tag{20}$$

where (i) $\mathbf{d} = d_1, \ldots, d_l$ is a tuple of integer constants, (ii) the first occurrence of $\mathtt{c}$ in $wr(\mathbf{a}, \mathtt{c}, \mathbf{t}(\mathtt{c}, \mathbf{d}, \mathbf{a}))$ stands for a tuple of terms all identical to $\mathtt{c}$, (iii) the guard ϕ_L contains the conjuncts $\mathtt{c} \neq d_i$ $(1 \leq i \leq l)$, and (iv)

$\phi_L, \mathbf{t}$ are purely arithmetical over $\mathtt{c}, \mathbf{d}, a_1[\mathtt{c}], \dots a_s[\mathtt{c}], a_1[d_1], \dots, a_s[d_l]$. Basically, simple+ local ground assignments differ from plain simple ones just because there are some 'idle' indices $\mathbf{d}$; in addition, the counter c can also be decremented.

The accelerated transition for (20) computed by Theorem 6.1 can be re-written as follows (we write $j \in [\mathtt{c}, \mathtt{c} \pm k]$ for $\mathtt{c} \le j \le \mathtt{c} + k$ or $\mathtt{c} - k \le j \le \mathtt{c}$, depending on whether we have increment or decrement in (20)):

$$\exists k \begin{pmatrix} k > 0 \wedge pc = l \wedge \forall j\, (j \in [\mathtt{c}, \mathtt{c} \pm k] \to \phi_L(j, \mathbf{d}, \mathbf{a})) \wedge pc' = l \wedge \mathbf{d}' = \mathbf{d} \wedge \\ \wedge\, \mathtt{c}' = \mathtt{c} \pm k \wedge \mathbf{a}' = \lambda j. (\mathtt{if}\ j \in [\mathtt{c}, \mathtt{c} \pm k]\ \mathtt{then}\ \mathbf{t}(j, \mathbf{d}, \mathbf{a})\ \mathtt{else}\ \mathbf{a}[j]) \end{pmatrix} \tag{21}$$

To show how acceleration and abstraction/refinement techniques can mutually benefit from each other, consider the procedure `allDiff`, represented by the *all diff 2* entry in Table1. This function tests whether all entries of the array a are pairwise different:

```
function allDiff ( int a[N] ):
1 r = true;
2 for (i = 1; i < N ∧ r; i++)
3   for (j = i-1; j ≥ 0 ∧ r; j--)
4     if (a[i] = a[j]) r = false;
5 assert (r → (∀x, y(0 ≤ x < y < N) → (a[x] ≠ a[y])))
```

This function is represented by the transition system specified below (in the specification, we omit identical updates to improve readability).

$$I(\mathbf{v}) = (r = 0 \wedge \mathtt{i} = 1 \wedge \mathtt{j} = 0 \wedge pc = l_1)$$

$$U(\mathbf{v}) = pc = l_4 \wedge \exists x, y. (0 \le x < y < \mathtt{a.Length} \wedge a[x] = a[y])$$

$$\tau_1 = \begin{pmatrix} pc = l_1 \wedge \mathtt{i} \le \mathtt{a.Length} \wedge r = 0 \wedge \\ pc' = l_2 \wedge \mathtt{j}' = \mathtt{i} - 1 \end{pmatrix}$$

$$\tau_2 = \begin{pmatrix} pc = l_1 \wedge \mathtt{i} > \mathtt{a.Length} \wedge \\ pc' = l_4 \end{pmatrix}$$

$$\tau_3 = \begin{pmatrix} pc = l_1 \wedge r = 1 \wedge \\ pc' = l_4 \end{pmatrix}$$

$$\tau_4 = \begin{pmatrix} pc = l_3 \wedge \\ pc' = l_1 \wedge \mathtt{i}' = \mathtt{i} + 1 \end{pmatrix}$$

$$\tau_5 = \begin{pmatrix} pc = l_2 \wedge r = 0 \wedge \mathtt{j} \ge 0 \wedge a[\mathtt{i}] = a[\mathtt{j}] \wedge \\ pc' = l_2 \wedge \mathtt{j}' = \mathtt{j} - 1 \wedge r = 1 \end{pmatrix}$$

$$\tau_6 = \begin{pmatrix} pc = l_2 \wedge r = 0 \wedge \mathtt{j} \ge 0 \wedge a[\mathtt{i}] \ne a[\mathtt{j}] \wedge \\ pc' = l_2 \wedge \mathtt{j}' = \mathtt{j} - 1 \end{pmatrix}$$

$$\tau_7 = \begin{pmatrix} pc = l_2 \wedge j < 0 \wedge \\ pc' = l_3 \end{pmatrix}$$

$$\tau_8 = \begin{pmatrix} pc = l_2 \wedge r = 1 \wedge \\ pc' = l_3 \end{pmatrix}$$

For this problem, the transition we want to accelerate is τ_6. Accelerating transition τ_6 is not sufficient to avoid divergence caused by the outer loop, though. On the other side, accelerating the inner loop simplifies the problem, which can be successfully verified by the model checker by exploiting abstraction/refinement techniques in 1.36 seconds (see Table1 for more details).

The acceleration of transition τ_6 requires simple+-assignements (implemented in the current release of MCMT). We follow MCMT implementation quite closely to explain what happens.

As a first observation, MCMT specification language requires that whenever two counters i and j both occur in array applications $a[\mathtt{i}], a[\mathtt{j}]$ (like in τ_6 above), the guard of the transition must contain either the literal $\mathtt{i} = \mathtt{j}$ or the literal $\mathtt{i} \ne \mathtt{j}$. Thus such transitions must be duplicated; in our case, the copy of τ_6 with

$\mathtt{i} = \mathtt{j}$ can be ignored because it has an inconsistent guard. The copy with $\mathtt{i} \neq \mathtt{j}$ in the guard satisfies the conditions for being a simple+-assignment. Thus, its acceleration, according to (21), can be written as

$$\exists k \left(\begin{array}{c} k > 0 \wedge \forall j\,(j \in [\mathtt{j}, \mathtt{j} \pm k] \rightarrow \mathtt{i} \neq j \wedge r = 0 \wedge j \geq 0 \wedge a[\mathtt{i}] \neq a[j]) \wedge \\ \wedge\, pc = 2 \wedge pc' = 2 \wedge \mathtt{i}' = \mathtt{i} \wedge r' = r \wedge \mathtt{j}' = \mathtt{j} \pm k \wedge \mathbf{a}' = \mathbf{a} \end{array} \right)$$

In the current release, MCMT is able to compute by itself the above accelerated transition and thus to certify safety of `allDiff` procedure.

C Experimental evaluation

Complete statistics for the experiments performed with MCMT are reported in Table1. Benchmarks have been taken from different sources:

- The benchmarks "filter test", "max in array test", "filter", "max in array 1", "max in array 2", "max in array 3" have been taken and/or adapted from programs on `http://proval.lri.fr/`.
- The "heap as array" program has been suggested by K. Rustan M. Leino and it is reported in Figure3.
- all the programs pN have been taken from "I. Dillig, T. Dillig, and A. Aiken. Fluid updates: Beyond strong vs. weak updates. In ESOP, pages 246-266, 2010.".
- The "bubble sort" example comes from the "Eureka" project `http://www.ai-lab.it/eureka` and has been used as a benchmark in the paper "A. Armando, M. Benerecetti, and J. Mantovani. Abstraction refinement of linear programs with arrays. In TACAS, pages 373-388, 2007."
- "all diff 1" and "all diff 2" have been suggested by Madhusudan Parthasarath and his group. They represent two different encoding of an algorithm that initializes an array to different values and then check if the array has been correctly initialized.
- "compare", "copy", "find 1", "find 2", "init", "init test", "partition" have been taken/adapted from "Krystof Hoder, Laura Kovács, Andrei Voronkov: Interpolation and Symbol Elimination in Vampire. In IJCAR, pages 188-195, 2010".
- The "linear search" program is used as a running example on the book "Aaron R. Bradley, Zohar Manna: The calculus of computation - decision procedures with applications to verification. Springer 2007, pp. I-XV, 1-366".
- "selection sort" example has been used in "M. N. Seghir, A. Podelski, and T. Wies. Abstraction Refinement for Quantified Array Assertions. In SAS, pages 3-18, 2009."
- "strcmp", "strcpy" and "strlen" have been adapted from the standard string C library.

The benchmarks named with " * test " refer to benchmarks with quantified assertions substituted by a for loop. For those programs, the postcondition does not have quantifiers: in these benchmarks it is even harder to come up with a quantified safe inductive invariant to prove that the program is correct. Thus, it is a remarkable fact that our tool can automatically synthetize such invariants.

Program	Status	No options	Abstraction	Acceleration	Accel. + Abstr.
filter test	safe	×	0.08	×	0.08
heap as array	safe	×	0.12	×	0.12
init test	safe	×	11.72	×	0.16
max in array test	safe	×	0.18	×	0.18
p01	safe	×	×	0.09	9.08
p02	safe	×	×	0.09	9.52
p03	safe	×	0.11	0.09	0.14
p08	safe	×	0.12	0.12	0.11
p09	safe	×	0.12	0.99	0.11
p14	safe	×	6.39	0.35	7.78
p17	safe	×	0.02	0.19	0.19
p04	unsafe	0.02	0.03	0.03	0.02
p10	unsafe	0.07	0.04	0.06	0.03
p11	unsafe	0.02	0.03	0.04	0.04
p15	unsafe	1.4	1.74	0.3	2.97
p16	unsafe	4.27	3.70	0.45	8.89
p18	unsafe	0.01	0.02	0.01	0.01
p19	unsafe	0.02	0.02	0.01	0.01
p20	unsafe	0.02	0.02	0.03	0.02
p22	unsafe	0.02	0.03	0.02	0.17
all diff 1	safe	×	×	0.08	0.13
all diff 2	safe	×	×	×	1.36
bubble sort	safe	×	1.23	×	1.23
compare	safe	×	0.04	×	0.04
copy	safe	×	0.03	0.03	0.03
filter	safe	×	0.11	×	0.11
find 1	safe	×	0.06	×	0.06
find 2	safe	×	0.07	0.06	0.17
init	safe	×	0.08	0.03	0.1
linear search	safe	×	0.04	0.05	0.02
max in array 1	safe	×	0.1	×	0.1
max in array 2	safe	×	0.11	×	0.13
max in array 3	safe	×	0.06	×	0.01
minusN	safe	×	×	0.77	1.4
partition	safe	×	0.05	×	0.03
selection sort	safe	×	7.87	×	45.07
strcat 1	safe	×	×	×	3.5
strcat 2	safe	×	×	×	3.62
strcmp	safe	×	0.04	0.06	0.02
strcpy	safe	×	0.03	0.02	0.01
strlen	safe	×	×	0.1	0.06
p01	safe	×	0.08	0.02	0.1
p02	safe	×	0.08	0.05	0.1
p03	safe	×	0.03	0.02	0.03
p08	safe	×	0.03	0.05	0.03
p09	safe	×	0.03	0.04	0.03
p18	safe	×	×	0.07	0.33
p20	safe	×	0.04	0.05	0.02
p04	unsafe	0.07	0.02	0.01	0.01
p11	unsafe	0.01	0.02	0.02	0.01
p14	unsafe	0.31	1.79	0.28	2.5
p15	unsafe	0.09	1.77	0.12	1.4
p16	unsafe	0.11	2.97	1.23	6.57
p17	unsafe	0.02	0.03	0.01	0.02
p19	unsafe	0.02	0.02	0.01	0.01

Table 1: Experimental results for different options. Time limit has been set to 60 seconds, and × denotes a timeout. Programs in the first part of the table are annotated with quantifier-free assertions, those in the second part have ∀-assertions. Notably, when abstraction and acceleration is combined MCMT is able to verify all the 55 programs.

```
var Heap: [int] int;
const unique F: int;  const unique G: int;
const F_final: int;  const G_final: int;
procedure HeapP ( )
  modifies Heap;
  requires F_final > 0 ∧ G_final > 0;
  ensures Heap[F] = F_final ∧ Heap[G] = G_final;
{
  Heap[F] := 0;  Heap[G] := G_final;
  while (Heap[F] < F_final)
    invariant Heap[F] ≤ F_final;
  {
    Heap[F] := Heap[F] + 1;
  }
}
```

Figure 3: The “heap as array” program.