Extracting Symbolic Transitions from $TLA+$ Specifications

Jure Kukovec, Thanh-Hai Tran, Igor Konnov

HAL Id: hal-01871131
https://inria.hal.science/hal-01871131

Submitted on 10 Sep 2018

Extracting Symbolic Transitions from TLA$^+$ Specifications [*]

Jure Kukovec[1], Thanh-Hai Tran[1], and Igor Konnov[1,2]✉

[1] TU Wien (Vienna University of Technology), Austria
{jkukovec,tran,konnov}@forsyte.at
[2] University of Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
igor.konnov@inria.fr

Abstract. In TLA$^+$, a system specification is written as a logical formula that restricts the system behavior. As a logic, TLA$^+$ does not have assignments and other imperative statements that are used by model checkers to compute the successor states of a system state. Model checkers compute successors either explicitly — by evaluating program statements — or symbolically — by translating program statements to an SMT formula and checking its satisfiability. To efficiently enumerate the successors, TLA's model checker TLC introduces side effects. For instance, an equality $x' = e$ is interpreted as an assignment of e to the yet unbound variable x.

Inspired by TLC, we introduce an automatic technique for discovering expressions in TLA$^+$ formulas such as $x' = e$ and $x' \in \{e_1, \ldots, e_k\}$ that can be provably used as assignments. In contrast to TLC, our technique does not explicitly evaluate expressions, but it reduces the problem of finding assignments to the satisfiability of an SMT formula. Hence, we give a way to slice a TLA$^+$ formula in symbolic transitions, which can be used as an input to a symbolic model checker. Our prototype implementation successfully extracts symbolic transitions from a few TLA$^+$ benchmarks.

1 Introduction

TLA is a general language introduced by Leslie Lamport for specifying temporal behavior of computer systems. It was later extended to TLA$^+$ [18], which provides the user with a concrete syntax for writing expressions over sets, functions, integers, sequences, etc. TLA$^+$ does not fix a model of computation, and thus it found applications in the design of concurrent and distributed systems, e.g., see [12,23,24,22,2].

A specification alone brings almost no guarantees of system correctness. As it is an untyped language, TLA$^+$ allows for expressions such as $1 \cup \{2\}$, which are considered ill-typed in statically-typed programming languages. To formally prove specification properties such as safety and liveness, one can use TLAPS — a proof

$$\text{MODULE } prodcons$$

VARIABLE S, $empty$
$Init \triangleq S = \{\} \land empty = \text{TRUE}$
$Produce \triangleq \land empty' = \text{FALSE}$
$\qquad\qquad \land \exists X \in \text{SUBSET } \{\text{"A", "B", "Z", "1", "8"}\} : S' = S \cup \{X\}$
$Consume \triangleq \neg empty \land S' \in \text{SUBSET } S \land empty' = (S' = \{\})$
$Next \triangleq Produce \lor Consume$

Fig. 1: A simple producer-consumer

system for TLA$^+$ [8]. Although progress towards proof automation was made in the last years [20], writing formal proofs is still a challenging task [23,24].

On the other side of the spectrum are model checkers that require little user effort to run. Indeed, TLA$^+$ users debug their specifications with TLC [26]. Beyond simple debugging, TLC found serious bugs in specifications of distributed algorithms [23]. Although TLC contains remarkable engineering solutions, its core techniques *enumerate* reachable states and inevitably suffer from state explosion.

Instead of enumerating states, software model checkers run SAT and SMT solvers in the background to reason about computations symbolically. To name a few, CBMC [15] and CPAChecker[3] implement bounded model checking [4] and CEGAR [9]. Domain-specific tools ByMC and Cubicle prove properties of parameterized distributed algorithms with SMT [10,14].

A simple example in Figure 1 illustrates the problems that one faces when developing a symbolic model checker for TLA$^+$. In this example, we model two processes: *Producer* that inserts a subset of $\{$"A", "B", "Z", "1", "8"$\}$ into the set S, and *Consumer* that removes from S its arbitrary subset. The system is initialized with the operator *Init*. A system transition is specified with the operator *Next* that is defined via a disjunction of operators *Produce* and *Consume*. Both Producer and Consumer maintain the state invariant $empty \Leftrightarrow (S = \emptyset)$. We notice the following challenges for a symbolic approach:

1. The specification does not have types. This is not a problem for TLC, since it constructs states on the fly and hence dynamically computes types. In the symbolic case, one can use type synthesis [20] or the untyped SMT encoding [21].
2. Direct translation of *Next* to SMT would produce a *monolithic* formula, e.g., it would not analyze *Produce* and *Consume* as independent actions. This is in sharp contrast to translation of imperative programs, in which variable assignments allow a model checker to focus only on the local state changes.

In this paper, we focus on the second problem. Our motivation comes from the observation on how TLC computes the successors of a given state [18, Ch. 14]. Instead of precomputing all potential successors — which would be anyway impossible without types — and evaluating *Next* on them, TLC explores subformulas of *Next*. The essential exploration rules are: (1) Disjunctions and conjunctions are evaluated from left to right, (2) an equality $x' = e$ assigns the value of e to x' if

x' is yet unbound, (3) if an unbound variable appears on the right-hand side of an assignment or in a non-assignment expression, TLC terminates with an error, and (4) operands of a disjunction assign values to the variables independently. In more detail, rule (4) means that whenever a disjunction $A \vee B$ is evaluated and x' is assigned a value in A, this value does not propagate to B; moreover, x' must be assigned a value in B.

In our example, TLC evaluates the actions *Produce* and *Consume* independently and assigns variables as prescribed by these formulas. As TLC is explicit, for each state, it produces at most 2^{2^5} successors in *Produce* as well as in *Consume*.

We introduce a technique to statically label expressions in a TLA$^+$ formula ϕ as assignments to the variables from a set V', while fulfilling the following:

1. For purely Boolean formulas, if one transforms ϕ into an equivalent formula $\bigvee_{1 \leq i \leq k} D_i$ in disjunctive normal form (DNF), then every disjunct D_i has *exactly one* assignment per variable from V'.
2. The assignments adhere the following partial order: if $x' \in V'$ is assigned a value in expression e, that uses a variable $y' \in V'$, then the assignment to y' precedes the assignment to x'.
3. In general, we formalize the above idea with the notion of a branch.

As expected, the following sequence of expressions is given as assignments in our example: (1) $empty' = \text{TRUE}$, (2) $S' = S \cup \{X\}$, (3) $S' \in \text{SUBSET}\, S$, and (4) $empty' = (S' = \emptyset)$. Using this sequence, our technique constructs two symbolic transitions that are equivalent to the actions *Produce* and *Consume*.

In general, finding assignments and slicing a formula into symbolic transitions is not as easy as in our example, because of quantifiers and IF-THEN-ELSE complicating matters. In this paper, we present our solution, demonstrate its soundness and report on preliminary experiments.

2   Abstract Syntax α-TLA$^+$

TLA$^+$ has rich syntax [18], which cannot be defined in this paper. To focus only on the expressions that are essential for finding assignments in a formula, we define abstract syntax for TLA$^+$ formulas below. In our syntax, the essential operators such as conjunctions and disjunctions are included explicitly, while the other non-essential operators are hidden under the star expression $\star$.

We assume predefined three infinite sets:

- A set $\mathcal{L}$ of *labels*. We use notation ℓ_i to refer to its elements, for $i \in \mathbb{N}$.
- A set *Vars'* of *primed variables* that are decorated with prime, e.g., x' and a'.
- A set *Bound* of *bound variables*, which are used by quantifiers.

$$Next \triangleq \ell_1 :: \Big(\ell_2 :: \big(\ell_3 :: empty' \in \ell_4 :: \star \wedge \ell_5 :: \exists\, X \in \ell_6 :: \star : \ell_7 :: S' \in \ell_8 :: \star \big)$$
$$\vee\, \ell_9 :: (\ell_{10} :: \star \wedge \ell_{11} :: S' \in \ell_{12} :: \star \wedge \ell_{13} :: empty' \in \ell_{14} :: \star(S')) \Big)$$

Fig. 2: The *Next* operator of producer-consumer in α-TLA$^+$

The abstract syntax α-TLA$^+$ is defined in terms of the following grammar:

$$expr ::= ex_\alpha \mid \ell :: \text{FALSE}$$
$$\mid \ell :: v' \in ex_\alpha \mid \ell :: expr \wedge \cdots \wedge expr \mid \ell :: expr \vee \cdots \vee expr$$
$$\mid \ell :: \exists x \in ex_\alpha :\ expr \mid \ell :: \text{IF}\ ex_\alpha\ \text{THEN}\ expr\ \text{ELSE}\ expr$$
$$ex_\alpha ::= \ell ::\ \star\,(v', \ldots, v')$$
$$\ell ::= \text{a unique label from the set}\ \mathcal{L}$$
$$v' ::= \text{a variable name from the set}\ Vars'$$
$$x ::= \text{a variable name from the set}\ Bound$$

A few comments on the syntax and its relation to TLA$^+$ expressions are in order. We require every expression to carry a unique label $\ell_i \in \mathcal{L}$. Although this is not a requirement in TLA$^+$, it is easy to decorate every expression with a unique label. The expressions of the form $\ell :: v' \in expr$ are of ultimate interest to us, as they are treated as assignment candidates. Under certain conditions, they can be used to assign to v' a value from the set represented by the expression *expr*. Perhaps somewhat unexpectedly, expressions such as $v' = e$ and UNCHANGED $\langle v_1, \ldots, v_k \rangle$ are not included in our syntax. To keep the syntax minimal, we represent them with $\ell :: v' \in expr$. Indeed, these expressions can be rewritten in an equivalent form: $v' = e$ as $v' \in \{e\}$, and UNCHANGED $\langle v_1, \ldots, v_k \rangle$ as $v'_1 \in \{v_1\} \wedge \cdots \wedge v'_k \in \{v_k\}$. Every non-essential TLA$^+$ expression e is presented in the abstract form $\ell :: \star(v'_1, \ldots, v'_k)$, where $v'_1, \ldots, v'_k$ are the names of the primed variables that appear in e. When no primed variable appears in an expression, we omit parenthesis and write $\ell :: \star$. TLA$^+$ expressions often refer to user-defined operators, which are not present in our abstract syntax. We simply assume that all non-recursive user-defined operators have been expanded, that is, recursively replaced with their bodies. All uses of recursive operators are hidden under $\star$; hence, recursive operator definitions are ignored when searching for assignment candidates.

It should be now straightforward to see how one could translate a TLA$^+$ expression to our abstract syntax. We write $\alpha(e)$ to denote the expression in α-TLA$^+$, that represents an expression e in the complete TLA$^+$ syntax. With γ we denote the reverse translation from α-TLA$^+$ to TLA$^+$ that has the property that $\gamma(\alpha(e)) = e$. Figure 2 shows the abstract expression $\alpha(Next)$ of the operator *Next* defined in Figure 1.

Discussions. Notice that α-TLA$^+$ is missing several fundamental constructs permitted in TLA$^+$, such as CASE expressions, universal quantifiers, and negations. They are all abstracted to $\star$. The primary purpose of α-TLA$^+$ is to allow us to

determine whether a given expression containing set inclusion — or equality — can be used as an assignment. If such an expression occurs under a universal quantifier, it is not clear which value should be used for an assignment. Hence, we abstract the expressions under universal quantifiers. For similar reason, we abstract the expressions under negation. The latter is consistent with TLC, which produces an error when given, for example, $Next == \neg(x' = 1)$. Finally, we abstract CASE, due to its semantics, which is defined in terms of the CHOOSE operator [18, Ch. 6]. In practice, there are no potential assignments under CASE in the standard TLA$^+$ examples.

3 Preliminary Definitions

Every TLA$^+$ specification declares a certain finite set of variables, which may appear in the formulas contained therein. Let ϕ be an α-TLA$^+$ expression. We assume, for the purposes of our analysis, that ϕ is associated with some finite set $Vars'(\phi)$, which is a subset of $Vars'$, containing all of the variables that appear in ϕ (and possibly additional ones). This is the set of variables declared by the specification in which $\gamma(\phi)$ appears.

Since the labels are unique, we write $\mathrm{lab}(\ell :: \psi)$ to refer to the expression label ℓ and $\mathrm{expr}(\ell)$ to refer to the expression that is labeled with ℓ. We refer to the set of all subexpressions of ϕ by $\mathrm{Sub}(\phi)$. See [16] for a formal definition.

The set $\mathrm{Sub}(\phi)$ allows us to reason about terms that appear inside an expression ϕ, at some unknown/irrelevant depth. We will often refer to the set of all labels appearing in ϕ, that is, $\mathrm{Labs}(\phi) = \{\mathrm{lab}(\psi) \mid \psi \in \mathrm{Sub}(\phi)\}$.

Of special interest to us are *assignment candidates*, i.e., expressions of the form $\ell :: v' \in \phi_1$. Given a variable $v' \in Vars'(\phi)$ and an α-TLA$^+$ expression ϕ, we write $\mathrm{cand}(v', \phi)$ to mean the set of labels that belong to assignment candidates for v' in subexpressions of ϕ. More formally, $\mathrm{cand}(v', \phi)$ is $\{\ell \mid (\ell :: v' \in \psi) \in \mathrm{Sub}(\phi)\}$. An exhaustive definition is included in [16]. We use the notation $\mathrm{cand}(\phi)$ to mean $\bigcup_{v' \in Vars'(\phi)} \mathrm{cand}(v', \phi)$.

Finally, we assign to each label ℓ in $\mathrm{Labs}(\phi)$ a set $\mathrm{frozen}_\phi(\ell) \subseteq Vars'(\phi)$. Intuitively, if a variable v' is in $\mathrm{frozen}_\phi(\ell)$, then no expression of the form $\hat{\ell} :: v' \in \psi$ can be treated as an assignment inside $\mathrm{expr}(\ell)$. Formally, for every $\ell \in \mathrm{Labs}(\phi)$ the set $\mathrm{frozen}_\phi(\ell)$ is defined as the minimal set satisfying all the constraints in Table 1.

The sets frozen_ϕ naturally lead to the dependency relations $\lhd_{v'}$ on $\mathrm{Labs}(\phi)$, where $v' \in Vars'(\phi)$. We will use $\ell_1 \lhd_{v'} \ell_2$ to mean that ℓ_1 is an assignment candidate for v', which also belongs to the frozen set of ℓ_2. Formally:

$$\ell_1 \lhd_{v'} \ell_2 \iff \ell_1 \in \mathrm{cand}(v', \phi) \land v' \in \mathrm{frozen}_\phi(\ell_2)$$

Intuitively, if $\ell_1 \lhd_{v'} \ell_2$ we want to make sure that $\mathrm{expr}(\ell_1)$ is evaluated before $\mathrm{expr}(\ell_2)$, if possible.

Example 1. Let us look at the following α-TLA$^+$ expression:

$$\ell_1 :: [\exists i \in [\ell_2 :: \star(y')] : \ell_3 :: x' \in [\ell_4 :: \star]]$$

Table 1: The constraints on frozen_ϕ

$\alpha\textbf{-TLA}^+$ **expression** ϕ	**Constraints on** frozen_ϕ
$\ell :: \star(v'_1, \ldots, v'_k)$	$\{v'_1, \ldots, v'_k\} \subseteq \text{frozen}_\phi(\ell)$
$\ell :: v' \in \phi_1$	$\text{frozen}_\phi(\ell) = \text{frozen}_\phi(\text{lab}(\phi_1))$
$\ell :: \bigwedge_{i=1}^s \phi_i$ or $\ell :: \bigvee_{i=1}^s \phi_i$	$\text{frozen}_\phi(\ell) \subseteq \text{frozen}_\phi(\text{lab}(\phi_i))$ for $i \in \{1, \ldots, s\}$
$\ell :: \exists x \in \phi_1 : \phi_2$	$\text{frozen}_\phi(\ell) \subseteq \text{frozen}_\phi(\text{lab}(\phi_1)) \subseteq \text{frozen}_\phi(\text{lab}(\phi_2))$
$\ell :: \text{IF } \phi_1 \text{ THEN } \phi_2 \text{ ELSE } \phi_3$	$\text{frozen}_\phi(\ell) \subseteq \text{frozen}_\phi(\text{lab}(\phi_1))$
	$\text{frozen}_\phi(\text{lab}(\phi_1)) \subseteq \text{frozen}_\phi(\text{lab}(\phi_i))$ for $i = 2, 3$

Take the subexpression $\ell_3 :: x' \in [\ell_4 :: \star]$, which we name ψ . By solving the constraints for $\text{frozen}_\psi(\ell_3)$ we conclude that $\text{frozen}_\psi(\ell_3) = \emptyset$. However, if we take the additional constraints for $\text{frozen}_\phi(\ell_3)$ into consideration, the empty set no longer satisfies all of them, specifically, it does not satisfy the condition imposed by the existential quantifier in ℓ_1. The additional requirement $\{y'\} \subseteq \text{frozen}_\phi(\ell_3)$ implies that $\text{frozen}_\phi(\ell_3) = \{y'\}$. This corresponds to the intuition that expressions under a quantifier, like ψ, implicitly depend on the bound variable and the expressions used to define it, which is $\text{expr}(\ell_2)$ in our example. $\triangleleft$

4 Formalizing Symbolic Assignments

As TLC evaluates formulas in a left-to-right order, there is a very clear notion of an assignment; the first occurrence of an expression $v' \in S$ is interpreted as an assignment to v'. In our work, we want to *statically* find expressions that can safely be used as assignments. If we were only dealing with Boolean formulas, we could transform the original TLA$^+$ formula to DNF, $\bigvee_{i=1}^s D_i$, and treat each D_i independently. However, we also need to find assignments, which may be nested under existential quantifiers. To transfer our intuition about DNF to the general case we first introduce a transformation boolForm, that captures the Boolean structure of the formula. Then, we introduce branches and assignment strategies to formalize the notion of assignments in the symbolic case.

Boolean structure of a formula and branches. The transformation boolForm maps an α-TLA$^+$ expression to a Boolean formula over variables from $\{b_\ell \mid \ell \in \mathcal{L}\}$. The definition of boolForm can be found in Table 2. As boolForm(ϕ) is a formula in Boolean logic, a model of boolForm(ϕ) is a mapping from $\{b_\ell \mid \ell \in \mathcal{L}\}$ to $\mathbb{B} = \{\text{true}, \text{false}\}$. Take $S \subseteq \mathcal{L}$. The set S naturally defines a model induced by S, denoted $\mathcal{M}[S]$, by the requirement that $\mathcal{M}[S] \vDash b_\ell$ if and only if $\ell \in S$.

The boolForm transformation allows us to formulate the central notion of a branch: A set $Br \subseteq \mathcal{L}$ is called a *branch* of ϕ if the following constraints hold:

(a) The set Br induces a model of boolForm(ϕ), i.e., $\mathcal{M}[Br] \vDash$ boolForm(ϕ), and
(b) The model $\mathcal{M}[Br]$ is minimal, that is, $\mathcal{M}[S] \nvDash$ boolForm(ϕ) for every $S \subset Br$.

Table 2: The definition of $\text{boolForm}(\phi)$

α-**TLA**$^+$ **expression** ϕ	$\text{boolForm}(\phi)$
$\ell :: \text{FALSE}$ or $\ell :: \star(v_1', \ldots, v_k')$ or $\ell :: v' \in \phi_1$	b_ℓ
$\ell :: \bigwedge_{i=1}^{s} \phi_i$	$\bigwedge_{i=1}^{s} \text{boolForm}(\phi_i)$
$\ell :: \bigvee_{i=1}^{s} \phi_i$	$\bigvee_{i=1}^{s} \text{boolForm}(\phi_i)$
$\ell :: \exists x \in \phi_1:\ \phi_2$	$\text{boolForm}(\phi_2)$
$\ell :: \text{IF } \phi_1 \text{ THEN } \phi_2 \text{ ELSE } \phi_3$	$\text{boolForm}(\phi_2) \vee \text{boolForm}(\phi_3)$

Then, $\text{Branches}(\phi)$ is the set of all branches of ϕ.

Example 2. Let us look the α-TLA$^+$ expression ϕ given by

$$\ell_1 :: [[\ell_2 :: x' \in \star] \wedge [\ell_3 :: [[\ell_4 :: x' \in \star] \vee [\ell_5 :: x' \in \star]]]]$$

We know that $\text{boolForm}(\phi) = b_{\ell_2} \wedge (b_{\ell_4} \vee b_{\ell_5})$. The set $S = \{\ell_2, \ell_4, \ell_5\}$ induces a model of $\text{boolForm}(\phi)$, but it is not a branch of ϕ because $\mathcal{M}[S]$ is not a minimal model. It is easy to see that ϕ has two branches $Br_1 = \{\ell_2, \ell_4\}$, and $Br_2 = \{\ell_2, \ell_5\}$. Therefore, we see that $\text{Branches}(\phi) = \{Br_1, Br_2\}$. $\triangleleft$

As our goal is to reason about the side-effects of variable assignments, the remainder of this section looks at how we can achieve this with the help of branches.

Assignment strategies. We want to statically mark some expressions as assignments, that is, pick a set $A \subseteq \text{Labs}(\phi)$. Below, we formulate the critical properties we require from such a set, which we will later call an assignment strategy.

Most obviously, we want to consider only assignment candidates:

Definition 1. *A set $H \subseteq \text{Labs}(\phi)$ is* homogeneous *if all the labels in H are assignment candidates. Formally, $H \subseteq \text{cand}(\phi)$.*

If we choose an arbitrary homogeneous set H, it might lack assignments on some branches or have multiple assignments to the same variable on others. Formally, we say that H has a *covering index* $d \in \mathbb{N}_0$ if there is a branch $Br \in \text{Branches}(\phi)$ and a variable $v' \in \text{Vars}'(\phi)$ for which $d = |Br \cap H \cap \text{cand}(v', \phi)|$. Now we define sets, that cover all branches with assignments:

Definition 2. *A homogeneous set C is a* covering *of ϕ, if it does not have 0 as a covering index. It is a* minimal covering *of ϕ, if it only has 1 as a covering index.*

Consider the TLA$^+$ formula $x' = y' \wedge y' = 2x'$. Its corresponding α-TLA$^+$ expression $\ell_0 :: (\ell_1 :: x' \in \ell_2 :: \star(y') \wedge \ell_3 :: y' \in \ell_4 :: \star(x'))$ has a minimal covering $\{\ell_1, \ell_3\}$. However, there is no way to order the assignments to x' and y'. To detect such cases, we define acyclic sets:

Definition 3. *A homogeneous set A is* acyclic*, if there is a strict total order $\prec_A$ on A, with the following property: For every variable $v' \in V$, every branch $Br \in$ Branches and every pair of labels ℓ_i and ℓ_j in $A \cap Br$ the relation $\ell_i \lhd_{v'} \ell_j$ implies $\ell_i \prec_A \ell_j$.*

Having defined homogeneous, minimal covering, and acyclic sets, we can formulate the notion of an *assignment strategy*.

Definition 4. *Let ϕ be an α-TLA$^+$ expression. A set $A \subseteq \mathcal{L}$ is an assignment strategy for ϕ, if it is an acyclic minimal covering.*

Static assignment problem. Given an α-TLA$^+$ expression ϕ, our goal is to find an assignment strategy, or prove that none exists.

5   Finding Assignment Strategies with SMT

For a given α-TLA$^+$ expression ϕ, we construct an SMT formula $\theta(\phi)$, that encodes the properties of assignment strategies. Technically, $\theta(\phi)$ is defined as $\theta_H(\phi) \wedge \theta_C(\phi) \wedge \theta_A(\phi)$, and consists of:

1. A Boolean formula $\theta_H(\phi)$, that encodes homogeneity.
2. A Boolean formula $\theta_C(\phi)$, that encodes the minimal covering property.
3. A formula $\theta_A(\phi)$, that encodes acyclicity. This formula requires the theories of linear integer arithmetic and uninterpreted functions (QF_UFLIA).

In the following, Propositions 1, 3, and 4 formally establish the relation between ϕ and its three SMT counterparts. Together, the propositions allows us to prove the following theorem:

Theorem 1. *For every α-TLA$^+$ formula ϕ and $A \subseteq$ Labs(ϕ), it holds that $\mathcal{M}[A] \vDash \theta(\phi)$ if and only if A is an assignment strategy for ϕ.*

5.1   Homogeneous Sets

We introduce a Boolean formula, whose models are exactly those induced by homogeneous sets. To this end, take the set of labels corresponding to expressions that are not assignment candidates, $\mathcal{N}(\phi)$, given by $\mathcal{N}(\phi) := $ Labs$(\phi) \setminus$ cand(ϕ). Then, we define the following:

$$\theta_H(\phi) := \bigwedge_{\ell \, \in \, \mathcal{N}(\phi)} \neg b_\ell$$

Proposition 1. *For every α-TLA$^+$ expression ϕ and $A \subseteq$ Labs(ϕ), it holds that $\mathcal{M}[A] \vDash \theta_H(\phi)$ if and only if A is homogeneous.*

Table 3: The definition of $\delta_{v'}(\phi)$

α-**TLA**$^+$ **expression** ϕ	$\delta_{v'}(\phi)$
$\ell :: \text{FALSE}$ or $\ell :: \star(v_1', \ldots, v_k')$	false
$\ell :: w' \in \phi_1$	$\begin{cases} b_\ell & ; w' = v' \\ \text{false} & ; \text{otherwise} \end{cases}$
$\ell :: \bigwedge_{i=1}^s \phi_i$	$\bigvee_{i=1}^s \delta_{v'}(\phi_i)$
$\ell :: \bigvee_{i=1}^s \phi_i$	$\bigwedge_{i=1}^s \delta_{v'}(\phi_i)$
$\ell :: \exists x \in \phi_1 : \phi_2$	$\delta_{v'}(\phi_2)$
$\ell :: \text{IF } \phi_1 \text{ THEN } \phi_2 \text{ ELSE } \phi_3$	$\delta_{v'}(\phi_2) \wedge \delta_{v'}(\phi_3)$

5.2 Minimal Covering Sets

Next we construct a Boolean formula $\theta_C^*(\phi)$, whose models are exactly those induced by covering sets. To this end, we define, for each $v' \in Vars'(\phi)$, the transformation $\delta_{v'}$ as shown in Table 3. Intuitively, $\delta_{v'}(\phi)$ is satisfiable exactly when there is an assignment to v' on every branch of ϕ. We then define

$$\theta_C^*(\phi) := \bigwedge_{v' \in Vars'(\phi)} \delta_{v'}(\phi)$$

Formally, the following proposition holds:

Proposition 2. *For every α-TLA$^+$ expression ϕ and $A \subseteq \text{Labs}(\phi)$, it holds that $\mathcal{M}[A] \models \theta_H(\phi) \wedge \theta_C^*(\phi)$ if and only if A is a covering set for ϕ.*

It is easy to restrict coverings to the minimal coverings. To do this, we define the set of collocated labels, denoted $\text{Colloc}(\phi)$, as

$$\text{Colloc}(\phi) := \{(\ell_1, \ell_2) \in \mathcal{L}^2 \mid \exists Br \in \text{Branches}(\phi) . \{\ell_1, \ell_2\} \subseteq Br\}$$

We can use this set to reason about minimal coverings: A minimal covering may contain, per variable, no more than one label from each pair of collocated assignments to that variable. We describe these labels by using the sets $\text{Colloc}_{v'}(\phi) := \text{Colloc}(\phi) \cap \text{cand}(v', \phi)^2$ and

$$\text{Colloc}_{Vars'}(\phi) := \bigcup_{v' \in Vars'(\phi)} \text{Colloc}_{v'}(\phi)$$

Then, the following SMT formula, in addition to $\theta_C^*(\phi)$, helps us find minimal covering sets:

$$\theta^{\exists!}(\phi) := \bigwedge_{\substack{(i,j) \in \text{Colloc}_{Vars'}(\phi) \\ i<j}} \neg(b_i \wedge b_j)$$

We denote by $\theta_C(\phi)$ the formula $\theta_C^*(\phi) \wedge \theta^{\exists!}(\phi)$.

Proposition 3. *For every α-TLA$^+$ expression ϕ and $A \subseteq \text{Labs}(\phi)$, it holds that $\mathcal{M}[A] \models \theta_H(\phi) \wedge \theta_C(\phi)$ if and only if A is a minimal covering set for ϕ.*

5.3 Acyclic Assignments

The last step is reasoning about acyclicity. Recall that, for $\ell_1, \ell_2 \in \mathcal{L}$, the relation $\ell_1 \lhd_{v'} \ell_2$ holds if and only if $\ell_1 \in \text{cand}(v', \phi) \wedge v' \in \text{frozen}_\phi(\ell_2)$. It is prudent to see that $\lhd_{v'}$ is not, in general, a strict total order (possibly not even irreflexive). However, the acyclicity property states that we can find a strict total order, which agrees with all relations $\lhd_{v'}$, on all branches.

Take $\text{Colloc}_\lhd(\phi)$ to be the filtering of $\text{Colloc}(\phi)$ by the relations $\lhd_{v'}$, i.e. the set $\{(i,j) \in \text{Colloc}(\phi) \cap \text{cand}(\phi)^2 \mid \exists v' \in \textit{Vars}'(\phi) \ . \ i \lhd_{v'} j\}$. The SMT formula describing acyclicity is as follows:

$$\theta_A^*(\phi) := \bigwedge_{(i,j) \ \in \ \text{Colloc}_\lhd(\phi)} b_i \wedge b_j \Rightarrow R(i) < R(j)$$

where R is an uninterpreted $\mathcal{L} \to \mathbb{N}$ function, capturing assignment order. In practice, we take $\mathcal{L} = \mathbb{N}$. Unlike the previous formulas, $\theta_A^*(\phi)$ extends beyond Boolean logic, requiring both linear integer arithmetic and uninterpreted functions. Thus, a model for $\theta_A^*(\phi)$ is a pair (M, r), where M models the Boolean part of the formula, i.e. assigns truth values to each b_i, and $r: \mathbb{N} \to \mathbb{N}$ is the interpretation of R.

To simplify the analysis, we force R to be injective, when it is restricted to $\text{Labs}(\phi)$. Otherwise we could always construct an injective function from R, which respects the required inequalities. The formula we therefore consider is as follows:

$$\theta_A(\phi) := \theta_A^*(\phi) \wedge \bigwedge_{\substack{\ell_1, \ell_j \ \in \ \text{Labs}(\phi) \\ \ell_i < \ell_j}} R(\ell_i) \neq R(\ell_j)$$

Proposition 4. *For every α-TLA^+ expression ϕ and $A \subseteq \text{Labs}(\phi)$, there is a function $r: \mathbb{N} \to \mathbb{N}$, for which $(\mathcal{M}[A], r) \vDash \theta_H(\phi) \wedge \theta_A(\phi)$ if and only if A is acyclic.*

6 Soundness of our Approach

In this section, we show the relation between assignment strategies and the original TLA^+ formulas. To this end, we introduce the notion of a slice. Together, branches allow us to rewrite a TLA^+ formula into an equivalent disjunction of slices.

In TLA^+, there are two kinds of variables: rigid variables that correspond to the variables declared with CONSTANT, and flexible variables whose values change during the course of an execution. Primed versions of the variables exist only for flexible variables and are used in transition formulas. Transition formulas in TLA^+ are first-order terms and formulas with flexible variables (unprimed and primed ones). We give the necessary definitions of TLA^+ semantics, whereas details can be found in [19]. An interpretation $\mathcal{I}$ defines a universe $|\mathcal{I}|$ of values and interprets each function symbol by a function and each predicate symbol by a relation. A state s is a mapping from unprimed flexible variables to values, and a state s' is a similar mapping for primed variables. A valuation ξ is a mapping from rigid variables to values. Given an interpretation $\mathcal{I}$, a pair of states (s, s'), and a valuation ξ, the

semantics of a TLA^+ transition formula E is the standard predicate logic semantics of E with respect to the extended valuation of s, s', ξ. With these definitions, $M = (\mathcal{I}, \xi, s, s')$ is a model for E, if E is equivalent to true under M. Let ϕ be a formula and $S \subseteq \mathcal{L}$. We define ϕ *sliced by* S, denoted $\text{slice}(\phi, S)$ in Table 4.

Table 4: The definition of $\text{slice}(\phi, S)$

α-TLA^+ **formula** ϕ	$\text{slice}(\phi, S)$
$\ell :: \text{FALSE}$	$\ell :: \text{FALSE}$
$\ell :: \star (v'_1, \dots, v'_1)$ or $\ell :: v' \in \phi_1$	$\begin{cases} \phi & ; \ell \in S \\ \ell :: \text{FALSE} & ; \text{otherwise} \end{cases}$
$\ell :: \bigwedge_{i=1}^{s} \phi_i$	$\ell :: \bigwedge_{i=1}^{s} \text{slice}(\phi_i, S)$
$\ell :: \bigvee_{i=1}^{s} \phi_i$	$\ell :: \bigvee_{i=1}^{s} \text{slice}(\phi_i, S)$
$\ell :: \exists x \in \phi_1 : \phi_2$	$\ell :: \exists x \in \phi_1 : \text{slice}(\phi_2, S)$
$\ell :: \text{IF } \phi_1 \text{ THEN } \phi_2 \text{ ELSE } \phi_3$	$\ell :: \text{IF } \phi_1 \text{ THEN } \text{slice}(\phi_2, S) \text{ ELSE } \text{slice}(\phi_3, S)$

Below, we show that the branches and slices induced by them naturally decompose a TLA^+ formula. Let ϕ be an α-TLA^+ expression and $\gamma(\phi)$ its corresponding TLA^+ formula. Then, the following holds:

Proposition 5. *Let ϕ be an α-TLA^+ expression and $M = (\mathcal{I}, \xi, s, s')$ a model of the TLA^+ formula $\gamma(\phi)$. There exists a branch Br of ϕ such that M is also a model of $\gamma(\text{slice}(\phi, Br))$.*

Proposition 6. *Let ϕ be an α-TLA^+ expression and $M = (\mathcal{I}, \xi, s, s')$ a model of the TLA^+ formula $\gamma(\text{slice}(\phi, Br))$. Then, M is also a model of $\gamma(\phi)$.*

Proposition 7. *Let ϕ be an α-TLA^+ expression. For every $S, T \subseteq \text{Labs}(\phi)$, every model M of the TLA^+ formula $\gamma(\text{slice}(\phi, S))$, is also a model of $\gamma(\text{slice}(\phi, S \cup T))$.*

It is easy to see that Proposition 7 does not hold in the other direction. For instance, take the empty set as S and $\text{Labs}(\phi)$ as T. This implies the following:

$$\gamma(\text{slice}(\phi, S)) = \text{FALSE and } \text{slice}(\phi, S \cup T) = \phi.$$

Obviously, FALSE cannot have a model, regardless of whether $\gamma(\phi)$ has one or not.

Since Propositions 5 and 6 hold, it would suffice to consider the set $\text{Branches}(\phi)$, together with an assignment strategy, to generate symbolic transitions. However, it is often the case that, for two distinct branches Br_1 and Br_2, the same assignments in A are chosen, that is, the intersections $Br_1 \cap A$ and $Br_2 \cap A$ are the same. We show that one can reduce the number of considered symbolic transitions, by analyzing how various branches intersect A.

An assignment strategy A naturally defines an equivalence relation $\sim_A$ on $\text{Branches}(\phi)$, given by $Br_1 \sim_A Br_2$ if and only if $Br_1 \cap A = Br_2 \cap A$. We use the notation $[Br]_A$ to refer to the equivalence class of Br by $\sim_A$, that is, the set $\{X \in \text{Branches}(\phi) \mid Br \sim_A X\}$.

Definition 5. *Let ϕ be an α-TLA$^+$ expression, A an assignment strategy for ϕ and Br a branch of ϕ. Using $X = [Br]_A$ and $Y = \bigcup_{Z \in X} Z$, we define the symbolic transition generated by Br and A to be $\mathrm{slice}(\phi, Y)$.*

Example 3. Let us look Example 2 again. The formula ϕ has two assignment strategies $A_1 = \{\ell_2\}$, and $A_2 = \{\ell_4, \ell_5\}$. If the first assignment strategy A_1 is chosen, we have that $Br_1 \cap A_1 = Br_2 \cap A_1 = \{\ell_2\}$. This implies that Br_1 and Br_2 are in the same equivalence class, or $Br_1 \sim_{A_1} Br_2$. Therefore, we have only one symbolic transition which is exactly ϕ. However, if A_2 is selected, branches Br_1 and Br_2 are not equivalent because $Br_1 \cap A_2 = \{\ell_4\}$ and $Br_2 \cap A_2 = \{\ell_5\}$. Therefore, we have two symbolic transitions:

$$T_1 = \ell_1 \ :: \ [[\ell_2 \ :: \ x' \in \star] \wedge [\ell_3 \ :: \ [[\ell_4 \ :: \ x' \in \star] \vee \ell_5 \ :: \ \text{FALSE}]]]$$
$$T_2 = \ell_1 \ :: \ [[\ell_2 \ :: \ x' \in \star] \wedge [\ell_3 \ :: \ [\ell_4 \ :: \ \text{FALSE} \vee [\ell_5 \ :: \ x' \in \star]]]]$$

The first assignment strategy A_1 seems to be better than A_2 because A_1 generates fewer symbolic transitions than A_2. However, in this paper, we do not introduce any metric, by which we could compare assignment strategies. In the implementation, we use any strategy found by the SMT solver. ◁

The equivalence relation $\sim_A$ allows us to define a counterpart to Proposition 7:

Proposition 8. *Let ϕ be an α-TLA$^+$ expression. For any selection $Br_1, \ldots, Br_k$ from the branches of ϕ, the following holds: If there exists a model M of the formula $\gamma(\mathrm{slice}(\phi, Br_1 \cup \cdots \cup Br_k))$, then M must be a model of $\gamma(\mathrm{slice}(\phi, Br))$, for some branch $Br \in \mathrm{Branches}(\phi)$. Additionally, if there is an assignment strategy A for ϕ, such that $Br_1, \ldots, Br_k$ all belong to the same equivalence class $[B]_A$, then M must be a model of $\gamma(\mathrm{slice}(\phi, Br))$, for some branch $Br \in [B]_A$.*

The following result allows us to use symbolic transitions, not individual branches:

Theorem 2. *Let ϕ be an α-TLA$^+$ expression and A an assignment strategy for ϕ. There is a model M of the TLA$^+$ formula $\gamma(\phi)$ if and only if there exists a $Br \in \mathrm{Branches}(\phi)$, such that M is a model of $\gamma(\psi)$, where ψ is the symbolic transition generated by Br and A.*

7 Preliminary Experiments and Potential Applications

Implementation and evaluation. Based on the theory presented in this paper, we have implemented a procedure to find assignment strategies and their corresponding symbolic transitions from TLA$^+$ specifications, or report that none exist. It uses Z3 as the background SMT solver.

We have chose specifications both from publicly available sources, e.g. EWD840 and Paxos from [1], and from a collection of algorithms we have encoded in TLA$^+$ ourselves. For each specification, we focus on the *Next* formula. We report the number of subexpressions in $\alpha(Next)$, that is, $|\mathrm{Sub}(\alpha(Next))|$, the number of assignments in the strategy found by our procedure, the number of symbolic transitions

Table 5: Experimental results

Specification	#subexpressions	size of strategy	#symbolic transitions	time (ms)
aba [6]	86	48	8	271
nbacg [13]	126	82	13	205
EWD840 [11]	47	16	4	25
prodcons (Fig. 1)	12	4	2	19
Paxos [17]	60	16	4	29
nbac [25]	47	15	14	26
bcastFolklore [7]	41	17	4	28

computed and the total runtime. The results are presented in Table 5. Note that the results for the specification in Fig. 1 are as expected; all assignment candidates must be part of the strategy and we find two symbolic transitions corresponding to *Produce* and *Consume*. We also see that the number of symbolic transitions is generally much smaller than the number of transitions an explicit-state model checker would find, as even simple specifications, like in Figure 1, would generate numerous transitions in explicit state model checking, but only two symbolic transitions.

Applications. We illustrate an application of our technique for bounded model checking [4] by the means of the example in Figure 3. In this example, three processes pass a unique token in one direction, with the goal of computing the largest process identifier.

Our technique extracts three symbolic transitions T_1, T_2, and T_3, each T_i being equivalent to $P(i) \wedge id' = id$ for $1 \leq i \leq 3$. As common in bounded model checking, with $[\![F]\!]_{i,i+1}$ we denote the SMT encoding of a transition by action F from an ith to an $(i+1)$-th state. For instance, $[\![Next]\!]_{0,1}$ and $[\![T_3]\!]_{0,1}$ encode the transitions from the state 0 to the state 1 by *Next* and T_3. Likewise, $[\![Init]\!]_0$ encodes SMT constraints by *Init* on the initial states. One can use the SMT encodings introduced in [20,21].

—— MODULE *max* ——

EXTENDS *Naturals*

VARIABLE *tok*, *max*, *id*

$Init \triangleq tok = 1 \wedge id \in [1..3 \rightarrow Nat] \wedge max = 0$

$P(i) \triangleq tok = i \wedge tok' = 1 + i \% 3 \wedge max' = \text{IF } id[i] > max \text{ THEN } id[i] \text{ ELSE } max$

$Next \triangleq (P(1) \vee P(2) \vee P(3)) \wedge id' = id$

Fig. 3: A distributed maximum computation in a ring of three processes in TLA^+

13

$$\llbracket Init \rrbracket_0 \qquad \llbracket Next \rrbracket_{0,1} \quad \llbracket Next \rrbracket_{1,2} \quad \llbracket Next \rrbracket_{2,3} \qquad \Bigg| \qquad \llbracket Init \rrbracket_0 \qquad \begin{array}{ccc} \llbracket T_1 \rrbracket_{0,1} & \llbracket T_1 \rrbracket_{1,2} & \llbracket T_1 \rrbracket_{2,3} \\ \llbracket T_2 \rrbracket_{0,1} & \llbracket T_2 \rrbracket_{1,2} & \llbracket T_2 \rrbracket_{2,3} \\ \llbracket T_3 \rrbracket_{0,1} & \llbracket T_3 \rrbracket_{1,2} & \llbracket T_3 \rrbracket_{2,3} \end{array}$$
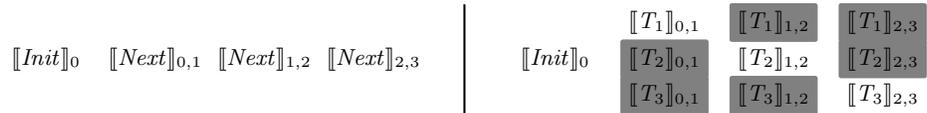
Fig. 4: SMT formulas that are constructed when checking the executions up to length 4: using the action *Next* (left), and using symbolic transitions (right). The gray formulas are excluded from the SMT context during the exploration.

Figure 4 shows the SMT formulas that are constructed by a bounded model checker when exploring executions up to length 4. (For the sake of space, we omit the formulas that check property violation.) On one hand, the monolithic encoding that uses only *Next* has to keep all the formulas in the SMT context. On the other hand, by incrementally checking satisfiability of the SMT context, the model checker can discover that some formulas — for example, $\llbracket T_2 \rrbracket_{0,1}$ and $\llbracket T_3 \rrbracket_{1,2}$ — lead to unsatisfiability and prune them from the SMT context. Similar approach improves efficiency of bounded model checking $\mathtt{C}$ programs [5][Ch. 16], hence, we expect it to be effective for the verification of TLA^+ specifications too.

8 Conclusions

We have introduced a technique to compute symbolic transitions of a TLA^+ specification by finding expressions that can be interpreted as assignments. Importantly, we designed the technique with soundness in mind. Detailed proofs can be found in the report [16]. We believe that our results can be used as a first preprocessing step when developing a symbolic model checker or a type checker for TLA^+.

As in the case of TLC, one can find TLA^+ specifications, for which no assignment strategy exists. However, TLA^+ users are systematically checking their specifications with TLC, in order to find simple errors. Hence, most of the benchmarks already come in a form compatible with TLC. Thus, we expect our approach to also work in practice. Based on these ideas, we are currently developing a symbolic model checker for TLA^+.

References

1. A collection of TLA^+ specifications. `https://github.com/tlaplus/Examples/`, [Online; accessed 21-October-2017]
2. Azmy, N., Merz, S., Weidenbach, C.: A rigorous correctness proof for pastry. In: Abstract State Machines, Alloy, B, TLA, VDM, and Z. pp. 86–101. Springer (2016)
3. Beyer, D., Keremoglu, M.E.: CPAchecker: A tool for configurable software verification. In: CAV. pp. 184–190 (2011)

4. Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic model checking without BDDs. In: TACAS. LNCS, vol. 1579, pp. 193–207 (1999)
5. Biere, A., Heule, M., van Maaren, H.: Handbook of satisfiability, vol. 185. IOS press (2009)
6. Bracha, G., Toueg, S.: Asynchronous consensus and broadcast protocols. Journal of the ACM 32(4), 824–840 (1985)
7. Chandra, T.D., Toueg, S.: Unreliable failure detectors for reliable distributed systems. Journal of the ACM 43(2), 225–267 (1996)
8. Chaudhuri, K., Doligez, D., Lamport, L., Merz, S.: The TLA$^+$ proof system: Building a heterogeneous verification platform. In: Theoretical aspects of computing. pp. 44–44. Springer-Verlag (2010)
9. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. J. ACM 50(5), 752–794 (2003)
10. Conchon, S., Goel, A., Krstic, S., Mebsout, A., Zaïdi, F.: Cubicle: A parallel smt-based model checker for parameterized systems - tool paper. In: CAV. pp. 718–724 (2012)
11. Dijkstra, E.W., Feijen, W.H., Van Gasteren, A.M.: Derivation of a termination detection algorithm for distributed computations. In: Control Flow and Data Flow: concepts of distributed programming, pp. 507–512. Springer (1986)
12. Gafni, E., Lamport, L.: Disk paxos. Distributed Computing 16(1), 1–20 (2003)
13. Guerraoui, R.: On the hardness of failure-sensitive agreement problems. Information Processing Letters 79(2), 99–104 (2001)
14. Konnov, I., Lazić, M., Veith, H., Widder, J.: A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. In: POPL. pp. 719–734 (2017)
15. Kroening, D., Tautschnig, M.: CBMC - C bounded model checker - (competition contribution). In: TACAS. pp. 389–391 (2014)
16. Kukovec, J., Tran, T.H., Konnov, I.: Extracting symbolic transitions from TLA$^+$ specifications (technical report 2018). `http://forsyte.at/wp-content/uploads/abz2018_full.pdf`, [Online; accessed 7-Feb-2018]
17. Lamport, L.: The part-time parliament. ACM TCS 16(2), 133–169 (1998)
18. Lamport, L.: Specifying systems: the TLA$^+$ language and tools for hardware and software engineers. Addison-Wesley Longman Publishing Co., Inc. (2002)
19. Merz, S.: The specification language TLA$^+$. In: Logics of specification languages, pp. 401–451. Springer (2008)
20. Merz, S., Vanzetto, H.: Automatic verification of TLA$^+$ proof obligations with SMT solvers. In: LPAR. vol. 7180, pp. 289–303. Springer (2012)
21. Merz, S., Vanzetto, H.: Harnessing SMT solvers for TLA+ proofs. ECEASST 53 (2012)
22. Moraru, I., Andersen, D.G., Kaminsky, M.: There is more consensus in egalitarian parliaments. In: SOSP. pp. 358–372. ACM (2013)
23. Newcombe, C., Rath, T., Zhang, F., Munteanu, B., Brooker, M., Deardeuff, M.: How Amazon web services uses formal methods. Comm. ACM 58(4), 66–73 (2015)
24. Ongaro, D.: Consensus: Bridging theory and practice. Ph.D. thesis, Stanford U. (2014)
25. Raynal, M.: A case study of agreement problems in distributed systems: Non-blocking atomic commitment. In: HASE. pp. 209–214 (1997)
26. Yu, Y., Manolios, P., Lamport, L.: Model checking TLA$^+$ specifications. In: Correct Hardware Design and Verification Methods, pp. 54–66. Springer (1999)