

# Translating B to TLA<sup>+</sup> for Validation with TLC

Dominik Hansen and Michael Leuschel

Institut für Informatik, Universität Düsseldorf\*\*  
Universitätsstr. 1, D-40225 Düsseldorf  
{hansen, leuschel}@cs.uni-duesseldorf.de

**Abstract.** The state-based formal methods B and TLA<sup>+</sup> share the common base of predicate logic, arithmetic and set theory. However, there are still considerable differences, such as the way to specify state transitions, the different approaches to typing, and the available tool support. In this paper, we present a translation from B to TLA<sup>+</sup> to validate B specifications using the model checker TLC. We provide translation rules for almost all constructs of B, in particular for those which are not built-in in TLA<sup>+</sup>. The translation also includes many adaptations and optimizations to allow efficient checking by TLC. Moreover, we present a way to validate liveness properties for B specifications under fairness conditions. Our implemented translator, TLC4B, automatically translates a B specification to TLA<sup>+</sup>, invokes the model checker TLC, and translates the results back to B. We use PROB to double check the counter examples produced by TLC and replay them in the PROB animator. We also present a series of case studies and benchmark tests comparing TLC4B and PROB.

**Keywords:** TLA<sup>+</sup>, B-Method, Tool Support, Model Checking, Animation.

## 1 Introduction and Motivation

B [1] and TLA<sup>+</sup> [8] are both state-based formal methods rooted in predicate logic, combined with arithmetic, set theory and support for mathematical functions. However, as already pointed out in [5], there are considerable differences:

- B is strongly typed, while TLA<sup>+</sup> is untyped. For the translation, it is obviously easier to translate from a typed to an untyped language than vice versa.
- The concepts of modularization are quite different.
- Functions in TLA<sup>+</sup> are total, while B supports relations, partial functions, injections, bijections, etc.
- B is limited to invariance properties, while TLA<sup>+</sup> also allows the specification of liveness properties.
- The structure of a B machine or development is prescribed by the B-method, while in a TLA<sup>+</sup> specification any formula can be considered as a system specification.

---

\*\* Part of this research has been sponsored by the EU funded FP7 project 287563 (ADVANCE) and the DFG funded project Gepavas.

As far as tool support is concerned,  $TLA^+$  is supported by the explicit state model checker TLC [13] and more recently by the TLAPS prover [3]. TLC has been used to validate a variety of distributed algorithms (e.g., [4]) and protocols. B has extensive proof support, e.g., in the form of the commercial product AtelierB [2] and the animator, constraint solver and model checker PROB [9,10]. Both AtelierB and PROB are being used by companies, mainly in the railway sector for safety critical control software. In an earlier work [5] we have presented a translation from  $TLA^+$  to B, which enabled applying the PROB tool to  $TLA^+$  specifications. In this paper we present a translation from B to  $TLA^+$ , this time with the main goal of applying the model checker TLC to B specifications. Indeed, TLC is a very efficient model checker for  $TLA^+$  with an efficient disk-based algorithm and support for fairness. PROB has an LTL model checker, but it does not support fairness (yet) and is entirely RAM-based. The model checking core of PROB is less tuned than  $TLA^+$ . On the other hand, PROB incorporates a constraint solver and offers several features which are absent from TLC, in particular an interactive animator with various visualization options. One feature of our approach is to replay the counter-examples produced by TLC within PROB, to get access to those features but also to validate the correctness of our translation. In this paper, we also present a thorough empirical evaluation between TLC and PROB. The results show that for lower-level, more explicit formal models, TLC fares better, while for certain more high-level formal models PROB is superior to TLC because of its constraint solving capabilities. The addition of a lower-level model checker thus opens up many new application possibilities.

## 2 Translation

The complete translation process from B to  $TLA^+$  and back to B is illustrated in Fig. 1.

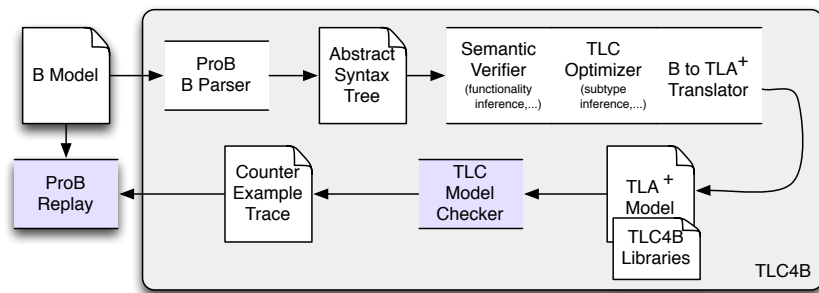


Fig. 1. The TLC4B Translation and Validation Process

Before explaining the individual phases, we will illustrate the translation with an example and explain the various phases based on that example. Translation rules of all data types and operators can be found in the extended version of our paper [6]. More specific implementation details will be covered in Section 4.

## 2.1 Example

Below we use a specification (adapted from [10]) of a process scheduler (Fig. 2). The specified system allows at most one process to be active. Each process can qualify for being selected by the scheduler by entering a FIFO queue. The specification contains two variables: a partial function *state* mapping each process to its state (a process must be created before it has a state) and a FIFO queue modeled as an injective sequence of processes. In the initial state no process is created and the queue is empty. The specification contains various operations to create (*new*), delete (*del*), or add a process to queue (*addToQueue*). Additionally, there are two operations to load a process into the processor (*enter*) and to remove a process from the processor (*leave*). The specification contains two safety conditions beside the typing predicates of the variables:

- At most one process should be active.
- Each process in the FIFO queue should have the state *ready*.

The translated TLA<sup>+</sup>- specification is shown in Fig. 3. At the beginning of the module some standard modules are loaded via the *EXTENDS* statement. These modules contain several operators used in this specification. The enumerated set *STATE* is translated as a constant definition. The definition itself is renamed (into *STATE\_1*) by the translator because *STATE* is keyword in TLA<sup>+</sup>. The invariant of the B specification is divided into several definitions in the TLA<sup>+</sup> module. This enables TLC to provide better feedback about which part of the invariant is violated. We translate each B operation as a TLA<sup>+</sup> action. Substitutions are translated as before-after predicates where a primed variable represents the variable in the next state. Unchanged variables must be explicitly specified. Note that a parameterized operation is translated as existential quantification. The quantification itself is located in the next-state relation *Next*, which is a disjunction of all actions. Some of the operations' guards appear in the *Next* definition rather than in the corresponding action. This is an instance of our translator optimizing the translation for the interpretation with TLC. The whole TLA<sup>+</sup> specification is described by the *Spec* definition. A valid behavior for the system has to satisfy the *Init* predicate for the initial state and then each step of the system must satisfy the next-state relation *Next*.

To validate the translated TLA<sup>+</sup> specification with TLC we have to provide an additional configuration file (Fig 4) telling TLC the main (specification) definition and the invariant definitions. Moreover, we have to assign values to all con-

```

MODEL Scheduler
SETS
  PROCESSES;
  STATE = {idle, ready, active}
VARIABLES
  state,
  queue
INVARIANT
  state ∈ PROCESSES ↔ STATE
  & queue ∈ iseq(PROCESSES)
  & card(state-1{active}) ≤ 1
  & !x.(x ∈ ran(queue) ⇒ state(x) = ready)
INITIALISATION
  state := {} || queue := []
OPERATIONS
  new(p) = SELECT p ∉ dom(state)
    THEN state := state ∪ {(p ↦ idle)} END;

  del(p) = SELECT p ∈ dom(state) ∧ state(p) = idle
    THEN state := {p} ⋄ state END;

  addToQueue(p) = SELECT p ∈ dom(state) ∧ state(p) = idle
    THEN state(p) := ready || queue := queue ← p END;

  enter = SELECT queue ≠ [] ∧ state-1{active} = {}
    THEN state(first(queue)) := active || queue := tail(queue) END;

  leave(p) = SELECT p ∈ dom(state) ∧ state(p) = active
    THEN state(p) := idle END
END

```

**Fig. 2.** MODEL Scheduler

stants of the module.<sup>1</sup> In this case we assign a set<sup>2</sup> of model values to the constant *PROCESSES* and single model values to the other constants. In terms of functionality, model values correspond to elements of an enumerated set in B. Model values are equal to themselves and unequal to all other model values.

<sup>1</sup> We translate a constant as variable if the axioms allow several solutions for this constant. All possible solution values will be enumerated in the initialization and the variable will be kept unchanged by all actions.

<sup>2</sup> The size of the set is a default number or can be specified by the user.

MODULE *Scheduler*

EXTENDS *Naturals, FiniteSets, Sequences, Relations, Functions,*  
*FunctionsAsRelations, SequencesExtended, SequencesAsRelations*

CONSTANTS *PROCESSES, idle, ready, active*

VARIABLES *state, queue*

$STATE\_1 \triangleq \{idle, ready, active\}$

$Invariant1 \triangleq state \in RelParFuncEleOf(PROCESSES, STATE\_1)$

$Invariant2 \triangleq queue \in ISeqEleOf(PROCESSES)$

$Invariant3 \triangleq Cardinality(RelImage(RelInverse(state), \{active\})) \leq 1$

$Invariant4 \triangleq \forall x \in Range(queue) : RelCall(state, x) = ready$

$Init \triangleq \wedge state = \{\}$   
 $\wedge queue = \langle \rangle$

$new(p) \triangleq \wedge state' = state \cup \{p, idle\}$   
 $\wedge UNCHANGED \langle queue \rangle$

$del(p) \triangleq \wedge RelCall(state, p) = idle$   
 $\wedge state' = RelDomSub(\{p\}, state)$   
 $\wedge UNCHANGED \langle queue \rangle$

$addToQueue(p) \triangleq \wedge RelCall(state, p) = idle$   
 $\wedge state' = RelOverride(state, \{p, ready\})$   
 $\wedge queue' = Append(queue, p)$

$enter \triangleq \wedge queue \neq \langle \rangle$   
 $\wedge RelImage(RelInverse(state), \{active\}) = \{\}$   
 $\wedge state' = RelOverride(state, \{Head(queue), active\})$   
 $\wedge queue' = Tail(queue)$

$leave(p) \triangleq \wedge RelCall(state, p) = active$   
 $\wedge state' = RelOverride(state, \{p, idle\})$   
 $\wedge UNCHANGED \langle queue \rangle$

$Next \triangleq \vee \exists p \in PROCESSES \setminus RelDomain(state) : new(p)$   
 $\vee \exists p \in RelDomain(state) : del(p)$   
 $\vee \exists p \in RelDomain(state) : addToQueue(p)$   
 $\vee enter$   
 $\vee \exists p \in RelDomain(state) : leave(p)$

$vars \triangleq \langle state, queue \rangle$

$Spec \triangleq Init \wedge \square [Next]_{vars}$

**Fig. 3.** Module Scheduler

```

SPECIFICATION Spec
INVARIANT Invariant1, Invariant2, Invariant3, Invariant4
CONSTANTS
PROCESSES = {PROCESSES1, PROCESSES2, PROCESSES3}
idle = idle
ready = ready
active = active

```

**Fig. 4.** Configuration file for module Scheduler

## 2.2 Translating Data Values and Functionality Inference

Due to the common base of B and  $TLA^+$ , most data types exist in both languages, such as sets, functions and numbers. As a consequence, the translation of these data types is almost straightforward.

$TLA^+$  has no built-in concept of Relations<sup>3</sup>, but  $TLA^+$  provides all necessary data types to define relations based on the model of the B-Method. We represent a relation in  $TLA^+$  as a set of tuples (e.g.  $\{\langle 1, TRUE \rangle, \langle 1, FALSE \rangle \langle 2, TRUE \rangle\}$ ). The drawback of this approach is that in contrast to B,  $TLA^+$ 's own functions and sequences are not based on the relations defined in this way. As an example, we cannot specify a  $TLA^+$  built-in function as a set of pairs; in B it is usual to do this as well as to apply set operators (e.g. the union operator as in  $f \cup \{2 \mapsto 3\}$ ) to functions or sequences. To support such a functionality in  $TLA^+$ , functions and sequences should be translated as relations if they are used in a “relational way”. It would be possible to always translate functions and sequences as relations. But in contrast to relations, functions and sequences are built-in data types in  $TLA^+$  and their evaluation is optimized by TLC (e.g. lazy evaluation). Hence we extended the B type-system to distinguish between functions and relations. Thus we are able to translate all kinds of relations and to deliver an optimized translation.

We use a type inference algorithm adapted to the extended B type-system to get the required type information for the translation. Unifying a function type with a relation type will result in a relation type (e.g.  $\mathbb{P}(\mathbb{Z} \times \mathbb{Z})$  for both sides of the equation  $\lambda x.(x \in 1..3|x + 1) = \{(1, 1)\}$ ). However, there are several relational operators preserving a function type if they are applied to operands with a function type (e.g. *ran*, *front* or *tail*). For these operators we have to deliver two translation rules (functional vs relational). Moreover the algorithm verifies the type correctness of the B specification (i.e. only values of the same type can be compared with each other).

## 2.3 Translating Operators

In  $TLA^+$  some common operators such as arithmetic operators are not built-in operators. They are defined in separate modules called standard modules that can

<sup>3</sup> Relations are not mentioned in the language description of [8]. In [7] Lamport introduces relations in  $TLA^+$  only to define the transitive closure.

be extended by a specification.<sup>4</sup> We reuse the concept of standard modules to include the relevant B operators. Due to the lack of relations in TLA<sup>+</sup> we have to provide a module containing all relational operators (Fig. 5).

MODULE <i>Relations</i>	
EXTENDS <i>FiniteSets, Naturals, TLC</i>	
$Relation(X, Y) \triangleq$	SUBSET $(X \times Y)$
$RelDomain(R) \triangleq$	$\{x[1] : x \in R\}$
$RelRange(R) \triangleq$	$\{x[2] : x \in R\}$
$RelInverse(R) \triangleq$	$\{(x[2], x[1]) : x \in R\}$
$RelDomRes(S, R) \triangleq$	$\{x \in R : x[1] \in S\}$ Domain restriction
$RelDomSub(S, R) \triangleq$	$\{x \in R : x[1] \notin S\}$ Domain subtraction
$RelRanRes(R, S) \triangleq$	$\{x \in R : x[2] \in S\}$ Range restriction
$RelRanSub(R, S) \triangleq$	$\{x \in R : x[2] \notin S\}$ Range subtraction
$RelImage(R, S) \triangleq$	$\{y[2] : y \in \{x \in R : x[1] \in S\}\}$
$RelOverride(R1, R2) \triangleq$	$\{x \in R : x[1] \notin RelDomain(R2)\} \cup R2$
$RelComposition(R1, R2) \triangleq$	$\{(u[1][1], u[2][2]) : u \in$ $\{x \in RelRanRes(R1, RelDomain(R2)) \times RelDomRes(RelRange(R1), R2) :$ $x[1][2] = x[2][1]\}$
⋮	

**Fig. 5.** An extract of the Module Relations

Moreover B provides a rich set of operators on functions such as all combinations of partial/total and injective/surjective/bijective. In TLA<sup>+</sup> we only have total functions. We group all operators on functions together in an additional module (Fig. 6). Sometimes there are several ways to define an operator. We choose the definition which can be best handled by TLC.<sup>5</sup>

Some operators exists in both languages but their definitions differ slightly. For example, the B-Method requires that the first operand for the modulo operator must be a natural number. In TLA<sup>+</sup> it can be also a negative number.

Operator	B-Method	TLA <sup>+</sup>
$a \text{ modulo } b$	$a \in \mathbb{N} \wedge b \in \mathbb{N}_1$	$a \in \mathbb{Z} \wedge b \in \mathbb{N}_1$

To verify B's well-definedness condition for modulo we use TLC's ability to check assertions. The special operator  $Assert(P, out)$  throws a runtime exception with the error message  $out$  if the predicate  $P$  is false. Otherwise,  $Assert$  will be evaluated to TRUE. The B modulo operator can thus be expressed in TLA<sup>+</sup> as

<sup>4</sup> TLC supports operators of the common standard modules Integers and Sequences in a efficient way by overwriting them with Java methods.

<sup>5</sup> Note that some of the definitions are based on the *Cardinality* operator that is restricted to finite sets.

```

EXTENDS FiniteSets
Range(f)  $\triangleq$  {f[x] : x ∈ DOMAIN f}
Image(f, S)  $\triangleq$  {f[x] : x ∈ S}
TotalInjFunc(S, T)  $\triangleq$  {f ∈ [S → T] :
    Cardinality(DOMAIN f) = Cardinality(Range(f))}
ParFunc(S, T)  $\triangleq$  UNION {[x → T] : x ∈ SUBSET S}
ParInjFunc(S, T)  $\triangleq$  {f ∈ ParFunc(S, T) :
    Cardinality(DOMAIN f) = Cardinality(Range(f))}
⋮

```

**Fig. 6.** An extract of the Module Functions

follows:

$$\text{Modulo}(a, b) \triangleq \text{IF } a \geq 0 \text{ THEN } a \% b \text{ ELSE } \text{Assert}(\text{FALSE}, \text{"WD ERROR"})$$

We also have to consider well-definedness conditions if we apply a function call to a relation:

$$\begin{aligned} \text{RelCall}(r, x) \triangleq & \text{IF } \text{Cardinality}(r) = \text{Cardinality}(\text{RelDom}(r)) \wedge x \in \text{RelDom}(r) \\ & \text{THEN } (\text{CHOOSE } y \in r : y[1] = x)[2] \\ & \text{ELSE } \text{Assert}(\text{FALSE}, \text{"WD ERROR"}) \end{aligned}$$

In summary, we provide the following standard modules for our translation:

- Relations
- Functions
- SequencesExtended (Some operators on sequences which are not included the standard module Sequences)
- FunctionsAsRelations (Defines all function operators on sets of pairs ensuring their well-definedness conditions)
- SequencesAsRelations (Defines all operators on sequences which are represented as sets of pairs.)
- BBuiltins (Miscellaneous operators e.g. modulo, min, max, sigma or pi)

## 2.4 Optimizations

**Subtype Inference** Firstly, we will describe how TLC evaluates expressions: In general TLC evaluates an expression from left to right. Evaluating an expression containing a bound variable such as an existential quantification ( $\exists x \in S : P$ ), TLC enumerates all values of the associated set and then substitutes them for the bound variable in the corresponding predicate. Due to missing constraint solving techniques, TLC is not able to evaluate another variant of the existential quantification without an associated set ( $\exists x : P$ ). This version is also a valid TLA<sup>+</sup> expression and directly corresponds to the way of writing an existential quantification in B ( $\exists x.(P)$ ). However, we confine our translations to the subset of TLA<sup>+</sup>



which is supported by TLC. Thus the translation is responsible for making all required adaptations to deliver an executable TLA<sup>+</sup> specification. For the existential quantification (or all other expressions containing bound variables), we use the inferred type  $\tau$  of the bound variable as the associated set ( $\exists x \in \tau_x : P.$ ) However, in most cases, it is not performant to let TLC enumerate over a type of a variable, in particular TLC aborts if it has to enumerate a infinite set. Alternatively, it is often possible to restrict the type of the bound variable based on a static analysis of the corresponding (typing) predicate. We use a pattern matching algorithm to find the following kind of expressions where  $x$  is a bound variable,  $e$  is an expression, and  $S$  is ideally a subset of the type<sup>6</sup>:  $x = e$ ,  $x \in S$ ,  $x \subseteq S$  or  $x \subset S$ .

If more than one of these patterns can be found for one variable, we build the intersection to keep the associated set as small as possible:

$$\begin{array}{ll} \text{B-Method} & \text{TLA}^+ \\ \exists x.(x = e \wedge x \in S_1 \wedge x \subseteq S_2 \wedge P) & \exists x \in (\{e\} \cap S_1 \cap \text{SUBSET } S_2) : P \end{array}$$

This reduces the number of times TLC has to evaluate the predicate  $P$ .<sup>7</sup>

**Lazy Evaluation** Sometimes TLC can use heuristics to evaluate an expression. For example TLC can evaluate  $\langle 1, 2, 1 \rangle \in \text{Seq}(\{1, 2\})$  to true without evaluating the infinite set of sequences. We will show how we can use these heuristics to generate an optimized translation. As mentioned before functions have to be translated as relations if they are used in a relational way in the B specification. How then should we translate the set of all total functions ( $S \rightarrow T$ )? The easiest way is to convert each function to a relation in TLA<sup>+</sup>:

$$\text{MakeRel}(f) \triangleq \{ \langle x, f[x] \rangle : x \in \text{DOMAIN } f \}$$

The resulting operator for the set of all total functions is:

$$\text{RelTotalFunctions}(S, T) \triangleq \{ \text{MakeRel}(f) : f \in [S \rightarrow T] \}$$

However this definition has a disadvantage, if we just want to check if a single function is in this set the whole set will be evaluated by TLC. Using the following definition TLC avoids the evaluation of the whole set:

$$\begin{aligned} \text{RelTotalFunctionsEleOf}(S, T) \triangleq \{ f \in \text{SUBSET}(S \times T) : \\ \wedge \text{Cardinality}(\text{RelDomain}(f)) = \text{Cardinality}(f) \\ \wedge \text{RelDomain}(f) = S \} \end{aligned}$$

In this case, TLC only checks if a function is a subset of the cartesian product (the whole Cartesian product will not be evaluated) and the conditions are checked

<sup>6</sup> The B language description in [2] requires that each (bound) variable must be typed by one of these patterns before use.

<sup>7</sup> In some cases, the associated set is still infinite and the user has to restrict the set to be finite.

only once. Moreover this definition fares well even if  $S$  or  $T$  are sets of functions (e.g.  $S \rightarrow V \rightarrow W$  in B). The advantage of the first definition is that it is faster when the whole set must be evaluated. As a consequence, we use both definitions for our translation and choose the first if TLC has to enumerate the set (e.g.  $\exists x \in \text{RelTotalFunctions}(S, T) : P$ ) and the second testing if a function belongs to the set (e.g.  $f \in \text{RelTotalFunctionsEleOf}(S, T)$  as an invariant).

### 3 Checking Temporal Formulas

One of the main advantages of  $\text{TLA}^+$  is that temporal properties can be specified directly in the language itself. Moreover the model checker TLC can be used to verify such formulas. But before we show how to write temporal formulas for a B specification we first have to describe a main distinction between the two formal methods. In contrast to B, the standard template of a  $\text{TLA}^+$  specification ( $\text{Init} \wedge \Box[\text{Next}]_{\text{vars}}$ ) allows stuttering steps at any time.<sup>8</sup> This means that a regular step of a  $\text{TLA}^+$  specification is either a step satisfying one of the actions or a stuttering step leaving all variables unchanged. When checking a specification for errors such as invariant violations it is not necessary to consider stuttering steps, because such an error will be detected in a state and stuttering steps only allow self transitions and do not add additional states. For deadlock checking stuttering steps are also not regarded by TLC, but verifying a temporal formula with TLC often ends in a counter-example caused by stuttering steps. For example, assuming we have a very simple specification of a counter in  $\text{TLA}^+$  with a single variable  $c$

$$\text{Spec} \triangleq c = 1 \wedge \Box[c' = c + 1]_c$$

We would expect that the counter will eventually reach 10 ( $\Diamond(c = 10)$ ). However TLC will report a counter-example, saying that at a certain state (before reaching 10), an infinite number of stuttering steps occurs and 10 will never be reached. From the B side we do not want to deal with these stuttering steps.  $\text{TLA}^+$  allows to add fairness conditions to the specification to avoid infinite stuttering steps. Adding weak fairness for the next-state relation ( $\text{WF}_{\text{vars}}(\text{Next})$ ) would prohibit an infinite number of stuttering steps if a step of the next-state relation is possible (i.e.  $\text{Next}$  is always enabled):

$$\text{WF}_{\text{vars}}(A) \triangleq \begin{aligned} &\vee \Box \Diamond (\langle A \rangle_{\text{vars}}) \\ &\vee \Box \Diamond (\neg \text{ENABLED } \langle A \rangle_{\text{vars}}) \end{aligned}$$

However this fairness condition is too strong: It asserts that either the action  $A$  will be executed infinitely often changing the state of the system ( $A$  must not be a stuttering step)

$$\langle A \rangle_{\text{vars}} \triangleq A \wedge \text{vars}' \neq \text{vars}$$

or  $A$  will be disabled infinitely often. Assuming weak fairness for the next state relation will also eliminate user defined stuttering steps. User defined stuttering

<sup>8</sup>  $\overline{[\text{Next}]_{\text{vars}}} \triangleq \text{Next} \vee \text{UNCHANGED vars}$

steps result from B operations which do not change the state of the system (e.g. skip or call operations). These stuttering steps may cause valid counter-examples and should not be eliminated. Hence, the translation should retain user defined stuttering steps in the translated TLA<sup>+</sup> specification and should disable stuttering steps which are implicitly included. In [12], Richards describes a way to distinguish between these two kinds of stuttering steps in TLA<sup>+</sup>. We use his definition of “Very Weak Fairness” applied to the next state relation ( $VWF_{vars}(Next)$ ) to disable implicit stuttering steps and allow user defined stuttering steps in the TLA<sup>+</sup> specification:

$$VWF_{vars}(A) \triangleq \begin{aligned} &\vee \square \diamond (\langle A \rangle_{vars}) \\ &\vee \square \diamond (\neg \text{ENABLED } \langle A \rangle_{vars}) \\ &\vee \square \diamond (\text{ENABLED } (A \wedge \text{UNCHANGED } vars)) \end{aligned}$$

The definition of  $VWF(A)$  is identical to  $WF(A)$  except for an additional third case allowing infinite stuttering steps if  $A$  is a stuttering action ( $A \wedge \text{UNCHANGED } vars$ ). We define the resulting template of the translated TLA<sup>+</sup> specification as follows:

$$Init \wedge \square [Next]_{vars} \wedge VWF_{vars}(Next)$$

We allow the B user to use following temporal operators to define liveness conditions for a B specification:<sup>9</sup>

- $\square f$  (Globally)
- $\diamond f$  (Finally)
- $ENABLED(op)$  (Check if the operation  $op$  is enabled)
- $\exists x.(P \wedge f)$  (Existential quantification)
- $\forall x.(P \Rightarrow f)$  (Universal quantification)
- $WF(op)$  (Weak Fairness will be translated to VWF)
- $SF(op)$  (Strong Fairness will be translated to “Almost Strong Fairness”<sup>10</sup>)

## 4 Implementation & Experiments

Our translator, called TLC4B, is implemented in Java and it took about six months to develop the initial version. Figure 1 in Section 2 shows the translation and validation process of TLC4B. After parsing the specification TLC4B performs some static analyses (e.g. type checking or checking the scope of the variables) verifying the semantic correctness of the B specification. Moreover, as explained in Section 2, TLC4B extracts required information from the B specification (e.g. subtype inference) to generate an optimized translation. Subsequently, TLC4B creates a TLA<sup>+</sup> module with an associated configuration file and invokes the model checker TLC. The results produced by TLC are translated back to B. For example, a goal predicate is translated as a negated invariant. If this invariant is violated, a “Goal found”

<sup>9</sup> We demonstrate the translation of a liveness condition with a concrete example in the extended version of this paper [6].

<sup>10</sup> Analogically Richards defines “Almost Strong Fairness” (ASF) as a weaker version of strong fairness (SF) reflecting the different kinds of stuttering steps

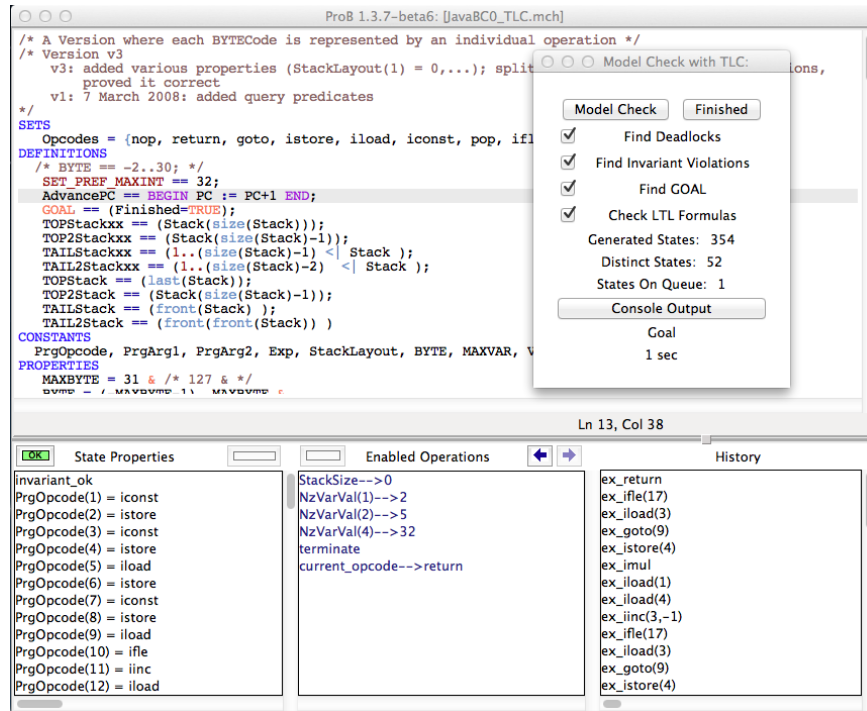


Fig. 7. PROB animator

message is reported. We expect TLC to find the following kinds of errors in the B specification:

- Deadlocks
- Invariant violations
- Assertion errors
- Violations of constants properties (i.e., axioms over the B constants are false)
- Well-definedness violations
- Violations of temporal formulas

For certain kinds of errors such as a deadlock or an invariant violation, TLC reports a trace leading to the state where the error occurs. A trace is a sequence of states where each state is a mapping from variables to values. TLC4B translates the trace back to B (as a list of B state predicates). TLC4B has been integrated into PROB as of version 1.3.7-beta10: The user needs no knowledge of TLA<sup>+</sup> because the translation is completely hidden. As shown in figure 7 counter-examples found by TLC are automatically replayed in the PROB animator (displayed in the history pane) to give the user an optimal feedback.

We have successfully validated several existing models from the literature (Fig. 8). The following examples show some fields of application of TLC4B. The experiments were all run on a Macbook Air with Intel Core i5 1,8 GHz processor, running TLC Version 2.05 and Prob version 1.3.7-beta9.

**Can-Bus** One example is a 314 line B specification of the Controller Area Network (CAN) Bus, containing 18 variables and 21 operations. The specification is rather low level, i.e., the operations consist of simple assignments of concrete values to variables (no constraint solving is required). TLC4B needs 1.5 seconds<sup>11</sup> to translate the specification to TLA<sup>+</sup> and less than 6 seconds for the validation of the complete state space composed of 132,598 states. PROB needs 192 seconds to visit the same number of states. Both model checkers report no errors. For this specification TLC benefits from its efficient algorithm for storing big state spaces.

**Invariant violations** We use a defective specification of a travel agency system (CarlaTravelAgencyErr) to test the abilities of TLC4B detecting invariant violations. The specification consists of 295 line of B code, 11 variables and 10 operations. Most of the variables are functions (total, partial and injective) which are also manipulated by relational operators. TLC4B needs about 3 seconds to translate the model and to find the invariant violation. 377 states are explored with the aid of the breadth first search and the resulting trace has a length of 5 states. PROB needs roughly the same time.

**Benchmarks** Besides the evaluation of real case studies, we use some specific benchmark tests comparing TLC4B and PROB. We use a specification of a simple counter testing TLC4B's abilities to explore a big (linear) state space. TLC4B needs 3 seconds to explore the state space with 1 million states. Comparatively, PROB takes 204 seconds. In another specification, the states of doors are controlled. The specification allows the doors to be opened and closed. We use two versions: In the first version the state of the doors are represented as a function (Doors\_Functions) and in the second as a relation (Doors\_Relations). The first version allows TLC4B to use TLA<sup>+</sup> functions for the translation and TLC needs 2 seconds to explore 32,768 states. For the second version TLC4B uses the newly introduced relations and takes 10 seconds. As expected, TLC can evaluate built-in operators faster than user defined operators. Hence the distinction TLC4B has between functions and relations can make a significant difference in running times. PROB needs about 100 seconds to explore the state space of both specifications. However, PROB needs less than a second using symmetry reduction.<sup>12</sup>

In summary, PROB is substantially better than TLC4B when constraint solving is required (NQueens, SumAndProduct, GraphIsomorphism<sup>13</sup>) or when naive enumeration of operation arguments is inefficient (GardnerSwitchingPuzzle). For some specifications (not listed in the table) TLC was not able to validate the translated TLA<sup>+</sup> specification because TLC had to enumerate an infinite set. On the other

---

<sup>11</sup> Most of this time is required to start the JVM and to parse the B specification.

<sup>12</sup> TLC's symmetry reduction does not scale for large symmetric sets.

<sup>13</sup> See <http://www.data-validation.fr/data-validation-reverse-engineering/> for larger industrial application of this type of task.

Model	Lines	Result	States	Transitions	PROB	TLC4B	$\frac{ProB}{Tlc4B}$
Counter	13	No Error	1000000	1000001	186.5	3.7	50.653
Doors_Functions	22	No Error	32768	983041	103.2	3.3	31.194
Can-Bus	314	No Error	132598	340265	191.8	7.2	26.624
KnightsTour <sup>(1)</sup>	28	Goal	508450	678084	817.5	34.1	23.998
USB_4Endpoints	197	NoError	16905	550418	72.5	5.7	12.632
Countdown	67	Inv. Viol.	18734	84617	31.4	2.8	11.073
Doors_Relations	22	No Error	32768	983041	103.3	11.6	8.926
Simpson_Four_Slot	78	No Error	46657	11275	33.7	4.3	7.874
EnumSetLockups	34	No Error	4375	52495	6.5	2.1	3.105
TicTacToe <sup>(1)</sup>	16	No Error	6046	19108	7.5	3.1	2.435
Cruise_finite1	604	No Error	1360	25696	6.2	3.2	1.954
CarlaTravelAgencyErr	295	Inv. Viol.	377	3163	3.3	3.1	1.069
FinalTravelAgency	331	No Error	1078	4530	4.7	4.4	1.068
CSM	64	No Error	77	210	1.4	1.6	0.859
SiemensMiniPilot_Abrial <sup>(1)</sup>	51	Goal	22	122	1.5	1.7	0.849
JavaBC-Interpreter	197	Goal	52	355	1.7	2.4	0.708
Scheduler	51	No Error	68	205	1.4	2.1	0.682
RussianPostalPuzzle	72	Goal	414	1159	1.7	2.8	0.588
Teletext_bench	431	No Error	13	122	1.8	3.7	0.496
WhoKilledAgatha	42	No Error	6	13	1.5	5.2	0.295
GardnerSwitchingPuzzle	59	Goal	206	502	2.5	11.7	0.213
NQueens_8	18	No Error	92	828	1.4	23.2	0.062
JobsPuzzle	66	Deadlock	2	2	1.6	29.3	0.053
SumAndProduct <sup>(1)</sup>	51	No Error	1	1	9.7	420.8	0.023
GraphIsomorphism	21	Deadlock	512	203	1.8	991.5	0.002

<sup>(1)</sup> Without Deadlock Check

**Fig. 8.** Empirical Results: Running times of Model Checking (times in seconds)

hand, TLC4B is substantially better than PROB for lower-level specifications with a large state space.

## 5 Correctness of the Translation

There are several possible cases where our validation of B models using TLC could be unsound: there could be a bug in TLC, there could be a bug in our TLA<sup>+</sup> library for the B operators, there could be a bug in our implementation of the translation from B to TLA<sup>+</sup>, there could be a fundamental flaw in our translation.

We have devised several approaches to mitigate those hazards. Firstly, when TLC finds a counter example it is replayed using PROB. In other words, every step of the counter example is double checked by PROB and the invariant or goal predicate is also re-checked by PROB. This makes it very unlikely that we produce incorrect counter examples. Indeed, PROB, TLC, and our translator have been

developed completely independently of each other and rely on different technology. Such independently developed double chains are often used in industry for safety critical tools.

The more tricky case is when TLC finds no counter example and claims to have checked the full state space. Here we cannot replay any counter example and we have the added difficulty that, contrary to PROB, TLC stores just fingerprints of states and that there is a small probability that not all states have been checked. We have no simple solution in this case, apart from re-checking the model using either PROB or formal proof. In addition, we have conducted extensive tests to validate our approach. For example, we use a range of models encoding mathematical laws to stress test our translation. These have proven to be very useful for detecting bugs in our translation and libraries (mainly bugs involving operator precedences). In addition, we have uncovered a bug in TLC relating to the cartesian product.<sup>14</sup> Moreover, we use a wide variety of benchmarks, checking that PROB and TLC produce the same result and generate the same number of states.

## 6 More Related Work, Discussion and Conclusion

Mosbahi et al. [11] were the first to provide a translation from B to TLA<sup>+</sup>. Their goal was to verify liveness conditions on B specifications using TLC. Some of their translation rules are similar to the rules presented in this paper. For example, they also translate B operations into TLA<sup>+</sup> actions and provide straightforward rules for operators which exist in both languages. However, there are also significant differences:

- Our main contribution is that we deliver translation rules for almost all B operators and in particular for those which are not build-in operators in TLA<sup>+</sup>. E.g., we specified the concept of relations including all operators on relations.
- Moreover, we also consider subtle differences between B and TLA<sup>+</sup> such as different well-definedness conditions and provide an appropriate translation.
- Regarding temporal formulas, we provide a way that a B user does not have to care about stuttering steps in TLA<sup>+</sup>.
- We restrict our translation to the subset of TLA<sup>+</sup> which is supported by the model checker TLC. Furthermore, we made many adaptations and optimizations allowing TLC to validate B specification efficiently.
- The implemented translator is fully automatic and does not require the user to know TLA<sup>+</sup>.

In the future, we would like to improve our automatic translator:

- Supporting modularization and the refinement techniques of B.
- Improving the performance of TLC by implementing Java modules for the new standard modules.
- Integrating TLC4B into Rodin and supporting Event-B specifications.

<sup>14</sup> TLC erroneously evaluates the expression  $\{1\} \times \{\} = \{\} \times \{1\}$  to *FALSE*.

In conclusion, by making TLC available to B models, we have closed a gap in the tool support and now have a range of complementary tools to validate B models: Atelier-B (or Rodin) providing automatic and interactive proof support, PROB being able to animate and model check high-level B specifications and providing constraint-based validation, and now TLC providing very efficient model checking of lower-level B specifications. The latter opens up many new possibilities, such as exhaustive checking of hardware models or sophisticated protocols. A strong point of our approach is the replaying of counter examples using PROB. Together with the work in [5] we have now constructed a two-way bridge between TLA<sup>+</sup> and B, and also hope that this will bring both communities closer together.

**Acknowledgements** We are grateful to Ivaylo Dobrikov for various discussions and support in developing the Tcl/Tk interface of TLC4B. We also would like to thank Stephan Merz and Leslie Lamport for very useful feedback concerning TLA<sup>+</sup> and TLC, and for giving us access to various specifications (like SumProduct). Finally, we are thankful to anonymous referees for their useful feedback.

## References

1. J.-R. Abrial. *The B-Book*. Cambridge University Press, 1996.
2. ClearSy. B language reference manual. [http://www.tools.clearsy.com/resources/Manrefb\\_en.pdf](http://www.tools.clearsy.com/resources/Manrefb_en.pdf). Accessed: 2013-11-10.
3. D. Cousineau, D. Doligez, L. Lamport, S. Merz, D. Ricketts, and H. Vanzetto. TLA+ proofs. In *FM*, pages 147–154, 2012.
4. E. Gafni and L. Lamport. Disk Paxos. *Distributed Computing*, 16(1):1–20, 2003.
5. D. Hansen and M. Leuschel. Translating TLA+ to B for validation with ProB. In *Proceedings iFM'2012*, LNCS 7321, pages 24–38. Springer, 2012.
6. D. Hansen and M. Leuschel. Translating B to TLA+ for validation with TLC. [http://www.stups.uni-duesseldorf.de/w/Special:Publication/HansenLeuschel\\_TLC4B\\_techreport](http://www.stups.uni-duesseldorf.de/w/Special:Publication/HansenLeuschel_TLC4B_techreport), 2013.
7. L. Lamport. The TLA+ hyperbook. <http://research.microsoft.com/en-us/um/people/lamport/tla/hyperbook.html>. Accessed: 2013-10-30.
8. L. Lamport. *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002.
9. M. Leuschel and M. Butler. ProB: A model checker for B. In K. Araki, S. Gnesi, and D. Mandrioli, editors, *FME 2003: Formal Methods*, LNCS 2805, pages 855–874. Springer-Verlag, 2003.
10. M. Leuschel and M. J. Butler. ProB: an automated analysis toolset for the B method. *STTT*, 10(2):185–203, 2008.
11. O. Mosbahi, L. Jemni, and J. Jaray. A formal approach for the development of automated systems. In J. Filipe, B. Shishkov, and M. Helfert, editors, *ICSOFT (SE)*, pages 304–310. INSTICC Press, 2007.
12. M. Reynolds. Changing nothing is sometimes doing something. Technical Report TR-98-02, Department of Computer Science, King's College London, February 1998.
13. Y. Yu, P. Manolios, and L. Lamport. Model checking TLA+ specifications. In L. Pierre and T. Kropf, editors, *Proceedings CHARME'99*, LNCS 1703, pages 54–66. Springer-Verlag, 1999.