

# Translating $TLA^+$ to B for Validation with PROB

Dominik Hansen and Michael Leuschel

Institut für Informatik, Universität Düsseldorf\*\*  
Universitätsstr. 1, D-40225 Düsseldorf

dominik.hansen@uni-duesseldorf.de, leuschel@cs.uni-duesseldorf.de

**Abstract.**  $TLA^+$  and B share the common base of predicate logic, arithmetic and set theory. However, there are still considerable differences, such as very different approaches to typing and modularization. There is also considerable difference in the available tool support. In this paper, we present a translation of the non-temporal part of  $TLA^+$  to B, which makes it possible to feed  $TLA^+$  specifications into existing tools for B. Part of this translation must include a type inference algorithm, in order to produce typed B specifications. There are many other tricky aspects, such as translating modules as well as let and if-then-else expressions. We also present an integration of our translation into PROB. PROB thus provides a complementary tool to the explicit state model checker TLC, with convenient animation and constraint solving for  $TLA^+$ . We also present a series of case studies, highlighting the complementarity to TLC. In particular, we highlight the sometimes dramatic difference in performance when it comes to solving complicated constraints in  $TLA^+$ .

**Keywords:** TLA, B-Method, Tool Support, Model Checking, Animation.

## 1 Introduction and Motivation

$TLA^+$  [5] and B [1] are both state-based formal methods rooted in predicate logic, combined with arithmetic and set theory. There are, however, considerable differences:

- $TLA^+$  is untyped, while B is strongly typed.
- the modularization concepts are very different (as we will see later in the paper).
- $TLA^+$  and B both support sets and functions. However, functions in  $TLA^+$  are total, while B supports relations, partial functions, injections, bijections, etc.
- $TLA^+$  has several constructs which are lacking in B, such as an IF-THEN-ELSE for expressions and predicates<sup>1</sup>, a LET construct or the choose operator. The latter enables one to define recursive functions over sets, which are awkward to define in B.

\*\* Part of this research has been sponsored by the EU funded FP7 project 214158: DEPLOY (Industrial deployment of advanced system engineering methods for high productivity and dependability).

<sup>1</sup> B only provides an IF-THEN-ELSE for substitutions.

As far as tool support is concerned, TLA<sup>+</sup> is supported by the explicit state model checker TLC [14], and more recently by the TLAPS prover [3, 2]. B has extensive proof support, e.g., in the form of the commercial product AtelierB [13]. The animator and model checker PROB [6] can also be applied to B specifications. Both AtelierB and PROB are being used by companies, mainly in the railway sector for safety critical control software. Some of the goals of our translation are to

- gain a better understanding of the common core and of the differences between TLA<sup>+</sup> and B
- gain an animator for TLA<sup>+</sup>
- gain a constraint solver for TLA<sup>+</sup>

Indeed, TLC is a very efficient model checker for TLA<sup>+</sup> with an efficient disk-based algorithm and support for fairness. PROB has an LTL model checker, but does not support fairness (yet) and is entirely RAM-based. The model checking core of PROB is less tuned than TLA<sup>+</sup>. However, PROB offers several features which are absent from TLC, notably an interactive animator with various visualization options. More importantly, the PROB kernel provides for constraint solving over predicate logic, set theory and arithmetic. PROB can also deal quite well with large data values. This has many applications, from constraint-based invariant or deadlock checking [4], over to test-case generation and on to improved animation because the user has to provide much less concrete values than with other tools. It also makes certain specifications “executable” which are beyond the reach of other tools such as TLC.

We suppose that the reader is familiar with either TLA<sup>+</sup> or B. Indeed, we hope that through our translation, TLA<sup>+</sup> constructs can be understood by B users and vice-versa. Below, in Sect. 2 we introduce the essentials of our translation on a simple example, while in Sect. 3 we present the translation more formally. (Some translation rules are relegated to an Appendix for referees; as are the sources of various examples.) In Sect. 4 we present case studies and experiments, and will also compare the tools PROB and TLC. We conclude with more related and future work in Sect. 5.

## 2 An example translation from TLA<sup>+</sup> to B

To allow B users to become familiar with TLA<sup>+</sup>, we present a variation of the well known HourClock example from Chapter 2 of [5]. Figure 1 shows the *MyHourClock* module, which avoids the IF-THEN-ELSE expression of the original at this point. The specification describes the typical behavior of a digital clock displaying only hours. The module starts with the MODULE clause followed by the name of the specification. The analogous clause of a B machine is MACHINE or MODEL. At the beginning of the module body, arithmetic operators such as + or “..” are loaded via EXTENDS from the standard module *Integers*. These operators are not built-in operators in TLA<sup>+</sup> and can either be defined by the user or imported with their usual meaning as here. The declaration of constants and variables is identical in

```

EXTENDS Integers
CONSTANTS c
ASSUME  $c \in 1..12$ 
VARIABLES hr
Init  $\triangleq hr = c$ 
add_1(p)  $\triangleq p + 1$ 
Inc  $\triangleq hr < 12 \wedge hr' = \text{add\_1}(hr)$ 
Reset  $\triangleq hr = 12 \wedge hr' = 1$ 
Next  $\triangleq \text{Inc} \vee \text{Reset}$ 

```

**Fig. 1.** Module *MyHourClock*

both languages. The ASSUME clause in TLA<sup>+</sup> corresponds to the PROPERTIES clause in B.

To understand the meaning of the other definitions in the module we need some additional information.<sup>2</sup> For our translation we use a configuration file, as TLC also uses, telling us the initial state and the next-state relation of the module. For this example we suppose *Init* to be the initial state predicate and *Next* to be the next-state relation. *Init* indicates that the variable *hr* has the value of the constant *c* in the initial state. *Next* is separated into two independent actions by the disjunction operator. An action is a before-after predicate describing a transition to a next-state with the aid of the prime operator('). A primed variable represents the variable in the next-state. The use of the additional *add\_1* operator may seem artificial here; its purpose is to demonstrate another aspect of our translation.

Figure 2 shows the translated B Machine of the *MyHourClock* example. We use the BECOMES/SUCH/THAT substitution under the INITIALIZATION clause to initialize the variables of the B machine. It assigns a value to the variable such that the predicate in the brackets is satisfied. The TLA<sup>+</sup> actions *Inc* and *Reset* are translated as separate B operations. We represent the prime operator in B by adding a local auxiliary variable for every variable.<sup>3</sup> The auxiliary variables (with suffix “\_n”) are generated in the ANY part of the ANY/WHERE/THEN substitution and get their value in the in the WHERE part. If the predicate in the WHERE part is not satisfiable the operation can not be executed. Finally, the values of the auxiliary variables are assigned to the corresponding global variables in the THEN part.

Operators such as *add\_1* are translated using B definitions. In B, definitions are a kind of macro and help to write frequently used expressions. Furthermore, they can have parameters. Using B definitions avoids to replace an operator call by the definition of the operator. The arithmetic operators are translated with the use of B’s built-in operators. Therefore, they do not appear in the DEFINITIONS clause.

<sup>2</sup> “What those definitions represent [...] lies outside the scope of mathematics and therefore outside the scope of TLA<sup>+</sup>” (see p. 21 of [5]).

<sup>3</sup> For technical reasons we cannot easily use the B’s \$0 construct.

```

MACHINE MyHourClock
DEFINITIONS  add_1(p) == p + 1
CONSTANTS  c
PROPERTIES  c ∈ 1..12
VARIABLES  hr
INVARIANT  hr ∈ ℤ
INITIALISATION  hr :(hr = c)
OPERATIONS
  Inc_Op = ANY hr_n
           WHERE hr < 12 ∧ hr_n = add_1(hr)
           THEN hr := hr_n
           END
  Reset_Op = ANY hr_n
            WHERE hr = 12 ∧ hr_n = 1
            THEN hr := hr_n
            END
END

```

**Fig. 2.** Machine MyHourClock

Finally, to obtain a correct B machine our translation has inferred and added the types of the variables in the INVARIANT clause.

### 3 The Translation from TLA<sup>+</sup> to B

#### 3.1 Type system

The basis of our translation is a mapping of TLA<sup>+</sup> values to B values. Due to the strict type system of B, every B value has to be associated with a type. Below we list the translations of the TLA<sup>+</sup> values and the resulting restrictions:

- Numbers: In B real numbers are not supported. Thus, only integers can be translated. They get the B type  $\mathbb{Z}$ .
- The boolean values TRUE and FALSE are identical in both languages. They get the B Type BOOL.
- The concepts of strings are different in both languages. In TLA<sup>+</sup> a string is a tuple of characters and a single character can be accessed. However, a string in B is atomic and has the base type STRING. For the translation we currently reject strings if they are used as tuples.
- A model value is none of TLA<sup>+</sup>'s own values but one of TLC's. But it is established to use TLA<sup>+</sup> together with TLC so we deliver a suitable translation. TLC allows to assign a model value or a set of model values to a constant in the configuration file. The equivalent of a model value is an element of an enumerated set in B. To make different model values comparable to each other we put them in the same enumerated set named: ENUM<sub>i</sub>. However if two model

values are never compared in the specification, we put them in different sets (such as  $\text{ENUM}_1$  and  $\text{ENUM}_2$ ). The B type of a model value is the name of the enumerated set containing it.

- There is one main difference between sets in  $\text{TLA}^+$  and B. In B all elements of a set must have the same type, i.e., a set has the B type  $\mathbb{P}\tau$  where  $\tau$  is the type of all elements.
- In both languages functions are a mapping from a domain to a range. The B type of a function is  $\mathbb{P}(\tau_1 \times \tau_2)$ , where  $\mathbb{P}\tau_1$  is the type of the domain and  $\mathbb{P}\tau_2$  is the type of the range.
- In  $\text{TLA}^+$  a record is a special case of a function whose domain is a set of strings (the field names) In B records have their own type  $\text{struct}(h_1 : \tau_1, \dots, h_n : \tau_n)$ , where  $h_1, \dots, h_n$  are the names of the fields and  $\tau_1, \dots, \tau_n$  the corresponding field types.
- Likewise, tuples are based on functions in  $\text{TLA}^+$ . The domain is a set of numbers from 1 to  $n$ , where  $n$  is the count of components. We translate tuples as sequences with the type  $\mathbb{P}(\mathbb{Z} \times \tau)$ . Thereby all components of a tuples get the same type  $\tau$ .

In B only values of the same type are comparable to each other and variables as well as constants must have fixed types. To verify these rules a type checking algorithm is required. Moreover we need a type inference algorithm to add missing type declarations to the translated B machine as shown in the example in Sect 2. Type checking and type inference are closely related and can be handled simultaneously.

We use an inference algorithm similar to [9], adapted to the B type system, where we add an extra type  $u$  representing an unspecified type. At the beginning each variable and constant have this type. The algorithm is based on the recursive method  $\text{eval}(e, \epsilon)$ , dealing with a  $\text{TLA}^+$  expression  $e$  and an expected type  $\epsilon$ . Evaluating an expression  $\text{eval}$  is applied recursively to its subexpressions. Moreover  $\text{eval}$  tries to unify the expected type with the type of the expression and returns the resulting type. The expected type of a subexpression is deduced from type informations of the operator calling this subexpression. Type informations of an operator arise from the translation. As an example the  $\text{TLA}^+$  operator  $+$  is translated by the B built-in operator  $+$  and its operands are assumed to be integers. There are polymorphic operators such as  $=$ , which only require that both operands have the same type. In this case the expected type for both operands is  $u$  but the resulting types of both sides have to be unified. Due to unification, the  $\text{eval}$  method is only once applied to each (used) expression of the  $\text{TLA}^+$  module. The algorithm fails either if a unification of two types fails or if a variable or constant still has a variable type  $u$  (or a type constructor containing  $u$  such as  $\mathbb{P}u$ ) at the end of algorithm.

### 3.2 Translation Rules

In this section we present translation rules for concepts which are different in  $\text{TLA}^+$  and B, or even missing in B. Translation rules for concepts which are almost identical can be found in Appendix B.

In contrast to  $TLA^+$ , B distinguishes between boolean values and predicates. The difference is already present at the syntactical level. Logical operators such as  $\wedge$  or  $\vee$  cannot be applied to boolean values. Similarly, variables or constants can not take a predicate as a value. Though, there is a way to convert from a predicate to boolean and vice versa. A predicate can be converted to a boolean value using the *bool* operator. The other way around, we can turn a boolean value into a predicate by comparing it with `TRUE`. The translation of the  $TLA^+$  predicate

$$TRUE = (TRUE \vee FALSE)$$

demonstrates both conversions:

$$TRUE = bool((TRUE = TRUE) \vee (FALSE = TRUE))$$

The `IF/THEN/ELSE` construct can be used in variety of ways in  $TLA^+$ . The two branches can consist of arbitrary expressions with or without primed variables. There is no general way to translate this construct with the `IF/THEN/ELSE` substitution of B. In order to make a translation to B possible, we first have to restrict both branches to the same type. In case the branches are predicates the construct

$$IF P THEN e_1 ELSE e_2$$

can be translated using two implications as

$$(P \Rightarrow e_1) \wedge (\neg(P) \Rightarrow e_2)$$

If  $e_1$  and  $e_2$  are expressions, we cannot use this scheme. Our solution is to create for both branches a lambda function, with respectively  $e_1$  and  $e_2$  as result expression. Moreover we choose `TRUE` as the sole dummy element of the domains. The “trick” is to add the condition  $P$  respectively its negation  $\neg(P)$  to the corresponding function. As a consequence, one of the functions is always empty. As already mentioned, B functions are sets and we can apply  $\cup$  to combine them (the result is still a function here). To get the value of the `IF/THEN/ELSE` construct, we just have to call the function with the value `TRUE` as argument:

$$(\lambda t.(t \in \{TRUE\} \wedge P|e_1) \cup \lambda t.(t \in \{TRUE\} \wedge \neg P|e_2))(TRUE)$$

The translation of the `CASE` construct is based on the same principle. However, every case is treated as single branch and only one case can be true at the same time.

The `LET d  $\hat{=}$  f IN e` construct allows to define a “local” operator  $d$  which can only be used in the expression  $e$ . This operator is treated as an ordinary operator and translated with the aid of a B definition; conserving the scope of the operator. In  $TLA^+$  operators of different `LET/IN` constructs could have the same name. We compensate this case by adding consecutive numbers as suffixes to multiple used names.

In  $TLA^+$  the `CHOOSE` operator is used to choose an arbitrary value of a set. The operator works in a deterministic way and chooses always the same value for

a given set. It is often combined with a recursive function such as determining the sum of a set. In B there is no way to express the general functionality of the CHOOSE operator<sup>4</sup> and recursive functions are still not (well) supported by B tools. Hence, we developed a way to handle frequently used operators which are based on the CHOOSE operator or on recursive functions. The principle is inspired by the way TLC overrides operators by its Java implementation: we create a new TLA<sup>+</sup> standard module (see Figure 3.2) with some useful operators, and during the translation these operators will be overridden by B built-in operators.

```

MODULE TLA2B
EXTENDS Integers, Sequences
MinOfSet(S)  $\triangleq$  CHOOSE p  $\in$  S :  $\forall n \in S : p \leq n$ 
MaxOfSet(S)  $\triangleq$  CHOOSE p  $\in$  S :  $\forall n \in S : p \geq n$ 
SetProduct(p)  $\triangleq$ 
  LET prod[S  $\in$  SUBSET Int]  $\triangleq$ 
    IF S = {} THEN 1
    ELSE LET q  $\triangleq$  CHOOSE pr  $\in$  S : TRUE
         IN q * prod[S \ {q}]
  IN prod[p]
SetSummation(p)  $\triangleq$ 
  LET sum[S  $\in$  SUBSET Int]  $\triangleq$ 
    IF S = {} THEN 0
    ELSE LET q  $\triangleq$  CHOOSE pr  $\in$  S : TRUE
         IN q + sum[S \ {q}]
  IN sum[p]
PermutedSequences(S)  $\triangleq$ 
  LET perms[ss  $\in$  SUBSET S]  $\triangleq$ 
    IF ss = {} THEN {⟨⟩}
    ELSE LET ps  $\triangleq$  [x  $\in$  ss  $\mapsto$ 
                  {Append(sq, x) : sq  $\in$  perms[ss \ {x}]]
         IN UNION {ps[x] : x  $\in$  ss}
  IN perms[S]

```

The concepts of modularization are different in TLA<sup>+</sup> and B. In B a machine is a closed system. Indeed, a machine can be included by another machine but variables can only be modified by its operations. As a result, a single machine of a compound system can be verified individually. A TLA<sup>+</sup> module does not need to satisfy this property. Hence, we translate a compound of TLA<sup>+</sup> modules as a single B machine:

- A module extending another module will be treated as a single module containing declarations and definitions (including local definitions of the extended module) of both modules. Otherwise, there are no further differences in comparison to a translation of a single module.

<sup>4</sup> Even though the operator does appear inside mathematical constructions of [1].

- The statement

$$I \triangleq \text{INSTANCE } M \text{ WITH } v_M \leftarrow v, c_M \leftarrow c$$

allows the specifier to use the definitions of the module  $M$ . Thereby, all variables and constants of  $M$  have to be overridden by variables and constants (or constant expressions) of the module instancing  $M$ . A definition  $d_M$  of Module  $M$  can be accessed via  $I!d_M$ . Also multiple instantiations of the same module are possible. We translate every definition  $d_M$  of  $M$  as an ordinary definition by only renaming it to  $I\_d_M$  and by overriding variables and constants as described in the statement.

In Sect. 2 we translated a  $\text{TLA}^+$  action from Fig. 1 to a B operation in Fig. 2, but we did not exactly define what  $\text{TLA}^+$  actions are and how they are extracted. An action is defined to be “an ordinary mathematical formula, except that it contains primed as well as unprimed variables” (see p. 16 of [5]). Following this definition we could handle the whole next state relation as a single action. However, this is not advisable, amongst others because of poor user feedback for animation, proof and model checking. Consequently we separate actions with the aid of the disjunction operator. If a disjunction of two actions occurs in a subdefinition of the next state relation, we also will treat them as separate actions unless the subdefinition has no parameter. Parameters indicate that a subdefinition can be used in different variations and multiple times; the translation should not dissolve this structure of a module. In this case the subdefinition is translated with a B definition.

A special translation is possible if an action contains an existential quantifier:

$$act \triangleq \exists x \in S : P(x)$$

The bounded variables of the quantification are handled as parameters of the resulting B operation:

$$act\_op(x) = \text{ANY} \dots \text{WHERE } x \in S \wedge P(x) \text{ THEN} \dots \text{END}$$

The advantage is that a user can choose a possible value for the parameter  $x$  during the animation process (values for  $x$  satisfying  $P(x)$  are generated by PROB).

## 4 Implementation and Experiments

The translator is implemented in Java and is called TLA2B. The frontend of TLA2B is based on SANY (cf., Chapter 12 of [5]) for parsing the module and performing a semantic analysis. Likewise, SANY serves as the frontend of the modelchecker TLC. Moreover, we reuse the configuration file parser of TLC. But the semantic analysis of the configuration file ist different: TLC requires a value for every constant of the corresponding module. In our case, a constant only has to be given a value if the type of the constant cannot be inferred from the module. Otherwise, values of constants can be chosen at a later point in time (PROB infers values for a



constant satisfying possible restrictions of the ASSUME clause). TLA2B can handle the clauses SPECIFICATION (temporal description of the specification), INVARIANT (an invariant holding in every state) and overriding of constants and definitions beside the already mentioned INIT (initial state) and NEXT (next state relation). Before inferring and checking types, we conduct a further analysis phase discarding the unused definitions of a module. As an example, temporal definitions are excluded from the translation. The remaining part of the TLA2B consists of implementations of the algorithms described in Sections 2 and 3. Finally, TLA2B creates a B machine file (.mch) containing the translated B machine.

TLA2B has been integrated into PROB as of version 1.3.5: opening a TLA<sup>+</sup> module ProB invokes TLA2B to translate the module. As can be seen in Fig. 3, the TLA<sup>+</sup> module is displayed in the editor while PROB runs the translated B machine in the background. The editor offers syntax highlighting and gives an easy way to modify the module.

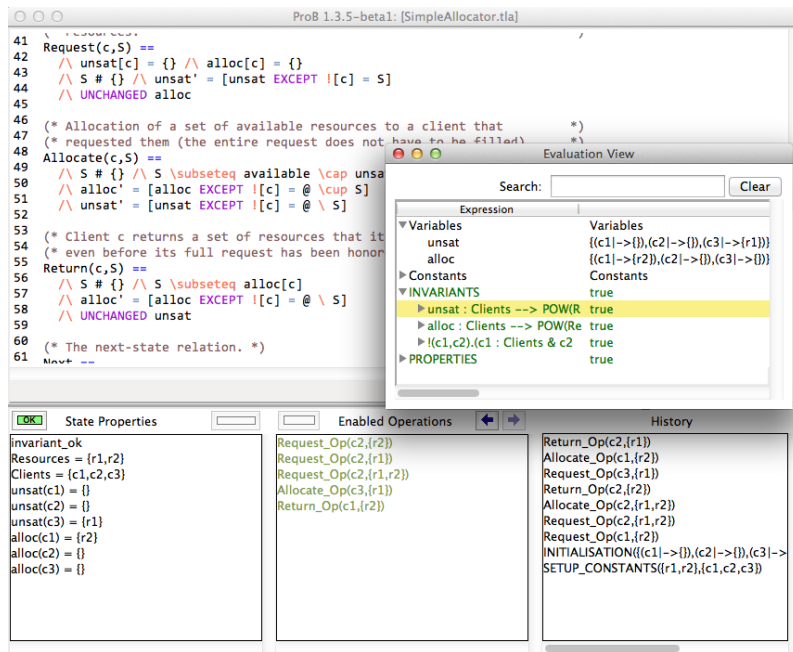


Fig. 3. PROB animator for the SimpleAllocator specification

The following examples show some fields of application of TLA2B in combination with PROB. It is not our intention to present a complete comparison between PROB and TLC. The experiments were all run on a system with Intel Core2 Duo 2 GHz processor, running Windows Vista 32 Bit, TLC2 2.03 and PROB 1.3.5-beta1.

Note that both PROB and TLC support symmetry, but in different ways. In TLC symmetries are provided by the user (e.g., in a configuration file) and do not seem to be checked, PROB identifies symmetries over given sets automatically.

**SimpleAllocator** As the first example we use the resource allocator case study from [8]. The purpose of the system is to manage a set of resources that are shared among a number of client processes. The first abstract specification of the system is the SimpleAllocator. TLA2B translates the module without the need for any modification (the TLA<sup>+</sup> module and the translated B machine are shown in Appendix A). Clients and resources are specified as sets of model values and allow TLC as well as PROB to use symmetry. Table 2 summarises the runtimes of model checking produced by TLC and PROB. Without symmetry TLC is superior to PROB, but for larger set sizes PROB’s symmetry outperforms TLC. It seems that TLC’s symmetry reduction cannot deal well with larger base set sizes and a lot of symmetrical states exist (incidentally, a situation where symmetry reduction could be particularly useful). This is actually to be expected, given the description of the symmetry reduction algorithm in [5]: when a state is added TLC checks for every permutation of it whether it already exists. This is expensive when there are many such permutations.

**Login** To specifically test this aspect of symmetry reduction, we have written the TLA specification Login which simply allows users to login and logout and deadlocks if all users have logged in (in Appendix C.4). Here, for 9 Users, TLC without symmetry reduction takes 1 second to find the deadlock, but did not terminate within 105 minutes with symmetry enabled. PROB takes 0.73 seconds without symmetry, and 0.04 using hash symmetry reduction [7]. For 21 users, TLC requires 141 seconds to find the deadlock without symmetry, and with symmetry an error message is generated.<sup>5</sup> PROB with hash symmetry takes 0.29 seconds to find the deadlock for 21 users. The constraint-based deadlock checking algorithm [4] finds a deadlock in less than 0.01 seconds for 21 users.

Clients	Resources	TLC (no symmetry)	TLC (symmetry)	ProB (no symmetry)	ProB (symmetry)
3	2	<1	<1	<2	<1
4	3	28	2	678	8
5	3	450	29	-	28
6	3	>4200	573	-	90

**Table 1.** SimpleAllocator: Runtimes of Model Checking (times in seconds)

**SchedulingAllocator** This is an advanced version of the SimpleAllocator from [8]. However, this time a small modification is required to be able to validate the specification using our tool. Indeed, the SchedulingAllocator contains the definition  $PERMSEQS(S)$  that is based on a recursive function and computes the set of

<sup>5</sup> “Attempted to construct a set with too many elements (>1000000)”. This error message already appears with 15 Users.

permutation sequences of the set  $S$ . To translate this to the  $B$  built-in operator, we have to override `PERMSEQS` with the `PERMUTEDSEQUENCES` definition provided by our `TLA2B` module. For this, we simply have to create a new module `MCSCHEDULINGALLOCATOR` (shown in Fig.4) and then add the override statement `PermSeqs <- PermutedSequences` to the configuration file. The results of model checking are comparable to the `SimpleAllocator` specification, and can be found in Table 2.

```

----- MODULE McSchedulingAllocator -----
EXTENDS SchedulingAllocator, TLA2B

```

**Fig. 4.** Module `McSchedulingAllocator`

Clients	Resources	TLC (symmetry)	ProB (symmetry)
3	2	1	2
4	3	70	165
5	3	>3600	1579

**Table 2.** `McSchedulingAllocator`: Runtimes of Model Checking (times in seconds)

**Producer-Consumer** Another example is the specification of a multi-threaded program by Charpentier taken from <http://www.cs.unh.edu/%7Echarpov/Teaching/TLA/>. The specification describes a system of threads working on a buffer. In case of a critical ratio between consumer and producer threads, the system can deadlock. After translation `PROB` reproduces the various deadlocks by model checking. For example, for 11 producers and 10 consumers a deadlock can be reached after 431 steps. Using the `AtelierB` provers we have also managed to prove the invariant of that model, i.e., that the buffer capacity is never exceeded and that the `waitSet` only contains valid participants. This required 8 interactive proofs and 5 automatic ones.

**Constraint Solving: GraphIso and N-Queens** One of the distinguishing features of `PROB` is its ability to solve complicated high-level constraints. For example, to find an isomorphism between two graphs (of out-degree exactly one) using the `TLA+` specification in Appendix C.3, `PROB` requires less than a second to find all solutions, while `TLC` requires over two hours to find the first solution.

As another example, consider the well-known `N-Queens` puzzle. We have experimented with two encodings of the puzzle (cf., Appendices C.1 and C.2): one<sup>6</sup> where we use the model checker to search for all valid placements of  $N$  queens on an  $N \times N$  chessboard and a more declarative encoding where we directly write a predicate describing all valid solutions (i.e., all solutions are generated in single set-builder rather than through an iterative algorithm). As can be seen in Table 3, the model checking approach can only deal with very small values of  $N$ . In contrast, `PROB` can

<sup>6</sup> The specification was written by S. Merz and is included in the `TLA+ Tools Distribution`.

handle values of N up to 13 for the declarative version of N-Queens. Furthermore, when one is interested in only one solution, PROB can, e.g., find a solution for N=50 in less than a second. (Restricting to single solutions does not make much of a performance difference for TLC however.)

N	Solutions	N-Queens (imperative)		N-Queens (declarative)	
		TLC	PROB	TLC'	PROB
4	2	1	<2	<1	<1
5	10	>3600	>3600	<1	<1
6	4	-	-	1	<1
7	40	-	-	16	<1
8	92	-	-	375	<1
9	352	-	-	2970	<1
10	724	-	-	-	<1
11	2,680	-	-	-	<1
12	14,200	-	-	-	9
13	73,712	-	-	-	41

**Table 3.** Finding *all* solutions for N-Queens (times in seconds)

We have also successfully animated several other existing models from the literature, and we are continuously extending the scope of our translation and the application of PROB to TLA<sup>+</sup> specifications. Otherwise, it should be emphasised that several specifications are rejected by TLA2B due to type conflicts or unsupported concepts such as real numbers. In summary, we are already able to animate a wide variety of TLA<sup>+</sup> specifications from the literature. The PROB constraint solving capabilities open up the way to animate and validate new kinds of specifications, which are outside the reach of TLC. TLC on the other hand is extremely valuable when it comes to explicit state model checking for large state spaces. However, PROB's symmetry reduction techniques seem to scale better than TLC's.

## 5 Discussion and Conclusion

The paper by Mokhtari and Merz [10] presents an animator and model checker for an executable subset of TLA<sup>+</sup>. The article clearly outlines the needs for an animator for TLA<sup>+</sup>; unfortunately the tool seems no longer to be available.

Mosbahi et al. [11] describe an approach of a translation from B to TLA<sup>+</sup>. In contrast to our translation they have to deal with concepts which are missing in TLA<sup>+</sup> such as partial functions. Moreover their main intention is to let TLC verify liveness properties on the translated TLA<sup>+</sup> specification, to overcome the restriction of the B-Method to invariance properties. Otherwise, [12] presents a LTL model checker, implemented inside PROB, that can verify liveness properties. So

<sup>7</sup> To display the trace leading to a deadlock TLC calculates each state of the trace again: this leads to a doubling of runtimes for this version of N-Queens in the table.

far, this model checker does not support fairness conditions, but an extension would give us the possibility to enlarge TLA2B to support the temporal part of TLA<sup>+</sup>.

In terms of features, we also plan to provide the graphical visualization features of PROB (available for B,Z and Event-B models) for TLA<sup>+</sup> as well. More work on translating various constructs effectively to B, such as the choose operator or recursive functions, is planned. Another important avenue of further work lies in improving our translation to B. In particular, we aim to generate various B style substitutions such as assignments or if-then-else constructs, rather than generic ANY substitutions. This makes the B translation more readable, but would also lead to noticeable performance improvements with PROB. E.g., in our experiments, this would lead, to a further 20 % runtime improvement for the SimpleAllocator example and up to a factor 2 for other examples.

We would also like to better exploit the symmetry reduction provided by PROB. While the SimpleAllocator, SchedulingAllocator and Login example worked well, the symmetry in the Producer-Consumer example could only be exploited by manually tweaking the B translation. We would like to automate this as much as possible, as it can lead to a considerable performance boost (after tweaking PROB with symmetry requires about a minute to find the 413 step counter example for the Producer-Consumer example; TLC requires more than three and a half hours to find a deadlock<sup>8</sup>). We are also interested in validation of our translation, maybe by exporting the state space computed by PROB to TLC and use TLC to check that it conforms to the original specification.

In conclusion, we have presented a translation from TLA<sup>+</sup> to B, which makes use of a type inference algorithm and effectively translates a large subset of TLA<sup>+</sup> to B. The complicated aspects of the translation are linked to the different modularization concepts, as well as to various operators which are not obvious to translate to B. The translation also identifies operations and parameters within the TLA<sup>+</sup> specification formula, in order to make the translation more readable as well as to enable effective application of B tools. In particular, by integrating our translation into the PROB validation tool, we obtain a new tool for TLA<sup>+</sup> specifications which is complementary to TLC, providing convenient animation, constraint solving and improved symmetry reduction. As our experiments show, TLC remains more effective for brute force explicit state model checking, at least for those specifications which do not require solving complicated constraints. As such it is very useful that both these tools can be applied to TLA<sup>+</sup> specifications. The translation itself is also human readable, and we hope that the paper also provides a bridge between the TLA<sup>+</sup> and B communities.

**Acknowledgements** We are grateful to Daniel Plagge for various discussions and helpful comments that helped in developing the translator. We also would like to thank Leslie Lamport and Stephan Merz for very useful feedback concerning TLA<sup>+</sup> and TLC and for giving us access to various specifications.

---

<sup>8</sup> When trying to use symmetry, the same error message occurs as in the Login example.

## References

1. J.-R. Abrial. *The B-Book*. Cambridge University Press, 1996.
2. K. Chaudhuri, D. Doligez, L. Lamport, and S. Merz. The TLA<sup>+</sup> proof system: Building a heterogeneous verification platform. In A. Cavalcanti, D. Déharbe, M.-C. Gaudel, and J. Woodcock, editors, *Proceedings ICTAC 2010*, LNCS 6255, page 44, 2010.
3. K. Chaudhuri, D. Doligez, L. Lamport, and S. Merz. Verifying safety properties with the TLA<sup>+</sup> proof system. In J. Giesl and R. Hähnle, editors, *Proceedings IJCAR 2010*, LNCS 6173, pages 142–148, 2010.
4. S. Hallerstede and M. Leuschel. Constraint-based deadlock checking of high-level specifications. *TPLP*, 11(4–5):767–782, 2011.
5. L. Lamport. *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002.
6. 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.
7. M. Leuschel and T. Massart. Efficient approximate verification of B via symmetry markers. *Annals of Mathematics and Artificial Intelligence*, 59(1):81–106, 2010.
8. S. Merz. TLA<sup>+</sup> Case Study: A Resource Allocator. Interne A04-R-101 — merz04a, 2004. Rapport interne.
9. S. Merz and H. Vanzetto. Automatic verification of tla<sup>+</sup> proof obligations with smt solvers. In N. Bjørner and A. Voronkov, editors, *Proceedings LPAR-18*, LNCS, Mérida, Venezuela, 2012. Springer.
10. Y. Mokhtari and S. Merz. Animating TLA specifications. In H. Ganzinger, D. A. McAllester, and A. Voronkov, editors, *Proceedings LPAR’99*, LNCS 1705, pages 92–110, 1999.
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. D. Plagge and M. Leuschel. Seven at a stroke: LTL model checking for high-level specifications in B, Z, CSP, and more. *STTT*, 11:9–21, 2010.
13. F. Steria, Aix-en-Provence. *Atelier B, User and Reference Manuals*, 2009. Available at <http://www.atelierb.eu/>.
14. Y. Yu, P. Manolios, and L. Lamport. Model checking TLA<sup>+</sup> specifications. In L. Pierre and T. Kropf, editors, *Proceedings CHARME’99*, LNCS 1703, pages 54–66. Springer-Verlag, 1999.

## A A more complicated translation: SimpleAllocator

---

MODULE *SimpleAllocator*

```
EXTENDS FiniteSets, TLC
CONSTANTS Clients, Resources
ASSUME IsFiniteSet(Resources)
VARIABLES unsat, alloc
TypeInvariant  $\triangleq$   $\wedge$  unsat  $\in$  [Clients  $\rightarrow$  SUBSET Resources]
 $\wedge$  alloc  $\in$  [Clients  $\rightarrow$  SUBSET Resources]
available  $\triangleq$  Resources \ (UNION {alloc[c] : c  $\in$  Clients})
Init  $\triangleq$   $\wedge$  unsat = [c  $\in$  Clients  $\mapsto$  {}]  $\wedge$  alloc = [c  $\in$  Clients  $\mapsto$  {}]
```

$$\begin{aligned}
\text{Request}(c, S) &\triangleq \wedge \text{unsat}[c] = \{\} \wedge \text{alloc}[c] = \{\} \\
&\wedge S \neq \{\} \wedge \text{unsat}' = [\text{unsat} \text{ EXCEPT } ![c] = S] \wedge \text{UNCHANGED } \text{alloc} \\
\text{Allocate}(c, S) &\triangleq \wedge S \neq \{\} \wedge S \subseteq \text{available} \cap \text{unsat}[c] \\
&\wedge \text{alloc}' = [\text{alloc} \text{ EXCEPT } ![c] = @ \cup S] \wedge \text{unsat}' = [\text{unsat} \text{ EXCEPT } ![c] = @ \setminus S] \\
\text{Return}(c, S) &\triangleq \wedge S \neq \{\} \wedge S \subseteq \text{alloc}[c] \\
&\wedge \text{alloc}' = [\text{alloc} \text{ EXCEPT } ![c] = @ \setminus S] \wedge \text{UNCHANGED } \text{unsat} \\
\text{Next} &\triangleq \exists c \in \text{Clients}, S \in \text{SUBSET } \text{Resources} : \\
&\quad \text{Request}(c, S) \vee \text{Allocate}(c, S) \vee \text{Return}(c, S) \\
\text{vars} &\triangleq \langle \text{unsat}, \text{alloc} \rangle \\
\text{SimpleAllocator} &\triangleq \wedge \text{Init} \wedge \square [\text{Next}]_{\text{vars}} \\
&\wedge \forall c \in \text{Clients} : \text{WF}_{\text{vars}}(\text{Return}(c, \text{alloc}[c])) \\
&\wedge \forall c \in \text{Clients} : \text{SF}_{\text{vars}}(\exists S \in \text{SUBSET } \text{Resources} : \text{Allocate}(c, S)) \\
\text{ResourceMutex} &\triangleq \forall c1, c2 \in \text{Clients} : c1 \neq c2 \Rightarrow \text{alloc}[c1] \cap \text{alloc}[c2] = \{\} \\
\text{Symmetry} &\triangleq \text{Permutations}(\text{Clients}) \cup \text{Permutations}(\text{Resources})
\end{aligned}$$


---

```

MACHINE SimpleAllocator
SETS ENUM1 = {r1, r2}; ENUM2 = {c1, c2, c3}
CONSTANTS Clients, Resources
PROPERTIES Clients = ENUM2 ∧ Resources = ENUM1
  ∧ ∃seq-.(seq- ∈ seq(Resources) ∧ ∀s.(s ∈ Resources ⇒
    ∃n.(n ∈ 1 .. size(seq-) ∧ seq-(n) = s)))
DEFINITIONS
  TypeInvariant == unsat ∈ Clients ⇒ ℙ(Resources)
    ∧ alloc ∈ Clients ⇒ ℙ(Resources);
  available == Resources - union(t-|∃c.(c ∈ Clients ∧ t- = alloc(c)));
  Init == unsat = λc.(c ∈ Clients| {}) ∧ alloc = λc.(c ∈ Clients| {});
  Request(c,S) == unsat(c) = {} ∧ alloc(c) = {}
    ∧ (S ≠ {} ∧ unsat-n = unsat ◁ {c ↦ S});
  Allocate(c,S) == S ≠ {} ∧ S ⊆ available ∩ unsat(c)
    ∧ alloc-n = alloc ◁ {c ↦ (alloc(c) ∪ S)}
    ∧ unsat-n = unsat ◁ {c ↦ (unsat(c) - S)};
  Return(c,S) == S ≠ {} ∧ S ⊆ alloc(c) ∧ alloc-n = alloc ◁ {c ↦ (alloc(c) - S)};
  ResourceMutex == ∀c1,c2.(c1 ∈ Clients ∧ c2 ∈ Clients ⇒
    (c1 ≠ c2 ⇒ alloc(c1) ∧ alloc(c2) = {}));
VARIABLES unsat, alloc
INVARIANT unsat ∈ ℙ(ENUM2 × ℙ(ENUM1))
  ∧ alloc ∈ ℙ(ENUM2 × ℙ(ENUM1)) ∧ TypeInvariant ∧ ResourceMutex
INITIALISATION unsat, alloc ∈ (Init)
OPERATIONS
  Request_Op(c, S) = ANY unsat-n
    WHERE c ∈ Clients ∧ S ∈ ℙ(Resources) ∧ Request(c, S)
    THEN unsat := unsat-n END;
  Allocate_Op(c, S) = ANY unsat-n, alloc-n
    WHERE c ∈ Clients ∧ S ∈ ℙ(Resources) ∧ Allocate(c, S)
    THEN unsat, alloc := unsat-n, alloc-n END;
  Return_Op(c, S) = ANY alloc-n
    WHERE c ∈ Clients ∧ S ∈ ℙ(Resources) ∧ Return(c, S)
    THEN alloc := alloc-n END;
END

```



## B Translations Rules (for referees)

### B.1 Logic

TLA <sup>+</sup>	B-Method
<i>TRUE</i>	<i>TRUE</i>
<i>FALSE</i>	<i>FALSE</i>
<i>BOOLEAN</i>	<i>BOOL</i>
$P \wedge Q$	$P \wedge Q$
$P \vee Q$	$P \vee Q$
$P \Rightarrow Q$	$P \Rightarrow Q$
$P \equiv Q$	$P \Leftrightarrow Q$
$\neg P$	$\neg P$
$\forall x \in S : P$	$\forall x.(x \in S \Rightarrow P)$
$\exists x \in S : P$	$\exists x.(x \in S \wedge P)$

### B.2 Sets

TLA <sup>+</sup>	B-Method
$\{x_1, \dots, x_n\}$	$\{x_1, \dots, x_n\}$
$\{x \in S : P\}$	$\{x   x \in S \wedge P\}$
$\{e : x \in S\}$	$\{t   \exists x.(x \in S \wedge t = e)\}$
$M_1 \otimes M_2$ for $\otimes \in \{=, \neq, \cup, \cap, \subseteq\}$	$M_1 \otimes M_2$
$e \in M$	$e \in M$
$e \notin M$	$e \notin M$
$e_1 \setminus e_2$	$e_1 - e_2$
SUBSET $S$	$\mathbb{P}(S)$
UNION $S$	$\text{union}(S)$
Cardinality( $S$ )	$\text{card}(S)$

### B.3 Functions

TLA <sup>+</sup>	B-Method
$[x \in S \mapsto e]$	$\lambda x.(x \in S   e)$
$f[e]$	$f(e)$
DOMAIN $f$	$\text{dom}(f)$
$[S \rightarrow T]$	$S \rightarrow T$
$[f \text{ EXCEPT } ![c_1] = e_1, \dots, ![c_n] = e_n]$	$f \triangleleft \{c_1 \mapsto e_1, \dots, c_n \mapsto e_n\}$

### B.4 Records

TLA <sup>+</sup>	B-Method
$r.h$	$r'h$
$[h_1 \mapsto e_1, \dots, h_n \mapsto e_n]$	$\text{rec}(h_1 : e_1, \dots, h_n : e_n)$
$[h_1 : S_1, \dots, h_n : S_n]$	$\text{struct}(h_1 : S_1, \dots, h_1 : S_n)$
$[r \text{ EXCEPT } !.b = e]$ for $r \in [a : S_1, b : S_2, c : S_3]$	$\text{rec}(a : r'a, b : e, c : r'c)$

## B.5 Integers

TLA <sup>+</sup>	B-Method
$e_1 \otimes e_2$	$e_1 \otimes e_2$
for $\otimes \in \{+, -, *, \div, \leq, \geq, <, >\}$	
$b^e$	$b^e$
$-e$	$-e$
$e_1..e_2$	$e_1..e_2$
<i>Nat</i>	NATURAL
<i>Int</i>	INTEGER

## B.6 Tuple

TLA <sup>+</sup>	B-Method
$\langle e_1, e_2, e_3 \rangle$	$[e_1, e_2, e_3]$
<i>Seq</i> ( <i>S</i> )	<i>seq</i> ( <i>S</i> )
<i>Len</i> ( <i>s</i> )	<i>size</i> ( <i>s</i> )
$s \circ t$	$s \hat{\ } t$
<i>Append</i> ( <i>s</i> , <i>e</i> )	$s \leftarrow e$
<i>Head</i> ( <i>s</i> )	<i>first</i> ( <i>s</i> )
<i>Tail</i> ( <i>s</i> )	<i>tail</i> ( <i>s</i> )
<i>SubSeq</i> ( <i>s</i> , <i>m</i> , <i>n</i> )	$(s \uparrow n) \downarrow (m - 1)$

## C Examples (for referees)

### C.1 N-Queens (imperative for model checking)

---

MODULE *Queens*

---

EXTENDS *Naturals*, *Sequences*

CONSTANT *N*

ASSUME  $N \in \text{Nat} \setminus \{0\}$

$\text{Attacks}(\text{queens}, i, j) \triangleq$   
 $\vee \text{queens}[i] = \text{queens}[j]$   
 $\vee \text{queens}[i] - \text{queens}[j] = i - j$   
 $\vee \text{queens}[j] - \text{queens}[i] = i - j$

$\text{IsSolution}(\text{queens}) \triangleq$   
 $\forall i \in 1.. \text{Len}(\text{queens}) - 1 : \forall j \in i + 1.. \text{Len}(\text{queens}) :$   
 $\neg \text{Attacks}(\text{queens}, i, j)$

$\text{Solutions} \triangleq \{\text{queens} \in [1..N \rightarrow 1..N] : \text{IsSolution}(\text{queens})\}$

VARIABLES *todo*, *sols*

$\text{Init} \triangleq \wedge \text{todo} = \{\}$   
 $\wedge \text{sols} = \{\}$

$\text{PlaceQueen} \triangleq \exists \text{queens} \in \text{todo} :$   
 LET  $\text{nextQ} \triangleq \text{Len}(\text{queens}) + 1$

$$\begin{aligned}
\text{cols} &\triangleq \{c \in 1..N : \neg \exists i \in 1..N : \text{Attacks}(\text{Append}(\text{queens}, c), i, \text{nextQ})\} \\
\text{exts} &\triangleq \{\text{Append}(\text{queens}, c) : c \in \text{cols}\} \\
\text{IN } &\text{IF } \text{nextQ} = N \\
&\quad \text{THEN } \wedge \text{todo}' = \text{todo} \setminus \{\text{queens}\} \\
&\quad \quad \wedge \text{sols}' = \text{sols} \cup \text{exts} \\
&\quad \text{ELSE } \wedge \text{todo}' = (\text{todo} \setminus \{\text{queens}\}) \cup \text{exts} \\
&\quad \quad \wedge \text{sols}' = \text{sols} \\
\text{vars} &\triangleq \langle \text{todo}, \text{sols} \rangle \\
\text{Spec} &\triangleq \text{Init} \wedge \square[\text{PlaceQueen}]_{\text{vars}} \wedge \text{WF}_{\text{vars}}(\text{PlaceQueen}) \\
\text{TypeInvariant} &\triangleq \\
&\quad \wedge \text{todo} \in \text{SUBSET } \text{Seq}(1..N) \wedge \forall s \in \text{todo} : \text{Len}(s) < N \\
&\quad \wedge \text{sols} \in \text{SUBSET } \text{Seq}(1..N) \wedge \forall s \in \text{sols} : \text{Len}(s) = N \\
\text{Invariant} &\triangleq \\
&\quad \wedge \text{sols} \subseteq \text{Solutions} \\
&\quad \wedge \text{todo} = \{\} \Rightarrow \text{Solutions} \subseteq \text{sols} \\
\text{Termination} &\triangleq \diamond(\text{todo} = \{\}) \\
\text{NoSolutions} &\triangleq \text{sols} = \{\}
\end{aligned}$$

## C.2 N-Queens (declarative encoding)

$$\begin{aligned}
&\text{MODULE } \text{QueensV2} \\
&\text{EXTENDS } \text{Naturals}, \text{FiniteSets} \\
&\text{CONSTANTS } N \\
&\text{VARIABLE } \text{solutions} \\
&\text{Init} \triangleq \text{solutions} = \{\} \\
&\text{Solve} \triangleq \text{solutions} = \{\} \wedge \text{solutions}' = \{\text{sol} \in [1..N \rightarrow 1..N] : \\
&\quad (\forall i \in 1..N : \forall j \in 2..N : i < j \Rightarrow \\
&\quad \quad \wedge \text{sol}[i] \neq \text{sol}[j] \\
&\quad \quad \wedge \text{sol}[i] + i - j \neq \text{sol}[j] \\
&\quad \quad \wedge \text{sol}[i] - i + j \neq \text{sol}[j])\} \\
&\text{Spec} \triangleq \text{Init} \wedge \square[\text{Solve}]_{\langle \text{solutions} \rangle}
\end{aligned}$$

## C.3 Simple Graph Isomorphism

$$\begin{aligned}
&\text{MODULE } \text{GraphIso} \\
&\text{EXTENDS } \text{Naturals} \\
&\text{VARIABLE } p, \text{solved} \\
&n \triangleq 9 \\
&g1 \triangleq [[i \in 1..n \mapsto i] \text{ EXCEPT } ![1] = 3, ![2] = 3,
\end{aligned}$$

$$\begin{aligned}
& ![3] = 6, ![4] = 6, ![5] = 6, ![8] = 9, ![9] = 8 \\
g2 & \triangleq [[i \in 1..n \mapsto i] \text{ EXCEPT } ![2] = 5, ![3] = 5, \\
& \quad ![4] = 5, ![6] = 4, ![7] = 4, ![1] = 9, ![9] = 1] \\
Init & \triangleq \wedge p = [i \in 1..n \mapsto 0] \\
& \quad \wedge solved = 0 \\
Solve & \triangleq \wedge solved = 0 \\
& \quad \wedge p' \in [1..n \rightarrow 1..n] \\
& \quad \wedge \forall i \in 1..n : (\exists j \in 1..n : p'[j] = i) \\
& \quad \wedge \forall i \in 1..n : (p'[g1[i]] = g2[p'[i]]) \\
& \quad \wedge solved' = 1 \\
Spec & \triangleq Init \wedge \square[Solve]_{(p, solved)}
\end{aligned}$$


---

## C.4 Login

---

MODULE *Login\_Deadlock*

---

EXTENDS *TLC*  
CONSTANT *Users*  
VARIABLES *logged\_in*  
*Init*  $\triangleq$  *logged\_in* = {}  
*Login*  $\triangleq$   $\exists u \in Users \setminus logged\_in : logged\_in' = logged\_in \cup \{u\}$   
*Logout*  $\triangleq$   $\exists u \in logged\_in :$   
 $\quad \wedge logged\_in \neq Users$       this causes a deadlock when *logged\_in* = *Users*  
 $\quad \wedge logged\_in' = logged\_in \setminus \{logged\_in\}$   
*Next*  $\triangleq$  *Login*  $\vee$  *Logout*  
*Spec*  $\triangleq$  *Init*  $\wedge$   $\square[Next]_{logged\_in}$   
*Symmetry*  $\triangleq$  *Permutations*(*User*)

---