

Validating B,Z and TLA⁺ Using PROB and Kodkod

Daniel Plagge and Michael Leuschel

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

Abstract. We present the integration of the Kodkod high-level interface to SAT-solvers into the kernel of PROB. As such, predicates from B, Event-B, Z and TLA⁺ can be solved using a mixture of SAT-solving and PROB's own constraint-solving capabilities developed using constraint logic programming: the first-order parts which can be dealt with by Kodkod and the remaining parts solved by the existing PROB kernel. We also present an empirical evaluation and analyze the respective merits of SAT-solving and classical constraint solving. We also compare to using SMT solvers via recently available translators for Event-B.

Keywords: B-Method, Z, TLA, Tool Support, SAT, SMT, Constraints.

1 Introduction and Motivation

TLA⁺ [10], B [1] and Z are all state-based formal methods rooted in predicate logic, combined with arithmetic and set theory. The animator and model checker PROB [12] can be applied to all of these formalisms and is being used by several companies, mainly in the railway sector for safety critical control software [13,14]. At the heart of PROB is a kernel dealing with the basic data types of these formalisms, i.e., integers, sets, relations, functions and sequences. An important feature of PROB is its ability to solve constraints; indeed constraints can arise in many situations when manipulating a formal specification: the tool needs to find values of constants which satisfy the stipulated properties, the tool needs to find acceptable initial values of a model, the tool has to determine whether an event or operation can be applied (i.e., is there a solution for the parameters which make the guard true) or whether a quantified expression is true or not. Other tasks involve more explicit constraint solving, e.g., finding counterexamples to invariant preservation or deadlock freedom proof obligations [7]. While PROB is good at dealing with large data structures and also at solving certain kinds of complicated constraints [7], it can fare badly on certain other constraints,

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

in particular relating to relational composition and transitive closure. (We will illustrate this later in the paper.)

Another state-based formalism is Alloy [8] with its associated tool which uses the Kodkod [21] library to translate its relational logic predicates into propositional formulas which can be fed into SAT solvers. Alloy can deal very well with complicated constraints, in particular those involving relational composition and transitive closure. Compared to B, Z and TLA⁺, the Alloy language and the Kodkod library only allow first-order predicates, e.g., they do not allow relations over sets or sets of sets.

The goal of this work is to integrate Kodkod into PROB, providing an alternative way of solving B, Z and TLA⁺ constraints. Note that we made sure that the animation and model checking engine as well as the user interface of PROB are agnostic as to how the underlying constraints are solved. Based on this integration we also conduct a thorough empirical evaluation of the performance of Kodkod compared to solving constraints with the existing constraint logic programming approach of PROB. As we will see later in the paper, this empirical evaluation provides some interesting insights. Our approach also ensures that the whole of B is covered, by delegating the untranslatable higher-order predicates to the existing PROB kernel.

2 B, Z, TLA⁺ and Kodkod in Comparison

PROB can support Z and TLA⁺ by translating those formalisms to B, because these formalisms have a common mathematical foundation. In the case of TLA⁺ a readable B machine is actually generated, whereas a Z specification is translated to PROB's internal representation because some Z constructs did not have a direct counterpart in B's syntax. In the next sections we refer only to B, but because all three notations share the same representation in PROB, all presented techniques can be applied likewise to the two other specification languages.

If we specify a problem in B, we basically have a number of variables, each of a certain type and a predicate. The challenge for PROB is then to find values for the variables that fulfil the predicate. For simplicity, we ignore B's other concepts like machines, refinement, etc.

Kodkod provides a similar view on a problem. We have to specify a number of relations (these correspond to our variables in B) and a formula (which corresponds to a predicate in B) and Kodkod tries to find solutions for the relations.

From this point of view, the main difference between B and Kodkod is the type system: Instead of having some basic types and operations like power set and Cartesian product to combine these, Kodkod has the concept of a universe consisting of atoms. To use Kodkod, we must define a list of atoms and for each relation we must specify a *bound* that determines a range of atoms that can be in the relation.

The bound mechanism can also be used to assign an exact value to a relation. This is later useful when we have already computed some values by PROB.

3 Architecture

3.1 Overview

We use a small example to illustrate the basic mechanism how Kodkod and SAT solving is integrated into PROB’s process to find a model for a problem. Details about the individual components are presented below after this overview.

Our small problem is taken from the “dragon book” [2] and formalised in B. The aim is to find loops in a control flow graph of a program (see Figure 1).

We model the basic blocks as an enumerated set *Blocks* with the elements $b_1, b_2, b_3, b_4, b_5, b_6, b_{entry}, b_{exit}$. The successor relation is represented by a variable *succs*, the set of the nodes that constitute the loop by *L* and the entry point of the loop by *lentry*. The problem is described by the B predicate:

$$\begin{aligned}
 & succs = \{b_{entry} \mapsto b_1, b_1 \mapsto b_2, b_2 \mapsto b_3, b_3 \mapsto b_3, \\
 & \quad b_3 \mapsto b_4, b_4 \mapsto b_2, b_4 \mapsto b_5, \\
 & \quad b_5 \mapsto b_6, b_6 \mapsto b_6, b_6 \mapsto b_{exit}\} \\
 & \wedge lentry \in L \\
 & \wedge succs^{-1}[L \setminus \{lentry\}] \subseteq L \\
 & \wedge \forall l.(l \in L \Rightarrow lentry \in (L \triangleleft succs \triangleright L)^+[\{l\}])
 \end{aligned}$$

In total, there are seven different solutions to this problem, for instance $L = \{b_2, b_3, b_4\}$ with $lentry = b_2$.

After parsing and type checking the predicate, we start a static analysis (the box “Analysis” in Fig. 2) to determine the integer intervals of all integer expressions. In our simple case, the analysis is not necessary. In Section 3.4 we describe how this analysis works and under which circumstances it is needed.

In the next phase, we try to translate the formula from B to Kodkod (“Translation” in Fig. 2). First we have a look at the used variables and their types: *succs* is of type $\mathbb{P}(Blocks \times Blocks)$, *L* of type $\mathbb{P}(Blocks)$ and *lentry* of type *Blocks*. *Blocks* is here the only basic type that is used. Thus we have to reserve 8 atoms in the Kodkod universe to represent this type; each atom in the universe corresponds directly to a block b_i . The variables can be represented by binary (*succs*) and unary (*L* and *lentry*) relations, where we have to keep in mind that the relation for *lentry* must contain exactly one element. The B predicate can be completely translated to a Kodkod problem. In Section 3.2 we will describe the translation in more detail. It can be useful to keep a part of the formula untranslated: since the part $succs = \{...\}$ is very easy to compute by PROB, we leave it untranslated. The translated formula has the form:

```

one lentry &&
lentry in L &&

```




Fig. 1. Control flow graph of a program

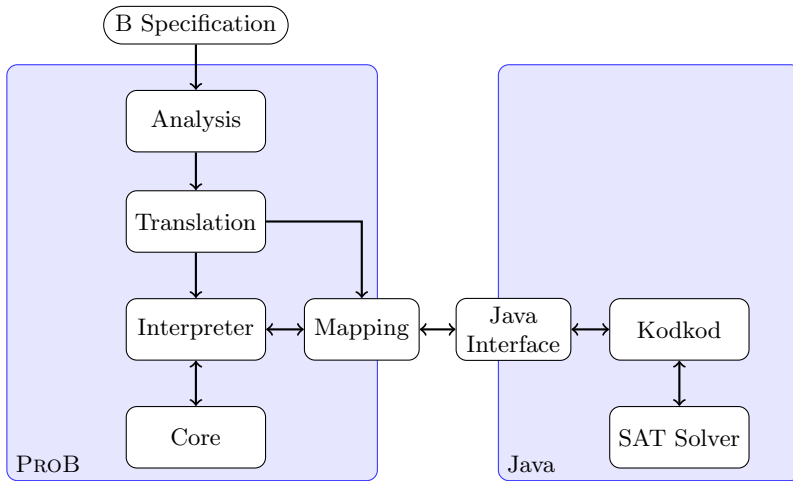


Fig. 2. Overview of the architecture

```

((L-lentry) . ~succs) in L &&
all l: one Blocks | (l in L =>
  lentry in (l.^((L->Blocks)&succs)&(Blocks->L))))
  
```

The translated description of the formula is then stored and a mapping between PROB’s internal representation and Kodkod’s representation of values is constructed (“Mapping” in Fig. 2). The message “new problem with following properties...” is sent to the Java process.

The Kodkod problem gets a unique identifier and the translated part of the B predicate is replaced by a reference to the problem, i.e., $succs = \{\dots\} \wedge kodkod(ID)$, and then given to the B interpreter of PROB.

When the PROB interpreter starts to evaluate the predicate, it prioritises which parts should be computed first. It chooses $succs = \{\dots\}$ because it can be computed deterministically by PROB’s core and finds a value for $succs$. Then a message “We have these values for $succs$, try to find values for the other variables” is sent to the Java process.

The Java process has now a complete description of the problem. It consists of the universe (with 8 atoms) and relations for the variables and the type *Blocks* itself. The bounds define the value of $succs$ and *Blocks* and ensure that all relations contain only atoms that match their corresponding types. This information is then given together with the formula to the solver of the Kodkod library (“Kodkod” in Fig. 2) that translates the problem into a SAT problem and passes this to the SAT solver.

The SAT solver finds solutions that are transformed by Kodkod to instances of the relations that fulfil the given formula. The values of the previously unknown relations that represent L and $lentry$ are sent back in an answer to the PROB process. The answer is then mapped to PROB’s internal representation of values.

The B interpreter can now continue with the found values. Now all predicates have been evaluated and the solutions can be presented to the user.

3.2 Translation

Representing Values. It turns out that the available data types in Kodkod are the main limitation when trying to translate a problem described in B. Let's first have a look at the available data types in B and how they can be translated to Kodkod. We have the basic data types:

Enumerated Sets. Enumerated sets can directly be translated to Kodkod. For each element of the set, we add an atom to the universe and create a unary relation that contains exactly that atom. The relation is needed in case the element is referred to in an expression. We create another unary relation for the whole set that contains exactly all atoms of the enumerated set.

Deferred Sets. Deferred Sets in B can have any number of elements that are not further specified. For animation, PROB chooses a fixed finite cardinality for the set, either by an analysis of the axioms or by using user preferences. Then we can treat deferred sets just like a special case of enumerated sets.

Booleans. The set of Booleans is a special case of an enumerated set with two elements TRUE and FALSE.

Integer. Integers in B represent mathematical numbers, they can be arbitrary large. It is possible to represent integer values in Kodkod, but the support is very limited and special care has to be taken. We describe the handling of integers in Section 3.3 in detail.

Thus, we can map a B variable of a basic data type to a Kodkod relation. Since Kodkod treats every relation as a set, we must ensure explicitly that the relations for such variables contain exactly one element.

Example. Let's assume that we use two types in our specification, an enumerated set $E = \{a, b, c\}$ and BOOL. Treating the Booleans as enumerated set $\text{BOOL} = \{\text{TRUE}, \text{FALSE}\}$, we have the following universe with five atoms:

	E			BOOL	
B value	a	b	c	TRUE	FALSE
atom	0	1	2	3	4

We can now represent a variable of type E by a unary relation $r1$ whose elements are bounded to be a subset of the atoms $0..2$. We also have to add the Kodkod formula `one r1`.

In B, two or more basic types can be combined with the Cartesian product. Variables of such a type can be represented by a relation.

Example. If we have a variable of type $(E \times E) \times \text{BOOL}$, we can represent it by a ternary relation $r2$ whose elements are bound to subsets of the atoms $0..2 \times 0..2 \times 3..4$. Like in the example above, we have to add the condition `one r2`.

We can construct the power set $\mathbb{P}(\alpha)$ for any type α in B. A variable of type $\mathbb{P}(\alpha)$ can be mapped to a Kodkod relation if α is itself not a set. A relation for $\mathbb{P}(\alpha)$ is defined exactly as a relation for α but without the additional restriction that it must contain exactly one element.

Finally, let's have a look at what we *cannot* translate. All “higher-order” data-types, i.e. sets of sets are not translatable. E.g a function $f \in A \mapsto \mathbb{P}(B)$ cannot be handled.

It turned out that unary and binary relations are handled very well. With relations of a higher arity we encounter the problem that many operators in Kodkod are restricted to binary relations. Thus it is not as easy to translate many properties using these data types.

Translating Predicates and Expressions. One of the central tasks in combining PROB and Kodkod is the translation of the B predicate that specifies the problem to a Kodkod formula. Many of B's most common operators can be directly translated to Kodkod, especially when basic set theory and relational algebra is used. It is not strictly necessary to cover all operators that B provides, because we always have the possibility to fall back to PROB's own constraint solving technique. Of course, we strive to cover as many operators as possible.

Operators on Predicates. The basic operators that act on predicates like conjunction, disjunction, etc. have a direct counterpart in Kodkod. This includes also universal and existential quantification.

Arithmetic Operators. Addition, subtraction and multiplication of integers can also directly be translated, whereas division is not supported by Kodkod. Other supported integer expressions are constant numbers and the cardinality of a set. If we want a variable to represent an integer, we have to convert explicitly between a relation that describes the value and an integer expression (see Section 3.3).

Relational Operators. Many operators that act on sets and relations can be translated easily to Kodkod.

Figure 3a shows a list of operators that have a direct counterpart in Kodkod. With $\mathcal{T}(A)$ we denote the translated version of the expression A . Please note that the expressions $A \in B$ and $A \subseteq B$ are translated to the same expression in Kodkod. This is due to the fact that single values are just a special case in Kodkod where a set contains just one element. The same effect can be found at the Cartesian product ($A \times B$) and a pair ($A \mapsto B$) and at the relational image ($A[B]$) and the function application ($A(B)$).

Other operators need a little bit more work. They can be expressed by combining other operators. Figure 3b shows a selection of such operators. In the table, we use an operator $\mathcal{A}(E)$ to denote the arity of the relation that represents the expression E . Again we can see that different operators in B (e.g. dom and prj_1) lead to the same result.

B	Kodkod	B	Kodkod
$A \in B$	$\mathcal{T}(A) \text{ in } \mathcal{T}(B)$	$\text{dom}(A)$	$\text{prj}[1:\mathcal{A}(\alpha)](\mathcal{T}(A))$
$A \subseteq B$	$\mathcal{T}(A) \text{ in } \mathcal{T}(B)$		with A being of type $\mathbb{P}(\alpha \times \beta)$
$A \times B$	$\mathcal{T}(A) \rightarrow \mathcal{T}(B)$	$\text{ran}(A)$	$\text{prj}[\mathcal{A}(\alpha)+1:\mathcal{A}(A)](\mathcal{T}(A))$
$A \mapsto B$	$\mathcal{T}(A) \rightarrow \mathcal{T}(B)$		with A being of type $\mathbb{P}(\alpha \times \beta)$
$A \cap B$	$\mathcal{T}(A) \& \mathcal{T}(B)$	$\text{prj}_1(A)$	$\mathcal{T}(\text{dom}(A))$
$A \cup B$	$\mathcal{T}(A) + \mathcal{T}(B)$	$\text{prj}_2(A)$	$\mathcal{T}(\text{ran}(A))$
$A \setminus B$	$\mathcal{T}(A) - \mathcal{T}(B)$	$A \triangleleft B$	$(\mathcal{T}(A) \rightarrow \mathcal{T}(B)) \& \mathcal{T}(B)$
$A[B]$	$\mathcal{T}(B). \mathcal{T}(A)$		with B being of type $\mathbb{P}(\alpha \times \beta)$
$A(B)$	$\mathcal{T}(B). \mathcal{T}(A)$	$A \triangleleft B$	$((\text{univ}-\mathcal{T}(A)) \rightarrow \mathcal{T}(\alpha)) \& \mathcal{T}(B)$
$A \triangleleft B$	$\mathcal{T}(A) ++ \mathcal{T}(B)$		with B being of type $\mathbb{P}(\alpha \times \beta)$
A^{-1}	$\sim \mathcal{T}(A)$	$\text{bool}(P)$	$\text{if } \mathcal{T}(P) \text{ then } \mathcal{T}(\text{TRUE}) \text{ else } \mathcal{T}(\text{FALSE})$
A^+	$\sim \mathcal{T}(A)$	$f \in A \mapsto B$	$\text{pfunc}(\mathcal{T}(f), \mathcal{T}(A), \mathcal{T}(B))$
		$f \in A \rightarrow B$	$\text{func}(\mathcal{T}(f), \mathcal{T}(A), \mathcal{T}(B))$
		$f \in A \mapsto B$	$\text{func}(\mathcal{T}(f), \mathcal{T}(A), \mathcal{T}(B)) \&\&$ $(\mathcal{T}(f). \sim \mathcal{T}(f)) \text{ in iden}$

(a) direct translation

(b) more complex rules

Fig. 3. Examples for translation rules

3.3 Integer Handling in Kodkod

Kodkod provides only a very limited support for integers. The reason for this is twofold. Since SAT solvers are used as the underlying technology, integers are encoded by binary numbers. Operations like addition then have to be encoded as boolean formulas. This makes the use of integers ineffective and cumbersome. Another reason is that the designers of Alloy – where Kodkod has its origin – argue [8] that integers are often not very useful and an indication of lack of abstraction when modeling systems.

Our intention is to make our tool applicable to as many specifications as possible, and many of the B specifications we tried contained some integer expressions. Indeed, integers are used to model sequences in B or multi-sets in Z.

When using Kodkod with integers, we have to specify the number of bits used in integer expressions. Integer overflows are silently ignored, e.g. the sum of two large naturals can be negative when the maximum integer size is exceeded. Thus we need to ensure that we use only integers in the specified range to prevent faulty results.

Kodkod distinguishes between relational and integer expressions. An integer expression is for example a constant integer or the sum of two integer expressions. Comparison of integer expressions like “less than” is also supported. In case we want a relation (i.e. a variable) that represents an integer, we must first assign values to some atoms. Figure 4 shows an example with a universe consisting of 9 atoms i_0, \dots, i_8 that represent integers.

	binary numbers										
	for integer sets										
	atom ...	i_0	i_1	i_2	i_3	i_4	i_5	i_6	i_7	i_8	...
associated integer value	-1	0	3	5	1	2	4	8	-16	...

Fig. 4. Mapping atoms to integer values

We have basically two options when we want represent integers by a relation:

- We can represent sets of integers in the interval $a..b$ by having an atom for each number in $a..b$. Then the relation simply represents the integers of its atoms.

E.g. with the universe in Figure 4, we can represent arbitrary subsets of $-1..5$ by using a relation that is bounded to the atoms i_0, \dots, i_4 .

The downside of this approach is that the number of atoms can become easily very large.

- Single integers can be represented more compactly by using a binary number. E.g. with the universe in Figure 4, we can represent a number of the interval $-16..15$ by using a relation that is bounded to the atoms $i_4..i_8$. A relation that consists of the atoms i_4, i_6, i_8 would represent the sum $1 + 4 + (-16) = -9$. Kodkod provides an operator to summarise the atoms of a relation, yielding an integer expression.

With this approach large numbers can be handled easily. The downside is that we cannot represent sets of numbers.

The atoms in the universe seen in Figure 4 are ordered in a way that we can use both approaches to represent integers in the same specification.

It can be seen that we need an exact knowledge of the possible size of integer expressions in the specification. To get the required information, a static analysis is applied to the specification before the translation. See below in Section 3.4 for details of the analysis.

Another problem that arises from having two kinds of integer representations, is that we have to ensure the consistency of formulas that use integer expressions. We briefly describe the problem in Section 3.5.

3.4 Predicate Analysis

In case that integers are used in the model, we need to know how large they can get in order to translate the expressions. To get this information we apply a static analysis on the given problem.

The first step of the analysis is that we create a graph that describes a constraint problem. For each expression in the abstract syntax tree, we create some of nodes depending on the expression's type that contain relevant information associated to the syntax node. E.g. this might be the possible interval for integers, the interval of the cardinality for sets or the interval in which all elements of a set lie.

By applying pattern matching on the syntax tree, we add rules that describe the flow of information in the graph. E.g. if we have a predicate $A \subseteq B$, we can propagate all information about elements of B to nodes that contain information about A . We evaluate all rules until a fixpoint or a maximum number of evaluation steps is reached.

Example. Let's take the predicate $A \subseteq 3..6 \wedge \text{card}(A) > 1$. For each integer node $(1, 3, 6, \text{card}(A))$ we create a node containing the integer range. For each of the sets A and $3..6$ we create two nodes: One contains the range of the set's cardinality, the other describes the integer range of the set's elements. Figure 5 shows the resulting graph. In the upper part of each node the expression is shown, in the lower part the kind of information that is stored in the node. The edges without any labels denote rules that pass information just from one node to another. Those which are labeled with $\leq, \geq, <, >$ express a relation between integer ranges of the source and target node. There is a special rule marked $i \rightarrow c$ ("interval to cardinality") which deduces a maximum cardinality from the allowed integers in a set. E.g. if all elements of a set I are in the range $x..y$, we know that $\text{card}(I) \leq y - x + 1$. The graph in Figure 5 shows the information we have about each node after the analysis. In particular, we know the bounds of all integer expressions.




Fig. 5. Constraint system for $A \subseteq 3..6 \wedge \text{card}(A) > 1$

Currently the analysis is limited to integer intervals and cardinality, because this was the concrete use case given by our translation to Kodkod. We plan to re-use the analysis for other aspects of PROB. E.g. if we can deduce the interval of a quantified integer variable, PROB can limit the enumeration of values to that range if it must test a predicate for all possible values of that variable.

Other types of information nodes are also of interest. For instance, we could infer the information if an expression is a function or sequence to assist PROB when evaluating predicates.

3.5 Integer Representations

We have seen above in Section 3.3 that we have two distinct forms of representing integers. Additionally we have Kodkod's integer expressions when we want to compare, add, subtract or multiply them. During the translation process we must ensure that the correct representation is chosen for each expression and that the representations are consistently used. If we take a simple equality $A = B$ with A, B being integers as an example, we must ensure that both sides use the same representation.

Internally, the result of the translation is an abstract syntax tree that describes the formula. Some expressions in this syntax tree have annotations about the needed integer representation. E.g., if the original expression in B is a set of integers, it has the annotation that the integer set representation and not the the binary number representation must be used. We now impose a kind of type checking on this syntax tree to infer if conversions between different integer representations have to be inserted in the formula.

3.6 Extent of the Translation: Partitioning

Theoretically we can take any translatable sub-predicate of a specification and replace it with a call to Kodkod. But usually, the overhead due to the communication between the processes can easily get so large that incorporating Kodkod has no advantage over using ProB alone.

A more sensible approach for a specification that is a conjunction of predicates $P_1 \wedge \dots \wedge P_n$ is to apply the translation to every P_i . All translatable predicates are then replaced by one single call to Kodkod. But even here we made the experience that the communication overhead can become large if not all predicates are translated.

Our current approach is to create a partition of the predicates P_1, \dots, P_n . Two predicates are then in the same set of the partition if they both use the same variable. We translate only complete partitions to keep the communication overhead small. There is one exception: We do not translate simple equations where one side is a variable and the other side an easy to compute constant. Such deterministic equations are computed first by the constraint solver, so the value for such a variable will be computed before the call to Kodkod is made. This keeps the translated formula small even for a large amount of data.

4 Experiments

We have chosen a number of problems to compare the performance of ProB's constraint solving technique and Kodkod's SAT solving approach. We have only used problems that can be completely translated. Completely translated models are still fully integrated into ProB, the results are converted to ProB's internal format and can be used for further animation and model checking. The results can be seen in Table 1.¹ All experiments were conducted on a dual-core Intel i7

¹ The source code of the examples are available in the technical report at:

http://www.stups.uni-duesseldorf.de/w/Special:Publication/PlaggeLeuschel_Kodkod2012 .

2.8 GHz processor running under Linux. MiniSat was used as a SAT solver. The measured times do not contain time for starting up PROB and for loading, parsing and type-checking the model. We measured the time to compute all solutions to each problem. For Kodkod, we measured two different times: The “total time” includes the translation of the problem, the communication between the two processes and the time needed by the solver to produce solutions. The “solver time” is the time that the Kodkod solver itself needs to find solutions, without the overhead of translation and communication between the processes.

Table 1. Comparing PROB and Kodkod (in milliseconds)

Model	PROB		Kodkod	
			total solver	
Who Killed Agatha?	177	123	12	
Crew Allocation	timeout*	297	112	
20-Queens	110	8223	8076	
Graph Colouring (integer sets)	50	2323	1859	
Graph Colouring (enumerated sets)	50	1037	818	
Graph Isomorphism	13	553	379	
Loop detection in control flow	23037	117	12	
SAT instance	11830	4143	588	
Send More Money	7	1773	1578	
Eratosthenes' sieve (1 step)	7	5833	5712	
Union of two sets (2000 elements)	33	4880	4659	
Requirements WRSPM model	timeout*	333	89	
BPEL deadlock check	20	337	68	

*: interrupted after 60 seconds

4.1 Analysis

Relational Operators. Let’s have a look at those problems where Kodkod is much faster than PROB. These are “crew allocation”, “loop detection” and “WR-SPM”. In all these problems we search for instances for sets or relations and the problem is described by relational operators and universal quantification. In this scenario, PROB sometimes starts to enumerate possible instances for the sets or relations which leads to a dramatic decrease of performance.

Arithmetic and Large Relations. The arithmetic problems (“Send More Money”, “Eratosthenes’ Sieve”) are solved by ProB much faster than by Kodkod. The first two problems deal with arithmetic. PROB uses internally a very efficient finite domain solver (CLP/FD) to tackle such problems. On the other side, arithmetic is one of the weaknesses of Kodkod, as we already pointed out.

Kodkod does not seem to scale well when encountering large relations (e.g. “union of two sets”). This has only been relevant for certain applications of PROB, such as the property verification on real data [13].

The graph colouring, graph isomorphism and 20-Queens problems are clearly faster solved by PROB. The structure of the problem is somehow fixed (by having e.g. total functions) and constraint propagation is very effective.

Room for Optimization. It can be seen that the graph colouring problem needs less than half the time when it is encoded with enumerated sets instead of integers. This indicates that the translation is not yet as effective as it should be. For the “SAT” problem, the translation and communication takes six times as long as the computation of the problem itself. This shows that we should investigate if we can optimize the communication.

4.2 SMT and other Tools

Very recently, an Event-B to SMT-Lib converter has become available for the Rodin platform [4]. This makes it possible to use SMT solvers (such as veriT, CVC and Z3 [3]; we used version 3.2 of the latter within the Rodin SMT Solvers Plug-in 0.8.0r14169 in our experiments below) on Event-B proof obligations. We have experimented with the translator on the examples from Table 1.² This is done by adding a theorem $1=2$ to the model: this generates an unprovable proof obligation which in turn produces a satisfiable SMT formula encoding the problem. For “Send More Money” from Table 1 Z3 initially reported “unknown”. After rewriting the model (making the inequalities explicit), Z3 was able to determine the solution after about 0.250 seconds. It is thus faster than Kodkod, but still slower than PROB. Surprisingly, Z3 was unable to solve the SMT-Lib translations for most of the other examples, such as the “Who killed Agatha” example, the “Set Union” example or the “Graph Colouring” example. Similarly, for the “Crew Allocation” example, Z3 was unable to find a solution already for three flights.³ Furthermore, for the constraint solving tasks related to deadlock checking, Z3 was not able to solve the translations of the simpler examples from [7]. It is too early for a conclusive result, but it seems that more work needs to be put into the B to SMT-Lib translator for this approach to be useful for model finding, animation or constraint-based checking.

Other tools for B are AnimB [17], Brama and BZTT [11]. They all have much weaker constraint-solving capabilities (see [13,14]) and are unable to solve most of the problems in Table 1. Another tool is TLC [24] for TLA⁺. It is very good at model checking, but constraints are solved by pure enumeration. As such, TLC is unable to solve, e.g., a 20 variable SAT problem, the NQueens problem for $N > 9$ and takes more than 2 hours for a variation of the graph isomorphism problem from Table 1.

5 More Related Work, Discussion and Conclusion

5.1 Alternative Approaches

Before starting our translation to Kodkod, we had experimented with several other alternate approaches to solve constraints in PROB. [22] offers the user a

² Apart from “loop” which cannot be easily translated to Event-B due to the use of transitive closure.

³ PROB solves this version in 0.06 seconds; Table 1 contains the problem for 20 flights.

Datalog-like language that aims to support program analysis. It uses BDDs to represent relations and compute queries on these relations. In particular, one has to represent a state of the model as a bit-vector and events have to be implemented as relations between two of those bit-vectors. These relations have to be constructed by creating BDDs directly with the underlying BDD library (JavaBDD) and storing them into a file. Soon after starting experimenting with BDDBDDB it became apparent that due to the lack of more abstract data types than bit vectors, the complexity of a direct translation from B to BDDBDDB was too high, even for small models, and this avenue was abandoned.

SAL [20] is a model-checking framework combining a range of tools for reasoning about systems. The SAL tool suite includes a state of the art symbolic (BDD-based) and bounded (SAT-based) model checkers. Some first results were encouraging for a small subset of the Event-B language, but the gap between B and SAL turned out to be too big in general and no realistic way was found to handle important B operators.⁴ More details about these experiments can be found in [19]. For Z, there is an ongoing attempt to use SAL for model checking Z specifications [5,6]. The examples presented in [5,6] are still relatively simple and pose no serious challenge in constraint solving. As the system is not publicly available, it is unclear how it will scale to more complicated specifications and constraints.

5.2 More Related Work

The first hand-translation of B to Alloy was undertaken in [18]. The paper [16] contains first experiments in translating Event-B to Alloy; but the work was also not pursued. Later, [15] presented a prototype Z to Alloy converter. The current status of this system is available at the website <http://homepages.ecs.vuw.ac.nz/~petra/zoy/>; the applicability seems limited by the lack of type inference and limited support for schemas. In contrast to these works, we translate directly to Kodkod and have a fully developed system, covering large subsets of B, Event-B, Z and TLA⁺ and delegating the rest to the PROB kernel.

A related system that translates a high-level logic language based on inductive definitions to SAT is IDP [23]. Another recent addition is Formula from Microsoft [9], which translates to the SMT solver Z3 [3].

5.3 Future Work

Currently our translation is only applicable for finding constraints satisfying the axioms as well as for constraint based deadlock checking. We are, however, working to also make it available for computing enabled events as well as for more general constraint-based testing and invariant checking.

Another avenue is to enlarge the area of applicability to some recurrent patterns of higher-order predicates. For example, many B specifications use total

⁴ Private communication from Alexei Iliasov and Ilya Lopatkin, March 6th, 2012.

functions of the type $f : \text{DOM} \rightarrow \text{POW}(\text{RAN})$ which cannot be translated as such to Kodkod. However, such functions can often be translated to relations of the form $fr : \text{DOM} \leftrightarrow \text{RAN}$ by adapting the predicates accordingly (e.g., translating $f(x)$ to $fr[\{x\}]$). More work is also needed on deciding automatically when to attempt the Kodkod translation and when predicates should be better left to the existing PROB kernel. Finally, inspired by the experiments, we also plan to improve the PROB kernel for better solving constraints over relational operators such as composition and closure.

5.4 Conclusion

After about three years of work and several attempts our translation to Kodkod is now mature enough to be put into practice and has been integrated into the latest version of the PROB toolset. The development required a considerable number of subsidiary techniques to be implemented. As our experiments have shown that the translation can be highly beneficial for certain kinds of constraints, and as such opens up new ways to analyze and validate formal specifications in B, Z and TLA⁺. However, the experiments have also shown that the constraint logic programming approach of PROB can be superior in a considerable number of scenarios; the translation to Kodkod and down to SAT is not (yet) the panacea. The same can be said of the existing translations from B to SMT. As such, we believe that much more research is required to reap the best of both worlds (SAT/SMT and constraint programming). An interesting side-effect of our work is that the PROB toolset now provides a double-chain (relying on technology developed independently and using different programming languages and paradigms) of validation for first-order predicates, which should prove relevant in high safety integrity level contexts.

Acknowledgements. We want to thank Sebastian Krings for helping us in the SMT experiments. We also thank Alexei Iliasov and Ilya Lopatkin for their earlier experiments involving translating Event-B to SAL and the insights gained in that work. We are also grateful for discussions with Daniel Jackson and Emina Torlak, in particular relating to encoding integers in Kodkod.

References

1. Abrial, J.-R.: *The B-Book*. Cambridge University Press (1996)
2. Aho, A.V., Lam, M.S., Sethi, R., Ullman, J.D.: *Compilers. Principles, Techniques, and Tools*, 2nd edn. Addison Wesley (2007)
3. de Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
4. Déharbe, D., Fontaine, P., Guyot, Y., Voisin, L.: SMT Solvers for Rodin. In: Derrick, J., Fitzgerald, J., Gnesi, S., Khurshid, S., Leuschel, M., Reeves, S., Riccobene, E. (eds.) ABZ 2012. LNCS, vol. 7316, pp. 194–207. Springer, Heidelberg (2012)

5. Derrick, J., North, S., Simons, A.J.H.: Z2SAL - Building a Model Checker for Z. In: Börger, E., Butler, M., Bowen, J.P., Boca, P. (eds.) ABZ 2008. LNCS, vol. 5238, pp. 280–293. Springer, Heidelberg (2008)
6. Derrick, J., North, S., Simons, A.J.H.: Z2SAL: a translation-based model checker for Z. *Formal Asp. Comput.* 23(1), 43–71 (2011)
7. Hallerstede, S., Leuschel, M.: Constraint-based deadlock checking of high-level specifications. *TPLP* 11(4–5), 767–782 (2011)
8. Jackson, D.: Alloy: A lightweight object modelling notation. *ACM Transactions on Software Engineering and Methodology* 11, 256–290 (2002)
9. Jackson, E.K., Tiham, Balasubramanian, D.: Reasoning about Metamodeling with Formal Specifications and Automatic Proofs. In: Whittle, J., Clark, T., Kühne, T. (eds.) MODELS 2011. LNCS, vol. 6981, pp. 653–667. Springer, Heidelberg (2011)
10. Lamport, L.: *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley (2002)
11. Legeard, B., Peureux, F., Utting, M.: Automated Boundary Testing from Z and B. In: Eriksson, L.-H., Lindsay, P.A. (eds.) FME 2002. LNCS, vol. 2391, pp. 21–40. Springer, Heidelberg (2002)
12. Leuschel, M., Butler, M.: ProB: A Model Checker for B. In: Araki, K., Gnesi, S., Mandrioli, D. (eds.) FME 2003. LNCS, vol. 2805, pp. 855–874. Springer, Heidelberg (2003)
13. Leuschel, M., Falampin, J., Fritz, F., Plagge, D.: Automated Property Verification for Large Scale B Models. In: Cavalcanti, A., Dams, D.R. (eds.) FM 2009. LNCS, vol. 5850, pp. 708–723. Springer, Heidelberg (2009)
14. Leuschel, M., Falampin, J., Fritz, F., Plagge, D.: Automated property verification for large scale B models with ProB. *Formal Asp. Comput.* 23(6), 683–709 (2011)
15. Malik, P., Groves, L., Lenihan, C.: Translating Z to Alloy. In: Frappier, M., Glässer, U., Khurshid, S., Laleau, R., Reeves, S. (eds.) ABZ 2010. LNCS, vol. 5977, pp. 377–390. Springer, Heidelberg (2010)
16. Matos, P.J., Marques-Silva, J.: Model Checking Event-B by Encoding into Alloy (Extended Abstract). In: Börger, E., Butler, M., Bowen, J.P., Boca, P. (eds.) ABZ 2008. LNCS, vol. 5238, pp. 346–346. Springer, Heidelberg (2008)
17. Métayer, C.: AnimB 0.1.1 (2010), <http://wiki.event-b.org/index.php/AnimB>
18. Mikhailov, L., Butler, M.: An Approach to Combining B and Alloy. In: Bert, D., Bowen, J.P., Henson, M.C., Robinson, K. (eds.) B 2002 and ZB 2002. LNCS, vol. 2272, pp. 140–161. Springer, Heidelberg (2002)
19. Plagge, D., Leuschel, M., Lopatkin, I., Romanovsky, A.: SAL, Kodkod, and BDDs for validation of B models. lessons and outlook. In: *Proceedings AFM 2009*, pp. 16–22 (2009)
20. Symbolic Analysis Laboratory (SAL) website, <http://sal.cs1.sri.com/>
21. Torlak, E., Jackson, D.: Kodkod: A Relational Model Finder. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 632–647. Springer, Heidelberg (2007)
22. Whaley, J., Lam, M.S.: Cloning-based context-sensitive pointer alias analysis using binary decision diagrams. In: Pugh, W., Chambers, C. (eds.) *Proceedings PLDI 2004*, pp. 131–144. ACM Press, New York (2004)
23. Wittocx, J., Mariën, M., Denecker, M.: Grounding FO and FO(ID) with bounds. *J. Artif. Intell. Res. (JAIR)* 38, 223–269 (2010)
24. Yu, Y., Manolios, P., Lamport, L.: Model Checking TLA⁺ Specifications. In: Pierre, L., Kropf, T. (eds.) CHARME 1999. LNCS, vol. 1703, pp. 54–66. Springer, Heidelberg (1999)