

# Synthesis of a simple self-stabilizing system

Adrià Gascón

SRI International  
California, USA

adriagascon@gmail.com

Ashish Tiwari

SRI International  
California, USA

tiwari@csl.sri.com

With the increasing importance of distributed systems as a computing paradigm, a systematic approach to their design is needed. Although the area of formal verification has made enormous advances towards this goal, the resulting functionalities are limited to detecting problems in a particular design. By means of a classical example, we illustrate a simple template-based approach to computer-aided design of distributed systems based on leveraging the well-known technique of bounded model checking to the synthesis setting.

## 1 Introduction

Consider a situation where a developer is trying to design and program a multi-agent distributed system to perform a certain task. The agents could be robots communicating with each other, sensors in a sensor network, processes in a multi-core machine, or processors connected to a bus in an avionics system. The task could involve achieving consensus, getting mutually exclusive access to some shared resource, computing some function of some sensed data, or something similar. The developer knows the underlying topology of the communication network, the (synchronous or asynchronous) communication model, as well as the nature of faults the network and the agents themselves can manifest. Does there exist a known algorithm that the developer could use in this particular scenario?

One path forward for the developer would be to study the literature on distributed algorithms [16]. It contains several impossibility results, as well as positive results and algorithms for several distributed problems, but all those results are accompanied with an applicability criterion. Does the condition for applicability hold for the developer's particular situation? Suppose that an impossibility result is applicable. Is there a workaround if certain system requirements are changed? Suppose the developer finds an algorithm that appears to be meaningful in his context, but the applicability criteria does not match. Can the discovered algorithm be massaged for his particular context?

All the questions raised above may be difficult to answer. The reason is that all intuitions about what works start to fail for distributed algorithms, and more so in the presence of faults. Reasoning about the correctness of an algorithm in the presence of faults is not only difficult, but also a surprisingly delicate task. The following quote, taken from the seminal paper that introduced the well-known Byzantine Generals Problem [13], talks about the correctness of an informally presented argument:

*“This argument may appear convincing, but we strongly advise the reader to be very suspicious of such nonrigorous reasoning. Although this result is indeed correct, we have seen equally plausible ‘proofs’ of invalid results. We know of no area in computer science or mathematics in which informal reasoning is more likely to lead to errors than in the study of this type of algorithm.”*

For these reasons, applying formal verification to distributed algorithms, as well as their fault-tolerant variants, has drawn considerable attention. In fact, several mechanized correctness proofs exist for some

classical distributed algorithms [24, 22, 14, 20], and such mechanization even lead to the detection of flaws in published results in some cases [22, 14]. In contrast to these examples of the application of verification technology, where the goal was to formally verify the correctness of a given algorithm, we are interested in *using formal methods to guide the human in the design process*, instead of just to verify its result.

Recently, there has been lots of excitement and progress in automatically synthesizing programs that satisfy some given requirement. This development was triggered by the observation that program synthesis starts becoming feasible if we start from a program sketch – rather than from a clean slate – and synthesize a program by completing the sketch so that it matches the given specification. This synthesis process has been effectively demonstrated for imperative programs [25, 26].

The main reason why distributed algorithms are specially suitable for a computer-aided design methodology is that the solutions to the kind of problems mentioned above are usually short and easy to describe, while their correctness (or impossibility) proofs can be very involved. This situation is ideal for synthesis tools, since their complexity is roughly the product of the size of the design space and the verification (checking) complexity. Consequently, most synthesis tools need either the design space to be small or the verification (checking) effort to be minimal.

However, sketches for distributed algorithms can not be written in imperative languages. A much richer language is needed (see [28] for related recent work on synthesis distributed protocols). The input language of formal verification tools, such as the SAL language, is a great option. It provides a very rich set of constructs – nondeterminism, synchronous and asynchronous composition, parametric module specifications, module instantiations, rich datatypes and rich expression language – that are needed for modeling the execution of distributed algorithms in presence of faults. However, the formal verification tools that run on SAL models, such as the SAL model checker, are just verification tools and hence they do not perform synthesis. If parts of the modeled system are not known, they can not help complete the algorithm in any way; though they can verify a manually completed sketch.

In this paper, we present computational techniques that can aid a human in exploring the design space of algorithms; that is, the field of *computer-aided synthesis*, with a focus in problems arising in *distributed systems*. Our proposed approach is based on using *synthesis-versions of popular formal verification techniques*. A general view of our approach to build computer-aided synthesis technology is shown in Figure 1. Just as SAT-based bounded model checking turns a verification problem into a search problem (over a large, but finite, search space), QBF-based *bounded model synthesis* turns a synthesis problem into a large, but finite, one-step  $\forall\exists$  game that can be solved using a QBF solver. Similarly, verification by  $k$ -induction can be lifted to  $k$ -inductive synthesis. There is a similar correspondence between the infinite state space versions of these techniques. In the present paper we focus on bounded model synthesis. Our approach is enabled by some impressive recent progress in the field of QBF solving and  $\exists\forall$  SMT solving [21, 29, 15, 19, 18, 11, 10, 4].

More concretely, in this paper we focus on leveraging the technique of bounded model checking to the template-based synthesis setting. Our templates are written in the SAL language, which is, as commented above, a suitable formalism to describe a distributed system. We also take advantage of the SAL model checker to construct a 2QBF formula that is afterwards sent to an off-the-shelf QBF solver. Our work indicates that the synthesis-extension of bounded model checking can be used to obtain surprising new algorithms, show non-existence of algorithms for certain classes of problems, and generate useful variants of known algorithms.

The rest of this paper is organized as follows. In the next section we present our running example, a problem inspired by Dijkstra’s paper [5]. In Section 3 we present our synthesis methodology and describe an opportunistic implementation using the SAL model checker. In Section 4 we describe our experience

Task	Inputs	Technique	Backend solver
Verification:	Model, Property	$\longrightarrow$ <i>bmc, k-ind</i>	$\longrightarrow$ SAT formula
Synthesis:	Partial Model, Property	$\longrightarrow$ <i>bms, k-inds</i>	$\longrightarrow$ $\forall\exists$ QBF formula
Verification:	Model, Property	$\longrightarrow$ <i>inf-bmc, k-ind</i>	$\longrightarrow$ SMT formula
Synthesis:	Partial Model, Property	$\longrightarrow$ <i>inf-bms, k-inds</i>	$\longrightarrow$ $\forall\exists$ SMT formula

Figure 1: Lifting some verification techniques to synthesis: Bounded model checking (bmc) and  $k$ -induction (k-ind) are generalized to bounded model synthesis (bms) and  $k$ -inductive synthesis (k-inds); and similarly, for the infinite state space versions of these techniques.

applying our approach to the running example, with references to all the SAL models implemented along the way. Finally, in Section 5 we provide some discussion and directions for further work.

## 2 Running example: Reaching Mutual Exclusion

In this section, we describe a simple example inspired by Dijkstra’s paper [5], which is a remarkable milestone in the study of fault tolerance. The example has the property that its solution is simple to describe, yet difficult to verify. Subsequently, we will use the same example to present our approach to computer-aided design of distributed algorithms. For another more complex example on fault-tolerant consensus, which also provides a proof of concept for our approach, the reader is referred to [8].

Consider a system with four machines  $m_0, m_1, m_2, m_3$ . Each machine  $m_i$  has two Boolean state variables  $A$  and  $B$ . The 4 machines are arranged in a ring topology in which every machine has read access to the state variables of its right and left neighbors; that is, machine  $m_i$  has read/write permission on its own state variables  $A$  and  $B$  and read permission on the state variables of machine  $m_{(i+1)\bmod 4}$ , which we denote as  $A_R$  and  $B_R$ , and the state variables of machine  $m_{(i-1)\bmod 4}$ , which we denote as  $A_L$ ,  $B_L$ . Each machine  $m_i$  updates its state according to a finite set of rules  $R_i$  of the form

IF privilege THEN make move ENDIF

where privilege is a Boolean condition on the state variables of the machine and its neighbors, and a move is an update to the values of  $A$  and  $B$ . We say that a rule is *enabled* at some step if its privilege evaluates to true in that step and its corresponding move *changes* the current state. At each step, a rule is arbitrarily selected from the set of enabled rules and executed. We say that the system is in a *legitimate* state if *exactly one* rule in  $\bigcup R_i$  is enabled. The problem is to find rules for each machine  $m_0, \dots, m_3$  such that:

- (a) At least one rule will always be enabled and the system is *guaranteed* to reach a legitimate state *regardless of its initial conditions* in a *finite* number of steps.
- (b) In each legitimate state, each possible move will bring the system again into a legitimate state.

Intuitively, the initial state is arbitrary and multiple machines can make a move, but eventually we want the machines to get a *mutually exclusive access* to make a move.

It is not at all obvious how to design local rules that will achieve convergence towards states satisfying (a) and (b). Note that the source of difficulty is that the initial state, as well as the subsequent moves of the system, are all picked nondeterministically. Another source of difficulty that we will consider later is requiring *fairness*, i.e. for every pair of machines  $m_i, m_j$ , there is a sequence of steps of the system going

from a legitimate state where a rule of  $m_i$  is enabled to a legitimate state where a rule of  $m_j$  is enabled. Note that there might be other reasonable definitions of fairness.

Some questions that may arise in the process of designing an algorithm for this problem might be: How many rules do we need? Is it useful to restrict the states space by fixing some variables to have a certain value? Is there a solution where all machines have the same set of rules? As stated by Dijkstra in the original paper, the discovery that the answer to the third question is “no” was crucial to obtaining an algorithm.

With our proposed approach, the third question can be automatically answered for a fixed (but reasonably big) number of rules of considerable complexity and for a fixed choice for the number of steps to achieving convergence.

Before going into the details of our approach to template-based synthesis of distributed systems in the next section, let us present a possible solution to our problem. We encourage the reader to think about the problem at this point.

In this solution,  $B_1$  is fixed to have value *false*,  $B_4$  is fixed to have value *true*, and the set of rules for each of the four machines is defined as follows:

$$\begin{aligned} R_0 &= \{\text{IF } (A = A_R) \wedge B_R \text{ THEN } A := \bar{A} \text{ ENDIF}\}, \\ R_3 &= \{\text{IF } A \neq A_L \text{ THEN } A := \bar{A} \text{ ENDIF}\} \\ R_1 = R_2 &= \{\text{IF } A \neq A_L \text{ THEN } A := \bar{A}, B := \textit{false} \text{ ENDIF}, \\ &\quad \text{IF } (A = A_R) \wedge B_R \text{ THEN } B := \textit{true} \text{ ENDIF}\} \end{aligned}$$

Note that every machine needs at most two different rules in this solution. We will not argue here about the correctness of this solution, which was obtained using the synthesis methodology described in the following section and later verified in SAL. (see [7] for the corresponding SAL model).

### 3 A Synthesis approach for $FG$ properties

Roughly speaking, any template-based program synthesis algorithm must traverse the space of possible instantiations of a given template and check if one of them satisfies the requirement, i.e. implements a solution to the given problem. Checking if a synthesized solution satisfies a requirement is a formal verification problem. Hence, synthesis can be simply performed as a loop over the formal verification tool. Our approach to synthesis is simpler: we merge the search and verify loop into just one constraint, as done in previous works such as [3, 6, 1, 27, 26, 25].

Our approach can be viewed as a generalization of the idea of bounded model checking to synthesis. Just as bounded model checking turns a verification problem into an *existential* constraint that encodes a *weaker* version of the verification problem, we turn synthesis into a *forall-exists* constraint that encodes a *weaker* version of the synthesis problem. The key step that makes automated synthesis effective is the step that defines the weaker version. A simpler version of the synthesis problem is obtained by

- (i) restricting the universe of possible algorithms that will be searched and
- (ii) replacing the verification step by an approximation step.

In our example, restriction (i) is achieved by fixing a template of the solution to be synthesized. That already restricts the search space for possible solutions to a finite (but possibly huge) set. As another example, along with fixing the number of processes to a constant, one can also fix the signature of the messages exchanged between processes (see [8]). We believe that the limitations derived from this kind of restrictions are harmless from the perspective of a system designer or researcher trying to gain insight

into a problem, although overcoming them in a general way is a challenging and important problem from the verification perspective, even specifically for distributed systems [12].

For restriction (ii), we modify the property that the solution has to satisfy. The modification is done so that the new property is easier to verify. Specifically, and analogously to how bounded model checking is used to check that a property is not violated in a fixed number of steps  $k$ , we replace our LTL property  $FG\phi$  by  $X^cG\phi$ , for some natural number  $c$ . Intuitively, that corresponds to relaxing a property of the form “eventually, it is always the case that  $\phi$  holds” to “after  $c$  steps, it is always the case that  $\phi$  holds”. The word “steps” might create some confusion here since it depends on the particular problem being analyzed. However, for distributed systems, regardless of their timing model, a notion of step always exist. Moreover, as we will see in our example, an adequate modeling of the problem might, in some cases, make the properties  $FG\phi$  and  $X^cG\phi$  equivalent for a suitable  $c$ .

However, the modified property  $X^cG\phi$  may still be too complex for our synthesis purposes. Hence, we can replace  $G\phi$  by just  $\phi$  or  $\phi \wedge X\phi$ . Thus, instead of meeting the requirement  $FG\phi$ , the synthesis tool may find a solution that satisfies a weaker requirement, say  $XXX(\phi)$ . For example, in our running example, one possible relaxed version of property (a) could be: At least one rule will always be enabled and the system is guaranteed to reach a legitimate state *regardless of its initial conditions* in 8 steps. Whether this property is enough to synthesize a solution depends on the provided template. Similarly, in our previous work on synthesis of distributed consensus algorithms, our relaxed synthesis property was that consensus must be achieved in at most 3 steps, instead of an arbitrarily large (but finite) number of steps.

Due to the modification of the requirement, a solution found automatically may not be sound with respect to the original requirements. It needs to be formally verified and hence, the synthesized solution is verified against the original property  $FG\phi$ . Since our approach leverages existing verification techniques to the synthesis setting, the final verification step does not need any extra encoding or translation work.

### 3.1 From the SAL model to the synthesis constraint

The modeling language of verification tools, such as SAL and NuSMV, just defines state transition systems, but provide powerful language constructs for this purpose that make it easy to model concurrent systems. Distributed algorithms, regardless of their timing model, can also be easily modeled as open (finite) state transition systems in these languages. Let  $\vec{x}$  denote all the state variables appearing in a model. Let  $I(\vec{x})$  be the predicate denoting the initial states and  $T(\vec{x}_1, \vec{x}_2)$  be the predicate denoting the transition relation (of the state transition system).

SAT-based (bounded) model checking is a powerful bug-detection technique that is available in many verification tools. Let us provide some details about bounded model checking. Given the transition system defined by  $I, T$ , the property  $G\phi$ , and a depth to search 3, a bounded model checker generates the following formula:

$$\exists \vec{x}_0, \vec{x}_1, \vec{x}_2, \vec{x}_3 : I(\vec{x}_0) \wedge T(\vec{x}_0, \vec{x}_1) \wedge T(\vec{x}_1, \vec{x}_2) \wedge T(\vec{x}_2, \vec{x}_3) \wedge \neg\phi(\vec{x}_3) \quad (1)$$

which states that there is 3-step execution of the system that violates the property  $G\phi$ .

Now consider the problem of synthesizing a transition system to satisfy  $F\phi$ . Let  $\vec{z}$  denote all the state variables appearing in a *template model/sketch* of the transition system. The set  $\vec{z}$  can be partitioned as  $\vec{x} \cup \vec{y}$ , where the  $\vec{y}$  are the (input) variables used to represent the synthesis search space and  $\vec{x}$  are the remaining regular (non-synthesis) variables (as in the verification case above).

Instead of synthesizing for  $F\phi$ , i.e. enforcing the LTL property  $F\phi$  in the resulting synthesized model, say we decide to satisfy the stronger requirement  $XXX\phi$ . Given the template transition system

defined by  $I, T$  with synthesis variables  $\vec{y}$  and non-synthesis variables  $\vec{x}$ , the property  $F\phi$ , and a depth for synthesis 3, a *bounded model synthesizer* generates the following formula:

$$\begin{aligned} \forall \vec{y}_0, \vec{y}_1, \vec{y}_2, \vec{y}_3 : \vec{y}_0 = \vec{y}_1 = \vec{y}_2 = \vec{y}_3 \Rightarrow \\ (\exists \vec{x}_0, \vec{x}_1, \vec{x}_2, \vec{x}_3 : I(\vec{z}_0) \wedge T(\vec{z}_0, \vec{z}_1) \wedge T(\vec{z}_1, \vec{z}_2) \wedge T(\vec{z}_2, \vec{z}_3) \wedge \neg\phi(\vec{z}_3)) \end{aligned} \quad (2)$$

where  $\vec{z}_i = \vec{x}_i \cup \vec{y}_i$  for all  $i$ .

Formula 2 says that “for every concrete instance of the state transition system (defined by assignment to  $\vec{y}_0$ ), there is an execution of that transition system that does not reach  $\phi$  in 3 steps”. If this formula is invalid, then it means that there is a concrete instantiation of the template that always reaches  $\phi$  in 3 steps. This indicates that synthesis is successful (for the requirement  $XXX\phi$ , and consequently for  $F\phi$ ). If the formula is valid, then it means that synthesis fails for the requirement  $XXX\phi$ . It is important to remark that this approach, as well as bounded model checking, assumes that the transition system of the modeled state machines is total, i.e. there are no deadlock states.

If the domains of all variables in Formula 2 have finite cardinality, then the formula can be written as a quantified ( $\forall\exists$ ) Boolean formula (QBF), which can be solved using off-the-shelf QBF solvers. The synthesized algorithm, if it exists, is obtained from the refutation of the formula generated by the QBF solver in form of a Herbrand model, i.e. a valuation for variables  $\vec{y}_0, \vec{y}_1, \vec{y}_2, \vec{y}_3$ .

Note that Formula 2 is not very different from Formula 1, which is generated by existing bounded model checkers. In the work presented in this paper, we modeled our template in SAL, and used the SAL bounded model checker to generate Formula 1, together with a mapping from variables of the SAL model to the corresponding arrays of Boolean variables occurring in Formula 1. Then, we used a simple script to convert Formula 1 into Formula 2. Specifically, our investigation was carried out by performing the following steps, described also in Figure 2:

1. We model the *template* of distributed algorithm in SAL [23, 17]. The model includes synthesis variables  $\vec{y}$  to define the transition relation.
2. We use the SAL bounded model checker to generate the SAT formula for the verification constraint (Formula 1). The SAT formula implicitly existentially quantifies all variables, including the synthesis variables  $\vec{y}$ .
3. We modify the SAT formula and convert it into a QBF formula by universally quantifying the synthesis variables. (This step uses the mapping from the original SAL variables to the Boolean SAT variables).
4. We use off-the-shelf QBF solvers (and a QBF preprocessor) to check validity of the  $\forall\exists$  formula (Formula 2). For the experiments reported in the next section, we used the QBF preprocessor Bloqer [2], followed by the QBF solver RareQS [9], although we have also experimented with DepQBF [15].
5. If the QBF solver returns *Unsat*, then the synthesis is declared *successful*, and if the QBF solver returns *Sat*, then the synthesis process is *unsuccessful*.
6. If synthesis is successful, the QBF solver outputs a valuation for the synthesis variables  $\vec{y}$  (a Herbrand model), which is used to obtain a concrete distributed algorithm.
7. The synthesized algorithm is formally verified: if the property was  $F\phi$ , there is nothing to verify; if the property was  $FG\phi$ , then the property that “after  $k$  steps, the property  $\phi$  is always true” is verified using  $k$ -induction or symbolic model checking.

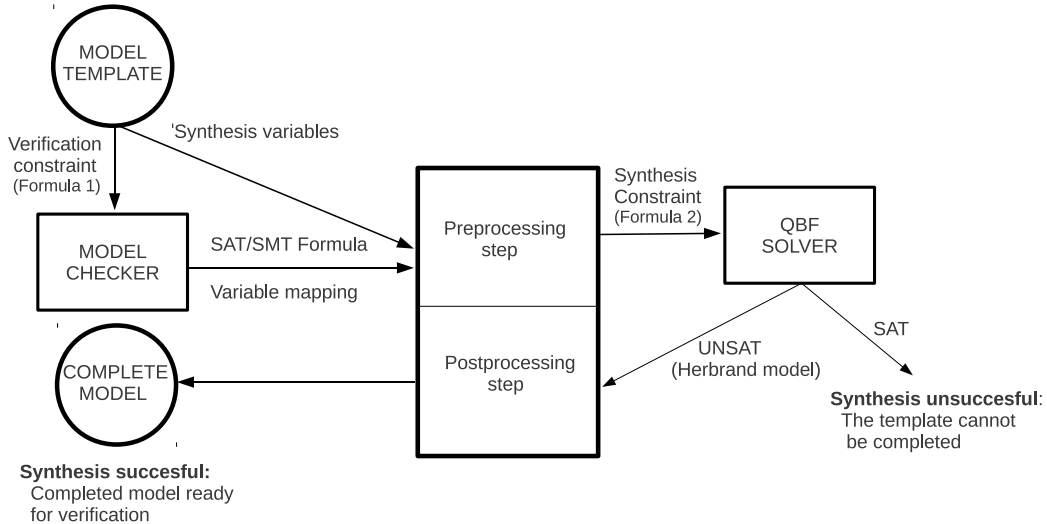


Figure 2: The bounded synthesis approach.

## 4 Synthesis of a self-stabilizing system with distributed control

In this section we present an example of our synthesis approach by finding a solution for the problem presented in Section 2. Instead of just presenting the template that was provided to our synthesis tool to produce a solution, we will illustrate one possible chain of interactions with the synthesis tool that leads to a solution. Our goal is to demonstrate how interacting with the synthesis tool is useful to get insight into the problem. To this end, we explain how Synthia, an imaginary character, used our approach to synthesize a solution to the problem of Section 2. We will not get into the modeling details, since all SAL models can be accessed at [7]. An advantage of our approach is that limited effort is needed to modify a template due to the expressivity of the SAL language.

### 4.1 How many rules?

The first question that came to Synthia’s mind was whether a really simple protocol would work. Is a single rule per process enough? What about the same rule for every process? That did not seem plausible but, to be sure, Synthia encoded the following simple solution template, which represents a finite family of possible solutions where (1) all 4 machines have the same rule set, (2) this rule set contains a single rule, (3) the privilege of that rule is a conjunction of two equality predicates comparing  $A$  and  $B$  to two, possibly negated, variables the machine can read. The corresponding SAL model is `single_rule.sal` in [7].

$$\begin{aligned}
 R_0 = R_1 = R_2 = R_3 &= \{\text{IF } c_A \wedge c_B \text{ THEN } A := v_A, B := v_B \text{ ENDIF}\}, \text{ where} \\
 c_A &\in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}, \quad \mathcal{D} = \{A, B, A_L, A_R, B_L, B_R, \text{false}, \text{true}\}, \\
 c_B &\in \{(B = b) \mid b \in \mathcal{D}\} \cup \{(B = \bar{b}) \mid b \in \mathcal{D}\}, \quad v_A, v_B \in \{A, \bar{A}, B, \bar{B}, \text{true}, \text{false}\}.
 \end{aligned}$$

To confirm her suspicion, Synthia asked the tool whether there is some instantiation of this template such that the system always reaches a legitimate state in four steps. Note that the interesting property is in fact  $FG(\textit{legitimate})$ , which gets transformed into  $X^4(\textit{legitimate})$ . In about a minute and a half the tool

told Synthia that there is no such instantiation. She also tried to synthesize a solution for  $X^8(\textit{legitimate})$  and  $X^{16}(\textit{legitimate})$ . As expected, the answer was again “no” in 2 and a half and 6 minutes, respectively.

Synthia was convinced that the rules had to implement some way of influencing the conditions of the rules of the neighbors. A possibility is preventing the left neighbor from making a move by having  $B_R$  as a condition of the rule. She asked the solver to complete the following small variation of the previous template. The corresponding SAL model is `single_rule_BR.sal` in [7].

$$R_0 = R_1 = R_2 = R_3 = \{\text{IF } c_A \wedge B_R \text{ THEN } A := v_A, B := v_B \text{ ENDIF}\}, \text{ where}$$

$$c_A \in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}, \quad \mathcal{D} = \{A, B, A_L, A_R, B_L, B_R, \textit{false}, \textit{true}\}, \text{ and}$$

$$v_A, v_B \in \{A, \bar{A}, B, \bar{B}, \textit{true}, \textit{false}\}.$$

After getting a negative answer for both  $X^4(\textit{legitimate})$  and  $X^{16}(\textit{legitimate})$  in less than 5 seconds, Synthia realized that the symmetry of the rules has to be broken somehow since otherwise the states where  $\forall i \in \{0, \dots, 3\} : B_i = \textit{false}$  would not have a successor. A possibility is to fix the value of  $B$  in machine 3 to *true*. She tried that and, additionally, fixing the value of  $B$  in machine 0 to value *false*, getting a negative answer in both cases.

To gain more intuition into the problem, Synthia tried to synthesize a solution assuming a particular initial state, changing the previous template to obtain the following (the corresponding SAL model is `single_rule_B_blocks_initialized.sal` in [7]):

$$B_0 \text{ initialized to } 1$$

$$A_0, A_1, A_2, A_3, B_1, B_2, B_3 \text{ initialized to } 0$$

$$R_0 = R_1 = R_2 = R_3 = \{\text{IF } c_A \wedge B_R \text{ THEN } A := v_A, B := v_B \text{ ENDIF}\}, \text{ where}$$

$$c_A \in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}, \quad \mathcal{D} = \{A, B, A_L, A_R, B_L, B_R, \textit{false}, \textit{true}\},$$

$$v_A, v_B \in \{A, \bar{A}, B, \bar{B}, \textit{true}, \textit{false}\}.$$

The enforced property was again  $X^4(\textit{legitimate})$ . The answer of the solver, in less than 2 seconds, was

$$B_0 \text{ initialized to } 1$$

$$A_0, A_1, A_2, A_3, B_1, B_2, B_3 \text{ initialized to } 0$$

$$R_0 = R_1 = R_2 = R_3 = \{\text{IF } B_R \text{ THEN } A := B_R, B := A \text{ ENDIF}\}$$

Synthia knew that this could not be generalized, since her former attempt to synthesize a solution had failed. She tried to verify the previous solution for the property  $X^4(\textit{legitimate})$  using symbolic model checking. It worked. The next step then, is to test  $FG(\textit{legitimate})$ . SAL returned a counterexample of length 10. Also, simulating by hand the execution of the previous complete model helped Synthia to get convinced that a solution where  $R_0 = R_1 = R_2 = R_3$  could not exist, although she did not worry about formally proving it.

## 4.2 Two rules per machine

Synthia extended the template to have two rules per machine. The corresponding SAL model is `two_rules_general.sal` in [7].

$$R_i = \{\text{IF } c_{A,i,1} \wedge c_{B,i,1} \text{ THEN } A := v_{A,i,1}, B := v_{B,i,1} \text{ ENDIF},$$

$$\text{IF } c_{A,i,2} \wedge c_{B,i,2} \text{ THEN } A := v_{A,i,2}, B := v_{B,i,2} \text{ ENDIF}\}, \text{ where}$$

$$c_{A,i,j} \in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}, \quad \mathcal{D} = \{A, B, A_L, A_R, B_L, B_R, \textit{false}, \textit{true}\},$$

$$c_{B,i,j} \in \{(B = b) \mid b \in \mathcal{D}\} \cup \{(B = \bar{b}) \mid b \in \mathcal{D}\}, \quad v_{A,i,j}, v_{B,i,j} \in \{A, \bar{A}, B, \bar{B}, \textit{true}, \textit{false}\}$$

$$\text{for every } i \in \{0, 1, 2, 3\}, j \in \{1, 2\}.$$



Once again, the enforced property was  $X^4(\textit{legitimate})$ , the bounded version of  $FG(\textit{legitimate})$ . The solver did not produce a solution in 10 minutes and Synthia lost her patience. It is important to remark here that the election for the value of the constant  $c$  in the strengtening  $X^c(\phi)$  of  $FG(\phi)$  may be crucial to obtain a solution. With this fact in mind, Synthia tried  $X^{12}(\textit{legitimate})$  with no success. The 2QBF  $\forall\exists$  instance corresponding to formula 2 for this system template and the property  $X^4(\textit{legitimate})$  has 128 universal and 23273 existential variables. The QBF preprocessor Bloqqer [2] reduced the number of clauses from 91714 to 15338.

Synthia knew that this template was too general, and thus many of its instances are either equivalent to some other instance or can be trivially discarded. The goal then, was to find a more restrictive template and reduce the number of universal variables in the resulting 2QBF problem. A simple option is considering the restriction of the previous template where  $R_1 = R_2$ . This requires trivial changes with respect to the previous template and is encoded in `two_rules_reduced.sal`. The resulting QBF formula for the property  $X^4(\textit{legitimate})$ , after being preprocessed with bloqqer, has 96 universal variables, 19059 existential variables, and 16099 clauses. As before, Synthia gives up after waiting for around 15 minutes.

After realizing that the tool will not give her all the answers, Synthia decided to go back to the idea of using  $B_R$  to block the left neighbor from making a move. Note that that case is not covered in the previous template. Additionally, she kept the restriction that  $R_1 = R_2$ . The resulting SAL model is `two_rules_reduced_BR.sal`.

$B_1$  is fixed to have value *false*, and the set of rules for each of the four machines is defined as follows:

$$\begin{aligned}
R_i &= \{ \text{IF } c_{A,k,1} \wedge B_R \text{ THEN } A := v_{A,k,1}, B := v_{B,k,1} \text{ ENDIF,} \\
&\quad \text{IF } c_{A,k,2} \wedge c_{B,k,2} \text{ THEN } A := v_{A,k,2}, B := v_{B,k,2} \text{ ENDIF} \}, \text{ where} \\
k &= 2 \text{ if } i = 3 \text{ and } k = i, \text{ otherwise and} \\
c_{A,k,j} &\in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}, \quad \mathcal{D} = \{A, B, A_L, A_R, B_L, B_R, \textit{false}, \textit{true}\}, \\
c_{B,k,2} &\in \{(B = b) \mid b \in \mathcal{D}\} \cup \{(B = \bar{b}) \mid b \in \mathcal{D}\}, \quad v_{A,k,j}, v_{B,k,j} \in \{A, \bar{A}, B, \bar{B}, \textit{true}, \textit{false}\} \\
&\text{for every } i \in \{0, 1, 2, 3\}, j \in \{1, 2\}.
\end{aligned}$$

After not obtaining a solution from the QBF solver in 10 minutes, Synthia decided to simplify her template even more, by restricting the domains of the conditions and the assignments of the rules as  $c_{A,k,j} \in \{(A = b) \mid b \in \mathcal{D}\} \cup \{(A = \bar{b}) \mid b \in \mathcal{D}\}$ , where  $\mathcal{D} = \{B, A_L, A_R\}$ , and  $v_{A,k,j}, v_{B,k,j} \in \{A, \bar{A}, \textit{true}, \textit{false}\}$ . Again, the enforced property is  $X^4(\textit{legitimate})$ . In 5 minutes the tool reported that there was no instance of the template satisfying the property.

Synthia was confused at this point. After using the tool to synthesize a solution for a particular case, Synthia realized that, not only the value of  $B$  in machine 0, but also the value of  $B$  in machine 3, must be fixed. The corresponding SAL model is `two_rules_reduced_BR_simpl_values.sal`. When enforcing  $X^4(\textit{legitimate})$ , the tool found an instance not satisfying  $FG(\textit{legitimate})$ , which was easily detected when trying to formally verify it. For the case of  $X^{12}(\textit{legitimate})$ , the tool found the following solution:

$B_1$  is fixed to have value *false*,  $B_4$  is fixed to have value *true*, and the set of rules for each of the four machines is defined as follows:

$$\begin{aligned}
R_0 &= \{ \text{IF } (A = B) \wedge B_R \text{ THEN } A := \bar{A} \text{ ENDIF,} \\
&\quad \text{IF } (A = B) \text{ THEN } A := \bar{A} \text{ ENDIF} \}, \\
R_3 &= \{ \text{IF } A \neq A_L \wedge B = A_R \text{ THEN } A := \bar{A} \text{ ENDIF} \} \\
R_1 = R_2 &= \{ \text{IF } A \neq A_r \wedge B_R \text{ THEN } A := \bar{A}, B := \textit{false} \text{ ENDIF} \}
\end{aligned}$$

The first thing that Synthia did was verifying  $FG(\textit{legitimate})$  to make sure that the synthesized solution preserves stabilization. This property could be proved by SAL using symbolic model checking in a few seconds. Also, the solution was checked for deadlock states. However, after inspecting the solution, Synthia realized that it is not fair. It was confirmed by trying to verify the more complex property  $FG(\textit{legitimate} \wedge M)$ , where  $M$  is a predicate that is satisfied iff every machine made a move at some step in the past, since symbolic model checking produced a counterexample.

Recall that the notion of fairness required in our example is that, for every pair of machines  $m_i, m_j$ , there is a sequence of steps of the system going from every legitimate state where a rule of  $m_i$  is enabled to a legitimate state where a rule of  $m_j$  is enabled. Note that this property requires the existence of an execution, and hence it intuitively corresponds to the  $E$  (Exists) temporal operator in Computational Tree Logic (CTL), and not the  $F$  operator in LTL. In the original problem presented by Dijkstra, the definition of enabled rule did not require a rule to change the current state to be enabled. However, note that every execution of the system with Dijkstra’s definition of enabled can be associated to an execution in our setting. Hence, the property  $FG(\textit{legitimate} \wedge M)$  correctly captures the original fairness condition.

Hence, Synthia used our tool to synthesize a solution for  $X^{12}(\textit{legitimate} \wedge M)$ , and obtained, in less than 30 seconds, the solution presented in Section 2.

The first question that came to Synthia’s mind was whether the synthesized solution could be generalized to  $n$  machines. However, before getting into that, Synthia asked one last question to the synthesis tool: is there any instantiation of the template satisfying  $X^{11}(\textit{legitimate})$ ? The tool quickly answered “no”. Synthia started wondering whether that bound holds for any algorithm satisfying the requirements. She then closed her laptop and grabbed pencil and paper.

## 5 Conclusion and further work

We have presented a practical approach to the synthesis of finite-state distributed systems based in bounded synthesis of LTL properties. Our approach can be seen as a natural first step in the extension of the capabilities of a model checker to synthesis and builds up on the fact that, while synthesizing a complex system from scratch is still unfeasible in practice, the recent progress in QBF solving enables synthesis from human-provided templates.

As further work, we plan to extend an existing model checker such as SAL to have synthesis capabilities. While the SAL language is very appropriate for the modeling of distributed systems, it does not provide specific constructs for describing templates. An important component of this task is the design and implementation of an extension of the SAL language to support definition of templates.

From another perspective, besides experimenting more with our approach, we are interested in leveraging it to the  $k$ -induction and infinite settings. The latter is enabled by the recent progress in  $\exists\forall$  SMT solving. However, more investigation is needed in finding decision procedures for that problem that are well suited for the instances that have to be solved in our setting.

## References

- [1] Rajeev Alur, Rastislav Bodík, Garvit Juniwal, Milo M. K. Martin, Mukund Raghothaman, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak & Abhishek Udupa (2013): *Syntax-guided synthesis*. In: *FMCAD*, pp. 1–17, doi:10.1109/FMCAD.2013.6679385.
- [2] A. Biere, F. Lonsing & M. Seidl (2011): *Blocked Clause Elimination for QBF*. In: *CADE*, pp. 101–115, doi:10.1007/978-3-642-22438-6\_10.

- [3] Roderick Bloem, Robert Könighofer & Martina Seidl (2014): *SAT-Based Synthesis Methods for Safety Specs*. In: *VMCAI*, pp. 1–20, doi:10.1007/978-3-642-54013-4\_1.
- [4] Chih-Hong Cheng, Natarajan Shankar, Harald Ruess & Saddek Bensalem (2013): *EFSMT: A Logical Framework for Cyber-Physical Systems*. *CoRR* abs/1306.3456. Available at <http://arxiv.org/abs/1306.3456>.
- [5] Edsger W. Dijkstra (1974): *Self-stabilizing Systems in Spite of Distributed Control*. *Commun. ACM* 17(11), pp. 643–644, doi:10.1145/361179.361202.
- [6] Bernd Finkbeiner & Sven Schewe (2013): *Bounded synthesis*. *STTT* 15(5-6), pp. 519–539, doi:10.1007/s10009-012-0228-z.
- [7] Adrià Gascón & Ashish Tiwari: *Sal modules of the running example*. Available at: [http://www.csl.sri.com/users/tiwari/software/synth\\_distributed](http://www.csl.sri.com/users/tiwari/software/synth_distributed).
- [8] Adria Gascón & Ashish Tiwari (2014): *A Synthesized Algorithm for Interactive Consistency*. In: *NASA Formal Methods*, pp. 270–284, doi:10.1007/978-3-319-06200-6\_23.
- [9] M. Janota, W. Klieber, J. Marques-Silva & E. M. Clarke (2012): *Solving QBF with Counterexample Guided Refinement*. In: *SAT*, pp. 114–128, doi:10.1007/978-3-642-31612-8\_10.
- [10] Mikolás Janota, Radu Grigore & Joao Marques-Silva (2013): *On QBF Proofs and Preprocessing*. In: *LPAR*, pp. 473–489. Available at 10.1007/978-3-642-45221-5\_32.
- [11] Mikolás Janota, William Klieber, João Marques-Silva & Edmund M. Clarke (2012): *Solving QBF with Counterexample Guided Refinement*. In: *SAT*, pp. 114–128, doi:10.1007/978-3-642-31612-8\_10.
- [12] Annu John, Igor Konnov, Ulrich Schmid, Helmut Veith & Josef Widder (2013): *Parameterized model checking of fault-tolerant distributed algorithms by abstraction*. In: *FMCAD*, pp. 201–209, doi:10.1109/FMCAD.2013.6679411.
- [13] Leslie Lamport, Robert E. Shostak & Marshall C. Pease (1982): *The Byzantine Generals Problem*. *ACM Trans. Program. Lang. Syst.* 4(3), pp. 382–401, doi:10.1145/357172.357176.
- [14] Patrick Lincoln & John M. Rushby (1993): *The Formal Verification of an Algorithm for Interactive Consistency under a Hybrid Fault Model*. In: *CAV*, pp. 292–304, doi:10.1007/3-540-56922-7\_24.
- [15] Florian Lonsing & Armin Biere (2010): *DepQBF: A Dependency-Aware QBF Solver*. *JSAT* 7(2-3), pp. 71–76. Available at [http://jsat.ewi.tudelft.nl/content/volume7/JSAT7\\_6\\_Lonsing.pdf](http://jsat.ewi.tudelft.nl/content/volume7/JSAT7_6_Lonsing.pdf).
- [16] Nancy A. Lynch (1996): *Distributed Algorithms*. Morgan Kaufmann.
- [17] Leonardo Mendonça de Moura, Sam Owre, Harald Rueß, John M. Rushby, Natarajan Shankar, Maria Sorea & Ashish Tiwari (2004): *SAL 2*. In: *CAV*, pp. 496–500, doi:10.1007/978-3-540-27813-9\_45.
- [18] Aina Niemetz, Mathias Preiner, Florian Lonsing, Martina Seidl & Armin Biere (2012): *Resolution-Based Certificate Extraction for QBF - (Tool Presentation)*. In: *SAT*, pp. 430–435, doi:10.1007/978-3-642-31612-8\_33.
- [19] Oswaldo Olivo & E. Allen Emerson (2011): *A More Efficient BDD-Based QBF Solver*. In: *CP*, pp. 675–690, doi:10.1007/978-3-642-23786-7\_51.
- [20] Sam Owre, John M. Rushby, Natarajan Shankar & Friedrich W. von Henke (1993): *Formal Verification for Fault-Tolerant Architectures: Some Lessons Learned*. In: *FME*, pp. 482–500, doi:10.1007/BFb0024663.
- [21] Darsh P. Ranjan, Daijue Tang & Sharad Malik (2004): *A Comparative Study of 2QBF Algorithms*. In: *SAT*. Available at <http://www.satisfiability.org/SAT04/programme/113.pdf>.
- [22] John M. Rushby & Friedrich W. von Henke (1993): *Formal Verification of Algorithms for Critical Systems*. *IEEE Trans. Software Eng.* 19(1), pp. 13–23, doi:10.1109/32.210304.
- [23] (2003): *The SAL intermediate language*. Computer Science Laboratory, SRI International, Menlo Park, CA. <http://sal.csl.sri.com/>.
- [24] Natarajan Shankar (1992): *Mechanical Verification of a Generalized Protocol for Byzantine Fault Tolerant Clock Synchronization*. In: *FTRTFT*, pp. 217–236, doi:10.1007/3-540-55092-5\_12.

- [25] Armando Solar-Lezama, Rodric M. Rabbah, Rastislav Bodík & Kemal Ebcioglu (2005): *Programming by sketching for bit-streaming programs*. In: *PLDI*, pp. 281–294, doi:10.1145/1065010.1065045.
- [26] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodík, Sanjit A. Seshia & Vijay A. Saraswat (2006): *Combinatorial Sketching for Finite Programs*. In: *ASPLOS*, pp. 404–415, doi:10.1145/1168857.1168907.
- [27] Saurabh Srivastava, Sumit Gulwani & Jeffrey S. Foster (2010): *From program verification to program synthesis*. In: *POPL*, pp. 313–326, doi:10.1145/1706299.1706337.
- [28] Abhishek Udupa, Arun Raghavan, Jyotirmoy V. Deshmukh, Sela Mador-Haim, Milo M. K. Martin & Rajeev Alur (2013): *TRANSIT: specifying protocols with concolic snippets*. In: *PLDI*, pp. 287–296, doi:10.1145/2462156.2462174.
- [29] Yinlei Yu & Sharad Malik (2005): *Validating the result of a Quantified Boolean Formula (QBF) solver: theory and practice*. In: *ASP-DAC*, pp. 1047–1051, doi:10.1145/1120725.1120821.