

# 1 Parameterized Synthesis of Self-Stabilizing 2 Protocols in Symmetric Rings

3 Nahal Mirzaie

4 University of Tehran, North Kargar st., Tehran, Iran  
5 mirzaienahal@ut.ac.ir

6 Fathiyeh Faghieh

7 University of Tehran, North Kargar st., Tehran, Iran  
8 f.faghieh@ut.ac.ir

9 Swen Jacobs<sup>1</sup>

10 CISPA Helmholtz Center i.G., Saarbrücken, Germany  
11 jacobs@cispa.saarland  
12  <https://orcid.org/0000-0002-9051-4050>

13 Borzoo Bonakdarpour<sup>2</sup>

14 Iowa State University, 207 Atanasoff Hall, Ames, IA 50011, USA  
15 borzoo@iastate.edu  
16  <https://orcid.org/0000-0003-1800-5419>

## 17 — Abstract —

---

18 *Self-stabilization* in distributed systems is a technique to guarantee *convergence* to a set of  
19 *legitimate states* without external intervention when a transient fault or bad initialization occurs.  
20 Recently, there has been a surge of efforts in designing techniques for automated synthesis of self-  
21 stabilizing algorithms that are correct by construction. Most of these techniques, however, are not  
22 parameterized, meaning that they can only synthesize a solution for a fixed and predetermined  
23 number of processes. In this paper, we report a breakthrough in parameterized synthesis of  
24 self-stabilizing algorithms in symmetric rings. First, we develop tight cutoffs that guarantee (1)  
25 closure in legitimate states, and (2) deadlock-freedom outside the legitimates states. We also  
26 develop a sufficient condition for convergence in *silent* self-stabilizing systems. Since some of  
27 our cutoffs grow with the size of local state space of processes, we also present an automated  
28 technique that significantly increases the scalability of synthesis in symmetric networks. Our  
29 technique is based on SMT-solving and incorporates a loop of synthesis and verification guided  
30 by counterexamples. We have fully implemented our technique and successfully synthesized  
31 solutions to maximal matching, three coloring, and maximal independent set problems.

32 **2012 ACM Subject Classification** Theory of computation → Logic and verification, Computer  
33 systems organization → Dependable and fault-tolerant systems and networks

34 **Keywords and phrases** Parameterized synthesis, Self-stabilization, Formal methods.

35 **Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2018.0

---

<sup>1</sup> This research was supported by the German Research Foundation (DFG) under the project ASDPS (JA 2357/2-1).

<sup>2</sup> This research has been partially supported by the NSF SaTC-1813388 and a grant from Iowa State University.



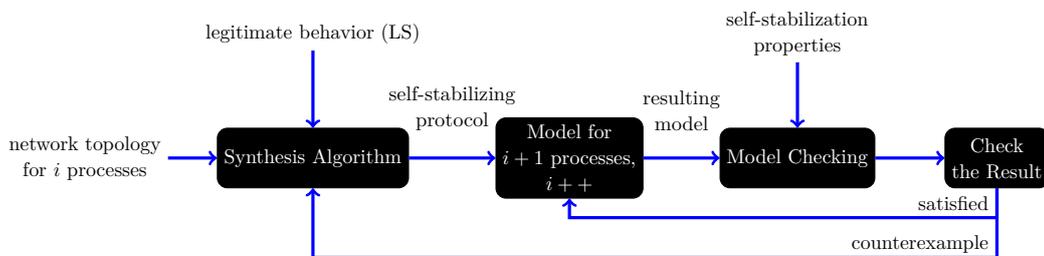
## 1 Introduction

Program *synthesis* (often called the “*holy grail*” of computer science) is the problem of automated generation of a computer program from a formally specified set of properties. The program generated in this fashion is guaranteed to be correct by construction. Program synthesis is known to be computationally intractable and, thus, is usually used to deal with small but intricate components of a system, e.g., concurrent/distributed algorithms that may exhibit obscure corner cases, where reasoning about their correctness is not straightforward.

Dijkstra [3] introduced the notion of *self-stabilization* in distributed systems, where the system always converges to a good behavior even if it is arbitrarily initialized or is subject to transient faults. Proof of self-stabilization is, however, often much more complex than what it initially seems like. Dijkstra himself published the proof of correctness of his seminal 3-state machine solution 12 years later [4]. This means that program synthesis can play a prime role in designing and reasoning about the correctness of self-stabilizing algorithms. In [8–12], we introduced a set of algorithms and tools for synthesizing self-stabilizing protocols. Our techniques take as input the network topology, timing model (asynchronous or synchronous), the good behavior of the protocol (either explicitly as a set of *legitimate states* or implicitly as a set of temporal logic formulas), type of symmetry, and type of stabilization (e.g., strong, weak, monotonic, ideal) and generate a set of first-order modulo theory (SMT) constraints. Then, an SMT-solver solves these constraints and, if satisfiable, produces a model that respects the input specification. Our tool ASSESS [10] has successfully synthesized complex algorithms such as Raymond’s distributed mutual exclusion [24], Dijkstra’s token ring [3] (for both three and four state machines), maximal matching [23], weak stabilizing token circulation in anonymous networks [2], and the three coloring problem [14]. Our algorithms are *complete* for a predetermined fixed number of processes; i.e., if they fail to find a solution to the synthesis problem, then there does not exist one. This completeness, however, comes at a big cost which is scalability. That is, for most instances, we could only synthesize solutions up to 5 processes at best.

In this paper, our goal is to address scalability as well as the shortcoming that the previous work can synthesize only a fixed and predetermined number of processes. To this end, we focus on automated synthesis of self-stabilizing protocols in *symmetric* and *parameterized* rings, where an unbounded number of processes exhibit identical behavior. We make two main contributions. First, we show how to solve the *parameterized synthesis problem* based on the notion of *cutoffs* [6] that can guarantee properties of distributed systems of arbitrary size by considering only systems of up to a certain fixed size  $c \in \mathbb{N}$ , and augmented by a sound but incomplete abstraction-based synthesis approach for properties that are known to be undecidable. In particular, we provide:

- cutoffs for the closure and deadlock-freedom properties, under the assumption that the



■ **Figure 1** SMT-based Counterexample-Guided Synthesis Technique

73 set of legitimate states is defined by a conjunction of predicates on the local state of  
 74 processes; we show that smaller cutoffs are possible under additional assumptions, and  
 75 that all our cutoffs are tight under the assumptions we consider.

76 ■ an abstraction-based method for the convergence property, which is known to be unde-  
 77 cidable in general [19, 21]; we show how, for the class of silent algorithms, a sufficient  
 78 condition for convergence of the parameterized system can be efficiently checked on a  
 79 finite system that over-approximates the behavior of systems of arbitrary size.

80 Note that these results can be used for both synthesis and verification. A drawback of our  
 81 cutoffs is that they are quadratic in the state space of a single process, so even with a cutoff,  
 82 we need synthesis methods that scale to a large number of processes. Thus, as our second  
 83 contribution, we propose a counterexample-guided synthesis technique that exploits our  
 84 symmetry assumption. More specifically, it consists of four steps (see Fig. 1):

- 85 1. First, we *synthesize* a solution for a small network of  $i$  processes using existing techniques;
- 86 2. Next, we trivially generalize this solution to a larger network with  $i + 1$  processes;
- 87 3. Then, we *verify* this solution using a model checker, and
- 88 4. If verification succeeds, we return to step 2 to attempt a larger network. Otherwise, we  
 89 obtain a counterexample that is added as a negative constraint for the synthesis algorithm,  
 90 and we return to step 1 for another round of synthesis with limited search space.

91 Using this approach and our cutoff results, we successfully synthesized parameterized self-  
 92 stabilizing protocols for well-known problems including three coloring, maximal matching,  
 93 and maximal independent set in less than 10 minutes. To our knowledge, this is the first  
 94 instance of such parameterized synthesis.

95 **Organization** The rest of the paper is organized as follows. Section 2 introduces the  
 96 preliminary concepts. In Section 3, we present the formal statement of our synthesis problem.  
 97 The parameterized correctness results are presented in Section 4, while our counterexample-  
 98 guided synthesis approach is presented in Section 5. Experimental results and case studies  
 99 are reported in Section 6. Related work is discussed in Section 7, and finally, we make  
 100 concluding remarks and discuss future work in Section 8.

## 101 2 Preliminaries

### 102 2.1 Distributed Programs

103 Most self-stabilizing algorithms are defined in the shared-memory model. Assume  $V$  to be  
 104 the set of all variables in the system, where each variable  $v \in V$  has a finite domain  $D_v$ . We  
 105 define a *state*  $s$  as a valuation of each variable in  $V$  by a value in its domain. The set of all  
 106 possible states is called the *state space*, and represented by  $S$ . A *transition* is defined as an  
 107 ordered pair  $(s_0, s_1)$ , where  $s_0, s_1 \in S$ . We denote the value of a variable  $v$  in state  $s$  by  $v(s)$ .

108 ► **Definition 1.** A *process*  $\pi$  is a tuple  $\langle R_\pi, W_\pi, T_\pi \rangle$ , where

- 109 ■  $R_\pi \subseteq V$  is the set of variables that  $\pi$  can read their values and is called the *read-set* of  $\pi$ ;
- 110 ■  $W_\pi \subseteq R_\pi$  is the set of variables that  $\pi$  can write to, and is called the *write-set* of  $\pi$ , and
- 111 ■  $T_\pi$  is the set of transitions of  $\pi$ , where for each transition  $(s_0, s_1) \in T_\pi$  and each variable  
 112  $v$ , such that  $v(s_0) \neq v(s_1)$ , we have  $v \in W_\pi$ . ◀

113 The third condition imposes the constraint that a process can only change the value of a  
 114 variable in its write-set, and the second condition states that this change cannot be blind. A  
 115 process  $\pi = \langle R_\pi, W_\pi, T_\pi \rangle$  is called *enabled* in state  $s_0$ , if there exists a state  $s_1$ , such that

116  $(s_0, s_1) \in T_\pi$ . The *local state space* of  $\pi$  is the set of all possible valuations of the variables  
 117 that  $\pi$  can read:  $S_\pi = \prod_{v \in R_\pi} D_v$

118 ► **Definition 2.** A *distributed program* is a tuple  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$ , where

119 ■  $P_{\mathcal{D}}$  is a set of processes over a common set  $V$  of variables, such that:

- 120 ■ for any two distinct processes  $\pi_1, \pi_2 \in P_{\mathcal{D}}$ , we have  $W_{\pi_1} \cap W_{\pi_2} = \emptyset$ ;
- 121 ■ for each process  $\pi \in P_{\mathcal{D}}$  and each transition  $(s_0, s_1) \in T_\pi$ , the following *read restriction*  
 122 holds:

$$123 \quad \forall s'_0, s'_1 : ((\forall v \in R_\pi : (v(s_0) = v(s'_0) \wedge v(s_1) = v(s'_1))) \wedge$$

$$124 \quad (\forall v \notin R_\pi : v(s'_0) = v(s'_1))) \implies (s'_0, s'_1) \in T_\pi \quad (1)$$

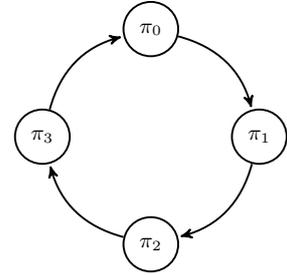
125 ■  $T_{\mathcal{D}}$  is the set of transitions and is the union of transitions of all processes:

$$T_{\mathcal{D}} = \bigcup_{\pi \in P_{\mathcal{D}}} T_\pi$$

126

127 Intuitively, the read restriction in Definition 2 imposes the constraint that for each process  $\pi$ ,  
 128 each transition in  $T_\pi$  depends only on the variables in the read-set of  $\pi$ . Thus, each transition  
 129 is an equivalence class in  $T_{\mathcal{D}}$ , which we call a *group* of transitions. The key consequence of  
 130 read restrictions is that during synthesis, if a transition is included (respectively, excluded)  
 131 in  $T_{\mathcal{D}}$ , then its corresponding group must also be included (respectively, excluded) in  $T_{\mathcal{D}}$ .

132 **Example.** We use the problem of distributed self-stabilizing  
 133 *one-bit maximal matching* as a running example to describe  
 134 the concepts throughout the paper. Consider a ring of 4 pro-  
 135 cesses (see Fig. 2), and let  $V = \{x_0, x_1, x_2, x_3\}$  be the set of  
 136 variables, where each  $x_i$  is a Boolean variable with domain  
 137  $\{\mathbf{F}, \mathbf{T}\}$ . Let  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  be a distributed program, where  
 138  $P_{\mathcal{D}} = \{\pi_0, \pi_1, \pi_2, \pi_3\}$ . Each process  $\pi_i$  ( $0 \leq i \leq 3$ ) can write to  
 139 the variable  $x_i$  (i.e.,  $W_{\pi_i} = \{x_i\}$ ), and read the variables of its  
 140 own and its neighbors ( $R_{\pi_i} = \{x_i, x_{(i+1) \bmod 4}, x_{(i-1) \bmod 4}\}$ ).  
 141 Notice that following Definition 2 and read/write restrictions of  
 142  $\pi_0$ , (arbitrary) transitions



143 ■ **Figure 2** One-bit maximal matching example.

143  $t_1 = ([x_0 = \mathbf{F}, x_1 = \mathbf{F}, x_2 = \mathbf{F}, x_3 = \mathbf{F}], [x_0 = \mathbf{T}, x_1 = \mathbf{F}, x_2 = \mathbf{F}, x_3 = \mathbf{F}])$

145  $t_2 = ([x_0 = \mathbf{F}, x_1 = \mathbf{F}, x_2 = \mathbf{T}, x_3 = \mathbf{F}], [x_0 = \mathbf{T}, x_1 = \mathbf{F}, x_2 = \mathbf{T}, x_3 = \mathbf{F}])$

146 are in the same group. The reason is that  $\pi_0$  cannot read  $x_2$ , and if, for example,  $t_1$  is  
 147 included in the set of transitions, while  $t_2$  is not, it implies that the execution in process  $\pi_0$   
 148 depends on the value of  $x_2$ , which is not possible.

► **Definition 3.** An *uninterpreted local function* for a process maps the *local state space* of a  
 process to a domain  $D_{lf}$ . The interpretation of an uninterpreted local function for a process  
 $\pi$  is a function:

$$S_\pi \rightarrow D_{lf}$$

149 where  $S_\pi$  is the local state space of  $\pi$ . 126

150 In the sequel, we use “uninterpreted functions” to refer to uninterpreted local functions.

**Example.** To formulate the requirements in the one-bit maximal matching example, we assume each process  $\pi_i$  is associated with an uninterpreted local function, called *match*, with the domain  $D_{match_i} = \{l, r, n\}$ , where  $l$ ,  $r$ , and  $n$  correspond to the cases where the process is matched to its left, right, and no neighbor (self-matched), respectively. The interpretation of  $match_i$  is a function:

$$(match_i)_I : \{F, T\} \times \{F, T\} \times \{F, T\} \rightarrow \{l, r, n\}$$

151 In other words, the value of  $match_i$  depends on the value of the process and its neighbors'  
152 Boolean variables.

► **Definition 4.** A *local predicate* of a process maps the *local state space* of a process to a Boolean:

$$S_\pi \rightarrow \{F, T\}$$

153

154 We use this definition to define legitimate states in Section 2.3.

### 155 2.1.1 Network Topology

156 A topology specifies the communication model of a distributed program.

157 ► **Definition 5.** A *topology* is a tuple  $\mathcal{T} = \langle V, |P_{\mathcal{T}}|, R_{\mathcal{T}}, W_{\mathcal{T}} \rangle$ , where  
158 ■  $V$  is a finite set of finite-domain discrete variables,  
159 ■  $|P_{\mathcal{T}}| \in \mathbb{N}_{\geq 1}$  is the number of processes,  
160 ■  $R_{\mathcal{T}}$  is a mapping  $\{0 \dots |P_{\mathcal{T}}| - 1\} \rightarrow 2^V$  from a process index to its read-set,  
161 ■  $W_{\mathcal{T}}$  is a mapping  $\{0 \dots |P_{\mathcal{T}}| - 1\} \rightarrow 2^V$  from a process index to its write-set, such that  
162  $W_{\mathcal{T}}(i) \subseteq R_{\mathcal{T}}(i)$ , for all  $i$  ( $0 \leq i \leq |P_{\mathcal{T}}| - 1$ ). ◀

163 **Example.** The topology of our maximal matching problem is a tuple  $\langle V, |P_{\mathcal{T}}|, R_{\mathcal{T}}, W_{\mathcal{T}} \rangle$ :

164 ■  $V = \{x_0, x_1, x_2, x_3\}$ , with domains  $D_{x_0} = D_{x_1} = D_{x_2} = D_{x_3} = \{T, F\}$ ,  
165 ■  $|P_{\mathcal{T}}| = 4$ ,  
166 ■  $R_{\mathcal{T}}(0) = \{x_0, x_1, x_3\}$ ,  $R_{\mathcal{T}}(1) = \{x_1, x_2, x_0\}$ ,  
167  $R_{\mathcal{T}}(2) = \{x_2, x_3, x_1\}$ ,  $R_{\mathcal{T}}(3) = \{x_3, x_0, x_2\}$ , and  
168 ■  $W_{\mathcal{T}}(0) = \{x_0\}$ ,  $W_{\mathcal{T}}(1) = \{x_1\}$ ,  $W_{\mathcal{T}}(2) = \{x_2\}$ , and  $W_{\mathcal{T}}(3) = \{x_3\}$ .

169 ► **Definition 6.** A distributed program  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  has topology  
170  $\mathcal{T} = \langle V, |P_{\mathcal{T}}|, R_{\mathcal{T}}, W_{\mathcal{T}} \rangle$  iff

171 ■ each process  $\pi \in P_{\mathcal{D}}$  is defined over  $V$   
172 ■  $|P_{\mathcal{D}}| = |P_{\mathcal{T}}|$   
173 ■ there is a mapping  $g : \{0 \dots |P_{\mathcal{T}}| - 1\} \rightarrow P_{\mathcal{D}}$  such that

$$\forall i \in \{0 \dots |P_{\mathcal{T}}| - 1\} : (R_{\mathcal{T}}(i) = R_{g(i)}) \wedge (W_{\mathcal{T}}(i) = W_{g(i)})$$

173

## 174 2.2 Symmetric Networks

175 Roughly speaking, a topology is symmetric, if the read-set and write-set of any two distinct  
 176 processes can be swapped (i.e., there is a bijection that maps read/write variables of a process  
 177 to another).

178 ► **Definition 7.** A topology  $\mathcal{T} = \langle V, |P_{\mathcal{T}}|, R_{\mathcal{T}}, W_{\mathcal{T}} \rangle$  is *symmetric*, iff for any distinct  
 179  $i, j \in \{0 \dots |P_{\mathcal{T}}| - 1\}$ , there exists

180 ■ a bijection  $f : R_{\mathcal{T}}(i) \rightarrow R_{\mathcal{T}}(j)$ , such that  $\forall v \in R_{\mathcal{T}}(i) : D_v = D_{f(v)}$ , and

181 ■ a bijection  $g : W_{\mathcal{T}}(i) \rightarrow W_{\mathcal{T}}(j)$ , such that  $\forall v \in W_{\mathcal{T}}(i) : D_v = D_{g(v)}$ . ◀

182 We call a symmetric topology a (bi-directional) *ring* (of size  $k = |P_{\mathcal{T}}|$ ) if for every  
 183  $i \in \{0 \dots |P_{\mathcal{T}}| - 1\}$ , we have  $R_{\mathcal{T}}(i) = W_{\mathcal{T}}(i - 1 \bmod k) \cup W_{\mathcal{T}}(i) \cup W_{\mathcal{T}}(i + 1 \bmod k)$ . Since  
 184 in this paper we only deal with rings, to simplify notation throughout the paper, arithmetic  
 185 on process indices is implicitly modulo the size of the ring.

186 **Example.** The topology of our one-bit maximal matching example is symmetric, and a ring  
 187 of size 4 (Fig. 2). For any two numbers  $i$  and  $j$ , function  $g$  is the mapping from  $x_i$  to a  $x_j$ ,  
 188 and function  $f$  maps  $x_i \mapsto x_j$ ,  $x_{(i+1)} \mapsto x_{(j+1)}$ , and  $x_{(i-1)} \mapsto x_{(j-1)}$ .

189 ► **Definition 8.** A distributed program  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  is called *symmetric* iff

190 ■ it has a symmetric topology, and

191 ■ for any two distinct processes  $\pi, \pi' \in P_{\mathcal{D}}$ , the following condition holds:

$$192 \quad \left( \forall (s_0, s_1) \in T_{\pi} : \exists (s'_0, s'_1) \in T_{\pi'} : \right. \\ \left. \left( \forall v \in R_{\pi} : (v(s_0) = f(v)(s'_0)) \right) \wedge \left( \forall v \in W_{\pi} : (v(s_1) = g(v)(s'_1)) \right) \right) \quad (2)$$

193 where  $f$  and  $g$  are the functions defined in Definition 7. ◀

194 In other words, in a symmetric distributed program the read- and write-sets of all processes  
 195 are identical up to renaming, and so are their transitions. Therefore, we also write  $\mathcal{T}^{\pi}$  for a  
 196 symmetric distributed program that has topology  $\mathcal{T}$  and where all processes are identical up  
 197 to renaming to  $\pi$ .

## 198 2.3 Self-Stabilization

199 Given a subset of the state space, called the set of *legitimate states* (denoted by  $LS$ ), a  
 200 *self-stabilizing* [3] program always recovers to a state in  $LS$  from any arbitrary state (e.g.,  
 201 due to bad initialization or occurrence of transient faults) in a finite number of steps, and  
 202 stays in  $LS$  thereafter.

203 ► **Definition 9.** A *computation* of  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  is an infinite sequence of states  $\bar{s} = s_0 s_1 \dots$ ,  
 204 such that: (1) for all  $i \geq 0$ , we have  $(s_i, s_{i+1}) \in T_{\mathcal{D}}$ , and (2) if a computation reaches a state  
 205  $s_i$ , from where there is no state  $\mathfrak{s} \neq s_i$ , such that  $(s_i, \mathfrak{s}) \in T_{\mathcal{D}}$ , then the computation stutters  
 206 at  $s_i$  indefinitely. Such a computation is called a *terminating computation*. ◀

207 ► **Definition 10.** A distributed program  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  is *self-stabilizing* for a set  $LS$  of  
 208 legitimate states iff

209 1. (*Convergence*) For any computation  $\bar{s} = s_0 s_1 \dots$ , there exists a state  $s_j \in \bar{s}$  ( $j \geq 0$ ),  
 210 such that  $s_j \in LS$ .

211 2. (*Closure*) For any transition  $(s_0, s_1) \in T_{\mathcal{D}}$ , if  $s_0 \in LS$ , then  $s_1 \in LS$ . ◀

212 ▶ **Definition 11.** A distributed program  $\mathcal{D} = \langle P_{\mathcal{D}}, T_{\mathcal{D}} \rangle$  is *silent* with respect to a given  $LS$   
213 if for any transition  $(s_0, s_1) \in T_{\mathcal{D}}$ , if  $s_0 \in LS$ , then  $s_1 = s_0$ . ◀

214 ▶ **Definition 12.** A set of legitimate states is *locally defined* if it can be defined by

$$215 \quad s \in LS \quad \text{if and only if} \quad \forall i \in \{0 \dots |P_{\mathcal{T}}| - 1\} : LS_i(s),$$

216 where  $LS_i$  is a predicate on the read-set of process  $\pi_i$ . ◀

217 **Example.** In a directed graph, a maximal matching is a maximal set of edges, in which  
218 no two edges share a common vertex. In a ring topology, each process can be matched  
219 to one of its two adjacent processes. To formulate this requirement, we assume each  
220 process  $\pi_i$  is associated with a local uninterpreted function, called  $match_i$ , with the domain  
221  $D_{match_i} = \{l, r, n\}$ .  $LS$  can be locally defined with

$$222 \quad LS_i = \{s \mid (match_{i-1}(\Pi_{R_{\pi_{i-1}}}(s)) = r \wedge match_i(\Pi_{R_{\pi_i}}(s)) = l \wedge match_{i+1}(\Pi_{R_{\pi_{i+1}}}(s)) = n) \vee$$

$$223 \quad (match_{i-1}(\Pi_{R_{\pi_{i-1}}}(s)) = n \wedge match_i(\Pi_{R_{\pi_i}}(s)) = r \wedge match_{i+1}(\Pi_{R_{\pi_{i+1}}}(s)) = l) \vee$$

$$224 \quad (match_{i-1}(\Pi_{R_{\pi_{i-1}}}(s)) = l \wedge match_i(\Pi_{R_{\pi_i}}(s)) = n \wedge match_{i+1}(\Pi_{R_{\pi_{i+1}}}(s)) = r) \vee$$

$$225 \quad (match_{i-1}(\Pi_{R_{\pi_{i-1}}}(s)) = l \wedge match_i(\Pi_{R_{\pi_i}}(s)) = r \wedge match_{i+1}(\Pi_{R_{\pi_{i+1}}}(s)) = l) \vee$$

$$226 \quad (match_{i-1}(\Pi_{R_{\pi_{i-1}}}(s)) = r \wedge match_i(\Pi_{R_{\pi_i}}(s)) = l \wedge match_{i+1}(\Pi_{R_{\pi_{i+1}}}(s)) = r)\}$$

The system is in its legitimate state, if and only if all processes are in their local legitimate states. For example, in a ring of size three with the set of processes  $P = \{\pi_0, \pi_1, \pi_2\}$ , the set of legitimate states can be formulated as the following:

$$\{s \mid LS_0(s) \wedge LS_1(s) \wedge LS_2(s)\}$$

228 Note how uninterpreted functions can be used to easily express  $LS$ . Without  $match_i$ , the  
229 user has to explicitly specify the cases where a process is matched to its left, right or itself,  
230 using the Boolean variables of its own and its adjacent processes (its read set).

### 231 3 Problem Statement

232 Our goal is to propose an automated method for parameterized synthesis of silent self-  
233 stabilizing protocols in symmetric ring networks. That is, we consider a problem where the  
234 size of the topology is a parameter, and we want to automatically synthesize the transition  
235 predicate and the interpretation of the uninterpreted function of each process, such that the  
236 resulting distributed program is silent self-stabilizing for any value of the parameter.

237 Formally, a *parameterized topology* is a sequence of symmetric topologies  $\mathcal{T}_1, \mathcal{T}_2, \dots$ , where  
238 for all  $n$  we have  $|P_{\mathcal{T}_n}| = n$  and bijections read-sets and write-sets, as required in Definition 7,  
239 also exist between process indices from different elements of the sequence. A *parameterized*  
240 *program* is a sequence of symmetric distributed programs  $\mathcal{D}_1, \mathcal{D}_2, \dots$  such that  $\mathcal{D}_i = \mathcal{T}_i^\pi$  for  
241 a parameterized topology  $\mathcal{T}_1, \mathcal{T}_2, \dots$ , and some process  $\pi$ .

242 The *parameterized synthesis problem* takes as input:

- 243 ■ a parameterized topology, and
- 244 ■ a set of locally defined legitimate states  $LS$ ,
- 245 and generates as output:

246 ■ a process  $\pi$  such that for every element  $\mathcal{T}_n$  of the topology, the program  $\mathcal{D}_n = \mathcal{T}_n^\pi$  is  
 247 self-stabilizing to  $LS$ .

248 ► **Definition 13.** For a given parameterized topology and a property under consideration, a  
 249 *cutoff* is a natural number  $c$  such that for any given process  $\pi$  and a locally defined  $LS$  the  
 250 following holds:  $\mathcal{D}_n = \mathcal{T}_n^\pi$  satisfies the property wrt.  $LS$  for all  $n \in \mathbb{N}$  iff  $\mathcal{D}_i = \mathcal{T}_i^\pi$  satisfies  
 251 the property wrt.  $LS$  for all  $i \in \{1 \dots c\}$ . ◀

252 Note that cutoffs can be used for both parameterized verification and synthesis. In Section 4,  
 253 we will present cutoffs for two properties: i) closure, and ii) the absence of deadlocks outside  
 254 of  $LS$ . Moreover, we will introduce an abstraction-based method that can be combined with  
 255 the cutoffs to solve the parameterized synthesis problem.

## 256 4 Parameterized Synthesis of Self-Stabilization in Symmetric Rings

257 In this section, we show how to reduce reasoning about parameterized programs to reasoning  
 258 about a finite number of finite programs. To prove self-stabilization, we need to prove that  
 259 the algorithm has the two properties of closure and convergence. We split the latter into two  
 260 properties: (1) the absence of deadlocks outside of  $LS$ , and (2) the absence of cycles outside  
 261 of  $LS$ . In the following, we provide tight cutoffs for closure and deadlocks outside of  $LS$ , as  
 262 well as a sound abstraction to prove the absence of cycles outside of  $LS$ . Finally, we provide  
 263 our main theorem that combines these results into a method for parameterized synthesis of  
 264 self-stabilizing algorithms in rings.

### 265 4.1 Cutoffs for Closure

266 Assume that the write-set of each process has  $l$  valuations. In other words, if process  $\pi_i$  has  
 267  $W_{\mathcal{T}}(i) = \{v_1, \dots, v_n\}$ , then  $l = |D_{v_1}| \times \dots \times |D_{v_n}|$ .

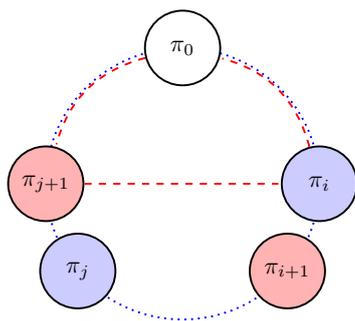
268 ► **Lemma 1.** For self-stabilizing algorithms on a ring topology, the following are cutoffs for  
 269 the closure property:

- 270 ■  $c = l^2 + 1$ , if  $LS$  is locally defined;
- 271 ■  $c = l + 1$ , if  $LS$  is locally defined and  $LS_i$  only depends on  $W_{\mathcal{T}}(i)$  and  $W_{\mathcal{T}}(i + 1)$ , and
- 272 ■  $c = 3$ , if  $LS$  is locally defined and  $LS_i$  only depends on  $W_{\mathcal{T}}(i)$ .

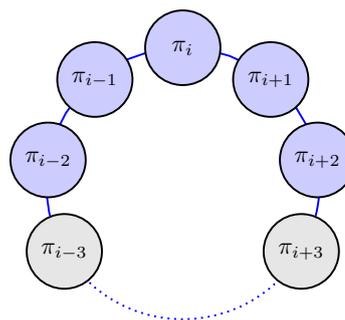
273 All of the cutoffs are tight under their respective assumptions.

274 **Proof.** We provide the proof idea for the first case. Consider a ring of size  $M > l^2 + 1$ .  
 275 Assume there exists a transition from  $s \in LS$  to  $s' \notin LS$ . WLOG, assume this action is taken  
 276 by  $\pi_0$ . Now, consider the  $M - 1$  pairs of consecutive processes  $(\pi_i, \pi_{i+1})$ , where  $i \in [0, M - 2]$ .  
 277 Note that  $\pi_0$  appears in just one tuple. Based on the pigeonhole principle, at least two  
 278 of these pairs of processes have the same valuation of their write-sets in  $s$ , since we have  
 279  $M - 1 \geq l^2 + 1$  tuples and only  $l^2$  possible valuations of the write-sets of a tuple. Assume that  
 280  $(\pi_i, \pi_{i+1})$  and  $(\pi_j, \pi_{j+1})$  have the same valuation of their write-sets. Then we can consider a  
 281 smaller ring composed of  $\pi_0, \dots, \pi_i, \pi_{j+1}, \dots, \pi_M$  with local valuations as in state  $s$  (Fig. 3).  
 282 Note that this construction does not change the valuations of the read-sets of the remaining  
 283 processes, and hence, the smaller ring will also be in a state in  $LS$ , and  $\pi_0$  can still take a  
 284 transition that leads to a state outside of  $LS$ . If necessary, we can repeat the removal of  
 285 processes until we arrive at a ring of size at most  $c = l^2 + 1$ .

286 To observe tightness, define a total order on pairs of valuations of write-sets, and assume  
 287 that  $LS$  is defined such that it requires the valuations along the ring to be compatible with



■ **Figure 3** Reducing a ring (blue dotted) to a smaller ring (red dashed).



■ **Figure 4** Blue processes act based on the synthesized algorithm and grey (with slanted lines) processes act randomly.

288 successor relation of that order. Then a state can only be in  $LS$  in a ring of size at least  
289  $c = l^2 + 1$ , and therefore violation of closure is not possible in any smaller ring.

290 For the second case, the same construction works even if we only consider single valuations  
291 of write-sets instead of pairs, except for one special case: if the processes with identical  
292 valuation are  $\pi_0$  and  $\pi_1$ , then  $\pi_1$  cannot be removed, as this might make the transition of  $\pi_0$   
293 impossible. Therefore, the cutoff is not  $l$ , but  $l + 1$ . Tightness follows by a similar argument  
294 as above.

295 Finally, in the third case we mainly need to ensure that  $\pi_0$  can still take its transition, as  
296 for every process  $LS_i$  only depends on its own valuation. Thus,  $c = 3$  is sufficient (and in  
297 general necessary, as otherwise the transition may not be possible). ◀

## 298 4.2 Cutoffs for Deadlock Detection

299 ▶ **Lemma 2.** For self-stabilizing algorithms on a ring topology, the following are cutoffs for  
300 the detection of deadlocks outside of  $LS$ :

- 301 ■  $c = l^2 + 1$ , if  $LS$  is locally defined;
- 302 ■  $c = l + 1$ , if  $LS$  is locally defined and transitions of processes only depend on  $W_{\mathcal{T}}(i)$  and  
303  $W_{\mathcal{T}}(i + 1)$  (i.e., the ring is uni-directional), and
- 304 ■  $c = 3$ , if  $LS$  is locally defined and transitions of processes only depend on  $W_{\mathcal{T}}(i)$  (i.e.,  
305 processes are completely independent).

306 All of the cutoffs are tight under their respective assumptions.

307 **Proof.** We only provide the proof idea for the first case, the other cases are variants similar  
308 to those seen above. Consider a ring of size  $M > l^2 + 1$  and assume that with the given  
309 program we have a deadlock state  $s \notin LS$ ; that is  $s \notin LS_i$  for at least one of the processes,  
310 and there is no active transition for any processes. We assume that  $i = 0$ , and by the same  
311 construction as in the proof of Lemma 1 we can obtain a smaller ring where every remaining  
312 process has the same valuation of its read-set. Therefore, the state in the smaller ring is not  
313 in  $LS_0$ , and no transitions are enabled. Tightness can be observed by a similar construction  
314 as above, except that now we define the transitions of processes such that no transition is  
315 enabled iff the order is respected. ◀

### 316 4.3 Process Abstraction for Convergence

317 As mentioned before, to prove self-stabilization of a parameterized program, we need to  
 318 prove closure and convergence. Closure can be proved based on Lemma 1, and Lemma 2  
 319 shows how to deal with deadlocks outside of  $LS$ . Thus, the missing part is a method to  
 320 prove that there are no cycles outside of  $LS$  that prevents a computation to eventually reach  
 321  $LS$ . In contrast to the two previous problems, we now consider infinite behaviors of the  
 322 system. Since parameterized verification and synthesis of symmetric self-stabilization in rings  
 323 is known to be undecidable [19,21], we cannot obtain cutoffs for this property. Therefore, we  
 324 resort to proving the absence of cycles based on a sound abstraction of the system behavior.

The basic idea is the following: we check whether there is a loop that starts and ends  
 in the same *local* state for an arbitrary process. If we can show that this is not possible,  
 then certainly no global loop is possible. Note that this is a stronger property than what we  
 want to prove; it proves that there could not be any loops in the protocol, neither inside nor  
 outside  $LS$ . It is obvious that this property can only be satisfied for silent protocols. To this  
 end, we fix five processes (see Fig. 4), and define the following property:

$$S \Rightarrow (\Diamond \Box S \vee \neg \Box \Diamond S),$$

325 where  $S$  is the local state of  $\pi_i$  (i.e., the valuation of its read-set),  $\Diamond$  is the ‘eventually’  
 326 operator and  $\Box$  is the ‘always’ operator in temporal logic. That is, given a local state in  
 327  $S$ , any future extension either reaches a state where  $S$  is continually true, or,  $S$  does not  
 328 become true infinitely often. Next, we attempt to prove the property in a ring of size 7,  
 329 where 5 processes behave according to the synthesized protocol, and the other two processes  
 330 have the same write-set, but can execute arbitrary transitions. The idea is that these two  
 331 processes over-approximate the possible behavior of all other processes. If we can prove the  
 332 property above in this abstraction of the system, then this implies that no loops are possible  
 333 in a concrete system in a ring of size  $\geq 5$ . Note that the precision of the abstraction can be  
 334 refined by increasing the number of processes that behave according to the protocol, or by  
 335 including the local state of additional processes into  $S$ . For the problems we considered in  
 336 our experiments (see Sect. 6), the fixed abstraction with  $5 + 2$  processes was sufficient.

### 337 4.4 Parameterized Self-Stabilization

338 Based on Lemmas 1 and 2, and the approach in Section 4.3, we obtain our main result.

339 ► **Theorem 14.** *Let  $\mathcal{T}_1, \mathcal{T}_2, \dots$  be a parameterized ring topology,  $\pi$  a process, and let  $LS$  be  
 340 locally defined by  $LS_i$ . Let  $c_1$  and  $c_2$  be cutoffs for closure and deadlock detection wrt.  $LS$ ,  
 341 respectively. If (1) closure holds in rings of size up to  $c_1$ , (2) deadlocks outside of  $LS$  are  
 342 impossible in rings of size up to  $c_2$ , and (3) the absence of cycles can be proven in rings of  
 343 up to size 4 and in an abstract system as above, then every instance of the parameterized  
 344 program  $\mathcal{D}_1 = \mathcal{T}_1^\pi, \mathcal{D}_2 = \mathcal{T}_2^\pi, \dots$  is self-stabilizing to  $LS$ . ◀*

## 345 5 SMT-based Counterexample-Guided Synthesis

### 346 5.1 General Idea

347 In [8, 9, 11], we introduced SMT-based methods to solve the synthesis problem for self-  
 348 stabilizing systems. In a nutshell, our techniques generate a set of SMT constraints from the  
 349 input synthesis instance and produce a model that represents a self-stabilizing protocol. In  
 350 order to scale up these technique to synthesize solutions up to the cutoff point efficiently, in

351 this section, we propose a method, where we find a solution for a larger topology using a  
 352 solution for a smaller topology. Let us first entertain a naïve idea, where we first synthesize  
 353 a protocol for a small topology and then simply use this solution for larger topologies with  
 354 the hope that since the protocol is symmetric, a small solution works in a larger network as  
 355 well. We now show that this approach is not conceivable even for very simple protocols.

356 **Example.** When applying our latest algorithm [11] to the one-bit maximal matching example,  
 357 the first synthesized solution for 4 processes is the following transition relation encoded by  
 358 guarded commands for each process  $\pi_i$ :

$$359 \quad \pi_i : \quad (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{F}) \quad \rightarrow \quad x_i := \mathbf{T}$$

$$360 \quad \quad \quad (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{T}) \quad \rightarrow \quad x_i := \mathbf{F}$$

362 and the following interpretation for uninterpreted function  $match_i$ :

$$363 \quad match_i : \quad (x_i = \mathbf{T}) \wedge ((x_{(i+1)} = \mathbf{T}) \vee (x_{(i-1)} = \mathbf{F})) \quad \mapsto \quad l$$

$$364 \quad \quad \quad (x_{(i+1)} = \mathbf{T}) \wedge (x_{(i-1)} = \mathbf{F}) \quad \mapsto \quad l$$

$$365 \quad \quad \quad (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{T}) \quad \mapsto \quad r$$

$$366 \quad \quad \quad (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{T}) \wedge (x_{(i-1)} = \mathbf{T}) \quad \mapsto \quad r$$

$$367 \quad \quad \quad (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{F}) \quad \mapsto \quad n$$

369 Now, if we trivially use the synthesized protocol on a topology with 5 processes, the resulting  
 370 protocol is incorrect. In particular, the following is a counterexample (i.e., a finite computation  
 371 that violates the specification) in terms of predicate  $match$ :

$$372 \quad \left( [match_0 = n, match_1 = n, match_2 = n, match_3 = n, match_4 = n], \right.$$

$$373 \quad \quad [match_0 = l, match_1 = n, match_2 = n, match_3 = n, match_4 = l],$$

$$374 \quad \quad \left. [match_0 = l, match_1 = r, match_2 = l, match_3 = n, match_4 = l] \right)$$

376 This computation violates convergence, as it reaches a deadlock state in  $\neg LS$ . This example  
 377 shows that a synthesized symmetric solution cannot be trivially extended to larger topologies.

## 378 5.2 The Counterexample-Guided Synthesis Algorithm

379 In order to limit the search space of SMT-solvers for a solution, we incorporate a synthesis-  
 380 verification loop guided by counterexamples. Our approach consists of the following steps:

- 381 1. Given a topology with  $i$  processes and a set of legitimate states, we use our existing  
 382 approach [8, 9, 11] to formulate the synthesis problem as an SMT instance.
- 383 2. We use an SMT solver to find a solution for the SMT instance, as a transition relation  
 384 and an interpretation for each uninterpreted function. Note that due to symmetry, the  
 385 transition relations and the interpretation functions are identical for all processes.
- 386 3. Next, we generalize the solution for a topology with  $i + 1$  processes and verify this solution  
 387 using a model checker.
- 388 4. If the result of verification is positive, then we go back to step 3 to check the properties for  
 389 a topology with  $i + 2$  processes. Otherwise, we transform the generated counterexample  
 390 into an SMT constraint and add it to the initial SMT instance (step 1) and return to  
 391 step 2.

392 For reasons of space, we do not include the details of our SMT-based synthesis technique [8,  
 393 9, 11]. We now analyze the nature of counterexamples. In the context of closure and  
 394 convergence, a model checker may generate a counterexample of the form  $\bar{s} = s_0 s_1 \cdots s_n$ .  
 395 Observe that  $\bar{s}$  is one of the following three types of counterexamples:

- 396 ■ If closure is violated, then  $\bar{s} = (s_0, s_1)$ , where  $s_0 \in LS$  and  $s_1 \notin LS$ .
- 397 ■ If convergence is violated by  $\bar{s} = s_0 s_1 \cdots s_n$ , where for all  $i \in [0, n]$ , we have  $s_i \notin LS$  and  
 398 either
  - 399 ■  $s_0 = s_n$ ; i.e., a loop exists outside the set of legitimate states, or
  - 400 ■ there does not exist a state  $\mathfrak{s}$ , where  $(s_n, \mathfrak{s})$  is a valid transition; i.e.,  $s_n$  is a *deadlock*  
 401 state outside the set of legitimate states.

402 For example, the counterexample presented in Section 5.1 is of the third type.

403 Dealing with the first type of counterexamples is pretty straightforward: we only add a  
 404 constraint to the SMT instance that disallows transition  $(s_0, s_1)$  in the transition relation. To  
 405 address deadlocks, we need to add a constraint to the SMT instance to enforce a change in the  
 406 resulting synthesized model, so that  $s_n$  is not a deadlock state. To this end, we propose two  
 407 sets of heuristics to change either the transition relation or the interpretation of uninterpreted  
 408 functions in Section 5.3. Dealing with loops is a bit more complicated. For example, one  
 409 can remove a transition from the loop to break it, but the choice of transition may involve a  
 410 combinatorial enumeration to find the right transition. This type of counterexamples is not  
 411 our focus in this paper and we leave it for future work. Interestingly, all of our case studies  
 412 in Section 6 do not involve loop counterexamples.

### 413 5.3 Heuristics Considering Transition Relations

414 The simplest method to resolve a deadlock is to formulate a constraint imposing the existence  
 415 of an outgoing transition from  $s_n$ . Since in this paper, our focus is on asynchronous systems,  
 416 a transition is the execution of one of the processes. We propose two strategies for selecting  
 417 a process to have an outgoing transition from a deadlock state.

**Progress Heuristic.** In this approach, we add a constraint stating that at least one of  
 the proceses should have an outgoing transition from  $s_n$ . More formally, assume that the  
 current topology includes  $i$  processes, where the read-set of each process has  $r$  variables, with  
 domains  $D_0, \dots, D_{r-1}$ , and the write-set of each process includes  $w$  variables, with domains  
 $D'_0, \dots, D'_{w-1}$ . Note that since the goal is to synthesize a symmetric program, all processes  
 execute similarly according to the function  $T_p$ :

$$T_p : \left( \prod_{j \in [0, r-1]} D_j \right) \rightarrow \left( \prod_{j \in [0, w-1]} D'_j \right)$$

and function  $f$  is of type:

$$f : \mathbb{N} \rightarrow \left( S \rightarrow \left( \prod_{j \in [0, r-1]} D_j \right) \right)$$

Then, the constraint to be added to the SMT instance can be written as:

$$\forall j \in [0, w-1] : \exists val_j \in D'_j : \bigvee_{k \in [0, i]} \left( (f(k)(s_n), [val_0, val_1, \dots, val_{w-1}]) \in T_p \right)$$

**Example.** Consider the counterexample mentioned in Section 5.1. Each process can read three Boolean variables and write to one Boolean variable and, hence,  $T_p$  is defined as follows:

$$T_p : \{F, T\} \times \{F, T\} \times \{F, T\} \rightarrow \{F, T\}$$

Note that for each process  $\pi_j$ ,  $f(j)$  returns  $[x_{(j-1)}, x_j, x_{(j+1)}]$ . In the counterexample we presented in the previous example, the last state where the deadlock happens is:

$$[x_0 = T, x_1 = F, x_2 = T, x_3 = F, x_4 = F]$$

418 Thus, we add the following constraint to the SMT instance:

$$419 \quad \exists val \in \{F, T\} : \left( ([F, T, F], val) \in T_p \vee ([T, F, T], val) \in T_p \vee ([F, T, F], val) \in T_p \vee \right. \\ 420 \quad \left. ([T, F, F], val) \in T_p \vee ([F, F, T], val) \in T_p \right) \\ 421$$

422 In the above constraint, the  $j$ th clause imposes a constraint on  $T_p$  to have an outgoing  
423 transition considering the local state of the  $j$ th process. (Note that the first and third clauses  
424 are the same, and we just put them for clarity.)

425 **Local LS Heuristic.** As mentioned in Section 2, we focus on sets  $LS$  that can be locally  
426 defined, i.e., the set of legitimate states can be described as a conjunction over local legitimate  
427 states of processes. In this case, a deadlock can be resolved by checking the local state of  
428 each process and imposing a constraint to have an outgoing transition for at least one of  
429 those processes that are not in their local legitimate states.

430 **Example.** For the counterexample of one-bit maximal matching with 4 processes, the local  
431 set of legitimate states is the following:

$$432 \quad (match_{(i-1)} = l \wedge match_i = n \wedge match_{(i+1)} = r) \vee \\ 433 \quad (match_{(i-1)} = r \wedge match_i = l \wedge match_{(i+1)} = n) \vee \\ 434 \quad (match_{(i-1)} = r \wedge match_i = l \wedge match_{(i+1)} = r) \vee \\ 435 \quad (match_{(i-1)} = l \wedge match_i = r \wedge match_{(i+1)} = l) \vee \\ 436 \quad (match_{(i-1)} = n \wedge match_i = r \wedge match_{(i+1)} = l) \\ 437$$

and the deadlock state is:

$$[x_0 = T, x_1 = F, x_2 = T, x_3 = F, x_4 = F]$$

For checking the local state of each process, we should first note the values of uninterpreted functions  $match_i$  in this state:

$$[match_0 = l, match_1 = r, match_2 = l, match_3 = n, match_4 = l]$$

438 Processes  $\pi_0$ ,  $\pi_1$ ,  $\pi_3$ , and  $\pi_4$  are not in a local legitimate state, and hence, the added  
439 constraint to the original SMT model will be as follows:

$$440 \quad \exists val \in \{F, T\} : \left( ([F, T, F], val) \in T_p \vee ([T, F, T], val) \in T_p \vee \right. \\ 441 \quad \left. ([T, F, F], val) \in T_p \vee ([F, F, T], val) \in T_p \right) \\ 442$$

443 Note that although this method seems more efficient than the progress approach in terms of  
444 having shorter constraints, it has the drawback of missing some solutions that the previous  
445 approach can find. More specifically, for a process being in a legitimate local state in a  
446 deadlock state, it may be the case that taking a transition by this process leads to a state,  
447 from which its neighbors can take other transitions that finally leads to a legitimate state.

448 **5.4 Heuristics Considering Uninterpreted Functions**

449 Our second class of heuristics focus on uninterpreted functions. That is, we impose a  
 450 constraint to change the interpretation function of at least one uninterpreted function in the  
 451 deadlock state. Similar to the heuristics introduced for transition relations, we introduce two  
 452 approaches for selecting at least one process to change the interpretation of its uninterpreted  
 453 function. Because of the similarity to the previous heuristics, we skip the details of this  
 454 heuristic.

455 **6 Case Studies and Experimental Results**

456 We used the model finder Alloy [15] and model checker NuSMV [7] to implement our  
 457 counterexample-guided synthesis approach. Our experimental platform is an 2.9 GHz Intel  
 458 Core i7 processor, with 16 GB of RAM. Our synthesis results are reported in Table 1.

459 **6.1 Three Coloring**

We consider the *three coloring problem* [14] on a ring, where each process  $\pi_i$  is associated with a variable  $c_i$  with domain  $\{0, 1, 2\}$ . Each value of the variable  $c_i$  represents a distinct color. A process can read and write its own variable. It can also read the variables of its neighbors. *LS* includes all states, where each process has a color different from its both neighbors. Thus, for a ring of 4 processes, *LS* is defined by the following predicate:

$$c_0(s) \neq c_1(s) \wedge c_1(s) \neq c_2(s) \wedge c_2(s) \neq c_3(s) \wedge c_3(s) \neq c_0(s)$$

460 Observe that the closure/deadlock-freedom cutoff point for this case study is  $3^2 + 1 = 10$   
 461 and, hence, we need to synthesize a solution for 10 processes. The synthesis time reported in  
 462 Table 1 is a bit smaller in the case of local *LS* heuristic, which is probably due to the smaller  
 463 constraints added in this case. The resulting protocols for the two heuristics are different.  
 464 Following is the one synthesized for the case of local LS:

$$\begin{aligned} 465 \quad \pi_i : \quad & (c_i = 2) \wedge (c_{i+1} = 2) \wedge (c_{i-1} \neq 0) \quad \rightarrow \quad c_i := 0 \\ 466 & (c_i \neq 0) \wedge (c_{i+1} = 1) \wedge (c_{i-1} = 2) \quad \rightarrow \quad c_i := 0 \\ 467 & (c_i = 1) \wedge (c_{i+1} = 1) \wedge (c_{i-1} \neq 2) \quad \rightarrow \quad c_i := 2 \\ 468 & (c_i = 0) \wedge (c_{i+1} \neq 2) \wedge (c_{i-1} = 0) \quad \rightarrow \quad c_i := 2 \\ 469 & (c_i \neq 1) \wedge (c_{i+1} = 2) \wedge (c_{i-1} = 0) \quad \rightarrow \quad c_i := 1 \\ 470 \end{aligned}$$

471 **6.2 One-Bit Maximal Matching**

472 This case study is the running example in this paper with cutoff point of  $2^2 + 1 = 5$  processes.  
 473 Note that using the heuristics considering transition relations, we could not synthesize a  
 474 protocol for this problem (Alloy reports unsatisfiability after adding the counterexample  
 475 constraints). The interesting point about this case study is that the progress heuristic has  
 476 better efficiency compared to the local *LS*. The reason may be due to the fact that the  
 477 constraints added in the local *LS* heuristic are too restrictive, and hence, Alloy needs to  
 478 search more in order to find a solution. The synthesized solutions using both heuristics are  
 479 the same for this case study, where the transition relation is the following:

$$\begin{aligned} 480 \quad \pi_i : \quad & (x_i = \mathbf{F}) \wedge (x_{i+1} = \mathbf{F}) \wedge (x_{i-1} = \mathbf{F}) \quad \rightarrow \quad x_i := \mathbf{T} \\ 481 & (x_i = \mathbf{T}) \wedge (x_{i+1} = \mathbf{T}) \quad \rightarrow \quad x_i := \mathbf{F} \\ 482 \end{aligned}$$

Problem	cutoff #	Heuristic	Synthesis Time	Model Checking Time
Three Coloring	10	Local <i>LS</i>	7m 3sec	16 msec
Three Coloring	10	Progress	9m 5sec	16 msec
One-Bit MM	5	Local <i>LS</i>	1m 48sec	27 msec
One-Bit MM	5	Progress	1m 44sec	33 msec
Maximal Matching	10	Local <i>LS</i>	7m 59sec	36 msec
Maximal Matching	10	Progress	4m 57sec	37 msec
Maximal Independent Set	5	Local <i>LS</i>	10sec	18 msec

■ **Table 1** Results for parameterized synthesis.

483 and the interpretation function for  $match_i$  is the following:

$$\begin{aligned}
484 \quad match_i : \quad & (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{T}) \wedge (x_{(i-1)} = \mathbf{T}) \mapsto l \\
& (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{F}) \mapsto l \\
485 & \\
486 & (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{T}) \mapsto r \\
487 & (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{T}) \mapsto r \\
488 & (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{T}) \wedge (x_{(i-1)} = \mathbf{F}) \mapsto n \\
489 & (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{T}) \mapsto n \\
490 &
\end{aligned}$$

### 491 6.3 Maximal Matching

492 In this case study, we used the same problem as in Section 6.2, but instead of using one  
493 Boolean variable for each process, we use a variable with three values  $\{l, r, n\}$  and, hence,  
494 we do not need the uninterpreted functions anymore. The resulting protocols for the two  
495 heuristics are different. As an example, the synthesized protocol for the case of local *LS* is  
496 the following:

$$\begin{aligned}
497 \quad \pi_i : \quad & (x_i = n) \wedge (x_{(i+1)} = n) \wedge (x_{(i-1)} = n) \rightarrow x_i := r \\
498 & (x_i \neq r) \wedge (x_{(i+1)} \neq r) \wedge (x_{(i-1)} = l) \rightarrow x_i := r \\
499 & (x_i \neq n) \wedge (x_{(i+1)} = r) \wedge (x_{(i-1)} = l) \rightarrow x_i := n \\
500 & (x_i = n) \wedge (x_{(i-1)} = r) \rightarrow x_i := l \\
501 & (x_i = r) \wedge (x_{(i+1)} = r) \wedge (x_{(i-1)} \neq l) \rightarrow x_i := l \\
502 &
\end{aligned}$$

### 503 6.4 Maximal Independent Set

An *independent set* in a graph is a subset of vertices in which no pair of vertices are adjacent. To synthesize a protocol that finds a maximal independent set, we consider a set of processes connected in a ring topology, where each process has a Boolean variable, the value of which shows whether or not it is included in the maximal independent set. The set of legitimate states include those states, where the processes whose variables have the true value form a maximal independent set. As an example, if  $c_i$  is the variable of the process  $\pi_i$ , then the set of legitimate states for the case of four processes is formulated by the following predicate:

$$(c_0(s) = \mathbf{T} \wedge c_1(s) = \mathbf{F} \wedge c_2(s) = \mathbf{T} \wedge c_3(s) = \mathbf{F}) \vee (c_0(s) = \mathbf{F} \wedge c_1(s) = \mathbf{T} \wedge c_2(s) = \mathbf{F} \wedge c_3(s) = \mathbf{T})$$

504 The point of this case study is that the resulting model for 4 processes worked for the size  
505 5 as well, and hence, no counterexample is found. Therefore, the result for both heuristics is

506 the same. The resulting protocol is the following:

$$\begin{array}{l}
 507 \quad \pi_i : \quad (x_i = \mathbf{F}) \wedge (x_{(i+1)} = \mathbf{F}) \wedge (x_{(i-1)} = \mathbf{F}) \quad \rightarrow \quad x_i := \mathbf{T} \\
 508 \quad \quad \quad (x_i = \mathbf{T}) \wedge (x_{(i+1)} = \mathbf{T}) \quad \rightarrow \quad x_i := \mathbf{F} \\
 509
 \end{array}$$

## 510 **7 Related Work**

511 Regarding the synthesis of self-stabilizing algorithms, one approach is to *add* self-stabilization  
 512 to a given algorithm. In contrast to our approach, the technique proposed by Ebneenasir  
 513 and Farahat [5] starts from a given non-stabilizing algorithm, it requires a more explicit  
 514 specification of the legitimate states, and it is not complete, i.e., it may fail to find a  
 515 solution even though there exists one. Klinkhammer and Ebneenasir show that adding strong  
 516 convergence is NP-complete in the size of the state space, which itself is exponential in the  
 517 size of variables of the protocol [18], and introduce a new method for adding self-stabilization  
 518 that is complete, but otherwise has the same limitations as mentioned above [20]. For ring  
 519 topologies, they have shown that parameterized verification of self-stabilization is undecidable  
 520 in uni-directional rings [19], while the parameterized synthesis problem is undecidable in bi-  
 521 directional rings, but surprisingly remains decidable in uni-directional rings [21]. Faghieh and  
 522 Bonakdarpour introduced an SMT-based synthesis technique for automatically synthesizing  
 523 self-stabilizing systems [8, 9] that is complete and not based on existing non-stabilizing  
 524 algorithms. An extension of this work [11] allows to symbolically specify the legitimate states  
 525 as a set of requirements, and supports the synthesis of ideal-stabilizing systems.

526 While these approaches are promising and can automatically synthesize a number of  
 527 well-known self-stabilizing systems, they all suffer from the problem of scalability, as the  
 528 complexity of the problem increases exponentially in the number of processes. For example, all  
 529 results reported by Faghieh and Bonakdarpour [8, 9, 11] correspond to automatically synthesis  
 530 of self-stabilizing systems with at most 5 processes. One way to address this scalability issue  
 531 in synthesis is to use a counterexample-guided synthesis method, as it has been proposed  
 532 for the completion of program sketches [25], for the lazy synthesis of reactive systems [13],  
 533 and for the synthesis of Byzantine-resilient systems [1]. The latter approach also supports  
 534 the synthesis of self-stabilizing systems, but counterexamples are only used to guide the  
 535 encoding of Byzantine-resilience, and the approach is limited to synchronous systems. In all  
 536 of these examples, a counterexample-guided approach can solve problems that are out of  
 537 reach for existing approaches. In our work, we for the first time used counterexamples to  
 538 guide synthesis for an increasing size of the topology, which allows us to scale the SMT-based  
 539 synthesis of self-stabilizing algorithms to systems with up to 200 processes.

540 Finally, the problem of scalability in the number of processes can be solved once and for  
 541 all by using a parameterized synthesis approach, as introduced by Jacobs and Bloem [16]  
 542 for (non-stabilizing) reactive systems. The approach relies on cutoff results, similar to the  
 543 ones we introduced in this work for closure and deadlock detection. Different techniques  
 544 are introduced in [17] to improve scalability of this approach in the complexity of the  
 545 specification, including the modular application of cutoff results in synthesis. An extension  
 546 of the approach [1] also supports the parameterized synthesis of self-stabilizing systems, but  
 547 only for synchronous systems, and not in all cases resulting in a completely symmetric system.  
 548 Finally, Lazic et al. [22] propose a method for synthesizing parameterized fault-tolerant  
 549 distributed algorithms. In contrast to our approach, synthesis is based on a sketch of an  
 550 asynchronous threshold-based fault-tolerant distributed algorithm, and the goal is to find  
 551 the right values for coefficients that may be missing in the guards.

## 8 Conclusion

In this paper, we proposed a new method for parameterized synthesis of self-stabilizing algorithms in symmetric rings using cutoff points. Furthermore, in order to scale the existing synthesis solutions [8–12] up to the cutoff point, we introduced an iterative loop of synthesis and verification guided by counterexamples. We demonstrated the effectiveness of our approach by synthesizing parameterized self-stabilizing protocols for well-known problems including self-stabilizing three coloring, maximal matching, and maximal independent set in less than 10 minutes. For future, we plan to work on asymmetric and dynamic networks as well as the case, where the protocol is live in the set of legitimate states.

## References

- 1 R. Bloem, N. Braud-Santoni, and S. Jacobs. Synthesis of self-stabilising and byzantine-resilient distributed systems. In *CAV*, pages 157–176, 2016.
- 2 S. Devismes, S. Tixeuil, and M. Yamashita. Weak vs. self vs. probabilistic stabilization. In *ICDCS*, pages 681–688, 2008.
- 3 E. W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11):643–644, 1974.
- 4 E. W. Dijkstra. A belated proof of self-stabilization. *Distributed Computing*, 1(1):5–6, 1986.
- 5 A. Ebneenasir and A. Farahat. A lightweight method for automated design of convergence. In *IPDPS*, pages 219–230, 2011.
- 6 E. A. Emerson and K. S. Namjoshi. On reasoning about rings. *International Journal on Foundations of Computer Science.*, 14(4):527–550, 2003.
- 7 A. Cimatti et. al. Nusmv 2: An opensource tool for symbolic model checking. In *CAV*, pages 359–364, 2002.
- 8 F. Faghih and B. Bonakdarpour. SMT-based synthesis of distributed self-stabilizing systems. In *SSS*, pages 165–179, 2014.
- 9 F. Faghih and B. Bonakdarpour. SMT-based synthesis of distributed self-stabilizing systems. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 10(3):21, 2015.
- 10 F. Faghih and B. Bonakdarpour. ASSESS: A tool for automated synthesis of distributed self-stabilizing algorithms. In *SSS*, pages 219–233, 2017.
- 11 F. Faghih, B. Bonakdarpour, S. Tixeuil, and S. Kulkarni. Specification-based synthesis of distributed self-stabilizing protocols. In *International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE)*, pages 124–141, 2016.
- 12 F. Faghih, B. Bonakdarpour, S. Tixeuil, and S. Kulkarni. Specification-based synthesis of distributed self-stabilizing protocols. *Logical Methods in Computer Science*, To appear.
- 13 B. Finkbeiner and S. Jacobs. Lazy synthesis. In *VMCAI*, 2012.
- 14 M. G. Gouda and H. B. Acharya. Nash equilibria in stabilizing systems. In *SSS*, pages 311–324, 2009.
- 15 D. Jackson. *Software Abstractions: Logic, Language, and Analysis*. MIT Press Cambridge, 2012.
- 16 S. Jacobs and R. Bloem. Parameterized synthesis. *Logical Methods in Computer Science*, 10(1), 2014.
- 17 A. Khalimov, S. Jacobs, and R. Bloem. Towards efficient parameterized synthesis. In *VMCAI*, volume 7737 of *Lecture Notes in Computer Science*, pages 108–127. Springer, 2013.
- 18 A. Klinkhamer and A. Ebneenasir. On the complexity of adding convergence. In *Fundamentals of Software Engineering (FSEN)*, pages 17–33, 2013.
- 19 A. Klinkhamer and A. Ebneenasir. Verifying livelock freedom on parameterized rings and chains. In *SSS*, pages 163–177, 2013.

- 599 **20** A. Klinkhamer and A. Ebneenasir. Synthesizing self-stabilization through superposition and  
600 backtracking. In *SSS*, pages 252–267, 2014.
- 601 **21** A. Klinkhamer and A. Ebneenasir. Synthesizing parameterized self-stabilizing rings with  
602 constant-space processes. In *Fundamentals of Software Engineering (FSEN)*, pages 100–  
603 115, 2017.
- 604 **22** Marijana Lazic, Igor Konnov, Josef Widder, and Roderick Bloem. Synthesis of distributed  
605 algorithms with parameterized threshold guards. In *OPODIS*, 2017.
- 606 **23** F. Manne, M. Mjeldel, L. Pilard, and S. Tixeuil. A new self-stabilizing maximal matching  
607 algorithm. *Theoretical Computer Science*, 410(14):1336–1345, 2009.
- 608 **24** Kerry Raymond. A tree-based algorithm for distributed mutual exclusion. *ACM Transac-*  
609 *tions on Computer Systems*, 7(1):61–77, 1989.
- 610 **25** Armando Solar-Lezama. Program sketching. *STTT*, 15(5-6):475–495, 2013.