

# 1 Verification of Randomized Distributed 2 Algorithms under Round-Rigid Adversaries

3 **Nathalie Bertrand** 

4 Univ. Rennes, Inria, CNRS, IRISA, Rennes, France

5 nathalie.bertrand@inria.fr

6 **Igor Konnov** 

7 University of Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

8 igor.konnov@inria.fr

9 **Marijana Lazić** 

10 TU Wien, Favoritenstraße 9–11, 1040 Vienna, Austria

11 lazic@forsyte.at

12 **Josef Widder** 

13 TU Wien, Favoritenstraße 9–11, 1040 Vienna, Austria

14 widder@forsyte.at

## 15 — Abstract —

---

16 Randomized fault-tolerant distributed algorithms pose a number of challenges for automated  
17 verification: (i) parameterization in the number of processes and faults, (ii) randomized choices and  
18 probabilistic properties, and (iii) an unbounded number of asynchronous rounds. This combination  
19 makes verification hard. Challenge (i) was recently addressed in the framework of threshold automata.

20 We extend threshold automata to model randomized algorithms that perform an unbounded  
21 number of asynchronous rounds. For non-probabilistic properties, we show that it is necessary and  
22 sufficient to verify these properties under round-rigid schedules, that is, schedules where processes  
23 enter round  $r$  only after all processes finished round  $r - 1$ . For almost-sure termination, we analyze  
24 these algorithms under round-rigid adversaries, that is, fair adversaries that only generate round-rigid  
25 schedules. This allows us to do compositional and inductive reasoning that reduces verification of  
26 the asynchronous multi-round algorithms to model checking of a one-round threshold automaton.  
27 We apply this framework and automatically verify the following classic algorithms: Ben-Or’s and  
28 Bracha’s seminal consensus algorithms for crashes and Byzantine faults, 2-set agreement for crash  
29 faults, and RS-Bosco for the Byzantine case.

30 **2012 ACM Subject Classification** Software and its engineering → Software verification; Software  
31 and its engineering → Software fault tolerance; Theory of computation → Logic and verification

32 **Keywords and phrases** threshold automata – counter systems – parameterized verification – ran-  
33 domized distributed algorithms – Byzantine faults

34 **Related Version** Full version with proofs is available at <https://hal.inria.fr/hal-01925533>.

## 35 **1** Introduction

36 Fault-tolerant distributed algorithms like Paxos and Blockchain recently receive much atten-  
37 tion. Still, these systems are out of reach with current automated verification techniques.  
38 One problem comes from the scale: these systems should be verified for a very large (ide-  
39 ally even an unbounded) number of participants. In addition, many systems (including  
40 Blockchain), provide probabilistic guarantees. To check their correctness, one has to reason  
41 about randomized distributed algorithms in the parameterized setting.

42 In this paper, we make first steps towards parameterized verification of fault-tolerant  
43 randomized distributed algorithms. We consider the algorithms that follow the ideas of Ben-  
44 Or [3]. Interestingly, these algorithms were analyzed in [15, 13] where probabilistic reasoning  
45 was done using the probabilistic model checker PRISM [14] for systems of 10-20 processes,



© N. Bertrand, I. Konnov, M. Lazić, J. Widder;

licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

```

1  bool v := input_value({0, 1});
2  int r := 1;
3  while (true) do
4    send (R,r,v) to all;
5    wait for n - t messages (R,r,*);
6    if received (n + t) / 2 messages (R,r,w)
7    then send (P,r,w,D) to all;
8    else send (P,r,?) to all;
9    wait for n - t messages (P,r,*);
10   if received at least t + 1
11     messages (P,r,w,D) then {
12     v := w;
13     if received at least (n + t) / 2
14       messages (P,r,w,D)
15     then decide w;
16   } else v := random({0, 1});
17   r := r + 1;
18 od

```

■ **Figure 1** Pseudo code of Ben-Or’s algorithm for Byzantine faults

46 while only safety was verified in the parameterized setting using Cadence SMV. From a  
 47 different perspective, these algorithms extend asynchronous threshold-guarded distributed  
 48 algorithms from [10, 9] with two features (i) a random choice (coin toss), and (ii) repeated  
 49 executions of the same algorithm until it converges (with probability 1).

50 A prominent example is Ben-Or’s fault-tolerant consensus algorithm [3] given in Figure 1.  
 51 It circumvents the impossibility of asynchronous consensus [7] by relaxing the termination  
 52 requirement to almost-sure termination, *i.e.*, termination with probability 1. Here processes  
 53 execute an infinite sequence of asynchronous loop iterations, which are called rounds  $r$ . Each  
 54 round consists of two stages where they first exchange messages tagged  $R$ , wait until the  
 55 number of received messages reaches a certain threshold (given as expression over parameters  
 56 in line 5) and then exchange messages tagged  $P$ . In the code,  $n$  is the number of processes,  
 57 among which at most  $t$  are Byzantine faulty (which may send conflicting information). The  
 58 correctness of the algorithm should be verified for all values of the parameters  $n$  and  $t$  that  
 59 meet a so-called resilience condition, *e.g.*,  $n > 3t$ . Carefully chosen thresholds  $(n - t, (n + t)/2$   
 60 and  $t + 1)$  on the number of received messages of a given type, ensure *agreement*, *i.e.*, that  
 61 two correct processes never decide on different values. At the end of a round, if there is no  
 62 “strong majority” for a value, *i.e.*, less than  $(n + t)/2$  messages were received (cf. line 13), a  
 63 process picks a new value randomly in line 16.

64 While these non-trivial threshold expressions can be dealt with using the methods in [9],  
 65 several challenges remain. The technique in [9] can be used to verify one iteration of the round  
 66 from Figure 1 only. However, consensus algorithms should prevent that there are no two  
 67 rounds  $r$  and  $r'$  such that a process decides 0 in  $r$  and another decides 1 in  $r'$ . This calls for a  
 68 compositional approach that allows one to compose verification results for individual rounds.  
 69 A challenge in the composition is that distributed algorithms implement “asynchronous  
 70 rounds”, *i.e.*, during a run processes may be in different rounds at the same time.

71 In addition, the combination of distributed aspects and probabilities makes reasoning  
 72 difficult. Quoting Lehmann and Rabin [16], “proofs of correctness for probabilistic dis-  
 73 tributed systems are extremely slippery”. This advocates the development of automated  
 74 verification techniques for probabilistic properties of randomized distributed algorithms in  
 75 the parameterized setting.

76 **Contributions.** We extend the threshold automata framework from [9] to round-based  
 77 algorithms with coin toss transitions. For the new framework we achieve the following:

- 78 1. For safety verification we introduce a method for compositional round-based reasoning.  
 79 This allows us to invoke a reduction similar to the one in [6]. We highlight necessary

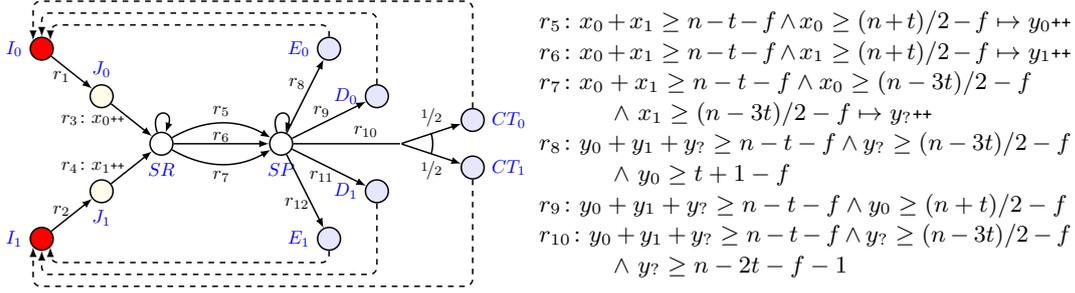


Figure 2 Ben-Or's algorithm as PTA with resilience condition  $n > 3t \wedge t > 0 \wedge t \geq f \geq 0$ .

80 fairness conditions on individual rounds. This provides us with specifications to be  
 81 checked on a one-round automaton.

- 82 2. We reduce probabilistic liveness verification to proving termination with positive prob-  
 83 ability within a fixed number of rounds. To do so, we restrict ourselves to round-rigid  
 84 adversaries, that is, adversaries that respect the round ordering. In contrast to existing  
 85 work that proves almost-sure termination for fixed number of participants, these are the  
 86 first parameterized model checking results for probabilistic properties.
- 87 3. We check the specifications that emerge from points 1. and 2. and thus verify challenging  
 88 benchmarks in the parameterized setting. We verify Ben-Or's [3] and Bracha's [5] classic  
 89 algorithms, and more recent algorithms such as 2-set agreement [19], and RS-Bosco [21].

## 90 2 Overview

91 We introduce probabilistic threshold automata to model randomized threshold-based algo-  
 92 rithms. An example of such an automaton is given in Figure 2. Nodes represent local states  
 93 (or locations) of processes, which move along the labeled edges or forks. Edges and forks  
 94 are called rules. Labels have the form  $\varphi \mapsto u$ , meaning that a process can move along the  
 95 edge only if  $\varphi$  evaluates to true, and this is followed by the update  $u$  of shared variables.  
 96 Additionally, each tine of a fork is labeled with a number in the  $[0, 1]$  interval, representing  
 97 the probability of a process moving along the fork to end up at the target location of the tine.  
 98 If we ignore the dashed arrows in Figure 2, a threshold automaton captures the behavior of  
 99 a process in one round, that is, a loop iteration in Figure 1.

100 The code in Figure 1 refers to numbers of received messages and, as is typical for  
 101 distributed algorithms, their relation to sent messages (that is the semantics of send and  
 102 receive) is not explicit in the pseudo code. To formalize the behavior, the encoding in the  
 103 threshold automaton directly refers to the numbers of sent messages, and they are encoded  
 104 in the shared variables  $x_i$  and  $y_i$ . The algorithm is parameterized:  $n$  is the number of  
 105 processes,  $t$  is the assumed number of faults and  $f$  is the actual number of faults. It should be  
 106 demonstrated to work under the resilience condition  $n > 3t \wedge t \geq f \wedge t > 0$ . For instance, the  
 107 locations  $J_0$  and  $J_1$  capture that a loop is entered with  $v$  being 0 and 1, respectively. Sending  
 108 an  $(R, r, 0)$  and  $(R, r, 1)$  message is captured by the increments on the shared variables  $x_0$   
 109 and  $x_1$  in the rules  $r_3$  and  $r_4$ , respectively; e.g., a process that is in location  $J_0$  uses rule  
 110  $r_3$  to go to location  $SR$  ("sent  $R$  message"), and increments  $x_0$  in doing so. Waiting for  $R$   
 111 and  $P$  messages in the lines 5 and 9, is captured by looping in the locations  $SR$  and  $SP$ .  
 112 In line 7 a process sends, e.g., a  $(P, r, 0, D)$  message if it has received  $n - t$  messages out  
 113 of which  $(n + t)/2$  are  $(R, r, 0)$  messages. This is captured in the guard of rule  $r_5$  where  
 114  $x_0 + x_1 \geq n - t - f$  checks the number of received messages in total, and  $x_0 \geq (n + t)/2 - f$

115 checks for the specific messages containing 0. The “ $-f$ ” term captures that  $f$  messages from  
 116 Byzantine faults may be received (we do not model faulty processes explicitly, that is, they  
 117 may not increase the shared variables  $x_i$  but  $f$  influences the evaluation of guards). The  
 118 branching at the end of the loop from lines 10 to 18 is captured by the rules outgoing of  $SP$ .  
 119 In particular rule  $r_{10}$  captures the coin toss in line 16. The non-determinism due to faults  
 120 and message delays is captured by multiple rules being enabled in the same configuration.

121 Liveness properties of distributed algorithms typically require fairness constraints, e.g.,  
 122 every message sent by a correct process to a correct process is eventually received. For  
 123 instance, this implies in Figure 1 that if  $n - t$  correct processes have sent messages of the  
 124 form  $(R, 1, *)$  and  $(n + t)/2$  correct processes have sent messages of the form  $(R, 1, 0)$  then  
 125 every correct process should eventually execute line 7, and proceed to line 9. We capture  
 126 this by the following fairness constraint: if  $x_0 + x_1 \geq n - t \wedge x_0 \geq (n + t)/2$ —that is, rule  
 127  $r_5$  is enabled without the help of the  $f$  faulty processes but by “correct processes alone”—  
 128 then the source location of rule  $r_5$ , namely  $SR$  should eventually be evacuated, that is, its  
 129 corresponding counter should eventually be 0.

130 The dashed edges, called round switch rules, encode how a process, after finishing a round,  
 131 starts the next one. The round number  $r$  serves as the loop iterator in Figure 2, and in each  
 132 iteration, processes send messages that carry  $r$ . To capture this, our semantics will introduce  
 133 fresh shared variables initialized with 0 for each round  $r$ . Because there are infinitely many  
 134 rounds, this means a priori we have infinitely many variables.

135 As parameterized verification of threshold automata is in general undecidable [12], we  
 136 consider the so-called “canonic” restrictions here, *i.e.*, only increments on shared variables,  
 137 and no increments of the same variable within loops. These restrictions still allow us to  
 138 model many threshold-based fault-tolerant distributed algorithms [9]. As a result, threshold  
 139 automata without probabilistic forks and round switching rules can be automatically checked  
 140 for safety and liveness [9]. Adding forks and round switches is required to adequately  
 141 model randomized distributed algorithms. Here we will use a convenient restriction that  
 142 requires that coin-toss transitions only appear at the end of a round, e.g., line 16 of Figure 1.  
 143 Intuitively, as discussed in Section 1, a coin-toss is only necessary if there is no strong  
 144 majority. Thus, all our benchmarks have this feature, and we exploit it in Section 7.

145 In order to overcome the issue of infinitely many rounds, we prove in Section 6 that we  
 146 can verify probabilistic threshold automata by analyzing a one-round automaton that fits in  
 147 the framework of [9]. We prove that we can reorder transitions of any fair execution such  
 148 that their round numbers are in an increasing order. The obtained ordered execution is  
 149 stutter equivalent with the original one, and thus, they satisfy the same  $LTL_X$  properties over  
 150 the atomic propositions describing only one round. In other words, our targeted concurrent  
 151 systems can be transformed to a sequential composition of one-round systems.

152 The main problem with isolating a one-round system is that consensus specifications  
 153 often talk about at least two different rounds. In this case we need to use round invariants  
 154 that imply the specifications. For example, if we want to verify agreement, we have to check  
 155 whether two processes decide different values, possibly in different rounds. We do this in two  
 156 steps: (i) we check the round invariant that no process changes its decision from round to  
 157 round, and (ii) we check that within a round no two processes disagree.

158 Finally, verifying almost-sure termination under round-rigid adversaries calls for distinct  
 159 arguments. Our methodology follows the lines of the manual proof of Ben Or’s consensus  
 160 algorithm by Aguilera and Toueg [1]. However, our arguments are not specific to Ben  
 161 Or’s algorithm, and we apply it to other randomized distributed algorithms (see Section 8).  
 162 Compared to their paper-and-pencil proof, the threshold automata framework required us to

163 provide a more formal setting and a more informative proof, also pinpointing the needed  
 164 hypothesis. The crucial parts of our proof are automatically checked by the model checker  
 165 ByMC [11]. Hence the established correctness stands on less slippery ground, which addresses  
 166 the mentioned concerns of Lehmann and Rabin.

### 167 **3 The Probabilistic Threshold Automata Framework**

168 A *probabilistic threshold automaton* PTA is a tuple  $(\mathcal{L}, \mathcal{V}, \mathcal{R}, RC)$ , where

- 169 ■  $\mathcal{L}$  is a finite set of locations, that contains the following disjoint subsets: *initial locations*  
 170  $\mathcal{I}$ , *final locations*  $\mathcal{F}$ , and *border locations*  $\mathcal{B}$ , with  $|\mathcal{B}| = |\mathcal{I}|$ ;
- 171 ■  $\mathcal{V}$  is a set of variables. It is partitioned in two sets:  $\Pi$  contains *parameter variables*, and  
 172  $\Gamma$  contains *shared variables*;
- 173 ■  $\mathcal{R}$  is a finite set of *rules*; and
- 174 ■  $RC$ , the *resilience condition*, is a formula in linear integer arithmetic over parameter  
 175 variables.

176 In the following we introduce rules in detail, and give syntactic restrictions on locations.  
 177 The resilience condition  $RC$  only appears in the definition of the semantics in Section 3.1.

178 A rule  $r$  is a tuple  $(from, \delta_{to}, \varphi, \vec{u})$  where  $from \in \mathcal{L}$  is the *source* location,  $\delta_{to} \in \text{Dist}(\mathcal{L})$  is  
 179 a probability distribution over the *destination* locations,  $\vec{u} \in \mathbb{N}_0^{|\Gamma|}$  is the *update vector*, and  $\varphi$   
 180 is a guard, *i.e.*, a conjunction of expressions of the form  $b \cdot x \geq \bar{a} \cdot \mathbf{p}^\top + a_0$  or  $b \cdot x < \bar{a} \cdot \mathbf{p}^\top + a_0$ ,  
 181 where  $x \in \Gamma$  is a shared variable,  $\bar{a} \in \mathbb{Z}^{|\Pi|}$  is a vector of integers,  $a_0, b \in \mathbb{Z}$ , and  $\mathbf{p}$  is the  
 182 vector of all parameters. If  $r.\delta_{to}$  is a Dirac distribution, *i.e.*, there exists  $\ell \in \mathcal{L}$  such that  
 183  $r.\delta_{to}(\ell) = 1$ , we call  $r$  a *Dirac rule*, and write it as  $(from, \ell, \varphi, \vec{u})$ . Destination locations of  
 184 non-Dirac rules are in  $\mathcal{F}$  (coin-toss transitions only happen at the end of a round).

185 Probabilistic threshold automata model algorithms with successive identical rounds.  
 186 Informally, a round happens between border locations and final locations. Then round switch  
 187 rules let processes move from final locations of a given round to border locations of the next  
 188 round. From each border location there is exactly one Dirac rule to an initial location, and  
 189 it has a form  $(\ell, \ell', \text{true}, \vec{0})$  where  $\ell \in \mathcal{B}$  and  $\ell' \in \mathcal{I}$ . As  $|\mathcal{B}| = |\mathcal{I}|$ , one can think of border  
 190 locations as copies of initial locations. It remains to model from which final locations to  
 191 which border location (that is, initial for the next round) processes move. This is done by  
 192 *round switch rules*. They can be described as Dirac rules  $(\ell, \ell', \text{true}, \vec{0})$  with  $\ell \in \mathcal{F}$  and  
 193  $\ell' \in \mathcal{B}$ . The set of round switch rules is denoted by  $\mathcal{S} \subseteq \mathcal{R}$ .

194 A location is in  $\mathcal{B}$  if and only if all the incoming edges are in  $\mathcal{S}$ . Similarly, a location is  
 195 in  $\mathcal{F}$  if and only if there is only one outgoing edge and it is in  $\mathcal{S}$ .

196 Figure 2 depicts a PTA with border locations  $\mathcal{B} = \{I_0, I_1\}$ , initial locations  $\mathcal{I} = \{J_0, J_1\}$ ,  
 197 and final locations  $\mathcal{F} = \{E_0, E_1, D_0, D_1, CT_0, CT_1\}$ . The only rule that is not Dirac rule  
 198 is  $r_{10}$ , and round switch rules are represented by dashed arrows.

#### 199 **3.1 Probabilistic Counter Systems**

200 A resilience condition  $RC$  defines the set of *admissible parameters*  $\mathbf{P}_{RC} = \{\mathbf{p} \in \mathbb{N}_0^{|\Pi|} : \mathbf{p} \models$   
 201  $RC\}$ . We introduce a function  $N : \mathbf{P}_{RC} \rightarrow \mathbb{N}_0$  that maps a vector of admissible parameters  
 202 to a number of modeled processes in the system. For instance, for the automaton in Figure 2,  
 203  $N$  is the function  $(n, t, f) \mapsto n - f$ , as we model only the  $n - f$  correct processes explicitly,  
 204 while the effect of faulty processes is captured in non-deterministic choices between different  
 205 guards as discussed in Section 2. Given a PTA and a function  $N$ , we define the semantics,  
 206 called *probabilistic counter system*  $\text{Sys}(\text{PTA})$ , to be the infinite-state MDP  $(\Sigma, I, \text{Act}, \Delta)$ ,

207 where  $\Sigma$  is the set of configurations for PTA among which  $I \subseteq \Sigma$  are initial, the set of actions  
 208 is  $\mathbf{Act} = \mathcal{R} \times \mathbb{N}_0$  and  $\Delta: \Sigma \times \mathbf{Act} \rightarrow \text{Dist}(\Sigma)$  is the probabilistic transition function.

209 **Configurations.** In a configuration  $\sigma = (\vec{\kappa}, \vec{g}, \mathbf{p})$ , a function  $\sigma.\vec{\kappa}: \mathcal{L} \times \mathbb{N}_0 \rightarrow \mathbb{N}_0$  defines values  
 210 of location counters per round, a function  $\sigma.\vec{g}: \Gamma \times \mathbb{N}_0 \rightarrow \mathbb{N}_0$  defines shared variable values per  
 211 round, and a vector  $\sigma.\mathbf{p} \in \mathbb{N}_0^{|\Pi|}$  defines parameter values. We denote the vector  $(\vec{g}[x, k])_{x \in \Gamma}$   
 212 of shared variables in a round  $k$  by  $\vec{g}[k]$ , and by  $\vec{\kappa}[k]$  we denote the vector  $(\vec{\kappa}[\ell, k])_{\ell \in \mathcal{L}}$  of  
 213 local state counters in a round  $k$ .

214 A configuration  $\sigma = (\vec{\kappa}, \vec{g}, \mathbf{p})$  is *initial* if all processes are in initial states of round 0, and  
 215 all global variables evaluate to 0, that is, if for every  $x \in \Gamma$  and  $k \in \mathbb{N}_0$  we have  $\sigma.\vec{g}[x, k] = 0$ ,  
 216 if  $\sum_{\ell \in \mathcal{B}} \sigma.\vec{\kappa}[\ell, 0] = N(\mathbf{p})$ , and finally if  $\sigma.\vec{\kappa}[\ell, k] = 0$ , for every  $(\ell, k) \in (\mathcal{L} \setminus \mathcal{B}) \times \{0\} \cup \mathcal{L} \times \mathbb{N}$ .

217 A threshold guard evaluates to true in a configuration  $\sigma$  for a round  $k$ , written  $\sigma, k \models \varphi$ ,  
 218 if for all its conjuncts  $b \cdot x \geq \bar{a} \cdot \mathbf{p}^\top + a_0$ , it holds that  $b \cdot \sigma.\vec{g}[x, k] \geq \bar{a} \cdot (\sigma.\mathbf{p}^\top) + a_0$  (and  
 219 similarly for conjuncts of the other form, *i.e.*,  $b \cdot x < \bar{a} \cdot \mathbf{p}^\top + a_0$ ).

220 **Actions.** An action  $\alpha = (r, k) \in \mathbf{Act}$  stands for the application of a rule  $r$  to a process in  
 221 round  $k$ . We write  $\alpha.from$  for  $r.from$ ,  $\alpha.\varphi$  for  $r.\varphi$ , etc.

222 An action  $\alpha = (r, k)$  is *unlocked* in a configuration  $\sigma$ , if its guard evaluates to true in its  
 223 round, that is  $\sigma, k \models \varphi$ . An action  $\alpha = (r, k)$  is *applicable* to a configuration  $\sigma$  if  $\alpha$  is unlocked  
 224 in  $\sigma$ , and  $\sigma.\vec{\kappa}[r.from, k] \geq 1$ . When an action  $\alpha$  is applicable to a configuration  $\sigma$ , and when  
 225  $\ell$  is a potential destination location for the probabilistic action  $\alpha$ , we write  $apply(\sigma, \alpha, \ell)$   
 226 for the resulting configuration: parameters are unchanged, shared variables are updated  
 227 according to the update function of  $r$ , and the values of counters are modified in a natural  
 228 way:  $(\vec{\kappa}[r.from, k])$  is decreased by 1 and  $\vec{\kappa}[\ell, k]$  is increased by 1). The probabilistic transition  
 229 function  $\Delta$  is defined by:  $\Delta(\sigma, \alpha)(\sigma') = \alpha.\delta_{to}(\ell)$  if  $apply(\sigma, \alpha, \ell) = \sigma'$ , and 0 otherwise.

### 230 3.2 Non-probabilistic Counter Systems

231 Non-probabilistic threshold automata are defined in [10], and they can be seen as a special  
 232 case of probabilistic threshold automata where all rules are Dirac rules.

233 With a PTA, one can naturally associate a non-probabilistic threshold automaton, by  
 234 replacing probabilities with non-determinism.

235 **► Definition 1.** *Given a PTA  $(\mathcal{L}, \mathcal{V}, \mathcal{R}, RC)$ , its (non-probabilistic) threshold automaton*  
 236 *is  $TA_{PTA} = (\mathcal{L}, \mathcal{V}, \mathcal{R}_{np}, RC)$  where the set of rules  $\mathcal{R}_{np}$  is defined as  $\{r_\ell = (from, \ell, \varphi, \vec{u}): r =$*   
 237  *$(from, \delta_{to}, \varphi, \vec{u}) \in \mathcal{R} \wedge \ell \in \mathcal{L} \wedge \delta_{to}(\ell) > 0\}$ .*

238 We write TA instead of  $TA_{PTA}$  when the automaton PTA is clear from the context. Every  
 239 rule from  $\mathcal{R}_{np}$  corresponds to exactly one rule in  $\mathcal{R}$ , and for every rule in  $\mathcal{R}$  there is at least  
 240 one corresponding rule in  $\mathcal{R}_{np}$  (and exactly one for Dirac rules).

241 If we understand a TA as a PTA where all rules are Dirac rules, we can define transitions  
 242 using the partial function  $apply$  in order to obtain an infinite (non-probabilistic) counter  
 243 system, which we denote by  $\text{Sys}_\infty(\text{TA})$ . Moreover, since in this case  $\mathcal{R} = \mathcal{R}_{np}$ , actions of  
 244 the PTA exactly match transitions of its TA. We obtain  $\sigma'$  by applying  $t = (r, k)$  to  $\sigma$ ,  
 245 and write this as  $\sigma' = t(\sigma)$ , if and only if for the destination location  $\ell$  of  $r$  holds that  
 246  $apply(\sigma, t, \ell) = \sigma'$ .

247 Also, starting from a PTA, one can define the counter system  $\text{Sys}(\text{PTA})$ , and consequently  
 248 its non-probabilistic counterpart  $\text{Sys}_{np}(\text{PTA})$ . As the definitions of  $\text{Sys}_{np}(\text{PTA})$  and  $\text{Sys}_\infty(\text{TA})$   
 249 are equivalent for a given PTA, we are free to choose one, and always use  $\text{Sys}_\infty(\text{TA})$ .

250 A (finite or infinite) sequence of transitions is called *schedule*, and it is often denoted  
 251 by  $\tau$ . A schedule  $\tau = t_1, t_2, \dots, t_{|\tau|}$  is applicable to a configuration  $\sigma$  if there is a sequence  
 252 of configurations  $\sigma = \sigma_0, \sigma_1, \dots, \sigma_{|\tau|}$  such that for every  $1 \leq i \leq |\tau|$  we have that  $t_i$  is  
 253 applicable to  $\sigma_{i-1}$  and  $\sigma_i = t_i(\sigma_{i-1})$ . A *path* is an alternating sequence of configurations  
 254 and transitions, for example  $\sigma_0, t_1, \sigma_1, \dots, t_{|\tau|}, \sigma_{|\tau|}$ , such that for every  $t_i$ ,  $1 \leq i \leq |\tau|$ , in the  
 255 sequence, we have that  $t_i$  is applicable to  $\sigma_{i-1}$  and  $\sigma_i = t_i(\sigma_{i-1})$ . Given a configuration  $\sigma_0$   
 256 and a schedule  $\tau = t_1, t_2, \dots, t_{|\tau|}$ , we denote by  $\text{path}(\sigma_0, \tau)$  a path  $\sigma_0, t_1, \sigma_1, \dots, t_{|\tau|}, \sigma_{|\tau|}$   
 257 where  $t_i(\sigma_{i-1}) = \sigma_i$ ,  $1 \leq i \leq |\tau|$ . Similarly we define an infinite schedule  $\tau = t_1, t_2, \dots$ ,  
 258 and an infinite path  $\sigma_0, t_1, \sigma_1, \dots$ , also denoted by  $\text{path}(\sigma_0, \tau)$ . An infinite path is *fair* if  
 259 no transition is applicable forever from some point on. Equivalently, when a transition is  
 260 applicable, eventually either its guard becomes false, or all processes leave its source location.

261 Since every transition in  $\text{Sys}_\infty(\text{TA})$  comes from an action in  $\text{Sys}(\text{PTA})$ , note that every  
 262 path in  $\text{Sys}_\infty(\text{TA})$  is a valid path in  $\text{Sys}(\text{PTA})$ .

### 263 3.3 Adversaries

264 As usual, the non-determinism is resolved by a so-called adversary. Let  $\text{Paths}$  be the set of  
 265 all finite paths in  $\text{Sys}(\text{PTA})$ . An *adversary* is a function  $\mathbf{a} : \text{Paths} \rightarrow \text{Act}$ , that given a finite  
 266 path  $\pi$  selects an action applicable to the last configuration of  $\pi$ . Given a configuration  $\sigma$   
 267 and an adversary  $\mathbf{a}$ , we generate a family of paths, depending on the outcomes of non-Dirac  
 268 transitions. We denote this set by  $\text{paths}(\sigma, \mathbf{a})$ . An adversary  $\mathbf{a}$  is *fair* if all paths in  $\text{paths}(\sigma, \mathbf{a})$   
 269 are fair. As usual, the MDP  $\text{Sys}(\text{PTA})$  together with an initial configuration  $\sigma$  and an  
 270 adversary  $\mathbf{a}$  induce a Markov chain, written  $\mathcal{M}_\mathbf{a}^\sigma$ . We write  $\mathbb{P}_\mathbf{a}^\sigma$  for the probability measure  
 271 over infinite paths starting at  $\sigma$  in the latter Markov chain.

272 We call an adversary a *round-rigid* if it is fair, and if every sequence of actions it produces  
 273 can be written as  $s_1 \cdot s_1^p \cdot s_2 \cdot s_2^p \dots$ , where for every  $k \in \mathbb{N}$ , we have that  $s_k$  contains only  
 274 Dirac actions of round  $k$ , and  $s_k^p$  contains only non-Dirac actions of round  $k$ . The set of all  
 275 round-rigid adversaries is denoted by  $\mathcal{A}^R$ .

### 276 3.4 Atomic Propositions and Stutter Equivalence

277 The atomic propositions we consider describe non-emptiness of the locations from  $\mathcal{L} \setminus \mathcal{B}$  in  
 278 a specific round. Formally, the set of all such propositions for a round  $k \in \mathbb{N}_0$  is denoted  
 279 by  $\text{AP}_k = \{p(\ell, k) : \ell \in \mathcal{L} \setminus \mathcal{B}\}$ . For every  $k$  we define a labeling function  $\lambda_k : \Sigma \rightarrow 2^{\text{AP}_k}$   
 280 such that  $p(\ell, k) \in \lambda_k(\sigma)$  iff  $\sigma.\vec{\kappa}[\ell, k] > 0$ , *i.e.*, if location  $\ell$  is nonempty in round  $k$  in  $\sigma$ . By  
 281 abusing notation, we write “ $\vec{\kappa}[\ell, k] > 0$ ” and “ $\vec{\kappa}[\ell, k] = 0$ ” instead of  $p(\ell, k)$  and  $\neg p(\ell, k)$ , resp.

282 We denote by  $\pi_1 \stackrel{\Delta}{\sim}_k \pi_2$  that the paths  $\pi_1$  and  $\pi_2$  are stutter equivalent [2] w.r.t.  $\text{AP}_k$ .  
 283 Two counter systems  $C_0$  and  $C_1$  are stutter equivalent w.r.t.  $\text{AP}_k$ , written  $C_0 \stackrel{\Delta}{\sim}_k C_1$ , if for  
 284 every path  $\pi$  from  $C_i$  there is a path  $\pi'$  from  $C_{1-i}$  such that  $\pi \stackrel{\Delta}{\sim}_k \pi'$ , for  $i \in \{0, 1\}$ .

## 285 4 Consensus Properties and their Verification

286 In probabilistic (binary) consensus every correct process has an initial value from  $\{0, 1\}$ .  
 287 It consists of safety specifications and an almost-sure termination requirement, which we  
 288 consider in its round-rigid variant:

289 **Agreement:** No two correct processes decide differently.

290 **Validity:** If all correct processes have  $v$  as the initial value, then no process decides  $1 - v$ .

291 **Probabilistic wait-free termination:** Under every round-rigid adversary, with probability 1  
 292 every correct process eventually decides.

293 We now discuss the formalization of these specifications in the context of Ben-Or's  
 294 algorithm whose threshold automaton is given in Figure 2.

295 **Formalization.** In order to formulate and analyze the specifications, we partition every  
 296 set  $\mathcal{I}$ ,  $\mathcal{B}$ , and  $\mathcal{F}$ , into two subsets  $\mathcal{I}_0 \uplus \mathcal{I}_1$ ,  $\mathcal{B}_0 \uplus \mathcal{B}_1$ , and  $\mathcal{F}_0 \uplus \mathcal{F}_1$ , respectively. For every  
 297  $v \in \{0, 1\}$ , the partitions satisfy the following:

298 (R1) The processes that are initially in a location  $\ell \in \mathcal{I}_v$  have the initial value  $v$ .

299 (R2) Rules connecting locations from  $\mathcal{B}$  and  $\mathcal{I}$  respect the partitioning, *i.e.*, they connect  $\mathcal{B}_v$   
 300 and  $\mathcal{I}_v$ . Similarly, rules connecting locations from  $\mathcal{F}$  and  $\mathcal{B}$  respect the partitioning.

301 We introduce two subsets  $\mathcal{D}_v \subseteq \mathcal{F}_v$ , for  $v \in \{0, 1\}$ . Intuitively, a process is in  $\mathcal{D}_v$  in a round  $k$   
 302 if and only if it decides  $v$  in that round. Now we can express the specifications as follows:

303 **Agreement:** For both values  $v \in \{0, 1\}$  the following holds:

$$304 \quad \forall k \in \mathbb{N}_0, \forall k' \in \mathbb{N}_0. \mathbf{A} \left( \mathbf{F} \bigvee_{\ell \in \mathcal{D}_v} \bar{\kappa}[\ell, k] > 0 \rightarrow \mathbf{G} \bigwedge_{\ell' \in \mathcal{D}_{1-v}} \bar{\kappa}[\ell', k'] = 0 \right) \quad (1)$$

305 **Validity:** For both  $v \in \{0, 1\}$  it holds

$$306 \quad \forall k \in \mathbb{N}_0. \mathbf{A} \left( \bigwedge_{\ell \in \mathcal{I}_v} \bar{\kappa}[\ell, 0] = 0 \rightarrow \mathbf{G} \bigwedge_{\ell' \in \mathcal{D}_v} \bar{\kappa}[\ell', k] = 0 \right) \quad (2)$$

307 **Probabilistic wait-free termination:** For every round-rigid adversary  $\mathbf{a}$

$$308 \quad \mathbb{P}_{\mathbf{a}} \left( \bigvee_{k \in \mathbb{N}_0} \bigvee_{v \in \{0, 1\}} \mathbf{G} \bigwedge_{\ell \in \mathcal{F} \setminus \mathcal{D}_v} \bar{\kappa}[\ell, k] = 0 \right) = 1 \quad (3)$$

309 Agreement and validity are non-probabilistic properties, and thus can be analyzed on the  
 310 non-probabilistic counter system  $\text{Sys}_{\infty}(\text{TA})$ . For verifying probabilistic wait-free termination,  
 311 we make explicit the following assumption that is present in all our benchmarks: all non-Dirac  
 312 transitions have non-zero probability to lead to an  $\mathcal{F}_v$  location, for both values  $v \in \{0, 1\}$ .

313 In Section 5 we formalize safety specifications and reduce them to single-round specifica-  
 314 tions. In Section 6 we reduce verification of multi-round counter systems to verification of  
 315 single-round systems. In Section 7 we discuss our approach to probabilistic termination.

## 316 **5 Reduction to Specifications with one Round Quantifier**

317 Agreement contains two round variables  $k$  and  $k'$ , and Validity considers rounds 0 and  $k$ .  
 318 Thus, both involve two round numbers. As ByMC can only analyze one round systems [9],  
 319 the properties are only allowed to use one round number. In this section we show how to  
 320 check formulas (1) and (2) by checking properties that refer to one round. Namely, we  
 321 introduce round invariants (4) and (5), and prove that they imply Agreement and Validity.

322 The first round invariant claims that in every round and in every path, once a process  
 323 decides  $v$  in a round, no process ever enters a location from  $\mathcal{F}_{1-v}$  in that round. Formally:

$$324 \quad \forall k \in \mathbb{N}_0. \mathbf{A} \left( \mathbf{F} \bigvee_{\ell \in \mathcal{D}_v} \bar{\kappa}[\ell, k] > 0 \rightarrow \mathbf{G} \bigwedge_{\ell' \in \mathcal{F}_{1-v}} \bar{\kappa}[\ell', k] = 0 \right) \quad (4)$$

325 The second round invariant claims that in every round in every path, if no process starts  
 326 a round with a value  $v$ , then no process terminates that round with value  $v$ . Formally:

$$327 \quad \forall k \in \mathbb{N}_0. \mathbf{A} \left( \mathbf{G} \bigwedge_{\ell \in \mathcal{I}_v} \bar{\kappa}[\ell, k] = 0 \rightarrow \mathbf{G} \bigwedge_{\ell' \in \mathcal{F}_v} \bar{\kappa}[\ell', k] = 0 \right) \quad (5)$$

328 The following proposition is proved using restriction (R2) and by an inductive argument  
329 over the round number.

330 ► **Proposition 2.** *If  $\text{Sys}_\infty(TA) \models (4) \wedge (5)$ , then  $\text{Sys}_\infty(TA) \models (1) \wedge (2)$ .*

## 331 **6 Reduction to Single-Round Counter System**

332 Given a property of one round, our goal is to prove that there is a counterexample to the  
333 property in the multi-round system iff there is a counterexample in a single-round system.  
334 This is formulated in Theorem 6, and it allows us to use ByMC on a single-round system.

335 The proof idea contains two parts. First, in Section 6.1 we prove that one can replace an  
336 arbitrary finite schedule with a round-rigid one, while preserving atomic propositions of a  
337 fixed round. We show that swapping two adjacent transitions that do not respect the order  
338 over round numbers in an execution, gives us a legal stutter equivalent execution, *i.e.*, an  
339 execution satisfying the same  $\text{LTL}_X$  properties.

340 Second, in Section 6.2 we extend this reasoning to infinite schedules, and lift it from  
341 schedules to transition systems. The main idea is to do inductive and compositional reasoning  
342 over the rounds. To do so, we need well-defined round boundaries, which is the case if  
343 every round that is started is also finished; a property we can automatically check for fair  
344 schedules. In more detail, regarding propositions for one round, we show that the multi-round  
345 transition system is stutter equivalent to a single-round transition system. This holds under  
346 the assumption that all fair executions of a single-round transition system terminate, and  
347 this can be checked using the technique from [9]. As stutter equivalence of systems implies  
348 preservation of  $\text{LTL}_X$  properties, this is sufficient to prove the main goal of the section.

### 349 **6.1 Reduction from arbitrary schedules to round-rigid schedules**

350 ► **Definition 3.** *A schedule  $\tau = (r_1, k_1) \cdot (r_2, k_2) \cdot \dots \cdot (r_m, k_m)$ ,  $m \in \mathbb{N}_0$ , is called round-rigid  
351 if for every  $1 \leq i < j \leq m$ , we have  $k_i \leq k_j$ .*

352 The following proposition shows that any finite schedule can be re-ordered into a round-  
353 rigid one that is stutter equivalent regarding  $\text{LTL}_X$  formulas over proposition from  $\text{AP}_k$ , for  
354 all rounds  $k$ . It is proved using arguments on the commutativity of transitions, similar to [6].

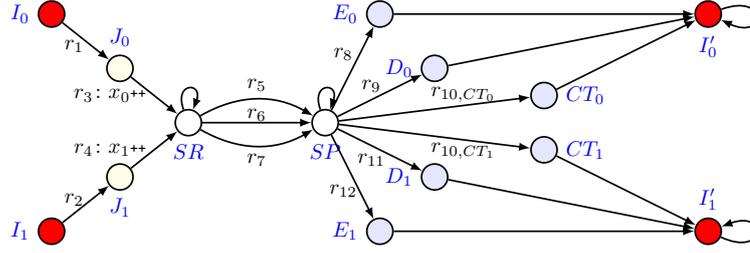
355 ► **Proposition 4.** *For every configuration  $\sigma$  and every finite schedule  $\tau$  applicable to  $\sigma$ , there  
356 is a round-rigid schedule  $\tau'$  such that the following holds:*

- 357 (a) *Schedule  $\tau'$  is applicable to configuration  $\sigma$ .*
- 358 (b)  *$\tau'$  and  $\tau$  reach the same configuration when applied to  $\sigma$ , *i.e.*,  $\tau'(\sigma) = \tau(\sigma)$ .*
- 359 (c) *For every  $k \in \mathbb{N}_0$  we have  $\text{path}(\sigma, \tau) \stackrel{\Delta}{=}_k \text{path}(\sigma, \tau')$ .*

360 Thus, instead of reasoning about all finite schedules of  $\text{Sys}_\infty(TA)$ , it is sufficient to reason  
361 about its round-rigid schedules. In the following section we use this to simplify the verification  
362 further, namely to a single-round counter system.

### 363 **6.2 From round-rigid schedules to single-round counter system**

364 For each PTA, we define a *single-round threshold automaton* that can be analyzed with the  
365 tools of [10] and [9]. Roughly speaking, we focus on one round, but also keep the border  
366 locations of the next round, where we add self-loops. Figure 3 represents the single-round  
367 threshold automaton associated with the PTA from Figure 2. We can prove that for specific  
368 fairness constraints, this automaton shows the same behavior as a round in  $\text{Sys}_\infty(TA)$ .



■ **Figure 3** The single-round threshold automaton  $\text{TA}^{\text{rd}}$  obtained from PTA in Figure 2

369 We restrict ourselves to fair schedules, that is, those where no transition is applicable  
 370 forever. We also assume that every fair schedule of a single-round system terminates, *i.e.*,  
 371 eventually every process reaches a state from  $\mathcal{B}'$ . Under the fairness assumption we check  
 372 the latter assumption with ByMC [11]. Moreover, we restrict ourselves to non-blocking  
 373 threshold automata, that is, we require that in each configuration each location has at least  
 374 one outgoing rule unlocked. As we use TAs to model distributed algorithms, this is no  
 375 restriction: locations in which no progress should be made unless certain thresholds are  
 376 reached, typically have self-loops that are guarded with `true` (e.g.  $SR$  and  $SP$ ). Thus for  
 377 our benchmarks one can easily check whether they are non-blocking using SMT (we have to  
 378 check that there is no evaluation of the variables such that all outgoing rules are disabled).

379 ► **Definition 5.** Given a PTA  $(\mathcal{L}, \mathcal{V}, \mathcal{R}, RC)$  or its TA  $(\mathcal{L}, \mathcal{V}, \mathcal{R}_{np}, RC)$ , we define a single-  
 380 round threshold automaton  $\text{TA}^{\text{rd}} = (\mathcal{L} \cup \mathcal{B}', \mathcal{V}, \mathcal{R}^{\text{rd}}, RC)$ , where  $\mathcal{B}' = \{\ell' : \ell \in \mathcal{B}\}$  are copies  
 381 of border locations, and  $\mathcal{R}^{\text{rd}} = (\mathcal{R}_{np} \setminus \mathcal{S}) \cup \mathcal{S}' \cup \mathcal{R}^{\text{loop}}$ , where  $\mathcal{R}^{\text{loop}} = \{(\ell', \ell', \text{true}, \vec{0}) : \ell' \in \mathcal{B}'\}$   
 382 are self-loop rules at locations from  $\mathcal{B}'$  and  $\mathcal{S}' = \{(\text{from}, \ell', \text{true}, \vec{0}) : (\text{from}, \ell, \text{true}, \vec{0}) \in$   
 383  $\mathcal{S} \text{ with } \ell' \in \mathcal{B}'\}$  consists of modifications of round switch rules. Initial locations of  $\text{TA}^{\text{rd}}$  are  
 384 the locations from  $\mathcal{B} \subseteq \mathcal{L}$ .

385 For a  $\text{TA}^{\text{rd}}$  and a  $k \in \mathbb{N}_0$  we define a counter system  $\text{Sys}^k(\text{TA}^{\text{rd}})$  as the tuple  $(\Sigma^k, I^k, R^k)$ .  
 386 A configuration is a tuple  $\sigma = (\vec{\kappa}, \vec{g}, \mathbf{p}) \in \Sigma^k$ , where  $\sigma.\vec{\kappa} : \mathcal{D} \rightarrow \mathbb{N}_0$  defines values of the  
 387 counters, for  $\mathcal{D} = (\mathcal{L} \times \{k\}) \cup (\mathcal{B}' \times \{k+1\})$ ; and  $\sigma.\vec{g} : \Gamma \times \{k\} \rightarrow \mathbb{N}_0$  defines shared variable  
 388 values; and  $\sigma.\mathbf{p} \in \mathbb{N}_0^{|\Pi|}$  is a vector of parameter values.

389 Note that by using  $\mathcal{D}$  in the definition of  $\sigma.\vec{\kappa}$  above, every configuration  $\sigma \in \text{Sys}^k(\text{TA}^{\text{rd}})$   
 390 can be extended to a valid configuration of  $\text{Sys}_\infty(\text{TA})$ , by assigning zero to all other counters  
 391 and global variables. In the following, we identify a configuration in  $\text{Sys}^k(\text{TA}^{\text{rd}})$  with its  
 392 extension in  $\text{Sys}_\infty(\text{TA})$ , since they have the same labeling function  $\lambda_k$ , for every  $k \in \mathbb{N}_0$ .

393 We define  $\Sigma_{\mathcal{B}}^k \subseteq \Sigma^k$ , for a  $k \in \mathbb{N}_0$ , to be the set of all configurations  $\sigma$  where every process  
 394 is in a location from  $\mathcal{B}$ , and all shared variables are set to zero in  $k$ , formally,  $\sigma.\vec{g}[x, k] = 0$  for  
 395 all  $x \in \Gamma$ , and  $\sum_{\ell \in \mathcal{B}} \sigma.\vec{\kappa}[\ell, k] = N(\mathbf{p})$ , and  $\sigma.\vec{\kappa}[\ell, i] = 0$  for all  $(\ell, i) \in \mathcal{D} \setminus (\mathcal{B} \times \{k\})$ . We call  
 396 these configurations *border configurations* for  $k$ . The set of initial states  $I^k$  is a subset of  $\Sigma_{\mathcal{B}}^k$ .

397 We define the transition relation  $R$  as in  $\text{Sys}_\infty(\text{TA})$ , *i.e.*, two configurations are in the  
 398 relation  $R^k$  if and only if they (or more precisely, their above described extensions) are in  $R$ .

399 For every  $k \in \mathbb{N}_0$  and every  $\sigma \in \Sigma_{\mathcal{B}}^k$ , there is a corresponding configuration  $\sigma' \in \Sigma_{\mathcal{B}}^0$   
 400 obtained from  $\sigma$  by renaming the round  $k$  to round 0. Let  $f_k$  be the renaming function, *i.e.*,  
 401  $\sigma' = f_k(\sigma)$ . Let us define  $\Sigma^u \subseteq \Sigma_{\mathcal{B}}^0$  to be the union of all renamed initial configurations from  
 402 all rounds, that is,  $\{f_k(\sigma) : k \in \mathbb{N}_0, \sigma \in I^k\}$ .

403 ► **Theorem 6.** Let TA be non-blocking, and let all fair executions of  $\text{Sys}^0(\text{TA}^{\text{rd}})$  w.r.t.  $\Sigma_{\mathcal{B}}^0$   
 404 terminate. Given a formula  $\varphi[i]$  from  $\text{LTL}_X$  over  $\text{AP}_i$ , for a round variable  $i$ , we have

- 405 (A)  $\text{Sys}^0(\text{TA}^{\text{rd}}) \models \mathbf{E}\varphi[0]$  w.r.t. initial configurations  $\Sigma^u$  if and only if  
 406 (B) there exists  $k \in \mathbb{N}_0$  such that  $\text{Sys}_\infty(\text{TA}) \models \mathbf{E}\varphi[k]$ .

407 **Proof sketch.** The theorem is proved using the following arguments. In statement (B), the  
 408 existential quantification over  $k$  corresponds to the definition of  $\Sigma^u$  as union, over all rounds,  
 409 of projections of all reachable initial configurations of that round.

410 Implication (B)  $\rightarrow$  (A) exploits the fact that all rounds are equivalent up to renaming of  
 411 round numbers (with the exception of possible initial configurations).

412 Implication (A)  $\rightarrow$  (B), note that an initial configuration in  $\Sigma^u$  is not necessarily initial in  
 413 round 0, so that one cannot *a priori* take  $k = 0$ . Let us explain how to extend an execution  
 414 of round  $k$  into an infinite execution in  $\text{Sys}_\infty(\text{TA})$ . By termination, all rounds up to  $k-1$   
 415 terminate, so that there is execution that reaches a configuration where all processes are in  
 416 initial locations of round  $k$ . The executions of round  $k$  mimick the ones of round 0 (modulo  
 417 the round number). Finally, the non-blocking assumption is required to be always able to  
 418 extend to infinite executions after round  $k$  is terminated. ◀

419 In Section 4 we showed how to reduce our specifications to formulas of the form  $\forall k \in$   
 420  $\mathbb{N}_0. \mathbf{A}\psi[k]$ . Theorem 6 deals with negations of such formulas, namely with existence of  
 421 a round  $k$  such that formula  $\mathbf{E}\varphi[k]$  holds. Therefore, we can check specifications on the  
 422 single-round system.

## 423 7 Probabilistic Wait-Free Termination

424 We start by defining two conditions that are sufficient to establish Probabilistic Wait-Free  
 425 Termination under round-rigid adversaries. Condition (C1) states the existence of a positive  
 426 probability lower-bound for all processes ending round  $k$  with equal final values. Condition  
 427 (C2) states that if all correct processes start round  $k$  with the same value, then they all will  
 428 decide on that value in that round.

429 (C1) There is a bound  $p \in (0, 1]$ , such that for every round-rigid adversary  $\mathbf{a}$ , and every  $k \in \mathbb{N}_0$ ,  
 430 and every configuration  $\sigma_k$  with parameters  $\mathbf{p}$  that is initial for round  $k$ , it holds that

$$431 \quad \mathbb{P}_{\mathbf{a}}^{\sigma_k} \left( \bigvee_{v \in \{0,1\}} \mathbf{G} \left( \bigwedge_{\ell \in \mathcal{F}_v} \vec{\kappa}[\ell, k] = 0 \right) \right) > p.$$

432 (C2) For every  $v \in \{0, 1\}$ ,  $\forall k \in \mathbb{N}_0. \mathbf{A} \left( \mathbf{G} \bigwedge_{\ell \in \mathcal{I}_{1-v}} \vec{\kappa}[\ell, k] = 0 \rightarrow \mathbf{G} \bigwedge_{\ell' \in \mathcal{F} \setminus \mathcal{D}_v} \vec{\kappa}[\ell', k] = 0 \right)$ .

433 Combining (C1) and (C2), under every round-rigid adversary, from any initial configuration  
 434 of round  $k$ , the probability that all correct processes decide before end of round  $k+1$  is at  
 435 least  $p$ . Thus the probability not to decide within  $2n$  rounds is at most  $(1-p)^n$ , which tends  
 436 to 0 when  $n$  tends to infinity. This reasoning follows the arguments of [1].

437 ▶ **Proposition 7.** *If  $\text{Sys}_\infty(\text{PTA}) \models (C1) \wedge (C2)$ , then  $\text{Sys}_\infty(\text{PTA}) \models (3)$ .*

438 Observe (C2) is a non-probabilistic property of the same form as (5), so that we can  
 439 check (C2) using the method of Section 6.

440 In the rest of this section, we detail how to reduce the verification of (C1), to a verification  
 441 task that can be handled by ByMC. First observe that (C1) contains a single round variable,  
 442 and recall that we restrict to round-rigid adversaries, so that it is sufficient to check them  
 443 (omitting the round variables) on the single-round system. We introduce analogous objects  
 444 as in the non-probabilistic case:  $\text{PTA}^{\text{rd}}$  (analogously to Definition 5), and its counter  
 445 system  $\text{Sys}(\text{PTA}^{\text{rd}})$ .

446 **7.1 Reducing probabilistic to non-probabilistic specifications**

447 Since probabilistic transitions end in final locations, they cannot appear on a cycle in  $\text{PTA}^{\text{rd}}$ .  
 448 Therefore, for fixed parameter valuation  $\mathbf{p}$ , any path contains at most  $N(\mathbf{p})$  probabilistic  
 449 transitions, and its probability is therefore uniformly lower-bounded. As a consequence:

450 **► Lemma 8.** *Let  $\mathbf{p} \in \mathbf{P}_{RC}$  be a parameter valuation. In  $\text{Sys}(\text{PTA}^{\text{rd}})$ , for every LTL formula  $\varphi$   
 451 over atomic proposition  $\text{AP}$ , the following two statements are equivalent:*

- 452 (a)  $\exists p > 0, \forall \sigma \in I_{\mathbf{p}}, \forall \mathbf{a} \in \mathcal{A}^R. \mathbb{P}_{\mathbf{a}}^{\sigma}(\varphi) > p,$   
 453 (b)  $\forall \sigma \in I_{\mathbf{p}}, \forall \mathbf{a} \in \mathcal{A}^R, \exists \pi \in \text{paths}(\sigma, \mathbf{a}). \pi \models \varphi.$

454 **7.2 Verifying (C1) on a non-probabilistic TA**

455 According to Lemma 8, in order to prove (C1), we only need to prove that in the sys-  
 456 tem  $\text{Sys}(\text{PTA}^{\text{rd}})$  it holds that

457 
$$\forall \sigma \in I_{\mathbf{p}}, \forall \mathbf{a} \in \mathcal{A}^R, \exists \pi \in \text{paths}(\sigma, \mathbf{a}). \pi \models \bigvee_{v \in \{0,1\}} \mathbf{G} \left( \bigwedge_{\ell \in \mathcal{F}_v} \bar{\kappa}[\ell] = 0 \right). \quad (6)$$

458 As in Section 6, it is possible to modify  $\text{PTA}^{\text{rd}}$  into a non-probabilistic TA, by replacing  
 459 probabilistic choices by non-determinism. Still, the quantifier alternation of (6) (universal  
 460 over initial configurations and adversaries vs. existential on paths) is not in the fragment  
 461 handled by ByMC [11]. Once an initial configuration  $\sigma$  and an adversary  $\mathbf{a}$  are fixed, the  
 462 remaining branching is induced by non-Dirac transitions. By assumption, these transitions  
 463 lead to final locations only, to both  $\mathcal{F}_0$  and  $\mathcal{F}_1$ , and under round-rigid adversaries, they are  
 464 the last transitions to be fired. To prove (6), it is sufficient to prove that all processes that  
 465 fire only Dirac transitions will reach final locations of the same type ( $\mathcal{F}_0$  or  $\mathcal{F}_1$ ). If this is the  
 466 case, then the existence of a path corresponds to all non-Dirac transitions being resolved in  
 467 the same way. This allows us to remove the non-Dirac transitions from the model as follows.  
 468 Given a  $\text{PTA}^{\text{rd}}$ , we define a threshold automaton  $\text{TA}^{\text{m}}$  with locations  $\mathcal{L}$  (without  $\mathcal{B}'$ ) such  
 469 that for every non-Dirac rule  $r = (\text{from}, \delta_{to}, \varphi, \vec{u})$  in  $\text{PTA}$ , all locations  $\ell$  with  $\delta_{to}(\ell) > 0$  are  
 470 merged into a new location  $\ell^{\text{mrg}}$  in  $\text{TA}^{\text{m}}$ . Note that this location must belong to  $\mathcal{F}$ . Naturally,  
 471 instead of a non-Dirac rule  $r$  we obtain a Dirac rule  $(\text{from}, \ell^{\text{mrg}}, \varphi, \vec{u})$ . Also we add self-loops  
 472 at all final locations. Paths in  $\text{Sys}(\text{TA}^{\text{m}})$  correspond to prefixes of paths in  $\text{Sys}(\text{PTA}^{\text{rd}})$ . In  
 473  $\text{Sys}(\text{TA}^{\text{m}})$ , from a configuration  $\sigma$ , an adversary  $\mathbf{a}$  yields a unique path, that is,  $\text{paths}(\sigma, \mathbf{a})$  is  
 474 a singleton set. Thus, the existential quantifier from (6) can be replaced by the universal one.

475 Finally, proving (6) on  $\text{PTA}^{\text{rd}}$  reduces to proving that  $\mathbf{A} \bigvee_{v \in \{0,1\}} \mathbf{G} (\bigwedge_{\ell \in \mathcal{F}_v} \bar{\kappa}[\ell] = 0)$   
 476 holds on  $\text{Sys}(\text{TA}^{\text{m}})$ , which can be checked automatically by ByMC.

477 **8 Experiments**

478 We have applied the approach presented in Sections 4–7 to five randomized fault-tolerant  
 479 distributed algorithms. (The benchmarks and the instructions on running the experiments  
 480 are available from: <https://forsyte.at/software/bymc/artifact42/>)

- 481 1. Protocol 1 for randomized consensus by Ben-Or [3], with two kinds of crashes: clean  
 482 crashes (**ben-or-cc**), for which a process either sends to all processes or none, and dirty  
 483 crashes (**ben-or-dc**), for which a process may send to a subset of processes. This algorithm  
 484 works correctly when  $n > 2t$ .  
 485 2. Protocol 2 for randomized Byzantine consensus (**ben-or-byz**) by Ben-Or [3]. This algorithm  
 486 tolerates Byzantine faults when  $n > 5t$ .

Label	Name	Automaton	Formula
S1	agreement_0	N	$\mathbf{A G}(\neg \text{EX}\{D0\}) \vee \mathbf{G}(\neg \text{EX}\{D1, E1\})$
S2	validity_0	N	$\mathbf{A ALL}\{V0\} \rightarrow \mathbf{G}(\neg \text{EX}\{D1, E1\})$
S3	completeness_0	N	$\mathbf{A ALL}\{V0\} \rightarrow \mathbf{G}(\neg \text{EX}\{D1, E1\})$
S4	round-term	N	$\mathbf{A fair} \rightarrow \mathbf{F ALL}\{D0, D1, E0, E1, CT\}$
S5	decide-or-flip	P	$\mathbf{A fair} \rightarrow \mathbf{F}(\text{ALL}\{D0, E0, CT\} \vee \text{ALL}\{D1, E1, CT\})$
S1'	sim-agreement	N	$\mathbf{A G}(\neg \text{EX}\{D0, E0\} \vee \neg \text{EX}\{D1, E1\})$
S1''	2-agreement	N	$\mathbf{A G}(\neg \text{EX}\{D0, E0\} \vee \neg \text{EX}\{D1, E1\} \vee \neg \text{EX}\{D2, E2\})$

■ **Table 1** Temporal properties verified in our experiments for value 0.

- 487 3. Protocol 2 for randomized consensus (*rab-c*) by Bracha [5]. It runs as a high-level  
488 algorithm together with a low-level broadcast that turns Byzantine faults into “little  
489 more than fail-stop (faults)”. We check only the high-level algorithm for clean crashes.
- 490 4.  $k$ -set agreement for crash faults (*kset*) by Raynal [19], for  $k = 2$ . This algorithm works in  
491 presence of clean crashes when  $n > 3t$ .
- 492 5. Randomized Byzantine one-step consensus (*rs-bosco*) by Song and van Renesse [21]. This  
493 algorithm tolerates Byzantine faults when  $n > 3t$ , and it terminates fast when  $n > 7t$  or  
494  $n > 5t$  and  $f = 0$ .

495 Following the reduction approach of Sections 4–7, for each benchmark, we have encoded  
496 two versions of one-round threshold automata: an N-automaton that models a coin toss  
497 by a non-deterministic choice in a coin-toss location, and is used for the non-probabilistic  
498 reasoning, and a P-automaton that never leaves the coin-toss location and which is used to  
499 prove probabilistic wait-free termination. Both automata are given as the input to Byzantine  
500 Model Checker (ByMC) [11], which implements the parameterized model checking techniques  
501 for safety [8] and liveness [9] of counter systems of threshold automata.

502 Both automata follow the pattern shown in Figure 2: They start in one of the initial  
503 locations (e.g.,  $V_0$  or  $V_1$ ), progress by switching locations and incrementing shared variables  
504 and end up in a location that corresponds to a decision (e.g.,  $D_0$  or  $D_1$ ), an estimate of a  
505 decision (e.g.,  $E_0$  or  $E_1$ ), or a coin toss (CT).

506 Table 1 summarizes the temporal properties that were verified in our experiments. Given  
507 the set of all possible locations  $\mathcal{L}$ , a set  $Y = \{\ell_1, \dots, \ell_m\} \subseteq \mathcal{L}$  of locations, and the designated  
508 crashed location  $\text{CR} \in \mathcal{L}$ , we use the shorthand notation:  $\text{EX}\{\ell_1, \dots, \ell_m\}$  for  $\bigvee_{\ell \in Y} \vec{\kappa}[\ell] \neq 0$   
509 and  $\text{ALL}\{\ell_1, \dots, \ell_m\}$  for  $\bigwedge_{\ell \in \mathcal{L} \setminus Y} (\vec{\kappa}[\ell] = 0 \vee \ell = \text{CR})$ . For *rs-bosco* and *kset*, instead of  
510 checking S1, we check S1' and S1''.

511 Table 2 presents the computational results of our experiments. The meaning of the  
512 columns is as follows: column  $|\mathcal{L}|$  shows the number of automata locations, column  $|\mathcal{R}|$  shows  
513 the number of automata rules, column  $|\mathcal{S}|$  shows the number of enumerated schemas (which  
514 depends on the structure of the automaton and the specification), column *time* shows the  
515 computation times — either in seconds or in the format HH:MM. For  $|\mathcal{R}|$ , we give the figures  
516 for the N-automata, since they have more rules in addition to the rules in P-automata. To  
517 save space, we omit the figures for memory use from the table: Benchmarks 1–5 need 30–170  
518 MB RAM, whereas *rs-bosco* needs up to 1.5 GB RAM per cluster node.

519 The benchmark *rs-bosco* presents a challenge for the schema enumeration technique of [9]:  
520 Its threshold automaton contains 12 threshold guards that can change their values almost in  
521 any order. Additional combinations are produced by the temporal formulas. ByMC reduces  
522 the number of combinations by analyzing dependencies between the guards. However, this

■ **Table 2** The experiments for first 5 rows were run on a single computer (Apple MacBook Pro 2018, 16GB). The experiments for last row (*rs-bosco*) were run in Grid5000 on 32 nodes (2 CPUs Intel Xeon Gold 6130, 16 cores/CPU, 192GB). Wall times are given.

Automaton			S1/S1'/S1''		S2		S3		S4		S5	
Name	$ \mathcal{L} $	$ \mathcal{R} $	$ \mathcal{S} $	Time								
<b>ben-or-cc</b>	10	27	9	1	5	0	5	0	5	0	5	0
<b>ben-or-dc</b>	10	32	9	1	5	1	5	0	5	0	5	1
<b>ben-or-byz</b>	9	18	3	1	2	0	2	0	2	0	2	1
<b>rabc-cr</b>	11	31	9	0	5	1	5	1	5	0	5	0
<b>kset</b>	13	58	65	3	65	17	65	12	65	39	65	40
<b>rs-bosco</b>	19	48	156M	3:21	156M	3:02	156M	3:21	n/a	n/a	156M	3:43

523 benchmark requires to enumerate between  $11!$  and  $14!$  schemas. To this end, we have run the  
 524 verification experiments for *rs-bosco* on 1024 CPU cores of the computing cluster Grid5000.  
 525 Table 2 presents the wall time results for *rs-bosco*, that is, the actual number of computation  
 526 hours on all the cores is the wall time multiplied by 1024.

527 For all the benchmarks in Table 2, ByMC has reported that the specifications hold. By  
 528 changing  $n > 3t$  to  $n > 2t$ , we found that *rabc-cr* can handle more faults (the original  $n > 3t$   
 529 was needed to implement the underlying communication structure which we assume given in  
 530 the experiments). In other cases, whenever we changed the parameters, that is, increased  
 531 the number of faults beyond the known bound, the tool reported a counterexample.

## 532 9 Conclusions

533 We lifted the threshold automata framework to multi-round randomized algorithms. We  
 534 proved a reduction that allows us to check  $LTL_X$  specifications over propositions for one round  
 535 in a single-round automaton so that the verification results transfer directly to the multi-  
 536 round counter system. Using round-based compositional reasoning, we have shown that this  
 537 is sufficient to check specifications that span multiple rounds, *e.g.*, agreement. Almost-sure  
 538 termination under round-rigid adversaries relies on a distinct reduction argument.

539 By experimental evaluation we showed that the verification conditions that came out of  
 540 our reduction can be automatically verified for several challenging randomized consensus  
 541 algorithms in the parameterized setting.

542 Our proof methodology for almost sure termination applies to round-rigid adversaries  
 543 only. As future work we shall prove that verifying almost-sure termination under round-rigid  
 544 adversaries is sufficient to prove it for more general adversaries. Transforming an adversary  
 545 into a round-rigid one while preserving the probabilistic properties over the induced paths,  
 546 comes up against the fact that, depending on the outcome of a coin toss in some step at  
 547 round  $k$ , different rules may be triggered later for processes in rounds less than  $k$ .

548 There are few contributions that address the automated verification of probabilistic  
 549 parameterized systems [20, 4, 18, 17]. In contrast with these, our processes are not finite-  
 550 state, due to the round numbers and parameterized guard expressions. The seminal work of  
 551 Pnueli and Zuck restricts to bounded shared variables and cannot use arithmetic thresholds  
 552 (different from 1 and  $n$ ) [20]. Algorithms based on well-structured transition systems [4]  
 553 do not directly apply to systems with several parameters, such as the ones generated by  
 554 probabilistic threshold automata. Regular model checking techniques [18, 17] cannot handle  
 555 arithmetic resilience conditions such as  $n > 3t$ , nor unbounded shared variables.

556 ——— **References** ———

- 557 **1** Marcos Aguilera and Sam Toueg. The correctness proof of Ben-Or’s randomized consensus  
558 algorithm. *Distributed Computing*, pages 1–11, 2012. online first.
- 559 **2** Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 560 **3** Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement  
561 protocols (extended abstract). In *PODC*, pages 27–30, 1983.
- 562 **4** Nathalie Bertrand and Paulin Fournier. Parameterized verification of many identical proba-  
563 bilistic timed processes. In *FSTTCS*, volume 24 of *LIPICs*, pages 501–513, 2013.
- 564 **5** Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143,  
565 1987.
- 566 **6** Tzilla Elrad and Nissim Francez. Decomposition of distributed programs into communication-  
567 closed layers. *Sci. Comput. Program.*, 2(3):155–173, 1982.
- 568 **7** Michael J. Fischer, Nancy A. Lynch, and M. S. Paterson. Impossibility of distributed consensus  
569 with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- 570 **8** Igor Konnov, Marijana Lazić, Helmut Veith, and Josef Widder. Para<sup>2</sup>: Parameterized path  
571 reduction, acceleration, and SMT for reachability in threshold-guarded distributed algorithms.  
572 *Formal Methods in System Design*, 51(2):270–307, 2017.
- 573 **9** Igor Konnov, Marijana Lazić, Helmut Veith, and Josef Widder. A short counterexample  
574 property for safety and liveness verification of fault-tolerant distributed algorithms. In *POPL*,  
575 pages 719–734, 2017.
- 576 **10** Igor Konnov, Helmut Veith, and Josef Widder. On the completeness of bounded model checking  
577 for threshold-based distributed algorithms: Reachability. *Information and Computation*, 252:95–  
578 109, 2017.
- 579 **11** Igor Konnov and Josef Widder. ByMC: Byzantine model checker. In *ISoLA (3)*, volume 11246  
580 of *LNCS*, pages 327–342. Springer, 2018.
- 581 **12** Jure Kukovec, Igor Konnov, and Josef Widder. Reachability in parameterized systems: All  
582 flavors of threshold automata. In *CONCUR*, pages 19:1–19:17, 2018.
- 583 **13** Marta Z. Kwiatkowska and Gethin Norman. Verifying randomized byzantine agreement. In  
584 *FORTE*, pages 194–209, 2002.
- 585 **14** Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: Verification of  
586 probabilistic real-time systems. In *CAV*, pages 585–591, 2011.
- 587 **15** Marta Z. Kwiatkowska, Gethin Norman, and Roberto Segala. Automated verification of a  
588 randomized distributed consensus protocol using Cadence SMV and PRISM. In *CAV*, pages  
589 194–206, 2001.
- 590 **16** Daniel J. Lehmann and Michael O. Rabin. On the advantages of free choice: A symmetric and  
591 fully distributed solution to the dining philosophers problem. In *POPL*, pages 133–138, 1981.
- 592 **17** Ondrej Lengál, Anthony Widjaja Lin, Rupak Majumdar, and Philipp Rümmer. Fair termination  
593 for parameterized probabilistic concurrent systems. In *TACAS*, volume 10205 of *LNCS*, pages  
594 499–517, 2017. doi:10.1007/978-3-662-54577-5\_29.
- 595 **18** Anthony Widjaja Lin and Philipp Rümmer. Liveness of randomised parameterised systems  
596 under arbitrary schedulers. In *CAV*, volume 9780 of *LNCS*, pages 112–133. Springer, 2016.  
597 doi:10.1007/978-3-319-41540-6\_7.
- 598 **19** Achour Mostéfaoui, Hamouma Moumen, and Michel Raynal. Randomized k-set agreement in  
599 crash-prone and Byzantine asynchronous systems. *Theor. Comput. Sci.*, 709:80–97, 2018.
- 600 **20** Amir Pnueli and Lenore D. Zuck. Verification of multiprocess probabilistic protocols. *Dis-  
601 tributed Computing*, 1(1):53–72, 1986. doi:10.1007/BF01843570.
- 602 **21** Yee Jiun Song and Robbert van Renesse. Bosco: One-step Byzantine asynchronous consensus.  
603 In *DISC*, volume 5218 of *LNCS*, pages 438–450, 2008.