

On automatic verification of self-stabilizing population protocols

Jun PANG^{1,2}, Zhengqin LUO³, Yuxin DENG (✉)³

¹ Department of Computer Science and Communications, Université du Luxembourg, Luxembourg 1359, Luxembourg

² State Key Laboratory of Novel Software Technology, Nanjing University, Nanjing 210093, China

³ Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China

© Higher Education Press and Springer-Verlag 2008

Abstract The population protocol model has emerged as an elegant computation paradigm for describing mobile *ad hoc* networks, consisting of a number of mobile nodes that interact with each other to carry out a computation. The interactions of nodes are subject to a fairness constraint. One essential property of population protocols is that all nodes must eventually converge to the correct output value (or configuration). In this paper, we aim to automatically verify self-stabilizing population protocols for leader election and token circulation in the Spin model checker. We report our verification results and discuss the issue of modeling strong fairness constraints in Spin.

Keywords distributed algorithms, model checking, population protocols, verification

1 Introduction

The field of distributed algorithms has enjoyed a rapid growth in the last two decades, due to the world-wide development and usage of mobile *ad hoc* networks. A great number of algorithms have been invented to solve hard problems in mobile *ad hoc* networks. However, these algorithms are only accessible to the distributed algorithms community, since their specifications and correctness arguments are often given at an informal level. This is insufficient to convince researchers outside the field of the validity of the arguments. If one wants to verify the correctness of some proofs, he has to prove substantial parts or entire sub-results, for which only informal arguments were given. This has been observed and illustrated in a recent paper [1] on formal reasoning about the correctness of a distributed consensus algorithm [2].

The last two decades have also seen an impressive amount of new techniques developed in the area of formal verification or model checking in particular. Model checking first builds a finite state space of a formal model of a system, and then verifies if a property, written in some temporal logic, about the system holds or not through an explicit state space search. Due to the finiteness of the state space, the search always terminates. Hence, model checking is largely automatic. It can produce an answer in a few minutes or even seconds for many models. A counterexample can be generated when the checked property fails to hold, which details why the formal model doesn't satisfy the property. Model checkers usually face a combinatorial blow up of the state-space, known as the state explosion problem. Techniques such as symbolic representation, symmetry reduction, and predicate abstraction, have been developed to deal with the state explosion problem and enhance the scalability of model checking. However, these techniques have not yet made impact on distributed algorithms, mainly because there have not yet been enough examples of non-trivial practical applications.

Clearly, both two fields, distributed algorithms vs. formal verification, can benefit from each other. On the one hand, formal verification can offer techniques to well-understand distributed algorithms. On the other hand, distributed algorithms can offer challenging examples to formal verification.

In this paper, we aim to automatically verify self-stabilizing population protocols for leader election and token circulation in the Spin model checker [3]. The population protocol model [4] has emerged as a new elegant computation paradigm for describing mobile *ad hoc* networks, consisting of a number of mobile nodes that interact with each other to carry out a computation. Each node has only a few states. One essential property of such protocols is that all nodes must eventually converge to the correct output value (or configuration), which is a typical

liveness property (something good will eventually happen) in terms of formal verification. To guarantee that such kind of property can be achieved, the interactions of nodes in population protocols are subject to a fairness constraint. The fairness condition is imposed on the adversary to ensure that the protocol makes progress. In population protocols, the required fairness condition will make the system behave nicely eventually, although it can behave arbitrarily for an arbitrarily long period [4]. That is why for population protocols correctness arguments are always rephrased as a property to be satisfied eventually. In formal verification, fairness is typically used to rule out some unrealistic runs due to non-determinism, i.e., it mainly concerns with a fair resolution of non-determinism in the models. So unsurprisingly, fairness has been a research topic to both communities, see e.g. Refs. [5–9].

In next section we review the basic population protocol model and the fairness conditions which are required for population protocols. The general framework for modeling population protocols in Spin is given in Section 3. In Section 4 and Section 5, we discuss experiment results on automatic verification of self-stabilizing leader election for complete graphs [10] and token circulation for directed rings [11], respectively. For leader election in complete graphs, we show that the algorithm also works under a weaker fairness condition. For token circulation in directed rings, the algorithm is model checked in a two-phase manner. We first show that under a particular activation order of nodes, satisfying the global fairness condition (see its definition in Section 2), some pre-defined safe configurations will be eventually reached. Then we show that from these safe configurations eventually token circulation is stabilized. In Section 6, we demonstrate that global fairness generally assumed for population protocols is necessary. We present counterexamples that we have observed in Spin to show that self-stabilizing token circulation and self-stabilizing leader election in directed rings cannot be achieved with local fairness (see its definition in Section 2). Finally, we discuss the difficulty of modeling global fairness in Spin and conclude the paper by pointing out some possible future work in Section 7.

Our contribution can be summarized as follows:

- We present a general framework for modeling population protocols in Spin;
- We successfully model check the self-stabilizing leader election in complete graphs under such a weaker fairness condition and the self-stabilizing token circulation in directed rings without explicitly encoding global fairness;
- We use Spin to automatically generate counterexamples to show why local fairness is insufficient for token circulation and leader election in directed rings.

2 The population protocol model

We briefly introduce the population protocol model in this section, and more details are available in Refs. [10,11].

2.1 Model and definitions

In our framework, the underlying network can be described by a directed graph $G = (V, E)$ without multi-edges and self-loops. Each vertex represents a simple finite-state sensing device, and each edge (u, v) means that u as an initiator could possibly interact with v as a responder.

A protocol is specified as a tuple $P(Q, \mathcal{C}, X, Y, O, \delta)$ which contains

- a finite set Q of states,
- a set \mathcal{C} of configurations,
- a finite set X of input symbols,
- a finite set Y of output symbols,
- an output function $O : Q \rightarrow Y$, and
- a transition function $\delta : (Q \times Y) \times (Q \times Y) \rightarrow 2^{Q \times Q}$.

If $(p', q') \in \delta((p, x), (q, y))$, then we write $((p, x), (q, y)) \rightarrow (p', q')$ and call it a transition. When δ always maps to a set that only contains a single pair of states, then we call the protocol and the transition function deterministic.

A configuration C is a mapping $C : V \rightarrow Q$ assigning to each node its internal state, and an input assignment $\alpha : V \rightarrow X$ specifies the input for each node. Let C and C' be configurations, α be an input assignment, and u, v be different nodes. If there is a pair

$$(C'(u), C'(v)) \in \delta((C(u), \alpha(u)), (C(v), \alpha(v))),$$

we say that C goes to C' via edge $e = (u, v)$ by transition

$$((C(u), \alpha(u)), (C(v), \alpha(v))) \rightarrow (C'(u), C'(v)),$$

abbreviated to $(C, \alpha) \xrightarrow{e} C'$. A pair of a transition r and an edge e constitutes an action $\sigma = (r, e)$. If C goes to C' via some edge, then C can go to C' in one step, written as $(C, \alpha) \rightarrow C'$.

An execution is an infinite sequence of configurations and assignments

$$(C_0, \alpha_0), (C_1, \alpha_1), \dots, (C_i, \alpha_i), \dots,$$

such that $C_0 \in \mathcal{C}$ and for each i , $(C_i, \alpha_i) \rightarrow C_{i+1}$.

2.2 Fairness conditions

In the following, we first summarize the fairness conditions for population protocols. Let

$$E = (C_0, \alpha_0), (C_1, \alpha_1), \dots, (C_i, \alpha_i), \dots$$

be an execution. Two different fairness conditions [10] are defined below:

Definition 2.1 (Global fairness) For every C, α , and C' such that $(C, \alpha) \rightarrow C'$, if $(C_i, \alpha_i) = (C, \alpha)$ for infinitely many i , then $(C_i, \alpha_i) = (C, \alpha)$ and $C_{i+1} = C'$ for infinitely many i . (Hence, the step $(C, \alpha) \rightarrow C'$ is taken infinitely many times in E .)

Definition 2.2 (Local fairness) For every action σ , if σ is enabled in (C_i, α_i) for infinitely many i , then $(C_i, \alpha_i) \xrightarrow{\sigma} C_{i+1}$ for infinitely many i . (Hence, the action σ is taken infinitely many times in E .)

It should be noticed that global fairness is strictly stronger than local fairness [10]. In population protocol model, steps specify how the whole protocol transforms from one configuration to another configuration, and actions specify the interactions between two nodes and only depend on the local states of the two interacting nodes. Global fairness requires that each step that can be taken infinitely often is actually taken infinitely often, while local fairness asserts that each action which is enabled infinitely often is actually taken infinitely often. Since one action can be enabled in different configurations, the global fairness condition insists that an action should be taken infinitely often in all such configurations, whereas the local fairness condition only requires that it occurs infinitely often in one of such configurations.

Let us illustrate the two different fairness conditions via an example. As shown in Fig. 1, the system contains two configurations C and C' , where action σ_1 is enabled in both configurations. An infinite execution containing the following loop

$$(C, \alpha) \xrightarrow{\sigma_1} (C, \alpha) \xrightarrow{\sigma_2} (C', \alpha') \xrightarrow{\sigma_4} (C', \alpha') \xrightarrow{\sigma_3} (C, \alpha)$$

does not satisfy global fairness condition, since σ_1 is enabled in configuration C' infinitely many times, but it is never performed in C' . However, the execution satisfies local fairness. On the other hand, an infinite execution containing the following loop meets both global and local fairness.

$$\begin{aligned} (C, \alpha) &\xrightarrow{\sigma_1} (C, \alpha) \xrightarrow{\sigma_2} (C', \alpha') \xrightarrow{\sigma_1} \\ (C', \alpha') &\xrightarrow{\sigma_4} (C', \alpha') \xrightarrow{\sigma_3} (C, \alpha) \end{aligned}$$




Fig. 1 Example of global and local fairness

In Ref. [10], two extra weak notions of fairness conditions are presented. The weak forms of fairness do not insist that particular steps occur infinitely often in E but only that the configurations that would result from those steps occur infinitely often. The relationship between these four kinds of fairness conditions are discussed thoroughly in Ref. [10].

As discussed before, in the area of formal verification, fairness is typically needed to prove liveness properties. It is concerned with a fair resolution of non-determinism, i.e., fairness conditions are used to rule out some unrealistic runs due to non-determinism. Usually, in formal verification a strong fairness condition states that if an activity is infinitely often enabled then it has to be executed infinitely often. This can be mapped into the population protocol model as global fairness and local fairness above, depending that the activity is either one step or one action. There is another notion of weak fairness in verification, stating that if an activity is continuously often enabled (no temporary disabling) then it has to be executed infinitely often. This notion of weak fairness is strictly weaker than those fairness conditions mentioned in Ref. [10].

Definition 2.3 (Weak fairness) For every action σ , if there exists i and for all $j > i$, σ is always enabled at (C_j, α_j) in E , then there exist C, α, C' such that $(C, \alpha) \xrightarrow{\sigma} C'$ occurs infinitely often.

This form of fairness is supported in Spin (see Section 3 for more details). Consider the example in Fig. 1 again, an infinite execution containing the following loop

$$(C, \alpha) \xrightarrow{\sigma_1} (C, \alpha) \xrightarrow{\sigma_2} (C', \alpha') \xrightarrow{\sigma_3} (C, \alpha)$$

satisfies only weak fairness, as σ_3 is enabled infinitely many times, but it is never performed. Moreover, an infinite execution containing the following loop does not meet any fairness condition mentioned in this paper.

$$(C, \alpha) \xrightarrow{\sigma_2} (C', \alpha') \xrightarrow{\sigma_3} (C, \alpha)$$

In particular, it is not weak fair since σ_1 is continuously often enabled, but it is never executed within the loop. Note that we cannot have a similar fairness condition for population protocols in which an action becomes enabled forever with respect to one concrete configuration, since the configuration of the system will be updated by other actions.

3 Automatic verification in Spin

The Spin model checker is a popular tool set for verification of concurrent systems [3,12]. A modeling language Promela (Process Meta Language) is used to specify a

concurrent system consisting of some processes that are the basic dynamic system components. Given a model described in Promela, Spin can either run random simulations to check the validation of functional behavior, or generate efficient C programs to verify the correctness with respect to some constraint conditions. Its verifier can find non-progress cycles, or verify general properties which are expressed by linear temporal logic (LTL) formulas. The verifier also provides an option for weak fairness among processes, that is, if a process is eventually permanently enabled in the run, then the process is executed infinitely often in the run. The reader is referred to Ref. [13] for more details.

In the population protocol model, one protocol consists of N nodes, numbered from 0 to $N - 1$. The protocol is usually described by a set of interaction rules (see e.g. Fig. 2). On the left hand side of each rule, the state and the input of the initiator and the responder should be matched by the rules. On the right hand side, the rule specifies the state of the initiator and the responder after the transition has been taken.

```

Rule 1.  ((L, *), (L, *)) → ((L), (-))
Rule 2.  ((-, F), (-, *)) → ((L), (-))
Rule 3.  ((-, T), (-, *)) → ((-), (-))
Each node outputs its own state.

```

Fig. 2 Algorithm for self-stabilizing leader election in complete graphs

Since the population protocols always depend on some kinds of fairness condition, such as global fairness and local fairness, we attempt to use the weak fairness condition (at process level) in Spin to model a fairness condition in population protocols. However, the fairness condition in population protocol model is related to actions/steps but not to nodes (processes). Thus, if we use a single process in Promela to model a single node of the population protocol model, then the weak fairness condition only guarantees that if a node is from some moment onward always enabled to interact with other nodes, then it will interact infinitely often. This obviously does not make any sense when verifying population protocols. Our strategy is to use a single process to represent an action which is related to an initiator, a responder and a transition rule. The process declaration for an action in Promela is described in the following way:

```
proctype Rulen(int i; int r)
```

The parameters i and r are identities of an initiator and a responder. The process name $Rulen$ corresponds to one concrete Rule n in the protocol. The state of each node is stored in some global variables. The process will check the global configuration to decide the executability of its own. For example, $Rule1(1,2)$ represents

an action between node 1 as an initiator and node 2 as a responder according to Rule 1. The entire system consists of all possible actions of every pair of nodes which can interact with each other.

```

run Rule1(0,1);
run Rule1(1,0);
run Rule1(1,2);
...
run Rule2(2,0);
...

```

The different communication patterns are determined by different network topologies, such as complete graphs and directed rings.

After modeling each possible action (between two nodes) as one process, if we use the weak fairness condition in the Spin model checker, we immediately obtain a condition strictly weaker than local fairness in the population protocol model. The condition is exactly weak fairness as discussed in Section 2. Otherwise we have to use a large LTL formula and auxiliary variables to characterize the strong (global and local) fairness conditions which will increase the complexity of the model (see more discussion in Section 7). Weak fairness only assumes that if an action is permanently enabled from some point, then it will be taken infinitely many times. Most self-stabilizing population protocols require either global or local fairness. However, we have found that some of them also work properly under this weaker fairness condition implemented in the Spin model checker. In the following sections, protocols will be verified under this condition.

Once a model has been built, we could define a bunch of propositions which refer to different protocol configurations. Finally, LTL formulas over these propositions can be used to specify some desired behaviors of the protocol. The LTL formulas will be translated into never claims in Prolema automatically, and verified by the Spin model checker.

4 Self-stabilizing leader election in complete graphs

A distributed system or a population protocol is said to be self-stabilizing [14] if it satisfies the following two properties:

- convergence: starting from an arbitrary configuration, the system is guaranteed to reach a correct configuration;
- closure: once the system reaches a correct configuration, it cannot become incorrect any more.

In this section, we show that self-stabilizing leader election in complete graphs can be achieved under the

weak fairness with the help of an eventually correct leader detector. The algorithm was originally given in Ref. [10]. Every node has one bit memory which represents two states, being a leader (L) or not ($-$). The leader detector gives each node an input true (T) or false (F) to indicate that whether there is a leader in the network. The detector may give wrong answers sometimes, but it will eventually return a correct answer permanently. A non-leader becomes a leader, when the leader detector signals the absence of a leader, and the responder is not a leader. When two leaders interact, the responder becomes a non-leader. Otherwise, no state change occurs.

The algorithm is described by the three interaction rules in Fig. 2. On the left hand side, the state and input of an initiator and a responder should be matched. The symbol “*” denotes that the input can always be matched. On the right hand side, the state of the two nodes would be updated by the rule.

In Ref. [10], it has already been shown that the algorithm implements self-stabilizing leader election in complete graphs under both global and local fairness, provided the existence of an eventual leader detector. Here, we have new verification result to show that the algorithm is correct even under the weak fairness condition.

4.1 Modeling leader election in Spin

The model for leader election in complete graphs follows the general paradigm of population protocol in Section 3, only with some additional definition issues.

The states of the whole system are represented by an array of bits $leader[N]$, in which N is the number of nodes in the network. When $leader[i]$ equals to 1, it indicates that node i claims to be a leader. Since we are modeling self-stabilizing protocols, we have to ensure that the protocol is correct starting from any arbitrary initial configuration. We employ atomic sequences and case selection in Promela to assign all possible values for every state variable in a single step.

```

atomic{
  if
    :: atomic{true -> leader[0]=0}
    :: atomic{true -> leader[0]=1}
  fi;
  ...
}
```

Thus, at the beginning of the protocol, every state variable could be assigned with all possible values. The verifier will check all these cases to ensure the self-stabilizing property of the system.

Besides, we have to model the eventually correct leader detector in the protocol. The detector is defined by two parts. First, there is a random process which randomly generates answers (encoded in the variable $detector$) when the detector is in “incorrect” state ($detectorcorrect == 0$):

```

proctype RandomDetector(){
  do
    :: (detectorcorrect == 0) -> detector=false;
    :: (detectorcorrect == 0) -> detector=true;
  od
}
```

Then, we define another process that can switch the detector’s state from “incorrect” to “correct” in a non-deterministic way ($detectorcorrect = 1$). The progress label ensures that the transition will finally occur.

```

proctype DetectorCorrect(){
  do
    :: (detectorcorrect == 0) ->
      progress : detectorcorrect=1;
  od
}
```

Once $detectorcorrect$ becomes true, the value of $detector$ will depend on whether the sum of $leader[i]$ ($0 \leq i < N$) is greater than 0.

Having defined the model, the LTL formula which specifies the desirable system behavior is relatively small. Under the weak fairness, the LTL formula for leader election in complete graph is simply as follows:

$$\langle \rangle [] oneLeader$$

This LTL formula says that along every path which satisfies the weak fairness condition a unique leader will eventually be elected. (Informally, $\langle \rangle$ reads as “eventually”, and $[]$ as “always”.) Here, $oneLeader$ is the proposition stating that the sum of all $leader[i]$ equals one. See Ref. [15] for the detailed model.

4.2 Verification results

It has been shown in Ref. [10] that the algorithm is valid under local fairness, and the fact that global fairness implies the local fairness condition yields that the algorithm is also valid under the global fairness condition. However, weak fairness is weaker than both of them. Thus it is interesting to see if the algorithm is still correct under such this weaker fairness condition. Surprisingly, the algorithm indeed implements self-stabilizing leader

Table 1 Leader election algorithm under weak fairness

	State size	Transition size	Time	Results
LE-3	558	92974	0.45 s	valid
LE-4	1661	629905	5.29 s	valid
LE-5	4856	3335330	41.71 s	valid
LE-6	13629	14810700	264.07 s	valid

election in complete graphs. We have verified the model with size up to six. The detailed results are given in Table 1.

4.3 Correctness under weak fairness

In this section, we show that the self-stabilizing leader election algorithm is correct for any number of nodes ($N \geq 2$), under the weak fairness condition. Our proof follows the scheme in Ref. [10]. We call a configuration of the protocol safe, if there is at least one node that has become leader.

Lemma 4.1 Let E be an execution of the algorithm starting from an arbitrary configuration. Then E contains a safe configuration.

Theorem 4.1 Given an eventually correct leader detector. Let E be a weak fair execution of the algorithm starting from an arbitrary configuration. Then eventually one unique leader will be elected.

The proof of Lemma 4.1 is the same as in Ref. [10]. In the following, we give the proof of Theorem 4.1, which is also similar to the one in Ref. [10].

Proof By Lemma 4.1, the algorithm can reach a safe configuration. By the rules of the algorithm, the number of leaders decreases by one only when two leaders interact via Rule 1. So there is always at least one leader in subsequent configurations. By the eventually correct leader detector, eventually all nodes will receive T , after which no more new leaders can be generated. Now we prove that the number of leaders eventually decreases using the weak fairness. We only consider the case when there is more than one leader in the configuration. Assume node i and node j are leaders. Since the graph is complete, the interaction between i and j is enabled via Rule 1. By subsequent steps in E ,

- either they cannot change the state of either i or j , then the interaction between i and j keeps enabled. By weak fairness, finally the interaction will take place, and the number of leaders decreases;
- or they can change the state of either i or j , then this must be done via Rule 1, since that is the only way to change the state of one node once the leader detector is correct. The interaction between i and j is disabled, and the number of leaders is decreased by one.

5 Token circulation in directed rings

The token circulation protocol in directed rings is proposed in Ref. [11]. The desired behavior of this protocol can be described as follows:

- there is only one node who holds the token.
- a node does not obtain again until every other node has obtained a token once.
- each node can have the token infinitely often.

The protocol is simple since we do not consider the case that some nodes are not willing to release the token. It is assumed that every node passes the token to next one right after it has got it. Furthermore, the protocol also requires the existence of a common leader. Informally, there is a static node with the leader mark L , and all other nodes have the non-leader mark N in every configuration. The state of each node is represented by a pair in $\{-, +\} \times \{0, 1\}$. $+$ means that the node is holding a token and $-$ means the opposite. The second part of a state of a node is called the label. The algorithm is given by the rules in Fig. 3.

$$\begin{aligned} \text{Rule 1. } & ((*b, N), (*b, L)) \rightarrow (-b), (+\bar{b}) \\ \text{Rule 2. } & ((*b, *), (*\bar{b}, N)) \rightarrow (-b), (+b) \end{aligned}$$

Fig. 3 Algorithm for token circulation in directed rings

The $*$ here still denotes an always-matched symbol. On the left hand side, the symbol b matches either 0 or 1 and \bar{b} is its complement. It should be noticed that different occurrences of b in a same rule refer to the same value. The input for each node informs them who is leader, which is unique in the network.

When two nodes interact, if the responder is the leader, it sets its label to the complement of the initiator's label; otherwise the responder copies the label from the initiator. If an interaction triggers a label change, a token is passed from the initiator to the responder. If a token is not present at the initiator, a new token is generated.

It has been proved in Ref. [16] that this algorithm implements a self-stabilizing token circulation in rings under global fairness, provided that there is a unique leader.

5.1 Modeling token circulation in Spin

The model for token circulation protocol is similar to the one in Section 4, only with some minor adaptations. The states of the whole system are represented by three arrays of bits $\text{leader}[N]$, $\text{token}[N]$ and $\text{label}[N]$, where N is the number of nodes in the network. Without loss of generality, we can assume that node 0 is always the leader. Therefore, we could simply set each node a fixed input ($\text{leader}[i]$) for leader election without considering complicated details of a dynamic leader election process, which we have analyzed in Section 4.

We still use each single process in Promela to model a single action between each possible initiator and responder. However, now the network topologies are directed rings instead of complete graphs. So each node can only be the responder of its predecessor in a ring. Thus, the system with four nodes is represented as follows:

```
run Rule1(0,1);
run Rule1(1,2);
run Rule1(2,3);
run Rule1(3,0);
run Rule2(0,1);
run Rule2(1,2);
run Rule2(2,3);
run Rule2(3,0);
```

The verification goal is represented by a conjunction of three LTL formulas, each for a goal of the protocol. For the first goal that there is only one token in the network, we use the LTL formula

$$\langle \rangle [] \text{oneToken}$$

where oneToken is the proposition stating the sum of all token[i] ($0 \leq i < N$) equals one.

For the second goal that a node does not obtain again until every other node has obtained a token once, it is obviously equivalent to the one that when a node is holding a token, nobody could obtain the token until its successor has obtained it once. For example, if the network has four nodes, then the assertion for node 2 can be specified by the LTL formula

$$\langle \rangle [] (\text{token}[2] \rightarrow (!\text{token}[0] \text{token}[1] \cup \text{token}[3])).$$

For the last goal that every node obtains the token infinitely often, we use a formula in the following form:

$$[] \langle \rangle \text{token}[0] \dots \text{left} [] \langle \rangle \text{token}[N-1]$$

We have done some experiments for the model under weak fairness, and the results are mostly negative. The protocol is correct only when the size of the network is three. When it comes to a size greater than three, the Spin verifier complained about some failure traces which satisfy the weak fairness condition (cf. Section 1). Therefore, weak fairness cannot guarantee the correctness of the protocol.

Since the token circulation in rings does not work properly under the weak fairness condition, we need to verify it

under global fairness. However, with the limitation of Spin model checker (see more discussion in Section 7), it seems infeasible to explicitly model global fairness. Here we use an alternative method to verify the protocol. The algorithm is model checked in a two-phase manner. We first show that under a particular activation order chosen by a scheduler (under the global fairness condition) safe configurations are eventually reached. Then we show from these safe configurations the token circulation is eventually stabilized (under the weak fairness condition). The idea follows the correctness proof of the protocol given in Ref. [16].

The safe configurations refer to those configurations in which all nodes have a same label. With a scheduler satisfying global fairness, the order of activation for nodes complies with the following sequence:

$$(0,1), (1,2), \dots, (N-2, N-1)$$

If there is no possible interaction between nodes i and $i+1$ while the scheduler selects the activation $(i, i+1)$, then the scheduler just turns to the next pair of nodes $(i+1, i+2)$.

In the corresponding model in Spin, we employ a special variable turn to record the current activation pair. When $\text{turn} = i$, it indicates that there should be an action between node i and $i+1$. Every transition process will check whether it is its turn to do the transition. If it is the case, the process will possibly be able to do the transition according to the states of corresponding nodes. After the transition having been done, the turn flag variable will be increased by one. Those transitions which are not selected by the scheduler will block themselves. A special watch-dog process is needed to handle the case when no transition is enabled.

```
:: timeout - >
```

```
if
```

```
:: turn = (N-2) - > turn = turn + 1;
```

```
fi;
```

The watchdog process uses a timeout to detect the block state of the entire system, and then increases turn by one. When turn reaches $N-1$, all processes (including the watchdog process) would be suspended, indicating the system has gone through this special activation order. Thus, by using this model, we can carry out the first part of the verification, checking whether a safe configuration is reachable from any initial configuration.

For the second part of the verification, we only need to show that each of the three protocol goals is satisfied from the same-label initial configurations. Thus we

only generate two possible initial configurations non-deterministically at the beginning of the system run. The modification is straightforward. See Ref. [15] for the detailed model.

Table 2 Verification results of token circulation algorithm (I)

	State size	Transition size	Time	Results
TC-4	203	1114	0.01 s	valid
TC-5	475	3466	0.01 s	valid
TC-6	1083	11434	0.02 s	valid

5.2 Verification results

For the first phase, we can see from Table 2 that it is indeed the case that some safe configuration is reachable under the particular activation order.

As to the second phase, the verification shows that weak fairness is enough to ensure the correctness of the protocol after reaching a safe configuration. The results are shown in Table 3. Note that global fairness is needed for the overall verification task.

Table 3 Verification results of token circulation algorithm (II)

	State size	Transition size	Time	Results
TC-4	1525	10121	0.02 s	valid
TC-5	7063	36831	0.13 s	valid
TC-6	19287	111535	0.42 s	valid

6 Counterexamples for local fairness

Counterexamples are of utmost importance in model checking. Model checking can normally produce an answer in a few minutes or even seconds for many models. A counterexample will be generated when the checked property fails to hold, which details why the formal model does not satisfy the property. In this section, we give counterexamples as evidences that self-stabilizing token circulation and leader election in directed rings require global fairness; local fairness cannot guarantee their correctness. For liveness properties, typically they contain loops in executions.

6.1 Token circulation in directed rings

In Section 5 we have shown that the algorithm is correct under weak fairness for three nodes. Hence, the algorithm is also correct under local fairness for a network of three nodes. In order to get counterexamples under local fairness, we need a network consisting of at least four nodes (and without using the particular activation order in Section 5).

We need to first present eight configurations which are involved in the counterexample, denoted by C_1, C_2, \dots, C_8 :

$$\begin{aligned}
 C_1 &= \begin{pmatrix} \text{Node}_0 : + & 1 \\ \text{Node}_1 : + & 0 \\ \text{Node}_2 : - & 1 \\ \text{Node}_3 : + & 1 \end{pmatrix} & C_2 &= \begin{pmatrix} \text{Node}_0 : + & 1 \\ \text{Node}_1 : - & 0 \\ \text{Node}_2 : + & 0 \\ \text{Node}_3 : + & 1 \end{pmatrix} \\
 C_3 &= \begin{pmatrix} \text{Node}_0 : - & 1 \\ \text{Node}_1 : + & 1 \\ \text{Node}_2 : + & 0 \\ \text{Node}_3 : + & 1 \end{pmatrix} & C_4 &= \begin{pmatrix} \text{Node}_0 : + & 0 \\ \text{Node}_1 : + & 1 \\ \text{Node}_2 : + & 0 \\ \text{Node}_3 : - & 1 \end{pmatrix} \\
 C_5 &= \begin{pmatrix} \text{Node}_0 : + & 0 \\ \text{Node}_1 : + & 1 \\ \text{Node}_2 : - & 0 \\ \text{Node}_3 : + & 0 \end{pmatrix} & C_6 &= \begin{pmatrix} \text{Node}_0 : + & 0 \\ \text{Node}_1 : - & 1 \\ \text{Node}_2 : + & 1 \\ \text{Node}_3 : + & 0 \end{pmatrix} \\
 C_7 &= \begin{pmatrix} \text{Node}_0 : - & 0 \\ \text{Node}_1 : + & 0 \\ \text{Node}_2 : + & 1 \\ \text{Node}_3 : + & 0 \end{pmatrix} & C_8 &= \begin{pmatrix} \text{Node}_0 : + & 1 \\ \text{Node}_1 : + & 0 \\ \text{Node}_2 : + & 1 \\ \text{Node}_3 : - & 0 \end{pmatrix}
 \end{aligned}$$

According to the interaction rules of the protocol, the actions enabled in each configuration are:

$$\begin{aligned}
 C_1 &: \text{Rule1}(3,0), \text{Rule2}(0,1), \text{Rule2}(1,2) \\
 C_2 &: \text{Rule1}(3,0), \text{Rule2}(0,1), \text{Rule2}(2,3) \\
 C_3 &: \text{Rule1}(3,0), \text{Rule2}(1,2), \text{Rule2}(2,3) \\
 C_4 &: \text{Rule2}(0,1), \text{Rule2}(1,2), \text{Rule2}(2,3) \\
 C_5 &: \text{Rule1}(3,0), \text{Rule2}(0,1), \text{Rule2}(1,2) \\
 C_6 &: \text{Rule1}(3,0), \text{Rule2}(0,1), \text{Rule2}(2,3) \\
 C_7 &: \text{Rule1}(3,0), \text{Rule2}(1,2), \text{Rule2}(2,3) \\
 C_8 &: \text{Rule2}(0,1), \text{Rule2}(1,2), \text{Rule2}(2,3)
 \end{aligned}$$

The counterexample contains the following loop starting from initial configuration C_1 , going through C_2 to C_8 , and then moving back to C_1 .

$$\begin{aligned}
 C_1 &\xrightarrow{\text{Rule2}(1,2)} C_2 \xrightarrow{\text{Rule2}(0,1)} C_3 \xrightarrow{\text{Rule1}(3,0)} C_4 \xrightarrow{\text{Rule2}(2,3)} \\
 C_5 &\xrightarrow{\text{Rule2}(1,2)} C_6 \xrightarrow{\text{Rule2}(0,1)} C_7 \xrightarrow{\text{Rule1}(3,0)} C_8 \xrightarrow{\text{Rule2}(2,3)} C_1
 \end{aligned}$$

We can observe that actions enabled in these configurations are $\text{Rule1}(3,0)$, $\text{Rule2}(0,1)$, $\text{Rule2}(1,2)$ and $\text{Rule2}(2,3)$. Since each configuration occurs infinitely many times in the trace, these four actions are also enabled infinitely often. Clearly, the given trace satisfies local fairness because these actions are actually taken infinitely often. (It does not satisfy global fairness.) However, there are more than one

token in the infinite execution persistently, which does not meet the requirement that eventually there is only one token in the ring. Note that in the trace we ignore the inputs for each node, since we fix the node 0 as the unique leader. This information is stored in the array leader[*i*].

6.2 Leader election in rings

We have modeled the algorithm for self-stabilizing leader election in rings from [10], which is more complicated than the one for complete graphs. An eventually correct leader detector is also needed. In this algorithm, each node has three types of memory slots for tokens: a bullet slot (*B*), a leader mark slot (*L*), and a shield slot (*S*). (–) represents an empty slot, and a full slot is denoted by its token. The order of slots in each node is (bullet, leader, and shield). The leader detector gives each node an input true (*T*) or false (*F*) to indicate that whether there is a leader in the network. The algorithm is described by the following rules.

When two nodes interact and the initiator’s input is false (*F*), a leader and a shield are created. At the same time, a bullet is fired (Rule 1). This is the only way for leaders and shields to be created. When the initiator’s input is true (*T*), the following rules apply: Shields move forward around the ring (Rules 2 and 3), and bullets move backward (Rule 5). Bullets are absorbed by any shield they encounter (Rules 2 and 3) but kill any leaders along the way (Rule 5). If a bullet moves into a node already containing a bullet, the two bullets merge into one. Similarly, when two shields meet, they merge into one. A leader fires a bullet whenever it is the initiator of an interaction (Rules 3 and 4).

It has been shown that leader election in rings does not work under the local fairness condition [10]. We used a model similar to the one in Section 4 and verified it in Spin. We have found several counterexamples which indicate that the algorithm does not work properly under local fairness. Here we present two of them.

6.2.1 Counterexample 1

First, we present three configurations for a network with three nodes, denoted by C_1 , C_2 and C_3 :

$$C_1 = \begin{pmatrix} \text{Node}_0 : B & L & - \\ \text{Node}_1 : - & - & S \\ \text{Node}_2 : - & L & S \end{pmatrix}$$

$$C_2 = \begin{pmatrix} \text{Node}_0 : - & L & S \\ \text{Node}_1 : - & - & S \\ \text{Node}_2 : B & L & - \end{pmatrix}$$

$$C_3 = \begin{pmatrix} \text{Node}_0 : - & L & S \\ \text{Node}_1 : - & - & - \\ \text{Node}_2 : - & L & S \end{pmatrix}$$

According to Fig. 4, the actions enabled in each configuration are:

$$C_1 : \text{Rule2}(1,2), \text{Rule3}(2,0), \text{Rule4}(0,1)$$

$$C_2 : \text{Rule2}(1,2), \text{Rule3}(0,1), \text{Rule4}(2,0)$$

$$C_3 : \text{Rule3}(0,1), \text{Rule3}(2,0)$$

- Rule 1. $((* * *, F), (* * *, *)) \rightarrow ((B L S), (* * *))$
- Rule 2. $((* - S, T), (* * *, *)) \rightarrow ((* - -), (- * S))$
- Rule 3. $((* L S, T), (* * *, *)) \rightarrow ((B L -), (- * S))$
- Rule 4. $((* L -, T), (- * *, *)) \rightarrow ((B L -), (- * *))$
- Rule 5. $((* * -, T), (B * *, *)) \rightarrow ((B - -), (- * *))$

Each node outputs its own state.

Fig. 4 Algorithm for self-stabilizing leader election in rings

The counterexample contains a loop starting from initial configuration C_1 , looping from C_1 to C_3 , and then going back to C_1 .

$$C_1 \xrightarrow{\text{Rule4}(0,1)} C_1 \xrightarrow{\text{Rule3}(2,0)} C_2 \xrightarrow{\text{Rule4}(2,0)} C_2 \xrightarrow{\text{Rule2}(1,2)} C_3 \xrightarrow{\text{Rule3}(0,1)} C_1$$

Since each configuration occurs infinitely many times in the trace, those five actions are also enabled infinitely often. Clearly, the given trace satisfies local fairness because these actions are actually taken infinitely often. However, there are two leaders in the infinite execution persistently. (In the trace, we ignore the inputs *T* for each node.)

6.2.2 Counterexample 2

Consider the following three configurations denoted by C_1 , C_2 and C_3 :

$$C_1 = \begin{pmatrix} \text{Node}_0 : B & L & - \\ \text{Node}_1 : - & L & S \\ \text{Node}_2 : B & L & - \end{pmatrix}$$

$$C_2 = \begin{pmatrix} \text{Node}_0 : B & L & - \\ \text{Node}_1 : B & L & - \\ \text{Node}_2 : - & L & S \end{pmatrix}$$

$$C_3 = \begin{pmatrix} \text{Node}_0 : - & L & S \\ \text{Node}_1 : B & L & - \\ \text{Node}_2 : B & L & - \end{pmatrix}$$


According to Fig. 4, the actions enabled in each configuration are as follows:

C_1 : Rule3(1,2), Rule4(0,1), Rule5(2,0)

C_2 : Rule3(2,0), Rule4(1,2), Rule5(0,1)

C_3 : Rule3(0,1), Rule4(2,0), Rule5(1,2)

The counterexample contains a loop starting from initial configuration C_1 , looping from C_1 to C_3 , and then going back to C_1 .



Note that this execution does not satisfy local fairness, since Rule5(2,0), Rule5(0,1) and Rule5(1,2) are enabled infinitely many times, but are never taken. However, it satisfies the weak fairness provided by Spin.

7 Concluding remarks

In this paper, we have reported our preliminary results on automatic verification of population protocols in Spin. We defined a weak form of fairness condition for population protocols, which is weaker than both global and local fairness as originally required for the population protocol model in Ref. [10]. Weak fairness can be supported by Spin. We have successfully model checked the self-stabilizing leader election in complete graphs under such a weaker fairness condition. Although global fairness is indeed necessary for the correctness of self-stabilizing token circulation in directed rings, we have managed to model check the algorithm in Spin by a two-phase approach without explicitly encoding global fairness. This approach follows the proof in Ref. [16]. More interestingly, counterexamples to show why local fairness is insufficient for token circulation and leader election in directed rings have been automatically generated in Spin. This has improved our understanding of the population protocols.

Direct encoding global fairness in Spin will require auxiliary variables in the model to characterize fairness situations, which will increase the complexity of the model. Furthermore, it will result in very large LTL formulas. Usually, a strong fairness condition in LTL has the following form:

$$[\]\langle \rangle\text{enabled } A = > [\]\langle \rangle A$$

in which A stands for an activity in the model and enabled is the proposition specifying the condition when A can be enabled. Global fairness for population protocols requires that each step that can be taken infinitely often is actually taken infinitely often. Since one step can be enabled in many different configurations, this will require the enabled proposition to encode all such configurations

when the step A can be enabled. Moreover, there are many more different steps in a population protocol. All these will end into a very large LTL formula. Spin cannot deal with large LTL formulas. The size of the formulas will also increase exponentially, when the number of nodes in the network increases. In Spin, the model checker transforms the negation of a given LTL formula into a Büchi automaton, then it builds the product of the automaton and the model to check emptiness of the product. The size of the Büchi automaton is exponential to the size of the LTL formula. Therefore, even we can model global fairness, this approach does not scale up for verifying population protocols.

Our study in this paper gives rise to two interesting open questions: (1) Possibly we can find a fairness condition which is weaker than global fairness but still strong enough to guarantee the correctness of population protocols. For such a fairness condition, hopefully the available model checkers can deal with efficiently. (2) For population protocols, we do not encounter the usual state explosion problem as in many other model checking exercises, since a node only has few states in the population protocol model. We in fact need efficient model checking algorithms to deal with large LTL formulas. The work reported in Ref. [17] is closely related, but it still cannot be applied to population protocols as we have checked in the paper. Our experience [18] of verification of self-stabilizing distributed algorithms in PVS suggests the possibility of using a theorem prover to check population protocols under global fairness, which is currently under our investigation. It is also possible to use probabilistic model checkers for the verification of population protocols, since if we associate probabilities with all enabled interactions in a protocol configuration, then an execution will be global fair with probability 1.0.

Acknowledgements Jun Pang wants to thank Rena Bakhshi and Wan Fokkink for bringing his attention to the paper [10] on self-stabilizing population protocols. Yuxin Deng would like to thank Jing Cao for interesting discussions on the Spin model checker. Yuxin Deng is supported by the National Natural Science Foundation of China (Grant No. 60703033). An extended abstract of this paper has appeared in the Proceedings of the 2nd IEEE International Symposium on Theoretical Aspects of Software Engineering (TASE 2008, IEEE Computer Society).

References

1. Fuzzati R. Merro M, Nestmann U. Distributed consensus, revisited. *Acta Informatica*, 2007, 44(6): 377–425
2. Chandra T D. Toueg S. Unreliable failure detector for reliable distributed systems. *Journal of the ACM*, 1996, 43(2): 225–267
3. Holzmann G J. The model checker Spin. *IEEE Transactions on Software Engineering*, 1997, 23(5): 279–295
4. Aspnes J. Ruppert E. An introduction to population protocols. *Bulletin of the European Association for Theoretical Computer Science, Distributed Computing Column*, 2007, 93: 98–117

5. Attie P C, Francez N, Grumberg O. Fairness and hyperfairness in multi-party interactions. *Distributed Computing*, 1993, 6: 245–254
6. Lamport L. Fairness and hyperfairness. *Distributed Computing*, 2000, 13(4): 239–245
7. Völzer H, Varacca D, Kindler E. Defining fairness. In: *Proceedings of the 14th Conference on Concurrency Theory*. Berlin: Springer, 2005, 3653: 458–472
8. Völzer H. On conspiracies and hyperfairness in distributed computing. In: *Proceedings of the 19th Conference on Distributed Computing*. Berlin: Springer, 2005, 3724: 33–47
9. Corradini F, Di Berardini M, Vogler W. Checking a mutex algorithm in a process algebra with fairness. In: *Proceedings of the 15th Conference on Concurrency Theory*. Berlin: Springer, 2006, 4137: 142–157
10. Fischer M J, Jiang H. Self-stabilizing leader election in networks of finite-state anonymous agents. In: *Proceedings of the 10th Conference on Principles of Distributed Systems*. Berlin: Springer, 2006, 4305: 395–409
11. Angluin D, Aspnes J, Fischer M J, Jiang H. Selfstabilizing population protocols. In: *Proceedings of the 9th Conference on Principles of Distributed Systems*. Berlin: Springer, 2005, 3974: 103–117
12. Holzmann G J. On-the-fly, LTL model checking with spin. <http://spinroot.com>
13. Holzmann G J. *The spin model checker: Primer and reference manual*. Addison-Wesley, 2003
14. Dijkstra E W. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 1974, 17(11): 643–644
15. Luo Z Q, Pang J, Deng Y X. Promela source codes of self-stabilizing population protocols. <http://basics.sjtu.edu.cn/~zhengqin/population>
16. Jiang H. *Distributed systems of simple interacting agents*. Yale University, 2007
17. Hammer M, Knapp A, Merz S. Truly on-the-fly LTL model checking. In: *Proceedings of the 11th Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Berlin: Springer, 2005, 3440: 191–205
18. Fokkink W J, Hoepman J H, Pang J. A note on K -state self-stabilization in a ring with $K = N$. *Nordic Journal of Computing*, 2005, 12(1): 18–26