Constant-Space Population Protocols for Uniform Bipartition

Hiroto Yasumi[1], Fukuhito Ooshita[2], Ken'ichi Yamaguchi[3], and Michiko Inoue[4]

1 **College of Information Science, National Institute of Technology, Nara College, Nara, Japan**
a0858@stdmail.nara-k.ac.jp
2 **Graduate School of Information Science, Nara Institute of Science and Technology, Nara, Japan**
f-oosita@is.naist.jp
3 **College of Information Science, National Institute of Technology, Nara College, Nara, Japan**
yamaguti@info.nara-k.ac.jp
4 **Graduate School of Information Science, Nara Institute of Science and Technology, Nara, Japan**
kounoe@is.naist.jp

Abstract

In this paper, we consider a uniform bipartition problem in a population protocol model. The goal of the uniform bipartition problem is to divide a population into two groups of the same size. We study the problem under various assumptions: 1) a population with or without a base station, 2) weak or global fairness, 3) symmetric or asymmetric protocols, and 4) designated or arbitrary initial states. As a result, we completely clarify constant-space solvability of the uniform bipartition problem and, if solvable, propose space-optimal protocols.

1998 ACM Subject Classification C.2.4 Distributed Systems, I.1.2 Algorithms

Keywords and phrases population protocol, uniform bipartition, distributed protocol

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2017.19

1 Introduction

1.1 The Background

A population protocol model [4] is an abstract model that represents computation on a network of low-performance devices. We refer to such devices as agents and a set of agents as a population. Agents can update their states by interacting with other agents, and proceed with computation by repeating the pairwise interactions. The population protocol model can be applied to many systems such as sensor networks and molecular robot networks. For example, one may construct sensor networks to monitor wild birds by attaching sensors to them. In this system, sensors collect and process data based on pairwise interactions when two sensors (or birds) come sufficiently close to each other. Another example is a system of low-performance molecular robots [22]. In this system, a large number of molecular robots compose a network inside a human body and discriminate the physical condition. To realize such systems, many protocols have been proposed as building blocks in the population protocol model [10]. For example, they include leader election protocols [1, 2, 8, 15, 17, 19, 23, 24, 25], counting protocols [9, 11, 12, 20], and majority protocols [1, 3, 6, 18].

21st International Conference on Principles of Distributed Systems (OPODIS 2017).
Editors: James Aspnes, Alysson Bessani, Pascal Felber, and João Leitão; Article No. 19; pp. 19:1–19:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper, we consider a uniform bipartition problem, which divides a population into two groups of the same size. The uniform bipartition problem is a special case of a group composition problem, which divides a population into multiple groups to satisfy some conditions. Some protocols for the group composition problem are developed as subroutines to realize fault-tolerant protocols [16] and periodic functions [21]. However, the complexity of the problem has not been studied deeply yet. For this reason, as the first step to study the complexity of the group composition problem, we focus on the space complexity of the uniform bipartition problem. Note that the uniform bipartition problem itself has some applications. For example, we can reduce energy consumption by switching on one group and switching off the other. In another example, we can assign a different task to each group and make agents execute multiple tasks at the same time. This can be regarded as differentiation of a population in the sense that initially identical agents are eventually divided into two groups and execute different tasks. In addition, by repeating uniform bipartition, we can divide a population into an arbitrary number of groups with almost the same size. For example, by repeating uniform bipartition four times, we can make sixteen groups of the same size. We can regroup the sixteen groups to three groups with almost the same size by partitioning them into five, five, and six groups.

1.2 Our Contributions

For the uniform bipartition problem, we clarify solvability and minimum requirement of agent space under various assumptions. More concretely, we consider four types of assumptions, 1) a population with or without a base station, 2) weak or global fairness, 3) symmetric or asymmetric protocols, and 4) designated or arbitrary initial states. A base station (BS) is a distinguishable agent with a powerful capability, and it is useful to realize good properties while it limits the range of application. Fairness is an assumption on interaction patterns of agents. While weak fairness assumes only that interaction occurs infinitely often between each pair of agents, global fairness makes a stronger assumption on the order of interactions (the definition is given in Section 2). Symmetric property of protocols is related to the power of symmetry breaking in the population. Asymmetric protocols may include transitions that make agents with the same states transit to different states. This requires a mechanism to break symmetry among agents and its implementation is sometimes difficult with low-performance agents such as molecular robots. Symmetric protocols do not include such transitions. The assumption of initial states is related to the requirement of initialization and the fault-tolerant property. If a protocol requires a designated initial state, we need to initialize all agents to execute protocols. On the other hand, when the protocol allows arbitrary initial states, initialization of agents other than the BS is not necessary. In addition, even if agents enter arbitrary states due to transient faults, the system can eventually reach the desired configuration by initializing the BS. If a protocol allows arbitrary initial states and does not require a BS, the protocol is self-stabilizing because it can work from arbitrary initial configurations.

For each combination of assumptions, we completely clarify constant-space solvability of the uniform bipartition problem and, if solvable, give a space-optimal protocol (except for protocols given in [16, 14]). The results are shown in Table 1. Each element of the table represents the minimum number of agent states (except for a BS) to solve the uniform bipartition problem on the setting. First, we consider protocols in the case of a single BS. If protocols assume designated initial states, we prove three states are necessary and sufficient. If protocols allow arbitrary initial states, four states are necessary and sufficient under global fairness, while no constant-space protocol exists under weak fairness. Next, we consider

Table 1 The minimum number of states to solve the uniform bipartition problem, where n is the number of agents.

BS	Fairness	Designated initial states		Arbitrary initial states	
		Asymmetric	Symmetric	Asymmetric	Symmetric
Single BS	Globally fair	3	3	4	4
	Weakly fair	3	3	$\Omega(n)$	$\Omega(n)$
No BS	Globally fair	3*	4**	Impossible	Impossible
	Weakly fair	3*	Impossible	Impossible	Impossible

* A protocol with three states is proposed in [16].

** A protocol with four states is obtained by a general transformer in [14].

protocols in the case of no BS. If protocols assume designated initial states, three states are necessary and sufficient for asymmetric protocols. However, if we focus on symmetric protocols, no protocol exists under weak fairness and four states are necessary and sufficient under global fairness. For the case of arbitrary initial states, we prove no protocol exists if we assume no BS. This implies that a BS is necessary for protocols with arbitrary initial states.

1.3 Related Works

The population protocol model was introduced by Angluin et al. [4, 7]. They regard initial states of agents as an input to the system, and resultant states of them as an output from the system. Following this definition, they clarified the class of computable predicates in the population protocol model.

In addition to such computability researches, many algorithmic problems have been considered in the population protocol model. For example, they include leader election [1, 2, 8, 15, 17, 19, 23, 24, 25], counting [9, 11, 12, 20], and majority [1, 3, 6, 18]. These problems are considered under various assumptions of a population with or without a base station, global or weak fairness, symmetric or asymmetric protocols, designated or arbitrary initial states. The leader election problem has been thoroughly studied for both designated and arbitrary initial states. For designated initial states, many researches aim to minimize the time and space complexity [1, 2, 17]. For arbitrary initial states, many papers have developed self-stabilizing and loosely-stabilizing protocols [8, 15, 19, 23, 24, 25]. Cai et al. [15] proposed a self-stabilizing leader election protocol with knowledge of n, and proved that knowledge of n is necessary to construct a self-stabilizing leader election protocol, where n is the number of agents. To overcome the requirement of knowledge of n, Sudo et al. [24] proposed a concept of loose stabilization and gave a loosely-stabilizing leader election protocol. The complexity and the requirement on communication graphs are improved later [19, 23, 25]. The counting problem aims to count the number of agents and it has been studied under assumptions of a single BS and arbitrary initial states. After the first protocol was proposed in [12], the space complexity was gradually minimized [11, 20]. In [9], a time and space optimal protocol was proposed. The majority problem is also a fundamental problem in the population protocol model. In this problem, each agent initially has a color x or y, and the goal is to decide which color gets a majority. For the majority problem, many protocols have been proposed [1, 3, 6, 18]. Recently an asymptotically space-optimal protocol for c colors ($c > 2$) has been proposed in [18].

As a similar problem to the uniform bipartition problem, a group composition problem is studied in [16, 21]. Delporte-Gallet et al. [16] proposed a protocol to divide a population into g groups of the same size. The protocol is asymmetric, assumes designated initial states, and works under global fairness in the model of no BS. When $g = 2$, the protocol solves the uniform bipartition problem with three states. However, the paper does not consider other setting. Lamani et al. [21] studied a problem that divides a population into groups of designated sizes. Although the proposed protocols assume arbitrary initial states, they also assume that $n/2$ pairs of agents make interactions at the same time and that agents know n. In addition, the protocol requires n states, that is, it is not a constant-space protocol.

2 Definitions

2.1 Population Protocol Model

A population A is defined as a collection of pairwise interacting agents. A protocol is defined as $P = (Q, \delta)$, where Q is a set of possible states of agents and δ is a set of transitions on Q. Each transition in δ is described in the form $(p, q) \to (p', q')$, which means that, when an agent in state p and an agent in state q interact, they change their states to p' and q', respectively. In this paper, only deterministic protocols are considered. If transition $(p, q) \to (p', q')$ satisfies $p = q$ and $p' \neq q'$, the transition is asymmetric; otherwise, the transition is symmetric. For protocol $P = (Q, \delta)$, P is symmetric if every transition in δ is symmetric, and P is asymmetric if every transition in δ is symmetric or asymmetric. Note that a symmetric protocol is also asymmetric.

A global state of a population is called a configuration. A configuration is defined as a vector of (local) states of all agents. We define $s(a, C)$ as the state of agent a at configuration C. When C is clear from the context, we simply write $s(a)$. If configuration C' is obtained from configuration C by a single transition of a pair of agents, we say $C \to C'$. For configurations C and C', if there is a sequence of configurations $C = C_0, C_1, \cdots, C_k = C'$ that satisfies $C_i \to C_{i+1}$ for any i $(0 \leq i < k)$, we say C' is reachable from C, denoted by $C \xrightarrow{*} C'$.

If an infinite sequence of configurations $E = C_0, C_1, C_2, \ldots$ satisfies $C_i \to C_{i+1}$ for any i $(i \geq 0)$, E is an execution of a protocol. An execution E is weakly fair if every pair of agents in A interacts infinitely often. An execution E is globally fair if, for every pair of configurations C and C' such that $C \to C'$, C' occurs infinitely often when C occurs infinitely often. Intuitively, global fairness guarantees that, if configuration C occurs infinitely often, every possible interaction at C occurs infinitely often. If C occurs infinitely often, C' satisfying $C \to C'$ occurs infinitely often, and consequently C'' satisfying $C' \to C''$ also occurs infinitely often. This implies that, under global fairness, if C occurs infinitely often, every configuration C^* reachable from C also occurs infinitely often.

In this paper, we consider two models, one with a single BS (base station) and one with no BS. In the model with a single BS, we assume that a single agent called a BS exists in A. The BS is distinguishable from other non-BS agents while non-BS agents are identical and cannot be distinguished. That is, state set Q is divided into state set Q_b of a BS and state set Q_p of non-BS agents. The BS can be as powerful as needed, in contrast with resource-limited non-BS agents. That is, we focus on the number of states $|Q_p|$ for non-BS agents and do not care the number of states $|Q_b|$ for the BS. In addition, even if we consider protocols with arbitrary initial states, we assume that the BS has a designated initial state while all non-BS agents have arbitrary initial states. If we consider protocols with designated initial states, all non-BS agents have the same designated initial states and the BS has another designated initial state. In the model with no BS, no BS exits and all agents are identical. In this case, they all have the same designated initial states or arbitrary initial states. In both models, no agent knows the total number of agents in the initial configuration.

2.2 Uniform Bipartition Problem

Let A_p be a set of all non-BS agents. Let $f : Q_p \to \{red, blue\}$ be a function that maps a state of a non-BS agent to *red* or *blue*. We define a color of $a \in A_p$ as $f(s(a))$. We say agent $a \in A_p$ is *red* if $f(s(a)) = red$ and agent $a \in A_p$ is *blue* if $f(s(a)) = blue$.

Configuration C is stable if there is a partition $\{R, B\}$ of A_p that satisfies the following condition: 1) $||R| - |B|| \leq 1$, and 2) for every C^* such that $C \xrightarrow{*} C^*$, each agent in R is *red* and each agent in B is *blue* at C^*.

An execution $E = C_0, C_1, C_2, \ldots$ solves the uniform bipartition problem if there is a stable configuration C_t in E. If each execution E of protocol P solves the uniform bipartition problem, we say protocol P solves the uniform bipartition problem. The main objective of this paper is to minimize the number of states for non-BS agents. Since the BS is powerful, we do not care the number of states for the BS. When protocol P requires x states for non-BS agents, we say P is a protocol with x states.

For simplicity, we use agents only to refer to non-BS agents in the following sections. To refer to the BS, we always use the BS (not an agent).

3 Uniform Bipartition Protocols with a Single BS

In this section, we consider the uniform bipartition problem under the assumption of a single BS. Recall that the BS is distinguishable from other non-BS agents, and we do not care the number of states for the BS.

3.1 Protocols with Designated Initial States

In this subsection, we consider protocols with designated initial states. We give a simple symmetric protocol with three states under global or weak fairness, and then prove that there exists no asymmetric protocol with two states under global or weak fairness. This implies that, in this case, three states are sufficient for asymmetric or symmetric protocols under global or weak fairness.

3.1.1 A protocol with three states

In this protocol, the state set of (non-BS) agents is $Q_p = \{initial, red, blue\}$, and we set $f(initial) = f(red) = red$ and $f(blue) = blue$. The designated initial state of all agents is *initial*. The idea of the protocol is to assign states *red* and *blue* to agents alternately when agents interact with the BS. To realize this, the BS has a state set $Q_b = \{b_{red}, b_{blue}\}$, and its initial state is b_{red}. The protocol consists of the following two transitions.

1. $(b_{red}, initial) \to (b_{blue}, red)$
2. $(b_{blue}, initial) \to (b_{red}, blue)$

That is, when the BS in state b_{red} (resp., b_{blue}) and a non-BS agent in state *initial* interact, the BS changes the state of the non-BS agent to *red* (resp., *blue*) and the state of itself to b_{blue} (resp., b_{red}). When two non-BS agents interact, no state transition occurs. Clearly, all non-BS agents evenly transit to state *red* or *blue*, and the difference in the numbers of *red* and *blue* agents is at most one. Note that the protocol contains no asymmetric transition and works correctly if every non-BS agent interacts with the BS. Therefore, we have the following theorem.

▶ **Theorem 1.** *In the model with a single BS, there exists a symmetric protocol with three states and designated initial states that solves the uniform bipartition problem under global or weak fairness.*

3.1.2 Impossibility with two states

Next, we show three states are necessary to construct an asymmetric protocol under global or weak fairness. This implies that, in this case, three states are necessary for asymmetric or symmetric protocols under global or weak fairness because a symmetric protocol is also asymmetric. That is, three states are necessary and sufficient in this case.

▶ **Theorem 2.** *In the model with a single BS, no asymmetric protocol with two states and designated initial states solves the uniform bipartition problem under global or weak fairness.*

Proof. We prove that such a protocol does not exist even if its execution satisfies both global and weak fairness. For contradiction, assume that such a protocol Alg exists. Without loss of generality, we assume $Q_p = \{s_1, s_2\}, f(s_1) = red, f(s_2) = blue$, and that the designated initial state of all agents is s_1. Let n is an even number that is at least four. We consider the following three cases.

First, for population A of a single BS and n (non-BS) agents $a_1, a_2, \ldots, a_n$, consider an execution $E = C_0, C_1, \ldots$ of Alg that satisfies both global and weak fairness. According to the definition, there exists a stable configuration C_t. That is, after C_t, the state of each agent does not change even if the BS and agents in states s_1 and s_2 interact in any order.

Next, for population A' of a single BS and $n+2$ agents $a_1, a_2, \ldots, a_{n+2}$, we define an execution $E' = C'_0, C'_1, \ldots, C'_t, C'_{t+1}, \ldots$ of Alg as follows.

- From C'_0 to C'_t, the BS and n agents $a_1, a_2, \ldots, a_n$ interact in the same order as the execution E.
- After C'_t, the BS and $n+2$ agents interact so as to satisfy both global and weak fairness.

Since the BS and agents $a_1, \ldots, a_n$ change their states similarly to E from C'_0 to C'_t, there are $n/2+2$ agents in state s_1 and $n/2$ agents in state s_2 at C'_t. Moreover, the state of the BS at C'_t is the same as the state of the BS at C_t. However, since the difference in the numbers of *red* and *blue* agents is two, C'_t is not a stable configuration. Consequently, after C'_t, some *red* or *blue* agent changes its state in execution E'.

Lastly, we consider execution E for population A again. Here, we consider interactions after stable configuration C_t, and apply interactions in E' to execution E. That is, we consider the following execution after C_t: 1) when the BS and an agent in state $s \in \{s_1, s_2\}$ interact at $C'_u \to C'_{u+1}$ ($u \geq t$) in E', the BS and an agent in state s interact at $C_u \to C_{u+1}$ in E, and 2) when two agents in states $s \in \{s_1, s_2\}$ and $s' \in \{s_1, s_2\}$ interact at $C'_u \to C'_{u+1}$ ($u \geq t$) in E', two agents in states s and s' interact at $C_u \to C_{u+1}$ in E. We can construct such an execution because, after stable configuration C_t, at least two agents are in s_1 and at least two agents are in s_2. In this execution E, since interactions occur similarly to E', some *red* or *blue* agent changes its state similarly to E' after C_t. This is a contradiction because C_t is a stable configuration. ◀

3.2 Protocols with Arbitrary Initial States

In this subsection, we consider protocols with arbitrary initial states. Recall that, since a BS is powerful, the BS can start the protocol from a designated initial state.

3.2.1 Under global fairness

Under global fairness, we give a symmetric protocol with four states, and prove impossibility of protocols with three states. That is, we show that four states are necessary and sufficient to construct a (symmetric or asymmetric) protocol in this case.

3.2.1.1 A symmetric protocol with four states

Here we show a symmetric protocol with four states under global fairness. In this protocol, each (non-BS) agent x has two variables $rb_x \in \{red, blue\}$ and $mark_x \in \{0, 1\}$. Variable rb_x represents the color of agent x. That is, for state s of agent x, $f(s) = red$ holds if $rb_x = red$ and $f(s) = blue$ holds if $rb_x = blue$. We define $\#red$ as the number of *red* agents and $\#blue$ as *blue* agents. We explain the role of variable $mark_x$ later.

The basic strategy of the protocol is that the BS counts *red* and *blue* agents by counting protocol *Count* [11] and changes colors of agents so that the numbers of *red* and *blue* agents become equal. Protocol *Count* is a symmetric protocol that counts the number of non-BS agents from arbitrary initial states under global fairness. Protocol *Count* uses only two states for each non-BS agent. In protocol *Count*, the BS has variable $Count.out$ that eventually outputs the number of agents. More concretely, $Count.out$ initially has value 0, gradually increases one by one, eventually equals to the number of agents, and stabilizes. The following lemma explains the characteristic of protocol *Count*.

▶ **Lemma 3** ([11]). *Let n be the number of non-BS agents. In the initial configuration, $Count.out = 0$ holds. When $Count.out < n$, $Count.out$ eventually increases by one under global fairness. When $Count.out = n$, $Count.out$ never changes and stabilizes.*

To count *red* and *blue* agents, the BS executes two instances of protocol *Count* in parallel to the main procedure of the uniform bipartition protocol. We denote by $Count_{red}$ and $Count_{blue}$ instances of protocol *Count* to count *red* and *blue* agents, respectively. The BS executes $Count_{red}$ when it interacts with a *red* agent. That is, the BS updates variables of $Count_{red}$ at the BS and the *red* agent by applying a transition of protocol $Count_{red}$. By this behavior, the BS executes $Count_{red}$ as if the population contains only *red* agents. Therefore, after the BS initializes its own variables of $Count_{red}$, it can correctly count the number of *red* agents by $Count_{red}$ (i.e., $Count_{red}.out$ eventually stabilizes to $\#red$) as long as a set of *red* agents does not change. Similarly, the BS executes $Count_{blue}$ when it interacts with a *blue* agent, and counts the number of *blue* agents. The straightforward approach to use the counting protocols is to adjust colors of agents after $Count_{red}.out$ and $Count_{blue}.out$ stabilize. However, the BS cannot know whether the outputs have stabilized or not. For this reason, the BS maintains estimated numbers of *red* and *blue* agents, and it changes colors of agents when the difference in the estimated numbers of *red* and *blue* agents is two. Note that, since the counting protocols assume that a set of counted agents does not change, the BS must restart $Count_{red}$ and $Count_{blue}$ from the beginning when the BS changes colors of some agents.

We explain the details of this procedure. The BS records the estimated numbers of *red* and *blue* agents in variables $C^*_{rb}[red]$ and $C^*_{rb}[blue]$, respectively. In the beginning of execution, these variables are identical to outputs of $Count_{red}$ and $Count_{blue}$. If the difference between $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ becomes two, the BS immediately changes colors of agents. At the same time, the BS updates $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ to reflect the change of colors. After the BS changes colors of some agents, it restarts $Count_{red}$ and $Count_{blue}$ from the beginning by initializing its own variables of the counting protocols. Since the counting protocols allow

Algorithm 1 Uniform bipartition protocol.

Variables at BS:

$C^*_{rb}[c](c \in \{red, blue\})$: the estimated number of c agents, initialized to 0

$Variables$: variables of $Count_c(c \in \{red, blue\})$

Variables at a mobile agent x:

$rb_x \in \{red, blue\}$: color of the agent, initialized arbitrarily

$mark_x \in \{0, 1\}$: a variable of $Count_c(c \in \{red, blue\})$, initialized arbitrarily

when a mobile agent x interacts with BS **do**
  update $mark_x$ and variables of $Count_{rb_x}$ at BS by applying a transition of $Count_{rb_x}$
  if $C^*_{rb}[rb_x] < Count_{rb_x}.out$ **then**
    $C^*_{rb}[rb_x] \leftarrow Count_{rb_x}.out$
  end if
  if $C^*_{rb}[rb_x] - C^*_{rb}[\overline{rb_x}] > 1$ **then**
    $C^*_{rb}[rb_x] \leftarrow C^*_{rb}[rb_x] - 1$
    $C^*_{rb}[\overline{rb_x}] \leftarrow C^*_{rb}[\overline{rb_x}] + 1$, $rb_x \leftarrow \overline{rb_x}$
    reset variables of $Count_{red}$ and $Count_{blue}$ at BS
  end if
end when

arbitrary initial states of non-BS agents, the BS can correctly count *red* and *blue* agents after that. Note that the BS does not initialize $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ because it knows such numbers of *red* and *blue* agents exist. If the output of $Count_{red}$ and $Count_{blue}$ exceeds $C^*_{rb}[red]$ and $C^*_{rb}[blue]$, the BS updates $C^*_{rb}[red]$ and $C^*_{rb}[blue]$, respectively. After that, if the difference between $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ becomes two, the BS changes colors of agents. By repeating this behavior, the BS adjusts colors of agents.

The pseudocode of this protocol is given in Algorithm 1. We define $\overline{red} = blue$ and $\overline{blue} = red$. Variable $mark_x$ is a two-state variable of counting protocols $Count_{red}$ and $Count_{blue}$. Since the BS restarts the counting protocols whenever it changes colors of agents, the BS keeps a set of *red* (resp., *blue*) agents unchanged until it restarts $Count_{red}$ (resp., $Count_{blue}$). In addition, each agent is involved in either $Count_{red}$ or $Count_{blue}$ at the same time. Hence it requires only a single variable $mark_x$ to execute $Count_{red}$ and $Count_{blue}$. When two non-BS agents interact, no state transition occurs in this protocol and counting protocols. When the BS and a *red* agent interact, they update $mark_x$ and variables of $Count_{red}$ at the BS by applying a transition of $Count_{red}$. This means that they execute $Count_{red}$ in parallel to the main procedure of the uniform bipartition protocol. After that, if $Count_{red}.out$ is larger than $C^*_{rb}[red]$, $C^*_{rb}[red]$ is updated with $Count_{red}.out$. If $C^*_{rb}[red]$ is larger than $C^*_{rb}[blue]$ by two or more, the *red* agent changes its color to *blue* and the BS updates $C^*_{rb}[red]$ and $C^*_{rb}[blue]$. After updating, the BS resets variables of $Count_{red}$ and $Count_{blue}$, and restarts counting. When the BS and a *blue* agent interact, they behave similarly.

In the following, we prove the correctness of Algorithm 1.

▶ **Lemma 4.** *In any configuration, $C^*_{rb}[red] \leq \#red$, $C^*_{rb}[blue] \leq \#blue$ and $|C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$ hold.*

Proof. We prove by induction on the index $k \geq 0$ of a configuration in an execution $C_0, C_1, C_2, \cdots, C_k, \cdots$. At the initial configuration C_0, the lemma holds. Let us assume that the lemma holds for configuration C_k and prove it for configuration C_{k+1}. From this assumption, $C^*_{rb}[red] \leq \#red$, $C^*_{rb}[blue] \leq \#blue$ and $|C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$ hold at C_k.

Assume that, when C_k transits to C_{k+1}, the BS and agent x interact. If $Count_{rb_x}.out$ becomes larger than $C^*_{rb}[rb_x]$, the BS updates $C^*_{rb}[rb_x]$ by $C^*_{rb}[rb_x] \leftarrow Count_{rb_x}.out$ (line 3). Note that, in this case, $C^*_{rb}[rb_x]$ increases by one from Lemma 3. In addition, $C^*_{rb}[red] \leq \#red$ and $C^*_{rb}[blue] \leq \#blue$ still hold. Recall that $|C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$ held before this update and $C^*_{rb}[rb_x]$ increases by one. Conseqently, at this moment (before line 5), $|C^*_{rb}[rb_x] - C^*_{rb}[\overline{rb_x}]| \leq 1$ or $C^*_{rb}[rb_x] - C^*_{rb}[\overline{rb_x}] = 2$ holds. Next, we consider lines 5 to 9. If $C^*_{rb}[rb_x] - C^*_{rb}[\overline{rb_x}] \leq 1$ at line 5, lines 6 to 8 are not executed, and thus $C^*_{rb}[red] \leq \#red$, $C^*_{rb}[blue] \leq \#blue$ and $|C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$ hold. If $C^*_{rb}[rb_x] - C^*_{rb}[\overline{rb_x}] = 2$ at line 5, agent x changes its color from rb_x to $\overline{rb_x}$, $C^*_{rb}[rb_x]$ decreases by one, and $C^*_{rb}[\overline{rb_x}]$ increases by one. This also preserves $C^*_{rb}[red] \leq \#red$, $C^*_{rb}[blue] \leq \#blue$ and $|C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$. Therefore, the lemma holds. ◀

▶ **Theorem 5.** *Algorithm 1 solves the uniform bipartition problem. That is, in the model with a BS, there exists a symmetric protocol with four states and arbitrary initial states that solves the uniform bipartition problem under global fairness.*

Proof. We define $phase = C^*_{rb}[red] + C^*_{rb}[blue]$. Initially, $phase = 0$ holds. We show that 1) $phase$ increases one by one if $phase < n$, and 2) Algorithm 1 solves the uniform bipartition problem if $phase = n$.

First consider the initial configuration. Since we assume global fairness, $Count_{red}.out$ or $Count_{blue}.out$ increases by one from Lemma 3 and at that time $phase$ increases by one.

Let us consider the transition $C \to C'$ such that $phase$ increases by one (i.e., line 4 is executed) and $phase < n$ holds at C'. We consider two cases.

- Case that lines 7 to 9 are not executed at $C \to C'$. In this case, since the BS does not change sets of *red* and *blue* agents, it can correctly continue to execute $Count_{red}$ and $Count_{blue}$. Since $phase < n = \#red + \#blue$ holds, either $\#red > C^*_{rb}[red]$ or $\#blue > C^*_{rb}[blue]$ holds. Consequently, from Lemma 3, either $Count_{red}.out > C^*_{rb}[red]$ or $Count_{blue}.out > C^*_{rb}[blue]$ holds eventually because we assume global fairness. At that time, $C^*_{rb}[red]$ or $C^*_{rb}[blue]$ increases by one and hence $phase$ increases by one.
- Case that lines 7 to 9 are executed at $C \to C'$. In this case, the BS changes sets of *red* and *blue* agents. At that time, the BS initializes its own variables of counting algorithms $Count_{red}$ and $Count_{blue}$. Since the counting algorithms work from arbitrary initial states of agents, the BS can correctly execute $Count_{red}$ and $Count_{blue}$ from the beginning under global fairness. Similarly to the first case, from Lemma 3, either $Count_{red}.out > C^*_{rb}[red]$ or $Count_{blue}.out > C^*_{rb}[blue]$ holds eventually. Then, $phase$ increases by one.

Lastly, consider the transition $C \to C'$ such that $phase$ increases by one and $phase = n$ holds at C'. From $phase = n$, $C^*_{rb}[red] + C^*_{rb}[blue] = n = \#red + \#blue$ holds, and consequently $C^*_{rb}[red] = \#red$ and $C^*_{rb}[blue] = \#blue$ hold from Lemma 4. This implies that $Count_{red}.out$ and $Count_{blue}.out$ never exceed $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ after that, respectively. Therefore, $C^*_{rb}[red]$ and $C^*_{rb}[blue]$ are never updated and consequently agents never change their colors any more. Since $|\#red - \#blue| = |C^*_{rb}[red] - C^*_{rb}[blue]| \leq 1$ holds from Lemma 4, we have the theorem. ◀

3.2.1.2 Impossibility with three states

Here we show the impossibility of asymmetric protocols with three states.

▶ **Theorem 6.** *In the model with a single BS, no asymmetric protocol with three states and arbitrary initial states solves the uniform bipartition problem under global fairness.*

Proof. For contradiction, assume that such a protocol Alg exists. Without loss of generality, we assume that the state set of agents is $Q_p = \{s_1, s_2, s_3\}$, $f(s_1) = f(s_2) = red$, and $f(s_3) = blue$. We consider the following three cases.

First, consider population $A = \{a_0, \ldots, a_n\}$ of a single BS and n agents such that n is even and at least 4. Assume that a_0 is a BS. Since each agent has an arbitrary initial state, we consider an initial configuration C_0 such that $s(a_i) = s_3$ holds for any $i (1 \le i \le n)$. Note that the BS a_0 has a designated initial state at C_0. From the definition of Alg, for any globally fair execution $E = C_0, C_1, \cdots$, there exists a stable configuration C_t. Hence, both the number of *red* agents and the number of *blue* agents are $n/2$ at C_t. After C_t, the color of agent a_i (i.e., $f(s(a_i))$) never changes for any $a_i (1 \le i \le n)$ even if the BS and agents interact in any order.

Next, consider population $A' = \{a'_0, \ldots, a'_{n+2}\}$ of a single BS and $n+2$ agents. Assume that agent a'_0 is a BS. We consider an initial configuration C'_0 such that $s(a'_i) = s_3$ holds for any i $(1 \le i \le n+2)$. From this initial configuration, we define an execution $E' = C'_0, C'_1, \cdots, C'_t, \cdots$ using the execution E as follows.

- For $0 \le u < t$, when a_i and a_j interact at $C_u \to C_{u+1}$, a'_i and a'_j interact at $C'_u \to C'_{u+1}$.
- For $t \le u$, an interaction occurs at $C'_u \to C'_{u+1}$ so that E' satisfies global fairness.

Since the BS and agents $a_1, \ldots, a_n$ change their states similarly to E from C'_0 to C'_t, $s(a'_i) = s(a_i)$ holds for $1 \le i \le n$. Hence, there exist $n/2$ *red* agents and $n/2 + 2$ *blue* agents at C'_t. Consequently C'_t is not a stable configuration. This implies that there exists a stable configuration $C'_{t'}$ for some $t' > t$. Clearly at least one *blue* agent becomes *red* from C'_t to $C'_{t'}$. That is, for some configuration $C'_{t^*} (t \le t^* < t')$, an agent in state s_3 transits to state s_1 or s_2 at $C_{t^*} \to C_{t^*+1}$. Assume that t^* is the smallest value that satisfies the condition.

Finally, for A we define an execution $E'' = C''_0, C''_1, \cdots$ using executions E and E' as follows.

- Let $C''_u = C_u$ for $0 \le u \le t$. That is, E'' reaches stable configuration C''_t in similarly to E.
- For $t \le u \le t^*$, we define an execution so that interaction at $C'_u \to C'_{u+1}$ also occurs at $C''_u \to C''_{u+1}$. Concretely, when a'_i and a'_j interact at $C'_u \to C'_{u+1}$, we define $a_{i'}$ and $a_{j'}$ as follows and they interact at $C''_u \to C''_{u+1}$. If $i \le n$, let $i' = i$. Otherwise, since $s(a'_i) = s_3$ holds at C'_u (because no agent in state s_3 changes its state from C'_t to C'_{t^*}), choose $i' (\le n)$ such that both $s(a_{i'}) = s_3$ and $i' \ne j$ hold. Similarly, if $j \le n$, let $j' = j$. Otherwise choose $j' (\le n)$ such that both $s(a_{j'}) = s_3$ and $j' \ne i'$ hold. Such i' and j' exist since at least two agents in state s_3 exist (because $n \ge 4$ holds and no agent in state s_3 changes its state from C'_t to C'_{t^*}).

Clearly, for $t \le u \le t^*$ and $i \le n$, $s(a_i)$ at C''_u is equal to $s(a'_i)$ at C'_u. Additionally, at $C''_{t^*} \to C''_{t^*+1}$, an agent in state s_3 transits to s_1 or s_2 as well as $C'_{t^*} \to C'_{t^*+1}$. This means that the agent changes its color at $C''_{t^*} \to C''_{t^*+1}$, which contradicts that C''_t is a stable configuration. ◀

▶ Remark. Recall that Section 3.1.1 gives a protocol with three states and designated initial states. In the protocol, the state set of agents is $Q_p = \{initial, red, blue\}$, we set $f(initial) = f(red) = red$ and $f(blue) = blue$, and the designated initial state is *initial*. The important point is that the designated initial state (i.e., *initial*) has the same color as one of other states (i.e., *red*).

In the proof of Theorem 6, we consider an execution such that all agents have the same initial state in the initial configuration. The difference from the above protocol is that the initial state does not have the same color as any other state. This means, even if we consider a protocol with three states and designated initial states, there exists no protocol such that the designated initial state does not have the same color as any other state. This fact holds even if the number of states is larger than three. ◀

3.2.2 Under weak fairness

Under weak fairness, we prove that no protocol with constant states solves the uniform bipartition problem. To prove this impossibility, we borrow techniques used in the impossibility proof for the counting problem [12]. This work shows that, in the model with a single BS, when the upper bound of the number of non-BS agents is n, no asymmetric protocol with $n-2$ states and arbitrary initial states solves the counting problem under weak fairness. We can apply the proof in [12] to the uniform bipartition problem in a straightforward manner.

▶ **Theorem 7.** *Let n be an even number that is at least four. In the model with a single BS, when the upper bound of the number of non-BS agents is n, no asymmetric protocol with $n-3$ states and arbitrary initial states solves the uniform bipartition problem under weak fairness.*

Proof. For contradiction, assume that such a protocol *Alg* exists. We consider the following two cases.

First, consider population $A = \{a_0, \dots, a_{n-2}\}$ of a single BS and $n-2$ agents such that a_0 is a BS. We consider an initial configuration C_0 such that initial states of $a_0, \dots, a_{n-2}$ are $s_0, \dots, s_{n-2}$ (s_0 is a designated initial state of the BS). Since the upper bound of the number of non-BS agents is n and agents do not know the number of agents, *Alg* should work correctly even if the number of non-BS agents is $n-2$. This implies that, for any execution $E = C_0, C_1, \cdots, C_t, \cdots$, there exists a stable configuration C_t. Since the number of states for non-BS agents is $n-3$, there exists y, a_p, and $a_{p'}$ such that configurations satisfying $y = s(a_p) = s(a_{p'})$ appear infinitely many times after C_t.

Next, consider population $A' = \{a'_0, \dots, a'_n\}$ of a single BS and n agents such that a'_0 is a BS. We consider an initial configuration C_0 such that initial states of $a'_0, \dots, a'_n$ are $s_0, \dots, s_{n-2}, y, y$, respectively. For A' we define an execution $E' = C'_0, C'_1, \cdots, C'_t, \cdots$ using the execution E as follow.

- For $0 \le u \le t-1$, when a_i and a_j interact at $C_u \to C_{u+1}$, a'_i and a'_j interact at $C'_u \to C'_{u+1}$.

Clearly, $s(a'_i) = s(a_i)$ holds at C'_t for any i $(0 \le i \le n-2)$. Since $s(a'_n) = s(a'_{n-1}) = y$ holds at C'_t, the difference in the numbers of *red* and *blue* agents remains two and consequently C'_t is not a stable configuration.

After C'_t, we define an execution as follows. This definition aims to make $n-2$ agents behave similarly to E and two agents keep state y.

- Until $y = s(a'_p) = s(a'_{p'})$ holds, if a_i and a_j interact at $C_u \to C_{u+1}$, a'_i and a'_j interact at $C'_u \to C'_{u+1}$.
- To define the remainder of E', we first define procedure $Proc(q, q')$, which creates a sub-execution from two indices q and q'. Procedure $Proc(q, q')$ can be applied to a configuration such that $y = s(a'_p) = s(a'_{p'}) = s(a'_{n-1}) = s(a'_n)$ holds. After that, $Proc(q, q')$ creates a sub-execution similar to E such that all agents in $A(q, q') = (A' - \{a'_p, a'_{p'}, a'_{n-1}, a'_n\}) \cup \{a'_q, a'_{q'}\}$ interact each other and the last configuration also satisfies the above condition. The concrete definition of $Proc(q, q')$ is as follows. When a_i and a_j interact at $C_u \to C_{u+1}$, a'_i and a'_j interact at $C'_u \to C'_{u+1}$ if $i, j \notin \{p, p'\}$. If $i = p$ or $j = p$, a'_q joins the interaction instead of a'_p. If $i = p'$ or $j = p'$, $a'_{q'}$ joins the interaction

instead of $a'_{p'}$. Procedure $Proc(q, q')$ continues these behaviors until all agents in $A(q, q')$ interact each other and satisfy $s(a'_q) = s(a'_{q'}) = y$.

By using $Proc(q, q')$, we define the remainder of E' to satisfy weak fairness as follows: Repeat $Proc(p, p')$, $Proc(p, n-1)$, $Proc(p, n)$, $Proc(n-1, p')$, $Proc(n, p')$, and $Proc(n, n-1)$.

Clearly, E' makes $n-2$ agents behave similarly to E and two agents keep state y. Hence, E' never converges to a stable configuration. Since E' is weakly fair, this is a contradiction. ◀

▶ Remark. Theorem 7 implies that no protocol with at most $n-4$ states solves the uniform bipartition problem under the same assumption. This is because, if a protocol with n_s states ($n_s \leq n-4$) is given, we can transform it to a protocol with $n-3$ states by adding $n-3-n_s$ dummy states. Hence, at least $n-2$ states are necessary to solve the uniform bipartition problem under this assumption.

On the other hand, the sufficient number of states to solve the uniform bipartition problem under this assumption is not known. To clarify the matching lower and upper bounds of the number of states is an open problem. ◀

4 Uniform Bipartition Protocols with No BS

In this section, we consider the uniform bipartition problem under the assumption of no BS. That is, all agents are identical.

4.1 Protocols with Designated Initial States

In this subsection, we consider protocols with designated initial states. Since we consider the model with no BS, all agents have the same initial state in the initial configuration.

4.1.1 Asymmetric protocols

First, we consider asymmetric protocols in this case. Since three states are necessary in the model with a BS from Theorem 2, three states are also necessary in the model with no BS. In addition, Delporte-Gallet et al. [16] gives a protocol with three states. This implies that three states are necessary and sufficient in this case.

Here, we briefly explain the protocol proposed in [16]. In this protocol, the state set of agents is $Q_p = \{initial, red, blue\}$, and we set $f(initial) = f(red) = red$ and $f(blue) = blue$. The designated initial state of all agents is $initial$. The protocol consists of a single asymmetric transition $(initial, initial) \rightarrow (red, blue)$. In this protocol, when two agents in state $initial$ interact, one agent transits to red and the other transits to $blue$. This implies that the number of agents in state red is always the same as the number of agents in state $blue$. Eventually all agents (possibly except one agent) transit to state red or $blue$. From $f(initial) = red$, the difference in the numbers of red and $blue$ agents is at most one. Note that the protocol works correctly if every pair of agents interacts once.

▶ **Theorem 8** ([16]). *In the model with no BS, there exists an asymmetric protocol with three states and designated initial states that solves the uniform bipartition problem under global or weak fairness.*

4.1.2 Symmetric protocols

Next, we consider symmetric protocols in this case. For this setting, we give three results: 1) a protocol with four states under global fairness, 2) impossibility with three states under global fairness, and 3) impossibility under weak fairness. These results show that, in this case, four states are necessary and sufficient to construct a symmetric protocol under global fairness, and no symmetric protocol exists under weak fairness.

4.1.2.1 A protocol with four states under global fairness

We can easily obtain a symmetric protocol with four states by a scheme proposed in [14]. The scheme transforms an asymmetric protocol with α states to a symmetric protocol with at most 2α states. By applying the scheme to an asymmetric protocol in Section 4.1.1 and deleting unnecessary states, we can obtain a symmetric protocol with four states.

For self-containment, we briefly explain the obtained protocol. Since no symmetric protocol solves the uniform bipartition problem for a population of two agents, we assume that a population consists of at least three agents. In this protocol, the state set of agents is $Q_p = \{initial, initial', red, blue\}$, and we set $f(initial) = f(initial') = f(red) = red$ and $f(blue) = blue$. The designated initial state of all agents is *initial*. The protocol consists of the following seven transitions.

1. $(initial, initial) \to (initial', initial')$
2. $(initial', initial') \to (initial, initial)$
3. $(initial, initial') \to (red, blue)$
4. $(initial, red) \to (initial', red)$
5. $(initial, blue) \to (initial', blue)$
6. $(initial', red) \to (initial, red)$
7. $(initial', blue) \to (initial, blue)$

The main behavior of the protocol is similar to the previous asymmetric protocol with three states. However, since asymmetric transition $(initial, initial) \to (red, blue)$ is not allowed in symmetric protocols, the scheme in [14] introduces a new state $initial'$. Transition 3 implies that, when agents in states *initial* and $initial'$ interact, they become *red* and *blue*, respectively. In addition, agents in states *initial* and $initial'$ become $initial'$ and *initial* respectively when they interact with some agents (except for interaction between one in state *initial* and one in state $initial'$). From global fairness, if at least two agents are in state *initial* or $initial'$, some two agents eventually enter states *initial* and $initial'$. After that, if the two agents interact, they enter states *red* and *blue*.

Figure 1 shows an example execution of the protocol for a population of four agents. Initially all agents are in state *initial* (Fig. 1 (a)). After interactions (a_1, a_2) and (a_3, a_4), all agents enter state $initial'$ (Fig. 1 (b)). Similarly, after interactions (a_1, a_4), (a_2, a_3), (a_1, a_3), and (a_2, a_4), all agents have the same state (Fig. 1 (c) and (d)). If these interactions happen infinite times, all agents keep the same state and never achieve the uniform bipartition. However, under the global fairness, such interactions do not happen infinite times. This is because, if some configuration C occurs infinite times, every configuration reachable from C should occur. This implies that, before a configuration in Fig. 1 (d) occurs infinite times, interactions (a_1, a_2) and (a_1, a_3) happen in this order from the configuration. Then, a_1 and a_3 enter states *red* and *blue*, respectively (Fig. 1 (e) and (f)). After that, in a similar way, the remaining agents eventually enter *red* and *blue* like Fig. 1 (g) and (h).

Theorem 8 and correctness of the scheme in [14] derives the following theorem.

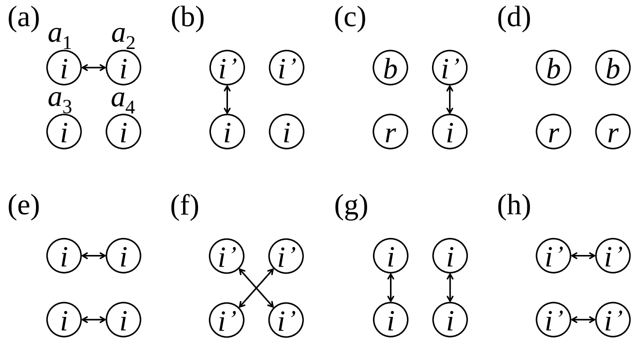

Figure 1 An example execution of the protocol. Symbols i, i', r, and b represent states *initial*, *initial*$'$, *red*, and *blue*, respectively. Arrows represent interactions of agents.

▶ **Theorem 9.** *In the model with no BS, when the number of agents is at least three, there exists a symmetric protocol with four states and designated initial states that solves the uniform bipartition problem under global fairness.*

4.1.2.2 Impossibility results

In the following, we show two impossibility results.

▶ **Theorem 10.** *In the model with no BS, no symmetric protocol with three states and designated initial states solves the uniform bipartition problem under global fairness.*

Proof. For contradiction, assume that such a protocol *Alg* exists. Without loss of generality, we assume that the state set of agents is $Q_p = \{s_1, s_2, s_3\}$, $f(s_1) = f(s_2) = red$, and $f(s_3) = blue$. Consider population $A = \{a_1, \ldots, a_n\}$ of n agents such that n is even and at least 6. First, assume that the designated initial state of all agents is s_3. Clearly, *Alg* has transition $(s_3, s_3) \to (s_i, s_i)$ for some $i \neq 3$. However, since $n/2$ agents in state s_3 exist at a stable configuration, some agents change their states from s_3 to s_i at the stable configuration. This implies that agents change their colors. Therefore, a designated initial state is s_1 or s_2.

Next, assume that the designated initial state of all agents is s_1 (Case of s_2 is the same). Since *Alg* is a symmetric protocol and all the initial states are s_1, *Alg* includes $(s_1, s_1) \to (s_i, s_i)$ for some $i \neq 1$. This implies that all agents can transit to state s_i from the initial configuration. Hence, *Alg* also includes $(s_i, s_i) \to (s_j, s_j)$ for some $j \neq i$. When $i = 3$, since $n/2$ *blue* agents exist at a stable configuration and they are in state s_3, the *blue* agents become *red* by transition $(s_3, s_3) \to (s_j, s_j)$. Therefore, $i \neq 3$ holds.

The remaining case is $i = 2$. If $j = 3$, that is, *Alg* includes $(s_2, s_2) \to (s_3, s_3)$, *red* agents (i.e., agents in state s_1 or s_2) change their colors at a stable configuration because *Alg* includes $(s_1, s_1) \to (s_2, s_2)$ and $(s_2, s_2) \to (s_3, s_3)$. This implies $j = 1$. In this case, *Alg* includes $(s_2, s_2) \to (s_1, s_1)$. Since some agents should transit to state s_3, *Alg* includes $(s_1, s_2) \to (s_k, s_l)$ such that k or l is 3. At a stable configuration, there exist $n/2$ agents with states s_1 or s_2. However, these agents can transit to state s_3 from transitions $(s_1, s_2) \to (s_k, s_l)$, $(s_2, s_2) \to (s_1, s_1)$, and $(s_1, s_1) \to (s_2, s_2)$. This is a contradiction. ◀

▶ **Theorem 11.** *In the model with no BS, no symmetric protocol with designated initial states solves the uniform bipartition problem under weak fairness.*

Proof. For contradiction, assume that such a protocol *Alg* exists. We assume that the state set of agents is $Q_p = \{s_1, s_2, \ldots\}$. Consider population $A = \{a_1, \ldots, a_n\}$ of n agents such that n is even and at least 2. Let s_{i_1} be the designated initial state of all agents, that is, $s(a_i) = s_{i_1}$ holds for any i $(1 \le i \le n)$ at the initial configuration. Clearly, symmetric protocol Alg has transition $(s_{i_1}, s_{i_1}) \to (s_{i_2}, s_{i_2})$ for some s_{i_2}. This implies that, if all pairs of two agents in state s_{i_1} interact, all agents transit to s_{i_2}. Similarly, if all pairs of two agents in state s_{i_2} interact, all agents transit to the same state (say s_{i_3}).

When the above execution is repeated, configurations such that all agents have the same state appear infinitely often. By changing pairs of two agents, we can make the above execution under weak fairness. If all agents are in the same state, such a configuration is not stable because the colors of all agents are the same. This is a contradiction. ◀

4.2 Protocols with Arbitrary Initial States

In this subsection, we consider protocols with arbitrary initial states. We show that, in this case, no protocol solves the uniform bipartition problem. That is, to allow agents to start from arbitrary initial states, a single BS is necessary.

▶ **Theorem 12.** *In the model with no BS, no asymmetric protocol with arbitrary initial states solves the uniform bipartition problem under global fairness*

Proof. For contradiction, assume that such a protocol Alg exists. Assume that n is even and at least 4. We consider the following two cases.

First, for population $A = \{a_1, \ldots, a_n\}$ of n agents, consider an execution $E = C_0, C_1, \cdots$ of Alg. From the definition of Alg, there exists a stable configuration C_t. Hence, both the number of *red* agents and the number of *blue* agents are $n/2$ at C_t. After C_t, the color of agent a_i (i.e.,$f(s(a_i))$) never changes for any a_i $(1 \le i \le n)$ even if agents interact in any order.

Next, for population $A' = \{a'_i | f(s(a_i, C_t)) = red\}$ of $n/2$ agents, consider an execution $E' = C'_0, C'_1, \cdots$ of Alg from the initial configuration C'_0 such that $s(a'_i, C'_0) = s(a_i, C_t)$ holds for any i $(1 \le i \le n/2)$. Since all agents are *red* at C'_0, some agents must change their colors to reach a stable configuration. This implies that, after C_t in execution E, agents change their colors if they interact similarly to E'. This is a contradiction. ◀

5 Conclusion

In this paper, we completely clarify constant-space solvability of the uniform bipartition problem and minimum requirement of agent space under various assumptions. This paper leaves many open problems:

- In the model of a single BS, how many states are necessary and sufficient to develop a uniform bipartition protocol with arbitrary initial states under weak fairness?
- Is it possible to extend our results to the uniform k-partition problem, which divides a population into k groups of the same size. Is it possible to construct a general protocol to solve the uniform k-partition problem? How many states are required to solve the problem?
- What is the relation between the uniform bipartition problem and other problems such as counting, leader election, and majority?
- What is the time complexity of the uniform bipartition problem under probabilistic fairness? The uniform bipartition problem has a close relationship to computation of function $f(n) = n/2$. The time complexity of $n/2$ computation has been studied in [5, 13]. Is it possible to derive the time complexity of the uniform bipartition problem from the results?

References

1 Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L Rivest. Time-space trade-offs in population protocols. In *Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2560–2579, 2017.

2 Dan Alistarh and Rati Gelashvili. Polylogarithmic-time leader election in population protocols. In *Proc. of the 42nd International Colloquium on Automata, Languages, and Programming*, pages 479–491, 2015.

3 Dan Alistarh, Rati Gelashvili, and Milan Vojnović. Fast and exact majority in population protocols. In *Proc. of the 2015 ACM Symposium on Principles of Distributed Computing*, pages 47–56, 2015.

4 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed computing*, 18(4):235–253, 2006.

5 Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, 2008.

6 Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 21(2):87–102, 2008.

7 Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.

8 Dana Angluin, James Aspnes, Michael J Fischer, and Hong Jiang. Self-stabilizing population protocols. In *International Conference On Principles Of Distributed Systems*, pages 103–117. Springer, 2005.

9 James Aspnes, Joffroy Beauquier, Janna Burman, and Devan Sohier. Time and space optimal counting in population protocols. In *Proc. of International Conference on Principles of Distributed Systems*, pages 13:1–13:17, 2016.

10 James Aspnes and Eric Ruppert. An introduction to population protocols. In *Middleware for Network Eccentric and Mobile Applications*, pages 97–120, 2009.

11 Joffroy Beauquier, Janna Burman, Simon Claviere, and Devan Sohier. Space-optimal counting in population protocols. In *Proc. of International Symposium on Distributed Computing*, pages 631–646, 2015.

12 Joffroy Beauquier, Julien Clement, Stephane Messika, Laurent Rosaz, and Brigitte Rozoy. Self-stabilizing counting in mobile sensor networks with a base station. In *Proc. of International Symposium on Distributed Computing*, pages 63–76, 2007.

13 Amanda Belleville, David Doty, and David Soloveichik. Hardness of computing and approximating predicates and functions with leaderless population protocols. In *Proc. of 44th International Colloquium on Automata, Languages, and Programming*, pages 141:1–141:14, 2017.

14 Olivier Bournez, Jérémie Chalopin, Johanne Cohen, and Xavier Koegler. Playing with population protocols. In *Proc. of the International Workshop on the Complexity of Simple Programs*, pages 3–15, 2008.

15 Shukai Cai, Taisuke Izumi, and Koichi Wada. How to prove impossibility under global fairness: On space complexity of self-stabilizing leader election on a population protocol model. *Theory of Computing Systems*, 50(3):433–445, 2012.

16 Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Eric Ruppert. When birds die: Making population protocols fault-tolerant. *Distributed Computing in Sensor Systems*, pages 51–66, 2006.

17 David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. In *Proc. of International Symposium on Distributed Computing*, pages 602–616, 2015.

18 Leszek Gasieniec, David Hamilton, Russell Martin, Paul G Spirakis, and Grzegorz Stachowiak. Deterministic population protocols for exact majority and plurality. In *Proc. of International Conference on Principles of Distributed Systems*, pages 14:1–14:14, 2016.

19 Taisuke Izumi. On space and time complexity of loosely-stabilizing leader election. In *Proc. of International Colloquium on Structural Information and Communication Complexity*, pages 299–312, 2015.

20 Tomoko Izumi, Keigo Kinpara, Taisuke Izumi, and Koichi Wada. Space-efficient self-stabilizing counting population protocols on mobile sensor networks. *Theoretical Computer Science*, 552:99–108, 2014.

21 Anissa Lamani and Masafumi Yamashita. Realization of periodic functions by self-stabilizing population protocols with synchronous handshakes. In *Proc. of International Conference on Theory and Practice of Natural Computing*, pages 21–33, 2016.

22 Satoshi Murata, Akihiko Konagaya, Satoshi Kobayashi, Hirohide Saito, and Masami Hagiya. Molecular robotics: A new paradigm for artifacts. *New Generation Computing*, 31(1):27–45, 2013.

23 Yuichi Sudo, Toshimitsu Masuzawa, Ajoy K Datta, and Lawrence L Larmore. The same speed timer in population protocols. In *Proc. of International Conference on Distributed Computing Systems*, pages 252–261, 2016.

24 Yuichi Sudo, Junya Nakamura, Yukiko Yamauchi, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. Loosely-stabilizing leader election in a population protocol model. *Theoretical Computer Science*, 444:100–112, 2012.

25 Yuichi Sudo, Fukuhito Ooshita, Hirotsugu Kakugawa, and Toshimitsu Masuzawa. Loosely-stabilizing leader election on arbitrary graphs in population protocols without identifiers nor random numbers. In *Proc. of International Conference on Principles of Distributed Systems*, pages 14:1–14:16, 2015.