

Large Flocks of Small Birds: On the Minimal Size of Population Protocols*

Michael Blondin¹, Javier Esparza², and Stefan Jaax³

- 1 Technische Universität München, Munich, Germany
blondin@in.tum.de
- 2 Technische Universität München, Munich, Germany
esparza@in.tum.de
- 3 Technische Universität München, Munich, Germany
jaax@in.tum.de

Abstract

Population protocols are a well established model of distributed computation by mobile finite-state agents with very limited storage. A classical result establishes that population protocols compute exactly predicates definable in Presburger arithmetic. We initiate the study of the minimal amount of memory required to compute a given predicate as a function of its size. We present results on the predicates $x \geq n$ for $n \in \mathbb{N}$, and more generally on the predicates corresponding to systems of linear inequalities. We show that they can be computed by protocols with $O(\log n)$ states (or, more generally, logarithmic in the coefficients of the predicate), and that, surprisingly, some families of predicates can be computed by protocols with $O(\log \log n)$ states. We give essentially matching lower bounds for the class of 1-aware protocols.

1998 ACM Subject Classification F.1.1 Models of Computation

Keywords and phrases Population protocols, Presburger arithmetic

Digital Object Identifier 10.4230/LIPIcs.STACS.2018.17

1 Introduction

Population protocols [4] are a model of distributed computation by anonymous, identical, and mobile finite-state agents. Initially introduced to model networks of passively mobile sensors, they also capture the essence of distributed computation in trust propagation or chemical reactions, the latter under the name of chemical reaction networks (see e.g. [18]). Structurally, population protocols can also be seen as a special class of Petri nets or vector addition systems [11].

Since the agents executing a protocol are anonymous and identical, its global state—called a *configuration*—is completely determined by the number of agents at each local state. In each computation step, a pair of agents, chosen by an adversary subject to a fairness condition stating that any repeatedly reachable configuration is eventually reached, interact and move to new states according to a joint transition function. In a closely related model, the adversary chooses the pair of agents uniformly at random.

A protocol computes a boolean value for a given initial configuration if in all fair executions all agents eventually agree to this value—so, intuitively, population protocols compute by reaching consensus. Given a set of initial configurations, the predicate computed by a protocol

* The full version of this paper can be found at <https://arxiv.org/abs/1801.00742>. M. Blondin was supported by the Fonds de recherche du Québec – Nature et technologies (FRQNT).



is the function that assigns to each configuration C the boolean value computed by the protocol starting from C .

Much research on population protocols has focused on their expressive power, i.e., the class of predicates computable by different classes of protocols (see e.g. [3, 6, 13, 16, 7]). In a famous result [6], Angluin et al. have shown that predicates computable by population protocols are exactly the predicates definable in Presburger arithmetic. There is also much work on complexity metrics for protocols. The main two metrics are the *runtime* of a protocol—defined for the model with a randomized adversary as the expected number of pairwise interactions until all agents have the correct output value—and its *state space size*, e.g. the number of states of each agent. In [5], Angluin et al. show that every Presburger predicate is computed with high probability by a population protocol with a leader—a distinguished auxiliary agent that assumes a specific state in the initial configuration irrespective of the input—in $O(n \log^4 n)$ interactions in expectation, where n is the number of agents of the initial configuration. Several recent papers study time-space trade-offs for specific tasks, like electing a leader [10], or for specific predicates, like majority [2, 1, 9].

In this paper we study the state space size of protocols as a function of the predicate they compute. In particular, we are interested in the minimal number of states needed to evaluate systems of linear constraints (a large subclass of the predicates computed by population protocols) as a function of the number of bits needed to describe the system. To the best of our knowledge, this question has not been considered so far. We study the question for protocols with and without leaders. Our results show that protocols with leaders can be exponentially more compact than leaderless protocols.

In order to introduce our results in the simplest possible setting, in the first part of the paper we focus on the family of predicates $\{x \geq n : n \in \mathbb{N}\}$. These predicates specify the well-known flock-of-birds problem [4], in which tiny sensors placed on birds have to reach consensus on whether the number of sick birds in a flock exceeds a given constant. The minimal number of states for computing $x \geq n$ formalizes a very natural question about emerging behavior: How many states must agents have in order to exhibit a “phase transition” when their number reaches n ? The standard protocol for the predicate $x \geq n$ (see Example 1) has $n + 1$ states. We are interested in protocols with at most $O(\log n)$ states, either leaderless or with at most $O(\log n)$ leaders. In the second part of the paper, we generalize our results to a much larger class of predicates, namely systems of linear inequalities $Ax \geq b$. Since $x \geq n$ is a (very) special case, our lower bounds for flock-of-birds protocols apply, while the upper bounds require new (and involved) constructions.

Protocol size for the flock-of-birds problem. In a first warm-up phase we exhibit a family of leaderless protocols with only $O(\log n)$ states. More precisely, we prove:

- (1) There exists a family $\{\mathcal{P}_n : n \in \mathbb{N}\}$ of leaderless population protocols such that \mathcal{P}_n has $O(\log_2 n)$ states and computes the predicate $x \geq n$ for every $n \in \mathbb{N}$.

We also give a lower bound:

- (2) For every family $\{\mathcal{P}_n : n \in \mathbb{N}\}$ of leaderless population protocols such that \mathcal{P}_n computes $x \geq n$, there exist infinitely many n such that \mathcal{P}_n has at least $(\log n)^{1/4}$ states.

However, this bound is only *existential* (“there exists infinitely many n ” instead of “for all n ”). Moreover, it follows from a counting argument that does not provide any information on the values of n realizing the bound. Is there a poly-logarithmic universal bound? We show that, surprisingly, the answer is negative:

- (3) There exists a family $\{\mathcal{P}_n : n \in \mathbb{N}\}$ of population protocols with two leaders, and values $c_0 < c_1 < \dots \in \mathbb{N}$, such that \mathcal{P}_n has $O(\log \log c_n)$ states and computes the predicate $x \geq c_n$ for every $n \in \mathbb{N}$.

Observe that in these protocols the “phase transition” occurs at $x = c_n$, even though no agent has enough memory to index a particular bit of c_n .

Can one go even further, and design $O(\log \log \log c_n)$ protocols? We show that the answer is negative for *1-aware* protocols. Both the standard protocol for $x \geq n$ and the families of (1) and (3) have the following, natural property: If the number of agents is greater than or equal to n , then the agents not only reach consensus 1, they also eventually *know* that they will reach this consensus. We say that these protocols are 1-aware.

We obtain lower bounds for 1-aware protocols that essentially match the upper bounds of (1) and (3):

- (4) Every leaderless, 1-aware population protocol computing $x \geq n$ has at least $\log_3 n$ states.
- (5) Every 1-aware protocol (leaderless or not) computing $x \geq n$ has at least $(\log \log(n)/151)^{1/9}$ states.

Protocols for systems of linear inequalities. In the second part of the paper we show that our results can be extended to other predicates. First, instead of the simple predicate $x \geq n$, we study the general linear predicate $a_1x_1 + a_2x_2 + \dots + a_kx_k \geq c$ for arbitrary integer coefficients $a_1, \dots, a_k, c \in \mathbb{Z}$. By means of a delicate construction we give protocols whose number of states grows only logarithmically in the size of the coefficients:

- (6) There is a protocol with at most $O(kn)$ states and $O(n)$ leaders that computes $a_1x_1 + \dots + a_kx_k \geq c$, where n is the size of the binary encoding of $\max(|a_1|, |a_2|, \dots, |a_k|, |c|)$.

Finally, in the most involved construction of the paper, we show that the same applies to arbitrary systems of linear inequalities. Note that the standard conjunction construction, which produces a protocol for $\varphi_1 \wedge \varphi_2$ from protocols computing predicates φ_1 and φ_2 , cannot be applied because it would lead to exponentially large protocols.

- (7) There is a protocol with at most $O((\log m + n)(m + k))$ states and $O(m(\log m + n))$ leaders that computes $A\mathbf{x} \geq \mathbf{c}$, where $A \in \mathbb{Z}^{m \times k}$ and n is the size of the largest entry in A and \mathbf{c} .

Structure of the paper. Section 2 introduces basic definitions, protocols with and without leaders, and a simple construction with an involved correctness proof showing how to simulate protocols with k -way interactions by standard protocols. Sections 3 to 5 present our bounds on the flock-of-birds predicates, and Section 6 the bounds on systems of linear inequalities. Due to space constraints, some proofs are deferred to the full version of this paper.

2 Preliminaries

Numbers. Let $n \in \mathbb{N}_{>0}$. The logarithm in base b of n is denoted by $\log_b n$. Whenever $b = 2$, we omit the subscript. We define $\text{bits}(n)$ as the set of indices of the bits occurring in the binary representation of n , e.g. $\text{bits}(13) = \{0, 2, 3\}$ since $13 = 1101_2$. The *size* of n , denoted $\text{size}(n)$, is the number of bits required to represent n in binary. Note that $|\text{bits}(n)| \leq \text{size}(n) = \lfloor \log n \rfloor + 1$.

Multisets. A *multiset* over a finite set E is a mapping $M: E \rightarrow \mathbb{N}$. The set of all multisets over E is denoted \mathbb{N}^E . For every $e \in E$, $M(e)$ denotes the number of occurrences of e in M , and for every $E' \subseteq E$ we define $M(E') \stackrel{\text{def}}{=} \sum_{e \in E'} M(e)$. The *support* and *size* of M are defined respectively as $\llbracket M \rrbracket \stackrel{\text{def}}{=} \{e \in E : M(e) > 0\}$ and $|M| \stackrel{\text{def}}{=} \sum_{e \in E} M(e)$. *Addition* and *comparison* are extended to multisets componentwise, i.e. $(M + M')(e) \stackrel{\text{def}}{=} M(e) + M'(e)$ for every $e \in E$, and $M \leq M' \stackrel{\text{def}}{\iff} M(e) \leq M'(e)$ for every $e \in E$. We define *multiset difference* as $(M \ominus M')(e) \stackrel{\text{def}}{=} \max(M(e) - M'(e), 0)$ for every $e \in E$. The empty multiset is denoted $\mathbf{0}$.

We sometimes denote multisets using a set-like notation, e.g. $\{f, 2 \cdot g, h\}$ is the multiset M such that $M(f) = 1$, $M(g) = 2$, $M(h) = 1$ and $M(e) = 0$ for every $e \in E \setminus \{f, g, h\}$.

Population protocols. We introduce a rather general model of population protocols, allowing for interactions between more than two agents and for leaders. A k -way population protocol is a tuple $\mathcal{P} = (Q, T, I, L, O)$ such that

- Q is a finite set of *states*,
- $T \subseteq \bigcup_{2 \leq i \leq k} Q^i \times Q^i$ is a set of *transitions*,
- $I \subseteq Q$ is a set of *initial states*,
- $L \in \mathbb{N}^Q$ is a set of *leaders*, and
- $O: Q \rightarrow \{0, 1\}$ is the *output mapping*.

We assume throughout the paper that agents can always interact, i.e., that for every pair of states (p, q) , there exists a pair of states (p', q') such that $((p, q), (p', q')) \in T$.

A *configuration* of \mathcal{P} is a multiset $C \in \mathbb{N}^Q$ such that $|C| > 0$. Intuitively, C describes a non empty collection containing $C(q)$ agents in state q for every $q \in Q$. We denote the set of configurations over $E \subseteq Q$ by $\text{Pop}(E)$. A configuration C is *initial* if $C = D + L$ for some $D \in \text{Pop}(I)$. So, intuitively, leaders are distinguished agents that are present in every initial configuration. The *number of leaders* of \mathcal{P} is $|L|$. We say that \mathcal{P} is *leaderless* if it has no leader, i.e. if $L = \mathbf{0}$. We discuss protocols with and without leaders later in this section.

Let $t = ((p_1, p_2, \dots, p_i), (q_1, q_2, \dots, q_i))$ be a transition. To simplify the notation, we denote t as $p_1, p_2, \dots, p_i \mapsto q_1, q_2, \dots, q_i$. Intuitively, t describes that i agents at states p_1, \dots, p_i may interact and move to states q_1, \dots, q_i . The *preset* and *postset* of t are respectively defined as $\bullet t \stackrel{\text{def}}{=} \{p_1, p_2, \dots, p_i\}$ and $t \bullet \stackrel{\text{def}}{=} \{q_1, q_2, \dots, q_i\}$. We extend presets and postsets to sets of transitions, e.g. $\bullet T \stackrel{\text{def}}{=} \bigcup_{t \in T} \bullet t$. The *pre-multiset* and *post-multiset* of t are respectively defined as $\text{pre}(t) \stackrel{\text{def}}{=} \{p_1, p_2, \dots, p_i\}$ and $\text{post}(t) \stackrel{\text{def}}{=} \{q_1, q_2, \dots, q_i\}$.

We say that t is *enabled* at $C \in \text{Pop}(Q)$ if $C \geq \text{pre}(t)$. If t is enabled at C , then it can *occur*, in which case it leads to the configuration $C' = (C \ominus \text{pre}(t)) + \text{post}(t)$. We denote this by $C \xrightarrow{t} C'$. We say that t is *silent* if $\text{pre}(t) = \text{post}(t)$. In particular, if t is silent and $C \xrightarrow{t} C'$, then $C = C'$. We write $C \rightarrow C'$ if $C \xrightarrow{t} C'$ for some $t \in T$. We write $C \xrightarrow{t_1 t_2 \dots t_k} C'$ if there exist $C_0, C_1, \dots, C_k \in \text{Pop}(Q)$ and $t_1, t_2, \dots, t_k \in T$ such that $C = C_0 \xrightarrow{t_1} C_1 \xrightarrow{t_2} \dots \xrightarrow{t_k} C_k = C'$. We write $C \xrightarrow{*} C'$ if $C \xrightarrow{\sigma} C'$ for some $\sigma \in T^*$. We say that C' is *reachable* from C if $C \xrightarrow{*} C'$. The *support* of a sequence $\sigma = t_1 t_2 \dots t_n \in T^*$ is $\llbracket \sigma \rrbracket \stackrel{\text{def}}{=} \{t_i : 1 \leq i \leq n\}$.

► **Example 1.** The flock-of-birds protocol mentioned in the introduction is formally defined as $\mathcal{P}_n = (Q, T, I, L, O)$ where $Q = \{0, 1, \dots, n\}$, $I = \{1\}$, $L = \mathbf{0}$, $O(a) = 1 \iff a = n$, and where T consists of the following transitions:

$$\begin{array}{ll} s_{a,b} : a, b \mapsto 0, \min(a + b, n) & \text{for every } 0 \leq a, b < n, \\ t_a : a, n \mapsto n, n & \text{for every } 0 \leq a \leq n. \end{array}$$

\mathcal{P}_n is 2-way and leaderless. Intuitively, it works as follows. Each agent stores a number. When two agents meet, one agent stores the sum of their values and the other one stores 0. Sums cap at n . Once an agent reaches n , all agents eventually get converted to n . To illustrate the above definitions, observe that: $\bullet s_{2,3} = \{2, 3\}$, $t_2 \bullet = \{n\}$, $\text{pre}(s_{2,3}) = \{2, 3\}$ and $\text{post}(t_2) = \{n, n\}$. Configuration $\{1, 1, 1\}$ is initial, but $\{1, 0, 2\}$ is not. We have $\{1, 1, 1\} \xrightarrow{s_{1,1}} \{1, 0, 2\} \xrightarrow{t_0} \{1, 2, 2\} \xrightarrow{t_1} \{2, 2, 2\}$, or more concisely $\{1, 1, 1\} \xrightarrow{\sigma} \{2, 2, 2\}$ where $\sigma = s_{1,1} t_0 t_1$. ◀

Computing with population protocols. An *execution* π is an infinite sequence of configurations $C_0 C_1 \dots$ such that $C_0 \rightarrow C_1 \rightarrow \dots$. We say that π is *fair* if for every

configuration D the following holds¹:

if $\{i \in \mathbb{N} : C_i \xrightarrow{*} D\}$ is infinite, then $\{i \in \mathbb{N} : C_i = D\}$ is infinite.

In other words, fairness ensures that a configuration cannot be avoided forever if it can be reached infinitely often along π . We say that a configuration C is a *consensus configuration* if $O(p) = O(q)$ for every $p, q \in \llbracket C \rrbracket$. If a configuration C is a consensus configuration, then its *output* $O(C)$ is the unique output of its states, otherwise it is \perp . An execution $\pi = C_0 C_1 \cdots$ *stabilizes* to $b \in \{0, 1\}$ if $O(C_i) = O(C_{i+1}) = \cdots = b$ for some $i \in \mathbb{N}$. The *output* of π is $O(\pi) \stackrel{\text{def}}{=} b$ if it stabilizes to b , and $O(\pi) \stackrel{\text{def}}{=} \perp$ otherwise. A consensus configuration C is *stable* if every configuration C' reachable from C is a consensus configuration such that $O(C') = O(C)$. It can easily be shown that a fair execution stabilizes to $b \in \{0, 1\}$ if and only if it contains a stable configuration whose output is b .

A population protocol $\mathcal{P} = (Q, T, I, L, O)$ is *well-specified* if for every initial configuration C_0 , there exists $b \in \{0, 1\}$ such that every fair execution π starting at C_0 has output b . If \mathcal{P} is well-specified, then we say that it *computes* the predicate $\varphi : \text{Pop}(I) \rightarrow \{0, 1\}$ if for every $D \in \text{Pop}(I)$, every fair execution starting at $D + L$ has output $\varphi(D)$.

► **Example 2.** Consider the protocol \mathcal{P}_2 defined in Example 1 (i.e., $n = 2$). We have $O(\uparrow 1, 1, 1) = 0$, $O(\uparrow 2, 2, 2) = 1$ and $O(\uparrow 1, 0, 2) = \perp$. The execution $\uparrow 1, 1, 1 \rightarrow \uparrow 1, 0, 2 \rightarrow \uparrow 1, 2, 2 \rightarrow \uparrow 2, 2, 2 \rightarrow \uparrow 2, 2, 2 \rightarrow \cdots$ is fair and its output is 1. However, the execution $\uparrow 1, 1, 1 \rightarrow \uparrow 1, 0, 2 \rightarrow \uparrow 1, 0, 2 \rightarrow \cdots$ is not fair since $\uparrow 1, 0, 2$ occurs infinitely often and can lead to $\uparrow 2, 2, 2$ which does not occur.

Leaders. Intuitively, leaders are extra agents present in every initial configuration. Allowing a large number of leaders may help to compute predicates with fewer states. To illustrate this, consider the leaderless protocol of Example 1. It computes $x \geq n$ with $n + 1$ states. We describe a 2-way protocol with only 4 states, but n leaders. It is an adaptation of the well-known basic majority protocol (see, e.g., [8]). Let $\mathcal{P}'_n = (Q, T, I, L_n, O)$ be the protocol where $Q \stackrel{\text{def}}{=} \{x, y, \bar{x}, \bar{y}\}$, $I \stackrel{\text{def}}{=} \{x\}$, $L_n \stackrel{\text{def}}{=} \uparrow n \cdot y$, $O(x) = O(\bar{x}) \stackrel{\text{def}}{=} 1$, $O(y) = O(\bar{y}) \stackrel{\text{def}}{=} 0$, and where T consists of the following transitions:

$$x, y \mapsto \bar{x}, \bar{y}, \quad x, \bar{y} \mapsto x, \bar{x}, \quad y, \bar{x} \mapsto y, \bar{y}, \quad \bar{x}, \bar{y} \mapsto \bar{x}, \bar{x}.$$

Informally, “active” agents in states x and y collide and become “passive” agents in states \bar{x} and \bar{y} . At some point, some active agents “win” and convert all passive agents to their output. It is known that this protocol is well-specified and computes the predicate $x \geq y$ when there are no leaders (i.e., if we set $L_n = \mathbf{0}$). So, by initially fixing n leaders in state y , \mathcal{P}'_n computes $x \geq n$.

Thus, the predicate $x \geq n$ can be computed either with $O(n)$ states and no leaders, or with 4 states and $O(n)$ leaders. This indicates a trade-off between states and leaders, and one should avoid hiding all of the complexity in one of them. For this reason, we make these two quantities explicit in all of our results.

The reason for considering protocols with leaders is that, as we shall see, even a constant number of leaders demonstrably leads to exponentially more compact protocols for some predicates. Other papers have made similar observations with respect to other resource measures (see e.g. [5, 14]).

¹ This definition of fairness differs from the original definition of Angluin et al. [4], but is equivalent.

From k -way to 2-way protocols. In our constructions it is very convenient to use k -way transitions for $k > 2$. The following lemma shows that k -way protocols can be transformed into 2-way protocols by introducing a few extra states. Intuitively, a k -way transition is simulated by a chain of 2-way transitions. The first part of the chain “collects” k participants one by one. First, two agents agree to participate, and one of them becomes “passive”, while the second “searches” for a third participant. This is iterated until k participants are collected. In the second part, the last collected agent “informs” all passive agents, one by one, that k agents have been collected; upon hearing this, the passive agents move to their destination states and become active again. To prevent faulty behavior when there are not enough agents, all transitions of the first part can be “reversed”, that is, the agent that is currently searching and the last collected agent can “repent” and “undo” the transition. While the construction is simple and intuitive, its correctness proof is very involved, because agents that reach their destination can engage in other interactions while other participants are still passive. The construction is presented in the full version of this paper.

► **Lemma 3.** *Let $\mathcal{P} = (Q, T, I, L, O)$ be a well-specified k -way population protocol. For every $3 \leq i \leq k$, let n_i be the number of i -way transitions of \mathcal{P} . There exists a 2-way population protocol \mathcal{P}' , with at most $|Q| + \sum_{3 \leq i \leq k} 3i \cdot n_i$ states, which is well-specified and computes the same predicate as \mathcal{P} .*

3 Leaderless protocols for $x \geq n$

In this section, we consider *leaderless* protocols for the predicate $x \geq n$. We first show that the number of states required to compute this predicate can be reduced from the known $O(n)$ bound to $O(\log n)$, using a similar binary encoding as in [1]. Then we show an existential lower bound of $O((\log n)^{1/4})$.

A protocol with $O(\log n)$ states. We describe a leaderless $\text{size}(n)$ -way protocol $\mathcal{P}_n = (Q_n, T_n, I_n, \mathbf{0}, O_n)$ with $\text{size}(n) + 3$ states that computes $x \geq n$. The states are $Q_n \stackrel{\text{def}}{=} \{\mathbf{0}, \mathbf{2}^0, \dots, \mathbf{2}^{\text{size}(n)}, \mathbf{n}\}$ and the sole initial state is $I_n \stackrel{\text{def}}{=} \{\mathbf{2}^0\}$. The output mapping is defined as $O_n(\mathbf{n}) \stackrel{\text{def}}{=} 1$ and $O_n(q) \stackrel{\text{def}}{=} 0$ for every state $q \neq \mathbf{n}$.

Before defining the set T_n of transitions, we need some preliminaries. For every state $q \in Q_n$, let $\text{val}(q)$ denote the number q stands for, i.e. $\text{val}(\mathbf{0}) = 0$, $\text{val}(\mathbf{n}) = n$ and $\text{val}(\mathbf{2}^i) = 2^i$ for every $0 \leq i \leq \text{size}(n)$. Moreover, for every configuration C , let $\text{val}(C) \stackrel{\text{def}}{=} \sum_{q \in Q_n} \text{val}(q) \cdot C(q)$. A configuration C is a *representation of m* if $\text{val}(C) = m$. For example, the configuration $\{ \mathbf{0}, \mathbf{2}^1, 5 \cdot \mathbf{2}^3 \}$ is a representation of $0 + 2^1 + 5 \cdot 2^3 = 42$. Observe that every initial configuration C_0 is a representation of $|C_0|$.

T_n is the union of two sets T_n^1 and T_n^2 . Intuitively, T_n^1 allows the protocol to reach from a representation of a number, say m , other representations of m . Formally, the transitions of T_n^1 are:

$$\begin{aligned} \mathbf{2}^i, \mathbf{2}^i &\mapsto \mathbf{2}^{i+1}, \mathbf{0} && \text{for every } 0 \leq i < \text{size}(n) \\ \mathbf{2}^{i+1}, \mathbf{0} &\mapsto \mathbf{2}^i, \mathbf{2}^i && \text{for every } 0 \leq i < \text{size}(n) \\ \{ \mathbf{2}^i : i \in \text{bits}(n) \} &\mapsto \mathbf{n}, \underbrace{\mathbf{0}, \dots, \mathbf{0}}_{|\text{bits}(n)|-1 \text{ copies}} \end{aligned}$$

The transitions of T_n^2 allow agents in state \mathbf{n} to “attract” all other agents to \mathbf{n} . Formally, they are:

$$\mathbf{n}, q \mapsto \mathbf{n}, \mathbf{n} \quad \text{for every } q \in Q_n.$$

Let us show that \mathcal{P}_n computes $x \geq n$. Let $C_0 = \{m \cdot \mathbf{2}^0\}$. If $m < n$, then $C(\mathbf{n}) = 0$ holds for every representation C of m . Therefore, every configuration C reachable from C_0

satisfies $C(\mathbf{n}) = 0$ and, since \mathbf{n} is the only state with output 1, the protocol stabilizes to 0. If $m \geq n$, then it is possible to reach a representation C of m satisfying $C(\mathbf{n}) > 0$, for example $C = \langle \mathbf{n}, (m - n) \cdot \mathbf{2}^0 \rangle$. Since for every transition $\mathbf{2}^i, \mathbf{2}^i \mapsto \mathbf{2}^{i+1}, \mathbf{0}$ the set T_n also contains the reverse transition $\mathbf{2}^{i+1}, \mathbf{0} \mapsto \mathbf{2}^i, \mathbf{2}^i$, every representation C of m satisfying $C(\mathbf{n}) = 0$ can reach a representation C' of m satisfying $C'(\mathbf{n}) > 0$. Let $\pi = C_0 C_1 C_2 \dots$ be a fair execution. By fairness, there is some $i \in \mathbb{N}$ such that $C_i(\mathbf{n}) > 0$. Again by fairness, and because of T_n^2 , there is also an index j such that $C_k = \langle m \cdot \mathbf{n} \rangle$ for every $k \geq j$, and so π stabilizes to 1.

Note that $|Q_n| = \text{size}(n) + 3$. Moreover, \mathcal{P}_n has one $|\text{bits}(n)|$ -way transition. Thus, by Lemma 3, we obtain the following theorem:

► **Theorem 4.** *There exists a family $\{\mathcal{P}_0, \mathcal{P}_1, \dots\}$ of leaderless and 2-way population protocols such that \mathcal{P}_n has at most $4\lceil \log n \rceil + 7$ states and computes the predicate $x \geq n$.*

An existential $(\log n)^{1/4}$ lower bound. We show that every family $\{\mathcal{P}_n\}_{n \in \mathbb{N}}$ of leaderless and 2-way protocols computing the family of predicates $\{x \geq n\}_{n \in \mathbb{N}}$ must contain infinitely many members of size $\Omega((\log n)^{1/4})$. We call this an existential lower bound, contrary to a universal lower bound, which would state that \mathcal{P}_n has size $\Omega((\log n)^{1/4})$ for every $n \geq 1$.

► **Theorem 5.** *Let $\{\mathcal{P}_0, \mathcal{P}_1, \dots\}$ be an infinite family of leaderless and 2-way population protocols such that \mathcal{P}_n computes the predicate $x \geq n$ for every $n \in \mathbb{N}$. There exist infinitely many indices n such that \mathcal{P}_n has at least $(\log n)^{1/4}$ states.*

Proof sketch. The proof boils down to bounding the number $d(m)$ of unary predicates computed by protocols with m states. The number of distinct sets of transitions, excluding silent ones, is bounded by $2^{m^4 - m^2}$. The number of possible initial states and output mappings are respectively m and 2^m . Altogether, we obtain:

$$d(m) \leq 2^{m^4 - m^2} \cdot m \cdot 2^m = 2^{m^4} \cdot \frac{2^m \cdot m}{2^{m^2}} \leq 2^{m^4}. \quad \blacktriangleleft$$

4 A $O(\log \log n)$ protocol with leaders for some $x \geq n$

The lower bound of Section 3 is not valid for every n , it only ensures that, for some values of n , protocols computing $x \geq n$ must have a logarithmic number of states. We prove that, surprisingly, there is an infinite sequence $n_1 < n_2 < \dots$ of values that break through the logarithmic barrier: The predicates $x \geq n_i$ can be computed by protocols with only $O(\log \log n_i)$ states and two leaders. So, loosely speaking, a flock of birds can decide if it contains at least n_i birds, even though no bird has enough memory to index a bit of n_i .

The result is based on a construction of [15]. In this paper, Mayr and Meyer study the word problem for commutative semigroup presentations. Given a finite set \mathcal{A} of generators, a presentation of a commutative semigroup generated by \mathcal{A} is a finite set of productions $\mathcal{S} = \{l_1 \rightarrow r_1, \dots, l_m \rightarrow r_m\}$, where $l_i, r_i \in \mathcal{A}^*$ for every $1 \leq i \leq m$, satisfying:

- Commutativity: $ab \rightarrow ba \in \mathcal{S}$ for every $a, b \in \mathcal{A}$;² and
- Reversibility: if $l \rightarrow r \in \mathcal{S}$, then $r \rightarrow l \in \mathcal{S}$.

Given $\alpha, \beta \in \mathcal{A}^*$, we say that β is *derived* from α in one step, denoted by $\alpha \rightarrow \beta$, if $\alpha = \gamma l \delta$ and $\beta = \gamma r \delta$ for some $\gamma, \delta \in \mathcal{A}^*$ and some $r \rightarrow l \in \mathcal{S}$. We say that β is *derived* from α if $\alpha \xrightarrow{*} \beta$, where $\xrightarrow{*}$ is the reflexive transitive closure of the relation induced by \rightarrow . Observe that, by reversibility, we have $\alpha \xrightarrow{*} \beta$ iff $\beta \xrightarrow{*} \alpha$. Further, by commutativity we have $\alpha \xrightarrow{*} \beta$ iff $\pi(\alpha) \xrightarrow{*} \pi'(\beta)$ for every permutation π of \mathcal{A} .

² In [15], the elements of S are written using uppercase letters. We use lowercase for convenience.

Mayr and Meyer study the following question: given a commutative semigroup presentation \mathcal{S} over \mathcal{A} , and initial and final letters $s, f \in \mathcal{A}$, what is the length of the shortest word α such that $s \xrightarrow{*} f\alpha$? They exhibit a family of presentations of size $O(n)$ for which the shortest α has double exponential length 2^{2^n} . More precisely, in [15, Sect. 6], they construct a family $\{\mathcal{S}_n\}_{n \geq 1}$ of presentations over alphabets $\{\mathcal{A}_n\}_{n \geq 1}$ satisfying the following properties:

- (1) $|\mathcal{A}_n| = 14n + 10$, $|\mathcal{S}_n| = 20n + 8$, and $\max\{|l|, |r| : l \rightarrow r \in \mathcal{S}_n\} = 5$.
- (2) $\{s_n, f_n, b_n, c_n\} \subseteq \mathcal{A}_n$ for every $n \geq 1$.
- (3) $s_n c_n \xrightarrow{*} f_n \alpha$ iff $\alpha = c_n b_n^{2^{2^n}}$ [15, Lemma 6 and 8].

To apply this result, for each $n \geq 1$ we construct a 5-way population protocol $\mathcal{P}_n = (Q_n, T_n, I_n, L_n, O_n)$ with two leaders as follows:

- $Q_n \stackrel{\text{def}}{=} \mathcal{A}_n \cup \{x\}$ for some $x \notin \mathcal{A}_n$.
- $T_n \stackrel{\text{def}}{=} T_n^1 \cup T_n^2$, where:
 - T_n^1 contains a transition $\text{pad}(p)$ for every production $p = l \rightarrow r$ of \mathcal{S}_n , obtained by “padding” p with x so that its left and right sides have the same length. For example, $\text{pad}(aab \rightarrow cd) = a, a, b \mapsto c, d, x$, and $\text{pad}(a \rightarrow bc) = a, x \mapsto b, c$,
 - $T_n^2 \stackrel{\text{def}}{=} \{f_n, q \mapsto f_n, f_n \mid q \in Q_n\}$,
- $I_n \stackrel{\text{def}}{=} \{x\}$,
- $L_n \stackrel{\text{def}}{=} \{c_n, s_n\}$, and
- $O_n(f_n) \stackrel{\text{def}}{=} 1$ and $O_n(q) \stackrel{\text{def}}{=} 0$ for every $q \neq f_n$.

Intuitively, T_n^1 allows \mathcal{P}_n to simulate derivations of \mathcal{S}_n : a step $C \xrightarrow{\text{pad}(p)} C'$ of \mathcal{P}_n simulates a one-step derivation of \mathcal{S}_n . We make this more precise. Given $\alpha \in \mathcal{A}_n^*$ and $m \geq |\alpha|$, let $C_{\alpha, m}$ be the configuration of \mathcal{P}_n defined as follows: $C_{\alpha, m}(x) = m$, and $C_{\alpha, m}(a) = |\alpha|_a$ for every $a \in \mathcal{A}_n$, where $|\alpha|_a$ is the number of occurrences of a in α . Further, given a configuration C of \mathcal{P}_n , let α_C be the element of \mathcal{S}_n given by $\alpha_C = a_1^{C(a_1)} \dots a_m^{C(a_m)}$, where a_1, \dots, a_m is a fixed enumeration of \mathcal{A}_n . We have:

- **Lemma 6.** *Let $\alpha, \beta \in \mathcal{A}_n^*$ and let C, C' be configurations of \mathcal{P}_n .*
- (a) *If $\alpha \xrightarrow{p_1 \dots p_k} \beta$ in \mathcal{S}_n , then for every $m \geq 4k$, $C_{\alpha, m} \xrightarrow{\text{pad}(p_1) \dots \text{pad}(p_k)} C_{\beta, m'}$ in \mathcal{P}_n for some $m' \geq 0$.*
 - (b) *If $C \xrightarrow{\text{pad}(p_1) \dots \text{pad}(p_k)} C'$ in \mathcal{P}_n , then $\alpha_C \xrightarrow{p_1 \dots p_k} \alpha_{C'}$ in \mathcal{S}_n .*

From Lemma 6, (1) and (3), the following can be shown:

► **Theorem 7.** *For every $n \in \mathbb{N}$, there is a 5-way protocol \mathcal{P}_n with at most $14n + 11$ states and at most $34n + 19$ transitions that computes the predicate $x \geq c_n$ for some number $c_n \geq 2^{2^n}$.*

Using Theorem 7 and Lemma 3, we obtain:

► **Corollary 8.** *There exists a family $\{\mathcal{P}_0, \mathcal{P}_1, \dots\}$ of 2-way protocols with two leaders and a family $\{c_0, c_1, \dots\}$ of natural numbers such that for every $n \in \mathbb{N}$ the following holds: $c_n \geq 2^{2^n}$ and protocol \mathcal{P}_n has at most $314 \log \log c_n + 131$ states and computes the predicate $x \geq c_n$.*

5 Universal lower bounds for 1-aware protocols

To the best of our knowledge, all the protocols in the literature for predicates $x \geq n$, including those of Section 3 and Section 4, share a very natural property: if the number of agents is greater than or equal to n , then the agents not only eventually reach consensus 1, they also eventually *know* that they will reach this consensus. Let us formalize this idea:

► **Definition 9.** A well-specified population protocol $\mathcal{P} = (Q, T, I, L, O)$ is *1-aware* if there is a set $Q_1 \subseteq Q \setminus (I \cup \llbracket L \rrbracket)$ of states such that for every initial configuration C_0 and every fair execution $\pi = C_0 C_1 \dots$

- (1) if π stabilizes to 0, then $C_i(Q_1) = 0$ for every $i \geq 0$, and
- (2) if π stabilizes to 1, then there is some $i \geq 0$ such that $C_j(Q \setminus Q_1) = 0$ for every $j \geq i$.

If in the course of an execution π an agent reaches a state of Q_1 , then π cannot stabilize to 0 by (1), and so, since \mathcal{P} is well-specified, it stabilizes to 1; intuitively, at this moment the agent “knows” that the consensus will be 1. Further, if an execution stabilizes to 1, then all agents eventually reach and remain in Q_1 by (2), and so eventually all agents “know”.³ Albeit seemingly restrictive, 1-aware protocols compute a significant subclass of predicates: monotonic Presburger predicates (see the full version of the paper for more details).

We say that a state q is *coverable* from a configuration C if $C \xrightarrow{*} C'$ for some configuration C' such that $C'(q) > 0$. The fundamental property of 1-aware protocols is that, loosely speaking, consensus reduces to coverability:

► **Lemma 10.** *Let $\mathcal{P} = (Q, T, \{x\}, L, O)$ be a 1-aware protocol computing a unary predicate φ . We have $\varphi(n) = 1$ if and only if some state of Q_1 is coverable from $\langle n \cdot x \rangle + L$.*

We show that for 1-aware protocols, the bounds of Sections 3 and 4 are essentially tight.

Leaderless protocols. We prove that a 1-aware, leaderless and 2-way protocol computing $x \geq n$ has at least $\log_3 n$ states. By Lemma 10, it suffices to show that some state of Q_1 is coverable from $\langle 3^k \cdot q \rangle$, where q is the initial state. Proposition 11 below is the key to the proof. It states that for every finite execution $C_1 \xrightarrow{\pi} C_2$, there is $C'_1 \xrightarrow{\pi'} C'_2$ such that C'_1 has the same support as C_1 and is not too large, and C'_2 contains a “record” of all states encountered during the execution of π (this is the set $\llbracket C_1 \rrbracket \cup \llbracket \pi \rrbracket^\bullet$).

Let us define the *norm* of a configuration C as $\|C\| \stackrel{\text{def}}{=} \max\{C(q) : q \in \llbracket C \rrbracket\}$. We obtain:

► **Proposition 11.** *Let $\mathcal{P} = (Q, T, I, L, O)$ be a k -way population protocol and let $C_1 \xrightarrow{\pi} C_2$ be a finite execution of \mathcal{P} . There exists a finite execution $C'_1 \xrightarrow{\pi'} C'_2$ such that (a) $\llbracket C'_1 \rrbracket = \llbracket C_1 \rrbracket$, (b) $\llbracket C'_2 \rrbracket = \llbracket C_1 \rrbracket \cup \llbracket \pi \rrbracket^\bullet$, and (c) $\|C'_1\| \leq (k + 1)^{|Q|}$.*

► **Theorem 12.** *Every 1-aware, leaderless and 2-way population protocol $\mathcal{P} = (Q, T, \{q_0\}, \mathbf{0}, O)$ computing $x \geq n$ has at least $\log_3 n$ states.*

Proof. Let $Q_1 \subseteq Q$ be the set of states from the definition of 1-awareness. Since $L = \mathbf{0}$, $C_0 = \langle n \cdot q_0 \rangle$ is the smallest initial configuration with output 1, and by Lemma 10 the smallest initial configuration from which some state $q_1 \in Q_1$ is coverable. Let $C_0 \xrightarrow{\pi} C \geq \langle q_1 \rangle$. Since $q_1 \neq q_0$, we have $q_1 \in \llbracket \pi \rrbracket^\bullet$. By Proposition 11, and since \mathcal{P} is 2-way, q_1 is also coverable from C'_0 satisfying $\llbracket C'_0 \rrbracket = \llbracket C_0 \rrbracket = \{q_0\}$ and $\|C'_0\| = 3^{|Q|}$. Thus, $C'_0 = \langle 3^{|Q|} \cdot q_0 \rangle$. By minimality of n , we get $n \leq 3^{|Q|}$, and thus $|Q| \geq \log_3 n$. ◀

Observe that the proof Theorem 12 uses the fact that \mathcal{P} is leaderless to conclude $C'_0 = \langle 3^{|Q|} \cdot q_0 \rangle$ from $\llbracket C'_0 \rrbracket = \llbracket C_0 \rrbracket$ and $\|C'_0\| = 3^{|Q|}$, which is not necessarily true with leaders.

Protocols with leaders. In the case of protocols with leaders we obtain a lower bound from Rackoff’s procedure for the coverability problem of vector addition systems [17].

³ We could also require the seemingly weaker property that eventually at least one agent “knows”. However, by adding transitions that “attract” all other agents to Q_1 , we can transform a protocol in which some agent “knows” into a protocol computing the same predicate in which all agents “know”.

17:10 Large Flocks of Small Birds: On the Minimal Size of Population Protocols

A *vector addition system* of dimension k (k -VAS) is a pair (A, \mathbf{v}_0) , where $\mathbf{v}_0 \in \mathbb{N}^k$ is an initial vector and $A \subseteq \mathbb{Z}^k$ is a set of vectors. An execution of a k -VAS is a sequence $\mathbf{v}_0 \mathbf{v}_1 \cdots \mathbf{v}_n$ of vectors of \mathbb{N}^k such that each $\mathbf{v}_{i+1} = \mathbf{v}_i + \mathbf{a}_i$ for some $\mathbf{a}_i \in A$. We write $\mathbf{v}_0 \xrightarrow{*} \mathbf{v}_n$ and say that the execution has *length* n . A vector \mathbf{v} is *coverable* in (A, \mathbf{v}_0) if $\mathbf{v}_0 \xrightarrow{*} \mathbf{v}'$ for some $\mathbf{v}' \geq \mathbf{v}$. The *size* of a vector $\mathbf{v} \in \mathbb{Z}^k$ is $\sum_{1 \leq i \leq k} \text{size}(\max(|\mathbf{v}(i)|, 1))$. The *size* of a set of vectors is the sum of the size of its vectors. In [17] Rackoff proves:

► **Theorem 13** ([17]). *Let $A \subseteq \mathbb{Z}^k$ be a set of vectors of size at most n and dimension $k \leq n$, and let $\mathbf{v}_0 \in \mathbb{N}^k$ be a vector of size n . For every $\mathbf{v} \in \mathbb{N}^k$, if \mathbf{v} is coverable in (A, \mathbf{v}_0) , then \mathbf{v} is coverable by means of an execution of length at most $2^{(3n)^n}$.*

Using a standard construction from the Petri net literature, it can be shown that every 2-way protocol \mathcal{P} with n states can be simulated by a VAS $\mathcal{V}_{\mathcal{P}}$ of size at most $12n^8$, where each execution of \mathcal{P} has a corresponding execution twice as long in $\mathcal{V}_{\mathcal{P}}$. Thus, by Theorem 13:

► **Proposition 14.** *Let $\mathcal{P} = (Q, T, I, L, O)$ be a 2-way population protocol and let $q \in Q$. For every configuration C , if q is coverable from C , then it is coverable by means of a finite execution of length at most $2^{(3m)^{m-1}}$ where $m = 12|Q|^8$.*

Using the above proposition, we derive:

► **Theorem 15.** *Let \mathcal{P} be a 1-aware and 2-way population protocol. For every $n \geq 2$, if \mathcal{P} computes $x \geq n$, then \mathcal{P} has at least $(\log \log(n)/151)^{1/9}$ states.*

6 Protocols for systems of linear inequalities

In Section 3, we have shown that the predicate $x \geq c$ can be computed by a leaderless protocol with $O(\log c)$ states. In this section, we will see that adding a few leaders allows to compute systems of linear inequalities. More formally, we show that there exists a protocol with $O((m+k) \cdot \log(dm))$ states and $O(m \cdot \log(dm))$ leaders computing the predicate $A\mathbf{x} \geq \mathbf{c}$, where $A \in \mathbb{Z}^{m \times k}$, $\mathbf{c} \in \mathbb{Z}^m$ and d is the the largest absolute value occurring in A and \mathbf{c} .

There are three crucial points that make systems of linear inequalities more complicated than flock-of-birds predicates: (1) variables have *coefficients*, (2) coefficients may be positive or *negative*, and (3) they are the *conjunction* of linear inequalities. We will explain how to address the two first points by considering the special case of linear inequalities. We will then discuss how to handle the third point.

Linear inequalities. Note that the predicate $\sum_{1 \leq i \leq k} a_i x_i \geq c$ is equivalent to $\sum_{1 \leq i \leq k} a_i x_i + (1-c) > 0$. Therefore, it suffices to describe protocols for predicates of the form $\sum_{1 \leq i \leq k} a_i x_i + c > 0$. In order to make the presentation more pleasant, we will first restrain ourselves to the predicate $ax - by + c > 0$ for some fixed $a, b \in \mathbb{N}$ and $c \in \mathbb{Z}$. Such a predicate admits the difficult aspects, i.e. coefficients and negative numbers. Moreover, as we will see, handling more than two variables is not an issue.

Let us now describe a protocol \mathcal{P}_{lin} for the predicate $ax - by + c > 0$. The idea is to keep a representation of $ax - by + c$ throughout executions of the protocol. Let $n \stackrel{\text{def}}{=} \text{size}(\max(\log |a|, \log |b|, \log |c|, 1))$. As in Section 3, we construct states to represent powers of two. However, this time, we also need states to represent negative numbers:

$$Q^+ \stackrel{\text{def}}{=} \{+2^i : 0 \leq i \leq n\} \quad \text{and} \quad Q^- \stackrel{\text{def}}{=} \{-2^i : 0 \leq i \leq n\}.$$

We also need states $X \stackrel{\text{def}}{=} \{\mathbf{x}, \mathbf{y}\}$ for the variables, and two additional states $R \stackrel{\text{def}}{=} \{+\mathbf{0}, -\mathbf{0}\}$. The set of all states of \mathcal{P}_{lin} is $Q \stackrel{\text{def}}{=} X \cup Q^+ \cup Q^- \cup R$, and the initial states are $I \stackrel{\text{def}}{=} X$.

Let us explain the purpose of R . Intuitively, we would like to have the transitions:

$$x \mapsto \wr +2^i : i \in \text{bits}(a) \wr \quad \text{and} \quad y \mapsto \wr -2^i : i \in \text{bits}(|b|) \wr.$$

This way, every agent in state \mathbf{x} (resp. \mathbf{y}) could be converted to the binary representation of a (resp. b). Unfortunately, this is not possible as these transitions produce more states than they consume. This is where leaders become useful. If R initially contains enough leaders, then R can act as a *reservoir* of extra states which allow to “pad” transitions. More formally, let $\text{rep}(z) : \mathbb{Z} \rightarrow \text{Pop}(Q \setminus X)$ be defined as follows:

$$\text{rep}(z) \stackrel{\text{def}}{=} \begin{cases} \wr +2^i : i \in \text{bits}(z) \wr & \text{if } z > 0, \\ \wr -2^i : i \in \text{bits}(|z|) \wr & \text{if } z < 0, \\ \wr -0 \wr & \text{if } z = 0. \end{cases}$$

For every $r \in R$, we add to \mathcal{P}_{lin} the following transitions:

$$\text{add}_{\mathbf{x},r} : \mathbf{x}, \underbrace{r, r, \dots, r}_{|\text{rep}(a)|-1 \text{ times}} \mapsto \text{rep}(a) \quad \text{and} \quad \text{add}_{\mathbf{y},r} : \mathbf{y}, \underbrace{r, r, \dots, r}_{|\text{rep}(b)|-1 \text{ times}} \mapsto \text{rep}(b).$$

We set the leaders to $L \stackrel{\text{def}}{=} \text{rep}(c) + \wr (4n+2) \cdot -0 \wr$. We claim that $4n+2$ reservoir states are enough, we will explain later why. Now, the key idea of the construction is that it is always possible to put $2n$ agents back into R . Thus, fairness ensures that the number of agents in X eventually decreases to zero, and then that the value represented over $Q^+ \cup Q^-$ is $ax - by + c$. We let the representations over Q^+ and Q^- “cancel out” until one side “wins”. If the positive (resp. negative) side wins, i.e. if $ax - by + c > 0$ (resp. $ax - by + c \leq 0$), then it signals all agents in R to move to $+0$ (resp. -0). To achieve this, for every $0 \leq i \leq n$, we add transition $\text{cancel}_i : +2^i, -2^i \mapsto +0, -0$ to the protocol. Since bits of the positive and negative numbers may not be “aligned”, we follow the idea of Section 3 and add further transitions to change representations to equivalent ones:

$$\begin{aligned} \text{up}_i^+ : +2^i, +2^i \mapsto +2^{i+1}, +0, & \quad \text{down}_{i+1,r}^+ : +2^{i+1}, r \mapsto +2^i, +2^i, \\ \text{up}_i^- : -2^i, -2^i \mapsto -2^{i+1}, -0, & \quad \text{down}_{i+1,r}^- : -2^{i+1}, r \mapsto -2^i, -2^i, \end{aligned}$$

where $0 \leq i < n$ and $r \in R$. Finally, for every $0 \leq i \leq n$, we add transitions to signal which side wins:

$$\begin{aligned} \text{signal}_i^+ : +2^i, -0 \mapsto +2^i, +0, & \quad \text{signal} : -0, +0 \mapsto -0, -0, \\ \text{signal}_i^- : -2^i, +0 \mapsto -2^i, -0. & \end{aligned}$$

Note that -0 “wins” over $+0$ because the predicate is false whenever $ax - by + c = 0$. It remains to specify the output mapping of \mathcal{P}_{lin} which we define as expected, i.e. $O(q) \stackrel{\text{def}}{=} 1$ if $q \in Q^+ \cup \{+0\}$, and $O(q) \stackrel{\text{def}}{=} 0$ otherwise.

Let us briefly explain why $4n+2$ reservoir states suffice. At any reachable configuration C , transitions of the form up_i^+ and up_i^- can occur until $C(\pm 2^i) \leq 1$ for every $0 \leq i < n$. Afterwards, at most $2n$ agents remain in these states. There can however be many agents in $S = \{+2^n, -2^n\}$. But, these two states represent numbers respectively larger and smaller than any coefficient, hence the number of agents in S can only grow by one each time a state from X is consumed. Overall, this means that $C \xrightarrow{*} C'$ for some C' such that $C'(R) \geq 2n$.

In order to handle more variables $\{x_1, x_2, \dots, x_k\}$, note that all we need to do is to set $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k\}$ instead, and add transitions $\text{add}_{\mathbf{x}_i,r}$ for every $1 \leq i \leq k$ and $r \in R$.

By applying Lemma 3 on \mathcal{P}_{lin} , we obtain:

► **Theorem 16.** *Let $a_1, a_2, \dots, a_k, c \in \mathbb{Z}$ and let $n = \text{size}(\max(|a_1|, |a_2|, \dots, |a_k|, |c|, 1))$. There exists a 2-way population protocol, with at most $10kn$ states and at most $5n + 2$ leaders, that computes the predicate $\sum_{1 \leq i \leq k} a_i x_i + c > 0$.*

Conjunction of linear inequalities. We briefly explain how to lift the construction for linear inequalities to systems of linear inequalities. The details of the formal construction and proofs are a bit involved, and are thus deferred to the full version of this paper. Let us fix some $A \in \mathbb{Z}^{m \times k}$ and $\mathbf{c} \in \mathbb{Z}^m$. We sketch a protocol \mathcal{P}_{sys} for the predicate $A\mathbf{x} + \mathbf{c} > \mathbf{0}$. For every $1 \leq i \leq m$, we construct a protocol \mathcal{P}_i for the predicate $\sum_{1 \leq j \leq k} A_{i,j} \cdot \mathbf{x}_j + \mathbf{c}_i > 0$. Protocol \mathcal{P}_i is obtained as presented earlier, but with some modifications. The largest power of two is picked as $n \stackrel{\text{def}}{=} \text{size}(d) + \lceil \log 2m^2 \rceil$ where

$$d \stackrel{\text{def}}{=} \max(1, \{|A_{i,j}| : 1 \leq i \leq m, 1 \leq j \leq k\}, \{|c_i| : 1 \leq i \leq m\}).$$

The reason for this modification is that the number of agents, in a largest power of two, should now increase by at most $1/m$ each time an initial state is consumed, as opposed to 1.

We also replace each *positive state* $q \in Q^+$ of \mathcal{P}_i by two states q_0 and q_1 , its *0-copy* and *1-copy*. The reason behind this is that positive states should not necessarily have output 1. Indeed, one linear inequality may be satisfied while the other ones are not. Therefore, $-\mathbf{0}$ and each *negative state* $q \in Q^-$ should be able to signal a 0-consensus to the positive states. The transitions of the form up_j^+ , down_j^+ and cancel_j are adapted accordingly.

Protocol \mathcal{P}_{sys} is obtained as follows. First, subprotocols $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$ are put side by side. Their initial (resp. reservoir) states are merged into a single set X (resp. R). For every $1 \leq j \leq k$, transitions $\text{add}_{\mathbf{x}_j, r}$ of the m subprotocols are replaced by a single transition consuming \mathbf{x}_j , and enough reservoir states, and producing $\text{rep}(A_{i,j})$ in each subprotocol \mathcal{P}_i , where $1 \leq i \leq m$. The signal mechanisms are replaced by these new ones:

- the 0-copy of state $+\mathbf{2}^0$ of *all* subprotocols can meet to convert $-\mathbf{0}$ to $+\mathbf{0}$,
- state $+\mathbf{0}$ can convert any positive state to its 1-copy,
- state $-\mathbf{0}$ or any negative state can convert $+\mathbf{0}$ to $-\mathbf{0}$, and any positive state to its 0-copy.

A careful analysis of the formal construction of \mathcal{P}_{sys} combined with Lemma 3 yields:

► **Theorem 17.** *Let $A \in \mathbb{Z}^{m \times k}$, $\mathbf{c} \in \mathbb{Z}^m$ and $n = \text{size}(\max(1, \{|A_{i,j}| : 1 \leq i \leq m, 1 \leq j \leq k\}), \{|c_i| : 1 \leq i \leq m\})$. There exists a 2-way population protocol, with at most $27(\log m + n)(m + k)$ states and at most $14m(\log m + n)$ leaders, that computes the predicate $A\mathbf{x} + \mathbf{c} > \mathbf{0}$.*

7 Conclusion and further work

We have initiated the study of the state space size of population protocols as a function of the size of the predicate they compute. Previous lower bounds were only for single predicates, like the majority predicate $x \leq y$, or for a variant of the model in which the number of states is a function of the number of agents.

There are many open questions. We conjecture that systems of linear inequalities can be computed by leaderless protocols with a polynomial number of states. A second, very intriguing question is whether the function $f(n)$ giving the minimal number of states of a two-leader protocol computing $x \geq n$ exhibits large gaps, i.e., if there are (families of) numbers c and $c + 1$ such that $f(c)$ is exponentially larger than $f(c + 1)$. A third question is whether there exist protocols with $O(\log \log \log n)$ states for the flock-of-birds predicates $x \geq n$. Such protocols cannot be 1-aware, but they might exist. Their existence is linked to the long standing question of whether the reachability problem for reversible VAS (a model equivalent to the commutative semigroup representations of [15]) has the same complexity as reachability for arbitrary VAS (see [12] for a brief introduction).

References

- 1 Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. Time-space trade-offs in population protocols. In *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2560–2579, 2017. doi:10.1137/1.9781611974782.169.
- 2 Dan Alistarh, Rati Gelashvili, and Milan Vojnović. Fast and exact majority in population protocols. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 47–56, 2015. doi:10.1145/2767386.2767429.
- 3 Dana Angluin, James Aspnes, Melody Chan, Michael J. Fischer, Hong Jiang, and René Peralta. Stably computable properties of network graphs. In *Proc. First International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 63–74, 2005. doi:10.1007/11502593_8.
- 4 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. In *Proc. 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 290–299, 2004. doi:10.1145/1011767.1011810.
- 5 Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, 2008. doi:10.1007/s00446-008-0067-z.
- 6 Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007. doi:10.1007/s00446-007-0040-2.
- 7 James Aspnes. Clocked population protocols. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 431–440, 2017. doi:10.1145/3087801.3087836.
- 8 James Aspnes and Eric Ruppert. An introduction to population protocols. In *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer Berlin Heidelberg, 2009. doi:10.1007/978-3-540-89707-1_5.
- 9 Andreas Bilke, Colin Cooper, Robert Elsässer, and Tomasz Radzik. Brief announcement: Population protocols for leader election and exact majority with $O(\log^2 n)$ states and $O(\log^2 n)$ convergence time. In *Proc. ACM Symposium on Principles of Distributed Computing, PODC*, pages 451–453, 2017. doi:10.1145/3087801.3087858.
- 10 David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. In *Proc. 29th International Symposium on Distributed Computing (DISC)*, pages 602–616, 2015. doi:10.1007/978-3-662-48653-5_40.
- 11 Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of population protocols. *Acta Informatica*, 54(2):191–215, 2017. doi:10.1007/s00236-016-0272-3.
- 12 Alain Finkel and Jérôme Leroux. Recent and simple algorithms for Petri nets. *Software and System Modeling*, 14(2):719–725, 2015. doi:10.1007/s10270-014-0426-0.
- 13 Rachid Guerraoui and Eric Ruppert. Names trump malice: Tiny mobile agents can tolerate byzantine failures. In *Proc. 36th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 484–495, 2009. doi:10.1007/978-3-642-02930-1_40.
- 14 Giuseppe Antonio Di Luna, Paola Flocchini, Taisuke Izumi, Tomoko Izumi, Nicola Santoro, and Giovanni Viglietta. Population protocols with faulty interactions: The impact of a leader. In *Proc. 10th International Conference Algorithms and Complexity (CIAC)*, pages 454–466, 2017. doi:10.1007/978-3-319-57586-5_38.
- 15 Ernst W. Mayr and Albert R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Advances in Mathematics*, 46(3):305–329, 1982. doi:10.1016/0001-8708(82)90048-2.
- 16 Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Mediated population protocols. *Theoretical Computer Science*, 412(22):2434–2450, 2011. doi:10.1016/j.tcs.2011.02.003.

17:14 Large Flocks of Small Birds: On the Minimal Size of Population Protocols

- 17 Charles Rackoff. The covering and boundedness problems for vector addition systems. *Theoretical Computer Science*, 6:223–231, 1978. doi:10.1016/0304-3975(78)90036-1.
- 18 David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008. doi:10.1007/s11047-008-9067-y.