



HAL
open science

Lower bounds on the state complexity of population protocols

Philipp Czerner, Javier Esparza, Jérôme Leroux

► **To cite this version:**

Philipp Czerner, Javier Esparza, Jérôme Leroux. Lower bounds on the state complexity of population protocols. *Distributed Computing*, 2023, 36 (3), pp.209-218. 10.1007/s00446-023-00450-4. hal-04286895

HAL Id: hal-04286895

<https://cnrs.hal.science/hal-04286895v1>

Submitted on 14 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Lower bounds on the state complexity of population protocols

Philipp Czerner¹ · Javier Esparza¹ · Jérôme Leroux²

Received: 29 November 2021 / Accepted: 25 April 2023 / Published online: 15 June 2023
© The Author(s) 2023

Abstract

Population protocols are a model of computation in which an arbitrary number of indistinguishable finite-state agents interact in pairs. The goal of the agents is to decide by stable consensus whether their initial global configuration satisfies a given property, specified as a predicate on the set of configurations. The state complexity of a predicate is the number of states of a smallest protocol that computes it. Previous work by Blondin et al. has shown that the counting predicates $x \geq \eta$ have state complexity $\mathcal{O}(\log \eta)$ for leaderless protocols and $\mathcal{O}(\log \log \eta)$ for protocols with leaders. We obtain the first non-trivial lower bounds: the state complexity of $x \geq \eta$ is $\Omega(\log \log \eta)$ for leaderless protocols, and the inverse of a non-elementary function for protocols with leaders.

Keywords Population protocols · State complexity

1 Introduction

Population protocols are a model of computation in which an arbitrary number of indistinguishable finite-state agents interact in pairs to decide if their initial global configuration satisfies a given property. Population protocols were introduced in [5, 6] to study the theoretical properties networks of mobile sensors with very limited computational resources, but they are also very strongly related to chemical reaction networks, a discrete model of chemistry in which agents are molecules that change their states due to collisions.

Population protocols decide a property by *stable consensus*. Each state of an agent is assigned a binary output (yes/no). In a correct protocol, all agents eventually reach the set of states whose output is the correct answer to the question “did our initial configuration satisfy the property?”, and stay in it forever. An example of a property decidable by population protocols is majority: initially agents are in one of two initial states, say A and B , and the property to be decided

is whether the number of agents in A is larger than the number of agents in B . In a seminal paper, Angluin et al. showed that population protocols can decide exactly the properties expressible in Presburger arithmetic, the first-order theory of addition [8].

Assume that at each step a pair of agents is selected uniformly at random and allowed to interact. The *parallel runtime* is defined as the expected number of interactions until a stable consensus is reached (i.e. until the property is decided), divided by the number of agents. Even though the parallel runtime is computed using a discrete model, under reasonable and commonly accepted assumptions the result coincides with the runtime of a continuous-time stochastic model. Many papers have investigated the parallel runtime of population protocols, and several landmark results have been obtained. In [6] it was shown that every Presburger property can be decided in $\mathcal{O}(n \log n)$ parallel time, where n is the number of agents, and [7] showed that population protocols with a fixed number of leaders can compute all Presburger predicates in polylogarithmic parallel time. (Loosely speaking, leaders are auxiliary agents that do not form part of the population of “normal” agents, but can interact with them to help them decide the property.) More recent results have studied protocols for majority and leader election in which the number of states grows with the number of agents, and shown that polylogarithmic time is achievable by protocols without leaders, even for very slow growth functions, see e.g. [2–4, 16, 19].

✉ Philipp Czerner
czerner@in.tum.de

Javier Esparza
esparza@in.tum.de

Jérôme Leroux
jerome.leroux@labri.fr

¹ Department of Computer Science, CIT School, Technical University of Munich, Munich, Germany

² LaBRI, University of Bordeaux, CNRS, Bordeaux, France

However, many protocols have a high number of states. For example, a quick estimate shows that the fast protocol for majority implicitly described in [7] has tens of thousands of states. This is an obstacle to implementations of protocols in chemistry, where the number of states corresponds to the number of chemical species participating in the reactions. The number of states is also important because it plays the role of memory in sequential computational models; indeed, the total memory available to a protocol is the logarithm of the number of states multiplied by the number of agents. Despite these facts, the *state complexity* of a Presburger property, defined as the minimal number of states of any protocol deciding the property, has received comparatively little attention¹. In [11, 12] Blondin et al. have shown that every predicate representable by a boolean combination of threshold and modulo constraints (every Presburger formula can be put into this form), with numbers encoded in binary, can be decided by a protocol with polynomially many states in the length of the formula. In particular, it is not difficult to see that every property of the form $x \geq \eta$, stating that the number of agents is at least η , can be decided by a leaderless protocol with $\mathcal{O}(\log \eta)$ states. A theorem of [11] also proves the existence of an infinite family of thresholds η such that $x \geq \eta$ can be decided by a protocol (with leaders) having $\mathcal{O}(\log \log \eta)$ states. However, to the best of our knowledge there exist no *lower* bounds on the state complexity, i.e. bounds showing that a protocol for $x \geq \eta$ needs $\Omega(f(\eta))$ states for some function f . This question, which was left open in [12], is notoriously hard due to its relation to fundamental questions in the theory of Vector Addition Systems.

In this paper we first show that every protocol, with or without leaders, needs a number of states that, roughly speaking, grows like the inverse Ackermann function, and then prove our main result: every leaderless protocol for $x \geq \eta$ needs $\Omega(\log \log \eta)$ states. The proof of the first bound relies on results on the maximal length of controlled antichains of \mathbb{N}^d , a topic in combinatorics with a long tradition in the study of Vector Addition Systems and other models, see e.g. [1, 9, 18, 23, 26]. The double logarithmic bound follows from Pottier's small basis theorem, a useful result of the theory of Diophantine equations [24].

The paper is organised as follows. Section 2 introduces population protocols, the state complexity function, and its inverse, the busy beaver function, which assigns to a number of states n the largest η such that a protocol with n states

decides $x \geq \eta$. Instead of lower bounds on state complexity, we present upper bounds on the busy beaver function for convenience. Section 3 presents some results on the mathematical structure of stable sets of configurations that are used throughout the paper. Section 4 shows an Ackermannian upper bound on the busy beaver function, valid for protocols with or without leaders, and explains why this very large bound might be optimal. Section 5 gives a triple exponential upper bound on the busy beaver function for leaderless protocols.

2 Population protocols and state complexity

2.1 Mathematical preliminaries

For sets A, B we write A^B to denote the set of functions $f: B \rightarrow A$. If B is finite we call the elements of \mathbb{N}^B *multisets* over B . We sometimes write multisets using set-like notation, e.g. $\{a, b, b\}$ and $\{a, 2 \cdot b\}$ denote the multiset m such that $m(a) = 1$, $m(b) = 2$ and $m(c) = 0$ for every $c \in B \setminus \{a, b\}$. Given a multiset $m \in \mathbb{N}^B$ and $B' \subseteq B$, we define $m(B') := \sum_{b \in B'} m(b)$. The *size* of m is $|m| := m(B)$; in other words, the total number of elements of m . The *support* of m is the set $\llbracket m \rrbracket = \{b \in B \mid m(b) > 0\}$. Abusing language we identify an element $b \in B$ with the one-element multiset containing it, i.e. with the multiset $m \in \mathbb{N}^B$ given by $m(b) = 1$ and $m(b') = 0$ for $b' \neq b$.

We call the elements of \mathbb{Z}^B *vectors* over B of dimension $|B|$. Observe that every multiset is also a vector. Arithmetic operations on vectors in \mathbb{Z}^B are defined as usual, extending the vectors with zeroes if necessary. For example, if $B' \subseteq B$, $u \in \mathbb{Z}^B$, and $v \in \mathbb{Z}^{B'}$, then $u + v \in \mathbb{Z}^B$ is defined by $(u + v)(b) = u(b) + v(b)$, where $v(b) = 0$ for every $b \in B \setminus B'$. For $u, v \in \mathbb{Z}^B$ we write $u \leq v$ if $u_i \leq v_i$ for all $i \in B$, and $u \not\leq v$ if $u \leq v$ and $u \neq v$. Given a vector $v \in \mathbb{Z}^k$, we define $\|v\|_1 = \sum_{i=1}^k |v_i|$ and $\|v\|_\infty = \max_{i=1}^k |v_i|$.

2.2 Population protocols

We recall the population protocol model of [6], with explicit mention of leader agents. A *population protocol* is a tuple $\mathcal{P} = (Q, T, L, X, I, O)$ where

- Q is a finite set of *states*;
- $T \subseteq Q_2 \times Q_2$ is a set of *transitions*, where Q_2 denotes the set of multisets over q of size 2;
- $L \in \mathbb{N}^Q$ is the *leader multiset*;
- X is a finite set of *input variables*;
- $I: X \rightarrow Q$ is the *input mapping*; and
- $O: Q \rightarrow \{0, 1\}$ is the *output mapping*.

¹ Notice that the time-space trade-off results of [2–4, 16, 19] refer to a more general model in which the number of states of a protocol grows with the number n of agents; in other words, a property is decided by a *family* of protocols, one for each value of n . Trade-off results bound the growth rate needed to compute a predicate within a given time. We study the minimal number of states of a *single* protocol that decides the property for *all* n .

We write $p, q \mapsto p', q'$ to denote that the pair $(\{p, q\}, \{p', q'\})$ is a transition. We assume that for every multiset $\{p, q\}$ there is at least one transition of the form $p, q \mapsto p', q'$.

Inputs and configurations An input to \mathcal{P} is a multiset $m \in \mathbb{N}^X$ such that $|m| \geq 2$. A configuration is a multiset $C \in \mathbb{N}^Q$ such that $|C| \geq 2$. Intuitively, a configuration represents a population of agents where $C(q)$ denotes the number of agents in state q . The initial configuration for input m is defined as

$$IC(m) := L + \sum_{x \in X} m(x) \cdot I(x).$$

When \mathcal{P} has a unique input x , i.e. $X = \{x\}$, we abuse language and write $IC(i)$ instead of $IC(i \cdot x)$ to denote the initial configuration for input $i \in \mathbb{N}$.

The output $O(C)$ of a configuration C is b if $C(q) \geq 1$ implies $O(q) = b$ for all $q \in Q$, and undefined otherwise. So a population has output b if all agents have output b .

Executions A transition $t = p, q \mapsto p', q'$ is enabled at a configuration C if $C \geq p + q$, and disabled otherwise. As $|C| \geq 2$ by the definition of configuration, every configuration enables at least one transition. If t is enabled at C , then it can be fired leading to configuration $C' := C - p - q + p' + q'$, which we denote $C \xrightarrow{t} C'$. We write $C \rightarrow C'$ if $C \xrightarrow{t} C'$ for some $t \in T$. Given a sequence $\sigma = t_1 t_2 \dots t_n$ of transitions, we write $C \xrightarrow{\sigma} C'$ if there exist configurations C_1, C_2, \dots, C_n such that $C \xrightarrow{t_1} C_1 \xrightarrow{t_2} C_2 \dots C_n \xrightarrow{t_n} C'$, and $C \xrightarrow{*} C'$ if $C \xrightarrow{\sigma} C'$ for some sequence $\sigma \in T^*$. For every set of transitions $T' \subseteq T$, we write $C \xrightarrow{T'} C'$ if $C \xrightarrow{t} C'$ for some $t \in T'$; we write $C \xrightarrow{T'^*} C'$, and say that C' is reachable from C , if $C \xrightarrow{\sigma} C'$ for some sequence $\sigma \in T'^*$. Given a set \mathcal{C} of configurations, $C \xrightarrow{*} \mathcal{C}$ denotes that $C \xrightarrow{*} C'$ for some $C' \in \mathcal{C}$.

An execution is a sequence of configurations $\sigma = C_0 C_1 \dots$ such that $C_i \rightarrow C_{i+1}$ for every $i \in \mathbb{N}$. The output $O(\sigma)$ of σ is b if there exist $i \in \mathbb{N}$ such that $O(C_i) = O(C_{i+1}) = \dots = b$, otherwise $O(\sigma)$ is undefined.

Executions have the *monotonicity property*: If $C_0 C_1 C_2 \dots$ is an execution, then for every configuration D the sequence $(C_0 + C) (C_1 + C) (C_2 + C) \dots$ is an execution too. We often say that a statement holds “by monotonicity”, meaning that it is a consequence of the monotonicity property.

Computations An execution $\sigma = C_0 C_1 \dots$ is fair if for every configuration C the following holds: if C is reachable from C_i for infinitely many $i \in \mathbb{N}$, then $C_j = C$ for infinitely many $j \in \mathbb{N}$. In other words, fairness ensures that an execution cannot avoid a reachable configuration forever. We say that a population protocol computes a predicate $\varphi: \mathbb{N}^X \rightarrow \{0, 1\}$ (or decides the property represented by the predicate) if for every $v \in \mathbb{N}^X$ every fair execution σ starting from $IC(v)$ satisfies $O(\sigma) = \varphi(v)$. Two protocols are equivalent if they compute the same predicate. It is known

that population protocols compute precisely the Presburger-definable predicates [8].

Example 2.1 Let $\mathcal{P}_k = (Q, T, 0, \{x\}, I, O)$ be the protocol where $Q := \{0, 1, 2, 3, \dots, 2^k\}$, $I(x) := 1$, $O(a) = 1$ iff $a = 2^k$, and the set T of transitions contains $a, b \mapsto 0, a + b$ if $a + b < 2^k$, and $a, b \mapsto 2^k, 2^k$ if $a + b \geq 2^k$ for every $a, b \in Q$. It is readily seen that \mathcal{P}_k computes $x \geq 2^k$ with $2^k + 1$ states. Intuitively, each agent stores a number, initially 1. When two agents meet, one of them stores the sum of their values and the other one stores 0, with sums capping at 2^k . Once an agent reaches 2^k , all agents eventually get converted to 2^k .

Now, consider the protocol $\mathcal{P}'_k = (Q', T', 0, \{x\}, I', O')$, where $Q' := \{0, 2^0, 2^1, \dots, 2^k\}$, $I'(x) := 2^0$, $O'(a) = 1$ iff $a = 2^k$, and T' contains $2^i, 2^i \mapsto 0, 2^{i+1}$ for each $0 \leq i < k$, and $a, 2^k \mapsto 2^k, 2^k$ for each $a \in Q'$. It is easy to see that \mathcal{P}'_k also computes $x \geq 2^k$, but more succinctly; while \mathcal{P}_k has $2^k + 1$ states, \mathcal{P}'_k has only $k + 1$ states.

Leaderless protocols A protocol $\mathcal{P} = (Q, T, L, X, I, O)$ has a multiset L of leaders. If $L = 0$, then the protocol is leaderless. Protocols with leaders and leaderless protocols compute the same predicates [8]. For $L = 0$ we have

$$IC(\lambda v + \lambda' v') = \lambda IC(v) + \lambda' IC(v')$$

for all inputs $v, v' \in \mathbb{N}^X$ and $\lambda, \lambda' \in \mathbb{N}$. In other words, any linear combination of initial configurations with natural coefficients is also an initial configuration.

2.3 State complexity of population protocols

Informally, the state complexity of a predicate is the minimal number of states of the protocols that compute it. We would like to define the state complexity function as the function that assigns to a number ℓ the maximum state complexity of the predicates of size at most ℓ . However, defining the size of a predicate requires to fix a representation. Population protocols compute exactly the predicates expressible in Presburger arithmetic [8], and so there are at least three natural representations: formulas of Presburger arithmetic, existential formulas of Presburger arithmetic, and semilinear sets [20]. Since the translations between these representations involve superexponential blow-ups, we focus on threshold predicates of the form $x \geq \eta$, for which the size of the predicate is the size of η , independently of the representation. We choose to encode numbers in unary, and so we define $STATE(\eta)$ as the number of states of the smallest protocol computing $x \geq \eta$.

The inverse of $STATE(\eta)$ is the function that assigns to a number n the largest η such that a protocol with n states computes $x \geq \eta$. Recall that the busy beaver function assigns to a number n the largest η such that a Turing machine with n states started on a blank tape writes η consecutive ones on the

tape and terminates. Due to this analogy, we call the inverse of the state complexity function the *busy beaver function*, and call protocols computing predicates of the form $x \geq \eta$ busy beaver protocols, or just *busy beavers*.

Definition 1 The *busy beaver function* $BB: \mathbb{N} \rightarrow \mathbb{N}$ is defined as follows: $BB(n)$ is the largest $\eta \in \mathbb{N}$ such that the predicate $x \geq \eta$ is computed by some leaderless protocol with at most n states. The function $BB_L(n)$ is defined analogously, but for general protocols, possibly with leaders.

In [12] Blondin et al. give lower bounds on the busy beaver function:

Theorem 2.2 ([12]) *For every number of states n : $BB(n) \in \Omega(2^n)$ and $BB_L(n) \in \Omega(2^{2^n})$.*

However, to the best of our knowledge no upper bounds have been given.

3 Mathematical structure of stable sets

We define the *stable configurations* of a protocol:

Definition 2 Let $b \in \{0, 1\}$. A configuration C of a protocol is *b-stable* if $O(C') = b$ for every configuration C' reachable from C . The set of *b-stable configurations* is denoted SC_b , and we let $SC = SC_0 \cup SC_1$.

It follows easily from the definitions that a population protocol *computes* a predicate $\varphi: \mathbb{N}^X \rightarrow \{0, 1\}$ iff for every input v and configuration C with $IC(v) \xrightarrow{*} C$ the condition $C \xrightarrow{*} SC_{\varphi(v)}$ holds.

Moreover, given a protocol computing φ , $\varphi(v) = b$ for an input v iff $IC(v) \xrightarrow{*} SC_b$.

A set \mathcal{C} of configurations is *downward closed* if $C \in \mathcal{C}$ and $C' \leq C$ implies $C' \in \mathcal{C}$. The sets SC_0, SC_1 , and SC are downward closed:

Lemma 3.1 *Let \mathcal{P} be a protocol with n states. For every $b \in \{0, 1\}$ the set SC_b is downward closed.*

Proof Assume $C \in SC_b$ and $C' \leq C$. We prove $C' \in SC_b$ by contradiction, so assume that $C' \xrightarrow{*} C''$ for some C'' such that $O(C'') \neq b$. By monotonicity, $C = C' + (C - C') \xrightarrow{*} C'' + (C - C')$, and since $O(C'') \neq b$ we have $O(C'' + (C - C')) \neq b$. So $C' \in SC_b$. \square

Given a downward closed set \mathcal{C} , a pair (B, S) , where B is a configuration and $S \subseteq Q$, is a *basis element* of \mathcal{C} if $B + \mathbb{N}^S \subseteq \mathcal{C}$. A *base* of \mathcal{C} is a finite set \mathcal{B} of *basis elements* such that $\mathcal{C} = \bigcup_{(B,S) \in \mathcal{B}} (B + \mathbb{N}^S)$. We define the *norm* of a basis element (B, S) as $\|(B, S)\|_\infty := \|B\|_\infty$, and the norm of a basis as the maximal norm of its elements.

It is well-known that every downward-closed set of configurations has a base. We prove a stronger result: the sets SC_0, SC_1 , and SC have bases of small norm.

Lemma 3.2 *Let \mathcal{P} be a protocol with n states. Every $\mathcal{C} \in \{SC_0, SC_1, SC\}$ has a basis of norm at most $2^{2(2n+1)!+1}$ with at most $\vartheta(n) := 2^{(2n+2)!}$ elements.*

Proof For the bound on the norm, let $\beta := 2^{2(2n+1)!}$ and fix a b -stable configuration C . Let $S := \{q \in Q \mid C(q) > 2\beta\}$, and define $B \leq C$ as follows: $B(i) := C(i)$ for $i \notin S$ and $B(i) := 2\beta$ for $i \in S$. Since $B \leq C$ and C is b -stable, so is B . We show that (B, S) is a basis element of SC_b , which proves the result for SC_0 and SC_1 . Assume the contrary. Then some configuration $C' \in B + \mathbb{N}^S$ is not b -stable. So $C' \xrightarrow{*} C''$ for some C'' satisfying $C''(q) \geq 1$ for some state $q \in Q$ with $O(q) \neq b$; we say that C'' covers q .

By Rackoff's Theorem [25], C'' can be chosen so that $C' \xrightarrow{\sigma} C''$ for a sequence σ of length $2^{2^{O(n)}}$; a more precise bound is $|\sigma| \leq \beta$ (see Theorem 3.12.11 in [17]). Since a transition moves at most two agents out of a given state, σ moves at most 2β agents out of a state. So, by the definition of B , the sequence σ is also executable from B , and also leads to a configuration that covers q . But this contradicts that B is b -stable. This concludes the proof for SC_0 and SC_1 . For SC , just observe that the union of the bases of SC_0 and SC_1 is a basis of SC .

To prove the bound on the number of elements of the bases, observe that the number of pairs (B, S) such that B has norm at most k and $S \subseteq Q$ is at most $(k + 2)^n$. Indeed, for each state q there are at most $k + 2$ possibilities: $q \in S$, or $q \notin S$ and $0 \leq B(q) \leq k$. So $\vartheta \leq (2^{2(2n+1)!+1} + 2)^n \leq 2^{(2n+2)!}$. \square

From now on we use the following terminology:

Definition 3 We call $\beta := 2^{2(2n+1)!+1}$ the *small basis constant* for the protocol \mathcal{P} . A *small basis* of SC_b or SC is a basis of norm at most β (guaranteed to exist by Lemma 3.2). Its elements are called *small basis elements*.

4 A general upper bound on the busy beaver function

We obtain a bound on the busy-beaver function $BB_L(n)$.

Fix a protocol $\mathcal{P}_n = (Q, T, L, \{x\}, I, O)$ with n states computing a predicate $x \geq \eta$. Observe that the unique input state is x , and so $IC(a) = a \cdot x + L$ for every input a .

Observe that for every input i we have $IC(i) \xrightarrow{*} C_i$ for some configuration $C_i \in SC$, and so $C_i \in B_i + \mathbb{N}^{S_i}$ for some basis element (B_i, S_i) of SC . If $i < \eta$, then $C_i \in SC_0$, and if $i \geq \eta$ then $C_i \in SC_1$. Lemma 4.1 below uses this observation to provide a sufficient condition for an input a to lie above η . The rest of the section shows that for a protocol with n states some number $a < f(n)$ satisfies the condition, where $f(n)$ is a function from the Fast Growing Hierarchy [18]. While the function $f(n)$ grows very fast, it is a recursive function. So, contrary to Turing machines, the busy-beaver function for

population protocols does not grow faster than any recursive function.

Lemma 4.1 *If there exist $a, b \in \mathbb{N}$, a basis element (B, S) of SC , and configurations $D_a, D_b \in \mathbb{N}^S$ satisfying*

1. $IC(a) \xrightarrow{*} B + D_a$, and
2. $b \cdot x \xrightarrow{*} D_b$,

then $\eta \leq a$.

Proof We first claim that $IC(a + \lambda b) \xrightarrow{*} B + D_a + \lambda D_b$ holds for every $\lambda \geq 0$. Observe that

$$IC(a + \lambda b) = (a + \lambda b) \cdot x + L = IC(a) + \lambda b \cdot x.$$

We have:

$$\begin{aligned}
 IC(a) + \lambda b \cdot x &\xrightarrow{*} B + D_a + \lambda b \cdot x && \text{by (1)} \\
 &\xrightarrow{*} B + D_a + \lambda D_b && \text{by (2)}
 \end{aligned}$$

and the claim is proved.

Assume now that $\eta > a$, i.e. \mathcal{P}_n rejects a . Since $D_a \in \mathbb{N}^S$, we have $B + D_a \in SC$, and so $B + D_a \in SC_0$ because \mathcal{P}_n rejects a . Since $D_b \in \mathbb{N}^S$, we have $B + D_a + \lambda D_b \in B + \mathbb{N}^S$, and so $B + D_a + \lambda D_b \in SC_0$ for every $\lambda \geq 0$. So \mathcal{P}_n rejects $a + \lambda b$ for every $\lambda \geq 0$, contradicting that \mathcal{P} computes $x \geq \eta$. \square

We now start our search for a number a satisfying the conditions of Lemma 4.1. First we identify a sequence of configurations $C_2, C_3, C_4 \dots$ of SC satisfying conditions close to 1. and 2. in Lemma 4.1.

Lemma 4.2 *There exists a sequence $C_2, C_3, C_4 \dots$ of configurations of SC satisfying:*

1. $IC(i) \xrightarrow{*} C_i$ for every $i \geq 2$, and
2. $C_i + j \cdot x \xrightarrow{*} C_{i+j}$ for every $j \geq 0$.

Proof Since \mathcal{P}_n computes $x \geq \eta$, for every $i \geq 2$ every fair run of \mathcal{P} starting at $IC(i)$ eventually reaches SC_0 or SC_1 , depending on whether $i < \eta$ or $i \geq \eta$, and stays there forever. We define C_2, C_3, C_4, \dots as follows. First, we let C_2 be any configuration of SC reachable from $IC(2)$. Then, for every $i \geq 2$, assume that C_i has already been defined and satisfies $IC(i) \xrightarrow{*} C_i$. Observe that $IC(i + 1) = IC(i) + x$. Since $IC(i) \xrightarrow{*} C_i$, we also have $IC(i + 1) = IC(i) + x \xrightarrow{*} C_i + x$. This execution can be extended to a fair run, which eventually reaches SC . We let C_{i+1} be any configuration of SC reachable from $C_i + x$.

Let us show that $C_2, C_3, C_4 \dots$ satisfies 1. and 2.. Property 1. holds for C_2 by definition, and for $i \geq 2$ because $IC(i + 1) = IC(i) + x \xrightarrow{*} C_i + I(x) \xrightarrow{*} C_{i+1}$. For property

2., by monotonicity and the definition of C_i we have for every $2 \leq i \leq k$:

$$\begin{aligned}
 C_i + j \cdot x &\xrightarrow{*} C_{i+1} + (j - 1) \cdot x \\
 &\xrightarrow{*} \dots \xrightarrow{*} C_{i+j-1} + x \xrightarrow{*} C_{i+j}
 \end{aligned}$$

\square

We can now easily prove the existence of a number a satisfying the conditions of Lemma 4.1. We start by recalling Dickson’s Lemma:

Lemma 4.3 (Dickson’s lemma) *For every infinite sequence v_1, v_2, \dots of vectors of the same dimension there is an infinite sequence $i_1 < i_2 < \dots$ of indices such that $v_{i_1} \leq v_{i_2} \leq \dots$*

By Dickson’s lemma, the sequence $C_2, C_3, C_4 \dots$ of configurations of SC constructed in Lemma 4.2 contains an ordered subsequence $C_{i_1} \leq C_{i_2} \leq C_{i_3} \dots$. Since SC has a finite basis, by the pigeonhole principle there exist numbers $k < \ell$ and a basis element (B, S) such that $C_{i_k}, C_{i_\ell} \in B + \mathbb{N}^S$. Since $C_k \leq C_\ell$, we have $C_k - C_\ell \in \mathbb{N}^S$, and so we can take:

$$a := k; b := \ell - k; D_a := C_k - B; D_b := C_\ell - C_k.$$

However, the proof of Dickson’s lemma is non-constructive, and gives no bound on the size of a . To solve this problem we observe that, in the terminology of [18], the sequence $C_2 C_3 \dots$ is *linearly controlled*: there is a linear control function $f: \mathbb{N} \rightarrow \mathbb{N}$ satisfying $|C_i| \leq f(i)$. Indeed, since $IC(i) \xrightarrow{*} C_i$, we have $|C_i| = |IC(i)| = |L| + i$, and so we can take $f(n) = |L| + n$. This allows us to use a result on linearly controlled sequences from [18]. Say a finite sequence v_0, v_1, \dots, v_s of vectors of the same dimension is *good* if there are two indices $0 \leq i_1 < i_2 \leq s$ such that $v_{i_1} \leq v_{i_2}$. The maximal length of good linearly controlled sequences has been studied in [9, 18, 23]. In particular, this lemma follows easily from results of [18]:

Lemma 4.4 [18] *For every $\delta \in \mathbb{N}$ and for every elementary function $g: \mathbb{N} \rightarrow \mathbb{N}$, there exists a function $F_{\delta,g}: \mathbb{N} \rightarrow \mathbb{N}$ at level \mathcal{F}_ω of the Fast Growing Hierarchy satisfying the following property: For every infinite sequence $v_0, v_1, v_2 \dots$ of vectors of \mathbb{N}^n satisfying $|v_i| \leq i + \delta$, there exist $i_0 < i_1 < \dots < i_{g(n)} \leq F_{\delta,g}(n)$ such that $v_{i_0} \leq v_{i_1} \leq \dots \leq v_{i_{g(n)}}$.*

We do not need the exact definition of the Fast Growing Hierarchy (see [18]); for our purposes it suffices to know that the level \mathcal{F}_ω contains functions that, crudely speaking, grow like the Ackermann function. From this lemma we obtain:

Theorem 4.5 *Let \mathcal{P}_n be a population protocol with n states and ℓ leaders computing a predicate $x \geq \eta$ for some $\eta \geq 2$. Then $\eta < F_{\ell,\vartheta}(n)$, where $\vartheta(n)$ is the function of Lemma 3.2.*

Proof By Lemma 4.4 there exist $\vartheta(n) + 1$ indices $i_0 < i_1 < \dots < i_{\vartheta(n)} \leq F_{\ell, \vartheta}(n)$ such that $C_{i_0} \leq C_{i_1} \leq \dots \leq C_{i_{\vartheta(n)}}$. By the definition of ϑ and the pigeonhole principle, there are indices $k < \ell$ and a basis element (B, S) of SC_0 such that $C_{i_k}, C_{i_\ell} \in B + \mathbb{N}^S$ and $C_{i_k} \leq C_{i_\ell}$. By Lemma 4.1, $\eta \leq i_k \leq F_{\ell, \vartheta}(n)$. \square

4.1 Is the bound optimal?

The function $F_{\ell, \vartheta}(n)$ grows so fast that one can doubt that the bound is even remotely close to optimal. However, recent results show that this would be less strange than it seems. If a protocol \mathcal{P} computes a predicate $x \geq \eta$, then η is the smallest number such that $IC(\eta) \xrightarrow{*} SC_1$. Therefore, letting $\mathbf{BBP}(n)$ denote the busy beaver protocols with at most n states, and letting $SC_1^{\mathcal{P}}$ and $IC^{\mathcal{P}}$ denote the set SC_1 and the initial mapping of the protocol \mathcal{P} , we obtain:

$$BB_L(n) = \max_{\mathcal{P} \in \mathbf{BBP}(n)} \min\{i \in \mathbb{N} \mid \exists C \in SC_1^{\mathcal{P}} : IC^{\mathcal{P}}(i) \xrightarrow{*} C\}$$

Consider now a deceptively similar function. Let All_1 be the set of configurations C such that $O(C) = 1$, i.e. all agents are in states with output 1. Further, let $\mathbf{PP}(n)$ denote the set of all protocols with alphabet $X = \{x\}$, possibly with leaders, and n states. Notice that we include also the protocols that do not compute any predicate. Define

$$f(n) = \max_{\mathcal{P} \in \mathbf{PP}(n)} \min\{i \in \mathbb{N} \mid \exists C \in All_1^{\mathcal{P}} : IC^{\mathcal{P}}(i) \xrightarrow{*} C\}$$

Using recent results in Petri nets and Vector Addition Systems [14, 15, 21, 22] it is easy to prove that $f(n)$ grows faster than any primitive recursive function.² However, a recent result [10] by Balasubramanian et al. shows $f(n) \in 2^{O(n)}$ for leaderless protocols.

These results suggest that a non-elementary bound on $BB_L(n)$ might well be optimal. However, in the rest of the paper we prove that this can only hold for population protocols with leaders. We show $BB(n) \in 2^{2^{O(n)}}$, i.e. leaderless busy beavers with n states can only compute predicates $x \geq \eta$ for numbers η at most double exponential in n .

² The paper [21] considers protocols with one leader, and studies the problem of moving from a configuration with the leader in a state q_{in} and all other agents in another state r_{in} , to a configuration with the leader in a state q_f and all other agents in state r_f . Combined with [14, 15, 22], this shows that the smallest number of agents for which this is possible grows faster than any elementary function in the number of states of the protocol.

5 An upper bound for leaderless protocols

Fix a *leaderless* protocol $\mathcal{P}_n = (Q, T, \emptyset, \{x\}, I, O)$ with $|Q| = n$ states computing a predicate $x \geq \eta$. Observe that the unique input state is x , and so $IC(a) = a \cdot x$ for every input a . We prove that $\eta \leq 2^{(2n+2)!} \in 2^{2^{O(n)}}$. We first introduce some well-known notions from the theory of Petri nets and Vector Addition Systems.

5.1 Potentially realisable multisets of transitions

The *displacement* of a transition $t = p, q \mapsto p', q'$ is the vector $\Delta_t \in \{-2, -1, 0, 1, 2\}^Q$ given by $\Delta_t := p' + q' - p - q$. Intuitively, $\Delta_t(q)$ is the change in the number of agents populating q caused by the execution of t . For example, if $Q = \{p, q, r\}$ and $t = p, q \mapsto p, r$ we have $\Delta_t(p) = 0$, $\Delta_t(q) = -1$, and $\Delta_t(r) = 1$. The *displacement* of a multiset $\pi \in \mathbb{N}^T$ is defined as $\Delta_\pi := \sum_{t \in T} \pi(t) \cdot \Delta_t$. We use the following notation:

$C \xrightarrow{\pi} C'$ denotes that $C' = C + \Delta_\pi$.

Intuitively, $C \xrightarrow{\pi} C'$ states that if C enables some sequence $t_1 t_2 \dots t_k \in T^*$ such that $\{t_1, \dots, t_k\} = \pi$, then the execution of σ leads to C' . However, such a sequence may not exist. We call the multiset $\{t_1, \dots, t_k\}$ the *Parikh mapping* of $t_1 t_2 \dots t_k$.

Say a configuration C is *j-saturated* if $C(q) \geq j$ for every $q \in Q$, i.e. if it populates all states with at least j agents. We have the following relations between $\xrightarrow{\sigma}$ and $\xrightarrow{\pi}$:

- Lemma 5.1** (i) If $C \xrightarrow{\sigma} C'$ then $C \xrightarrow{\pi} C'$, where π is the *Parikh mapping* of σ
- (ii) If $C \xrightarrow{\pi} C'$ and C is $2|\pi|$ -saturated, then $C \xrightarrow{\sigma} C'$ for any σ with *Parikh mapping* π .

Proof (i): Easy induction on $|\sigma|$.

(ii): By induction on $|\pi|$. The basis case $\pi = \emptyset$ is trivial. Otherwise, let σ be any sequence with *Parikh mapping* π . Since this sequence is non empty, it can be decomposed as $t\sigma'$ where t is a transition and σ' is a sequence with *Parikh mapping* π' defined by $\pi'(t) = \pi(t) - 1$ and $\pi'(t') = \pi(t')$ for every $t' \neq t$. As C is $2|\pi|$ -saturated, we have $C \xrightarrow{t} C''$ for some configuration C'' . Further, C'' is $2|\pi'|$ -saturated and $C'' \xrightarrow{\pi'} C'$. By induction hypothesis $C'' \xrightarrow{\sigma'} C'$. It follows that $C \xrightarrow{\sigma} C'$, and we are done. \square

We introduce the set of potentially realisable multisets of transitions of a protocol:

Definition 4 A multiset π of transitions is *potentially realisable* if there are $i \in \mathbb{N}$ and $C \in \mathbb{N}^Q$ such that $IC(i) \xrightarrow{\pi} C$.

The reason for the name is as follows. If π is not potentially realisable, then by Lemma 5.1(i) no sequence $\sigma \in T^*$ with

Parikh mapping π can be executed from any initial configuration, and so π cannot be “realised”. In other words, potential realisability is a necessary but not sufficient condition for the existence of an execution that “realises” π .

5.2 Structure of the proof

We can now give a high-level view of the proof of the bound $\eta \leq 2^{(2n+2)!}$. The starting point is a version of Lemma 4.1 in which, crucially, the condition $IC(b) \xrightarrow{*} D_b$ is replaced by the weaker $IC(b) \xrightarrow{*} D_b$.

Lemma 5.2 *If there exist $a, b \in \mathbb{N}$, a basis element (B, S) of SC , configurations $D_a, D_b \in \mathbb{N}^S$, and a configuration D satisfying*

- (i) $IC(a) \xrightarrow{*} D \xrightarrow{*} B + D_a$, and
- (ii) $IC(b) \xrightarrow{\pi} D_b$ for some $\pi \in \mathbb{N}^T$ such that D is $2|\pi|$ -saturated,

then $\eta \leq a$.

Proof We first claim that

$$IC(a + \lambda b) \xrightarrow{*} B + D_a + \lambda D_b$$

holds for every $\lambda \geq 0$. To prove this, observe first that $IC(b) \xrightarrow{\pi} D_b$ implies $D + IC(b) \xrightarrow{\pi} D + D_b$. Since D is $2|\pi|$ -saturated so is $D + IC(b)$, and, by Lemma 5.1(ii), we have $D + IC(b) \xrightarrow{\sigma} D + D_b$ where σ is any sequence with Parikh mapping π . With an induction on λ , we immediately derive

$$D + \lambda IC(b) \xrightarrow{\sigma^\lambda} D + \lambda D_b. \tag{*}$$

Since \mathcal{P}_n is leaderless, $IC(a + \lambda b) = IC(a) + \lambda IC(b)$ holds, and:

$$\begin{aligned} IC(a) + \lambda IC(b) &\xrightarrow{*} D + \lambda IC(b) && \text{by (i)} \\ &\xrightarrow{*} D + \lambda D_b && \text{by (*)} \\ &\xrightarrow{*} B + D_a + \lambda D_b && \text{by (i)}. \end{aligned}$$

This proves the claim.

Assume now that $\eta > a$, i.e. \mathcal{P}_n rejects a . Since $D_a, D_b \in \mathbb{N}^S$, we have $B + D_a + \lambda D_b \in SC$, and so $B + D_a + \lambda D_b \in SC_0$ for every $\lambda \geq 0$. So, by the claim, \mathcal{P}_n rejects $a + \lambda b$ for every $\lambda \geq 0$, contradicting that it computes $x \geq \eta$. \square

In the next sections we show that $a := 2^{(2n+2)!}$ satisfies the conditions of Lemma 5.2. We proceed in three steps:

- (a) Section 5.3 proves that for every $j \in \mathbb{N}$ and input $a \geq j3^n$ the initial configuration $IC(a)$ can reach a j -saturated configuration D . We remark that this result is only true for leaderless protocols.

- (b) Let $S \subseteq Q$. Section 5.4 proves that $IC(a) \xrightarrow{*} B + D_a$, for $a \in \mathbb{N}$, $D_a \in \mathbb{N}^S$ and a configuration B , implies $IC(b) \xrightarrow{*} D_b$ for some $b \in \mathbb{N}$, $D_b \in \mathbb{N}^S$, provided that a is “large” relative to $|B|$.
- (c) Section 5.5 puts everything together, and gives the final bound.

5.3 Reaching j -saturated configurations

Recall our assumption that for every state $q \in Q$ there exists an input i_q such that $IC(i_q) \xrightarrow{*} C_q$ for some configuration C_q such that $C_q(q) > 0$. By monotonicity, we have $IC(i) \xrightarrow{*} C$ for the input $i := \sum_{q \in Q} i_q$ and the 1-saturated configuration $C := \sum_{q \in Q} C_q$. We show that we can choose $i < 3^n$ and that $IC(i) \xrightarrow{\sigma} C$ for some σ such that $|\sigma| \leq 3^n$. It follows that for every $j \in \mathbb{N}$ the input $j3^n$ can reach a j -saturated configuration by executing j times σ .

Lemma 5.3 *Let C be a configuration satisfying $x \in \llbracket C \rrbracket \subseteq Q$. There exists a transition $p, q \mapsto p', q'$ such that $\{p, q\} \subseteq \llbracket C \rrbracket$ and $\{p', q'\} \not\subseteq \llbracket C \rrbracket$.*

Proof Since $\llbracket C \rrbracket$ is strictly included in Q , there exist $i \in \mathbb{N}$, a word $\sigma \in T^*$, a configuration C' such that $IC(i) \xrightarrow{\sigma} C'$ and $\llbracket C' \rrbracket \not\subseteq \llbracket C \rrbracket$. Assume w.l.o.g. that σ has minimal length. If $\sigma = \epsilon$ then $\llbracket C' \rrbracket \subseteq \{x\}$ contradicting $\llbracket C' \rrbracket \not\subseteq \llbracket C \rrbracket$. So $\sigma = \sigma' t$ for some transition $t = p, q \mapsto p', q'$, and $IC(i) \xrightarrow{\sigma'} C'' \xrightarrow{t} C'$ for some configuration C'' . From $C'' \xrightarrow{t} C'$ we derive $\{p, q\} \subseteq \llbracket C'' \rrbracket$ and $\llbracket C' \rrbracket \subseteq \llbracket C'' \rrbracket \cup \{p', q'\}$. By minimality of σ and $|\sigma'| < |\sigma|$, we deduce that $\llbracket C'' \rrbracket \subseteq \llbracket C \rrbracket$. In particular $\{p, q\} \subseteq \llbracket C \rrbracket$, and $\llbracket C' \rrbracket \subseteq \llbracket C \rrbracket \cup \{p', q'\}$. As $\llbracket C' \rrbracket \not\subseteq \llbracket C \rrbracket$, we deduce that $\{p', q'\} \not\subseteq \llbracket C \rrbracket$. \square

Lemma 5.4 *There exists a word $\sigma \in T^*$, a 1-saturated configuration C and a sequence σ of length at most 3^n such that $IC(3^n) \xrightarrow{\sigma} C$.*

Proof We build a sequence $\sigma_0, \sigma_1 \dots$ of words in T^* and a non decreasing sequence $C_0, C_1 \dots$ of configurations such that for every $j \geq 0$ we have (with the convention $C_{-1} = 0$)

- $IC(3^j) \xrightarrow{\sigma_j} C_j$;
- $\llbracket C_{j-1} \rrbracket$ is saturated, or $\llbracket C_{j-1} \rrbracket \subset \llbracket C_j \rrbracket$, and
- $|\sigma_j| = (3^j - 1)/2$.

Since \mathcal{P}_n has n states, some C_j with $0 \leq j \leq n$ is saturated.

The sequence is built inductively. Choose $C_0 = IC(1)$ and $\sigma_0 = \epsilon$. Assume that $\sigma_0, \dots, \sigma_k$ has been built. If C_k is saturated, then choose $C_{k+1} := 3C_k$ and $\sigma_{k+1} := \sigma_k^3$. Assume C_k is not saturated. Since $x \in \llbracket C_1 \rrbracket$ we have $x \in \llbracket C_k \rrbracket \subset Q$. By Lemma 5.3 there exists a transition $t = p, q \mapsto p', q'$ such that $p, q \in \llbracket C_k \rrbracket$ and $\{p', q'\} \not\subseteq \llbracket C_k \rrbracket$. Since $IC(3^k) \xrightarrow{\sigma_k} C_k$ and \mathcal{P}_n is leaderless, we have $IC(3^{k+1}) \xrightarrow{\sigma_k^3} 3C_k$. Since $p, q \in \llbracket C_k \rrbracket$, transition t is enabled at $2C_k$. We choose

$C_{k+1} := C_k + C'_k$ where C'_k is the configuration satisfying $2C_k \xrightarrow{t} C'_k$. Now, just set $\sigma_{k+1} := \sigma_k^3 t$. \square

5.4 Reaching ϵ -concentrated stable configurations

This section contains the core of the proof. We want to find some $b \in \mathbb{N}$ and a configuration $D_b \in \mathbb{N}^S$ with $IC(b) \xrightarrow{*} D_b$, where $S \subseteq Q$ is given by a basis element (B, S) . We start by noting that any sequence $IC(a) \xrightarrow{*} B + D_a$, again with $a \in \mathbb{N}$ and $D_a \in \mathbb{N}^S$, already fulfils this condition *approximately* if a is large and (B, S) is a small basis element, i.e. B is “small” relative to D_a . The next lemma formalises this notion.

Definition 5 Let $0 < \epsilon \leq 1$ and $S \subseteq Q$. A configuration C is ϵ -concentrated in S if $C(S) \geq (1 - \epsilon)|C|$ or, equivalently, $C(Q \setminus S) \leq \epsilon|C|$.

Lemma 5.5 Let $k \geq 1$, let $a := kn\beta$, where β is the constant of Definition 3, and let D denote a configuration with $IC(a) \xrightarrow{*} D$. There exists a small basis element (B, S) of SC , and a $D_a \in \mathbb{N}^S$ s.t.

1. $IC(a) \xrightarrow{*} D \xrightarrow{*} B + D_a$, and
2. $B + D_a$ is $\frac{1}{k}$ -concentrated in S .

Proof As D is reachable from an input configuration, it can reach a configuration $B + D_a$, where $D_a \in \mathbb{N}^S$ and (B, S) is a small basis element of SC . The latter implies $\|B\|_\infty \leq \beta$, so we have $|B| \leq n\beta = a/k = |B + D_a|/k$, and $B + D_a$ is $\frac{1}{k}$ -concentrated in S . \square

In the remainder of this section we use another small basis theorem, Pottier’s small basis theorem for Diophantine equations, to extend the above result: not only can we reach ϵ -concentrated configurations for arbitrarily small $\epsilon > 0$, we can also *potentially* reach a 0-concentrated configuration, i.e. a configuration in which *all* agents populate S .

For this, we observe that the potentially realisable multisets π (see Definition 4) are precisely the solutions of the system of $|Q| - 1$ Diophantine equations over the variables $\{\pi(t)\}_{t \in T}$ given by

$$\sum_{t \in T} \pi(t) \Delta_t(q) \geq 0 \quad \text{for } q \in Q \setminus \{x\}.$$

Let $A \cdot y \geq 0$ be a homogeneous system of linear Diophantine inequalities, i.e. a solution is a vector m over the natural numbers such that $A \cdot m \geq 0$. A set \mathcal{B} of solutions is a basis if every solution is the sum of a multiset of solutions of \mathcal{B} . Formally, \mathcal{B} is a *basis* if for every solution m , there exists a multiset $M \in \mathbb{N}^{\mathcal{B}}$ such that $m = \sum_{b \in \mathcal{B}} M(b) \cdot b$. It is easy to see that every system has a finite basis. Pottier’s theorem shows that it has a small basis:

Theorem 5.6 ([24]) Let $A \cdot y \geq 0$ be a system of e linear Diophantine equations on v variables. There exists a basis $\mathcal{B} \subseteq \mathbb{N}^v$ of solutions such that for every $m \in \mathcal{B}$:

$$\|m\|_1 \leq \left(1 + \max_{i=1}^e \sum_{j=1}^v |a_{ij}| \right)^e.$$

Since the potentially realisable multisets are the solutions of a system of Diophantine equations, by applying Pottier’s theorem we obtain:

Corollary 5.7 Let $\xi := 2(2|T| + 1)^{|Q|}$. There exists a basis of potentially realisable multisets such that every element π of the basis satisfies $|\pi| \leq \xi/2$. Moreover, $IC(i) \xrightarrow{\pi} C$ for some $i \leq \xi$ and some C such that $C(Q) \leq \xi$ and $C(x) = 0$.

Proof The potentially realisable multisets are precisely the solutions of a system of $|Q| - 1$ Diophantine equations on $|T|$ variables. The matrix A of the system satisfies $|a_{ij}| = |\Delta_j(q_i)|$, and so $|a_{ij}| \leq 2$ for every $1 \leq i < |Q|$ and every $1 \leq j \leq |T|$. The result follows from Theorem 5.6.

For $C(x) = 0$, note that $IC(i) \xrightarrow{\pi} C$ implies $IC(i - C(x)) \xrightarrow{\pi} C - x \cdot C(x)$. \square

Recall that Definition 3 introduced the small basis constant β . Analogously, we now introduce the Pottier constant:

Definition 6 We call $\xi := 2(2|T| + 1)^{|Q|}$ the *Pottier constant* for the protocol \mathcal{P} .

Remark 1 For a deterministic population protocol (i.e. a protocol such that for every pair of states there is at most one transition) we could instead use $\xi = 2(|Q| + 2)^{|Q|}$.

The following lemma is at the core of our result. Intuitively, it states that if some potentially realisable multiset leads to a $(1/\xi)$ -concentrated configuration in S , then some *small* potentially realisable multiset leads to a 0-concentrated configuration in S .

Lemma 5.8 Let $IC(i) \xrightarrow{*} C$ and $S \subseteq Q$. If C is $(1/\xi)$ -concentrated in S , then $IC(j) \xrightarrow{\theta} C'$ for some number j , potentially realisable multiset $\theta \in \mathbb{N}^T$, and configuration C' such that $|\theta| \leq \xi/2$, $j \leq \xi$ and C' is 0-concentrated on S , i.e. $C' \in \mathbb{N}^S$.

Proof If $x \notin S$ then we take θ as the empty multiset and are done. Otherwise let $l := C(x)$ and set $i^* := i - l$ and $C^* := C - l \cdot x$. Clearly, $C^*(S) < i^*/\xi$ follows from $C(S) < i/\xi$. It therefore suffices to prove the lemma in the case of $C^* = C$ and $i^* = i$, i.e. $l = 0$.

Since $IC(i) \xrightarrow{*} C$, we have $IC(i) \xrightarrow{\pi} C$ for some potentially realisable multiset π . By Corollary 5.7 there exist potentially realisable multisets $\pi_1, \dots, \pi_k \in \mathbb{N}^T$ s.t. $\pi = \pi_1 + \dots + \pi_k$, numbers $i_1, \dots, i_k \in \mathbb{N}$ and configurations C_1, \dots, C_k (not necessarily distinct) such that

- $i = \sum_{j=1}^k i_j, \pi = \sum_{j=1}^k \pi_j$, and $C = \sum_{j=1}^k C_j$.
- $IC(i_j) \xrightarrow{\pi_j} C_j$.
- $i_j \leq \xi$ and $C_j(Q) \leq \xi$ for $j \in \{1, \dots, k\}$.

Let $J = \{j \in \{1, \dots, k\} \mid i_j > 0\}$. We have

$$i = \sum_{j=1}^k i_j = \sum_{j \in J} i_j \leq \xi \cdot |J|.$$

Since $i > \xi \cdot C(S)$ by hypothesis, we deduce that $|J| > C(S)$. Assume that $C_j(S) > 0$ for every $j \in J$. Then $C(S) = \sum_{j \in J} C_j(S) \geq |J|$, a contradiction. So $C_j(S) = 0$ for some $j \in J$, and we take $\theta := \pi_j$. \square

5.5 The upper bound

We have now obtained all the results needed for our final theorem.

Theorem 5.9 *Let \mathcal{P}_n be a leaderless population protocol with n states, and let β and ξ be the constants of Definition 3 and Definition 6. If \mathcal{P}_n computes a predicate $x \geq \eta$, then*

$$\eta \leq \xi n \beta 3^n \leq 2^{(2n+2)!}.$$

Proof Assume \mathcal{P}_n computes $x \geq \eta$. We first prove that $\eta \leq a$ holds for $a := \xi n \beta 3^n$. It suffices to show that a satisfies the conditions of Lemma 5.2 for a suitable choice of $b, (B, S), D_a, D_b, D$, and π .

Lemma 5.4 shows that a 1-saturated configuration C with $IC(3^n) \rightarrow C$ exists. Choose $D := \xi n \beta \cdot C$. Observe that D is $(\xi n \beta)$ -saturated and reachable from $IC(a)$.

Now choose (B, S) and D_a as the configurations and basis element of Lemma 5.5, where $k := \xi 3^n$ and D is chosen as above. With this choice of k Lemma 5.5 yields:

- 1 $IC(a) \xrightarrow{*} D \xrightarrow{*} B + D_a$; and
- 2 $B + D_a$ is $(1/\xi)$ -concentrated in S .

(Regarding 2., note that $(1/\xi 3^n)$ -concentrated implies $(1/\xi)$ -concentrated.) Property 1. coincides with condition (i) of Lemma 5.2. Finally, we choose $b, D_b \in \mathbb{N}^S$, and π so that condition (ii) of Lemma 5.2 holds as well. In particular, b, D_b , and π must satisfy:

$$IC(b) \xrightarrow{\pi} D_b \text{ for some } \pi \text{ s.t. } D_b \text{ is } 2|\pi| \text{-saturated.}$$

By 1. we have $IC(a) \xrightarrow{*} B + D_a$. Applying Lemma 5.8 with $i := a$ and $C := B + D_a$, we conclude that $IC(j) \xrightarrow{\theta} C'$ for some $0 < j \leq \xi$, some multiset θ such that $|\theta| \leq \xi/2$, and some configuration C' 0-concentrated in S , i.e. satisfying $C'(Q \setminus S) = 0$.

Using $b := j, \pi := \theta$, and $D_b := C'$, we have $IC(b) \xrightarrow{\pi} D_b$, which fulfils condition (ii) of Lemma 5.2 as well. Further, since D is $(n\beta\xi)$ -saturated and $n\beta\xi \geq 2(\xi/2) \geq 2|\pi|$, the configuration D is $2|\pi|$ -saturated. Applying the lemma we get $\eta \leq a$, concluding the first part of our proof.

We now show that $\xi n \beta 3^n \leq 2^{(2n+2)!}$. Since $\xi = 2(2|T| + 1)^{|Q|}$ and $\beta = 2^{2^{(2n+1)!+1}}$, we deduce (for $n \geq 2$):

$$\begin{aligned} \eta &\leq 2(2|T| + 1)^n n 2^{2^{(2n+1)!+1}} 3^n \\ &\leq 4(2n^4 + 1)^n n 2^{2^{(2n+1)!}} 2^{2n} \\ &\leq 4 \cdot 3^n n^{4n} n 2^{2^{(2n+1)!}} 2^{2n} \\ &\leq 4 \cdot 3^n n^{4n+1} 2^{2^{(2n+1)!}} 2^{2n} \end{aligned}$$

and so

$$\begin{aligned} 2^n &\leq 2 + 2n + (4n + 1) \log_2 n + 2(2n + 1)! + 2n \\ &\leq 2(2n + 1)! + (2n + 1)(2 + 2 \log_2 n) \\ &\leq (2n + 2)! \end{aligned}$$

Therefore $\eta \leq 2^{(2n+2)!}$. \square

6 Conclusion

We have obtained the first non-trivial lower bounds on the state complexity of population protocols computing counting predicates of the form $x \geq \eta$, a fundamental question about the model. The obvious open problems are to close the gap between the $\Omega(\log \log \eta)$ lower bound and the $\mathcal{O}(\log \eta)$ upper bound for the leaderless case, and the even larger gap between $\mathcal{O}(\log \log \eta)$ and (roughly speaking), the $\Omega(\alpha(\eta))$ lower bound for protocols with leaders, where $\alpha(\eta)$ is the inverse of the Ackermann function.

Acknowledgements This work was supported by an ERC Advanced Grant (787367: PaVeS), by the Research Training Network of the Deutsche Forschungsgemeinschaft (DFG) (378803395: ConVeY), and by the grant ANR-17-CE40-0028 of the French National Research Agency ANR (project BRAVAS). A previous version of this paper appeared in the proceedings of PODC 2021 [13].

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copy-

right holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abriola, S., Figueira, S., Senno, G.: Linearizing well quasi-orders and bounding the length of bad sequences. *Theor. Comput. Sci.* **603**, 3–22 (2015)
2. Alistarh, D., Aspnes, J., Eisenstat, D., Gelashvili, R., Rivest, R.L.: Time-space trade-offs in population protocols. In: *SODA*, pp. 2560–2579. SIAM, Philadelphia (2017)
3. Alistarh, D., Aspnes, J., Gelashvili, R.: Space-optimal majority in population protocols. In: *SODA*, pp. 2221–2239. SIAM, Philadelphia (2018)
4. Alistarh, D., Gelashvili, R.: Recent algorithmic advances in population protocols. *SIGACT News* **49**(3), 63–73 (2018)
5. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. In: *PODC*, pp. 290–299. ACM (2004)
6. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M.J., Peralta, R.: Computation in networks of passively mobile finite-state sensors. *Distrib. Comput.* **18**(4), 235–253 (2006)
7. Angluin, D., Aspnes, J., Eisenstat, D.: Fast computation by population protocols with a leader. *Distrib. Comput.* **21**(3), 183–199 (2008)
8. Angluin, D., Aspnes, J., Eisenstat, D., Ruppert, E.: The computational power of population protocols. *Distrib. Comput.* **20**(4), 279–304 (2007)
9. Balasubramanian, A.R.: Complexity of controlled bad sequences over finite sets of \mathbb{N}^d . In: *LICS*, pp. 130–140. ACM (2020)
10. Balasubramanian, A.R., Esparza, J., Raskin, M.A.: Finding cut-offs in leaderless rendez-vous protocols is easy. In: *FoSSaCS*. Lecture Notes in Computer Science, vol. 12650. Springer, New York (2021)
11. Blondin, M., Esparza, J., Genest, B., Helfrich, M., Jaax, S.: Succinct population protocols for Presburger arithmetic. In: *STACS*, vol. 154, pp. 40:1–40:15. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
12. Blondin, M., Esparza, J., Jaax, S.: Large flocks of small birds: On the minimal size of population protocols. In: *STACS*, vol. 96, pp. 16:1–16:14. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018)
13. Czermer, P., Esparza, J.: Lower bounds on the state complexity of population protocols. In: *PODC*, pp. 45–54. ACM (2021)
14. Czerwinski, W., Lasota, S., Lazic, R., Leroux, J., Mazowiecki, F.: The reachability problem for Petri nets is not elementary. *J. ACM* **68**(1), 7:1–7:28 (2021)
15. Czerwinski, W., Orlikowski, L.: Reachability in vector addition systems is Ackermann-complete. *CoRR*, [arxiv:2104.13866](https://arxiv.org/abs/2104.13866) (2021)
16. Elsässer, R., Radzik, T.: Recent results in population protocols for exact majority and leader election. *Bull. EATCS* **126** (2018)
17. Esparza, J.: Petri nets lecture notes (2019). <https://archive.model.in.tum.de/um/courses/petri/SS2019/PNSkript.pdf>
18. Figueira, D., Figueira, S., Schmitz, S., Schnoebelen, P.: Ackermannian and primitive-recursive bounds with Dickson’s lemma. In: *LICS*, pp. 269–278. IEEE Computer Society (2011)
19. Gąsieniec, L., Stachowiak, G.: Enhanced phase clocks, population protocols, and fast space optimal leader election. *J. ACM* **68**(1), 1–12 (2020)
20. Haase, C.: A survival guide to Presburger arithmetic. *ACM SIGLOG News* **5**(3), 67–82 (2018)
21. Horn, F., Sangnier, A.: Deciding the existence of cut-off in parameterized rendez-vous networks. In: *CONCUR*, vol. 171, pp. 46:1–46:16. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
22. Leroux, J.: The reachability problem for Petri nets is not primitive recursive. *CoRR* [arxiv:2104.12695](https://arxiv.org/abs/2104.12695) (2021)
23. McAloon, K.: Petri nets and large finite sets. *Theor. Comput. Sci.* **32**, 173–183 (1984)
24. Pottier, L.: Minimal solutions of linear diophantine systems: bounds and algorithms. In: *RTA*. Lecture Notes in Computer Science, vol. 488, pp. 162–173. New York, Springer (1991)
25. Rackoff, C.: The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.* **6**, 223–231 (1978)
26. Schmitz, S.: Complexity hierarchies beyond elementary. *ACM Trans. Comput. Theory* **8**(1), 3:1–3:36 (2016)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.