



HAL
open science

State Complexity of Protocols with Leaders

Jérôme Leroux

► **To cite this version:**

Jérôme Leroux. State Complexity of Protocols with Leaders. PODC '22: ACM Symposium on Principles of Distributed Computing, Jul 2022, Salerno, Italy. pp.257-264, 10.1145/3519270.3538421 . hal-03863775

HAL Id: hal-03863775

<https://cnrs.hal.science/hal-03863775v1>

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.

State Complexity of Protocols With Leaders

JÉRÔME LEROUX, LaBRI, CNRS, Univ. Bordeaux, France

Population protocols are a model of computation in which an arbitrary number of anonymous finite-memory agents are interacting in order to decide by stable consensus a predicate. In this paper, we focus on the counting predicates that asks, given an initial configuration, whether the number of agents in some initial state i is at least n . In 2018, Blondin, Esparza, and Jaax shown that with a fix number of leaders and interaction-width, there exists infinitely many n for which the counting predicate is stably computable by a protocol with at most $O(\log \log(n))$ states. We provide in this paper a matching lower-bound (up to a square root) that improves the inverse-Ackermannian lower-bound presented at PODC in 2021.

1 INTRODUCTION

Population protocols were introduced by Angluin, Aspnes, Diamadi, Fischer, and Peralta in [1, 2] to study the computational power of networks of resource-limited mobile agents. In this model, each agent has a state in a finite set of states. When agents interact, their states are updated accordingly to a finite interaction table. This table corresponds intuitively to a conservative Petri net (the Petri net is conservative since the number of agents is preserved by each transition) where each line of the interaction table is matched by a transition of the Petri net. In this model, an agent may accept or reject depending only on its own state. A population protocol is said to be stably computing a predicate, if for any initial configurations, eventually and forever, under some natural fairness conditions, either all agents accept or all agents reject. Moreover, this outcome should only depend on the initial configuration and not on the way interactions are performed.

Deciding if a protocol stably computes some unknown predicate is a problem called the *well-specification problem*. This problem was proved to be decidable in [9, 10] by observing that well-specification problem is equivalent to the reachability problem for Petri nets up to elementary reductions. Since this last problem was recently proved to be Ackermannian-complete [8, 11], it means that deciding the well-specification problem is Ackermannian-complete. Intuitively, population protocols maybe intrinsically very complicated.

Despite this Ackermannian complexity result, in [4], Angluin, Aspnes, Eisenstat, and Ruppert have shown that predicates stably computable by population protocols cannot be more complicated than the one definable in the Presburger arithmetic. Combined with [1, 2], it follows that predicates stably computable by population protocols are exactly the predicates definable in the Presburger arithmetic.

Since deciding if a population protocol is stably computing some Presburger predicate is Ackermannian-complete, a natural question is the conciseness of population protocols. Intuitively, is it possible to define a population protocol computing predicates that are very complex compared to the number of states of the protocol? This problem is related to the so-called *state complexity* of a Presburger predicate intuitively defined as the minimal number of states of a population protocols deciding it.

State complexity upper-bounds are obtained thanks to algorithms computing from predicates protocols stably computing it with a number of states as small as possible. In [5], by revisiting the construction of population protocols deciding Presburger predicates, some improvement on state complexity upper-bounds was derived. On the other side, state complexity lower-bounds is also a difficult task since such a bound requires to prove that there is no way to stably compute a predicate with a given amount of states. In this context, focusing on the state complexity of simple Presburger predicates is a natural question. The simplest non trivial Presburger predicates are clearly the counting predicates that corresponds to the set of configurations such that the number of agents in a given state is larger than

or equal to some positive number n . In 2018, Blondin, Esparza, and Jaax shown in [6] that with a fix number of leaders and interaction-width (the number of agents that can interact at each interaction step), there exists infinitely many n for which the counting predicate is stably computable by a protocol with at most $O(\log \log(n))$ states.

This state complexity upper-bound was recently completed by a state complexity lower-bound in [7]. In that paper, Czerner and Esparza shown that the number of states of a population protocol deciding a counting predicate with a bounded number of leaders and a bounded interaction-width is at least $\Omega(A^{-1}(n))$ where A is some Ackermannian function, leaving a gap between the $O(\log \log(n))$ upper-bound and the $\Omega(A^{-1}(n))$ lower-bound.

Main result. In this paper, we follow the model of protocols introduced by Dana Angluin, James Aspnes, and David Eisenstat in 2006 that allows agents creations and destructions [3]. We slightly extends that model to allow leaders. We close the previously mentioned state complexity gap by proving that for any $h < \frac{1}{2}$, and under a fix number of leaders and a bounded number of interaction-width, any protocol stably computing a counting predicate requires at least $\Omega((\log \log(n))^h)$ states.

Outline. In Section 2 we recall some basic definitions and results about protocols. In Section 3, we introduce the notion of communication-width of a protocol and show that protocol with finite communication-width are naturally related to the model of Petri nets. In Section 4 we introduce the state complexity problem and show that counting the number of states of protocols without taking into account the number of leaders or the communication-width is not relevant. In Section 5 we introduce the notions of stabilized configurations and show that those configurations are characterized by their small values. Those results are obtained thanks to Rackoff's techniques originally introduced in the context of the coverability problem for Petri nets. Section 6 contains the central technical lemma. It is a lemma about Petri nets that intuitively shows that from any initial configuration we can reach with short executions kind of bottom configurations. Section 7 recalls the model of Petri net with control-states and provides a result on small cycles satisfying some properties. This last result is obtained by introducing a linear system and by applying Pottier's techniques [12] in order to obtain small solutions for that linear system. Results of the previous sections are combined in Section 8 to obtain our state complexity lower-bound. Some related open problems and future work are presented in Section 9.

2 PROTOCOLS

In this section, we introduce our model of protocols that slightly differs from the one introduced in [3] by allowing leaders. We extend the definition of stably computing a predicate with such a protocol. This definition is a straightforward extension of the one given in [3] that generalizes several definitions of *fair computations* in particular in the context of unconservative protocols.

Let P be a finite set of elements called *states*. A P -*configuration* (or just a *configuration* if P is clear from the context) is a mapping in \mathbb{N}^P . Given a configuration ρ , the number $|\rho| \stackrel{\text{def}}{=} \sum_{p \in P} \rho(p)$ is called the *number of agents* in ρ . Let Q be another finite set of states. We associate with a P -configuration ρ the Q -configuration $\rho|_Q$ defined for every $q \in Q$ by $\rho|_Q(q) = \rho(q)$ if $q \in P$ and zero otherwise. Notice that Q is not necessarily a subset of P . Given $p \in P$ we simply denote by $p|_P$ (or just p when P is clear from the context) the mapping in \mathbb{N}^P that maps p on 1 and the other states on zero. The sum $\alpha + \beta$ of two configurations α, β and the product $n \cdot \rho$ where $n \in \mathbb{N}$ and ρ is a configuration are defined component-wise as expected.

Let R be a binary relation on P -configurations for some finite set of states P . We say that R is *additive* if $(\alpha, \beta) \in R$ implies $(\alpha + \rho, \beta + \rho) \in R$ for every configuration α, ρ, β . As usual, R is called a *preorder* if it is *reflexive* and *transitive*. We also say that R is *conservative* if $|\alpha| = |\beta|$ for every $(\alpha, \beta) \in R$.

A *protocol* is a tuple $(P, \xrightarrow{*}, \rho_L, I, \gamma)$ where P is a finite set of *states*, $\xrightarrow{*}$ is an additive preorder on the P -configurations, ρ_L is a P -configuration called the *configuration of leaders*, $I \subseteq P$ is the set of *initial states*, and $\gamma : P \rightarrow \{0, \star, 1\}$ is the *output function*. The value $|\rho_L|$ is called the number of *leaders*. A protocol is said to be *leaderless* when this number is zero. A configuration of the form $\rho_L + \rho|_P$ with $\rho \in \mathbb{N}^I$ is called an *initial configuration*. A protocol is said to be *conservative* if $\xrightarrow{*}$ is conservative. The function γ is extended on any $\rho \in \mathbb{N}^P$ by:

$$\gamma(\rho) = \{j \in \{0, \star, 1\} \mid \exists p \in P \rho(p) > 0 \wedge \gamma(p) = j\}$$

For $j \in \{0, 1\}$, we introduce the following sets S_j called the j -output stable configurations. Notice that S_0 and S_1 are not defined exactly the same way in order to manage the zero configuration. With our definition we interpret the output of the zero configuration as 0. Notice that the definition of 0-output stable configurations introduced in [3] does not care about the zero configuration since protocols are semantically restricted to non zero configurations. We do not introduce a set S_\star since intuitively the element \star in the image of γ is interpreted as an undetermined output.

$$S_0 \stackrel{\text{def}}{=} \{\alpha \in \mathbb{N}^P \mid \forall \beta \alpha \xrightarrow{*} \beta \Rightarrow \gamma(\beta) \subseteq \{0\}\}$$

$$S_1 \stackrel{\text{def}}{=} \{\alpha \in \mathbb{N}^P \mid \forall \beta \alpha \xrightarrow{*} \beta \Rightarrow \gamma(\beta) = \{1\}\}$$

A *predicate* is a mapping $\phi : \mathbb{N}^I \rightarrow \{0, 1\}$. We say that a protocol stably computes ϕ if for every $\rho \in \mathbb{N}^I$ and for every $\alpha \in \mathbb{N}^P$ such that $\rho_L + \rho|_P \xrightarrow{*} \alpha$, there exists $\beta \in S_{\phi(\rho)}$ such that $\alpha \xrightarrow{*} \beta$. A predicate ϕ is *stably computable* if there exists a protocol that stably computes it.

REMARK 1. In [3], predicates that are stably computable by leaderless protocols restricted to functions γ such that $\gamma(P) \subseteq \{0, 1\}$ are proved to be exactly the predicates definable in the Presburger arithmetic. We think that this property can be extended to any protocol with leaders and using the \star element in the definition of γ . We left this problem open since it is outside of the scope of this paper.

3 INTERACTION-WIDTH

The class of additive preorders is central for defining protocols as shown in the previous section. Additive preorders are a natural generalization of the Petri net reachability relations used in classical population protocols. For parameterized complexity purposes, we introduce in this section the notion of interaction-width that intuitively limits the number of agents that can communicate together in a single interaction step. As expected additive preorders with finite interaction-width are exactly the reachability relations of Petri nets.

A P -*transition* (or simply a *transition* when the finite set P is clear from the context) is a pair $t \stackrel{\text{def}}{=} (\alpha_t, \beta_t)$ of P -configurations. Given such a transition, we introduce $|t| \stackrel{\text{def}}{=} \max\{|\alpha_t|, |\beta_t|\}$ called the *interaction-width* of t . We associate with a transition t the binary relation \xrightarrow{t} over the configurations defined by $\alpha \xrightarrow{t} \beta$ if there exists a configuration ρ such that $\alpha = \alpha_t + \rho$ and $\beta = \beta_t + \rho$. Notice that \xrightarrow{t} is the minimal for the inclusion additive binary relation that contains t . Given a word $\sigma = t_1 \dots t_k$ of transitions, we introduce the binary relation $\xrightarrow{\sigma}$ over the configurations defined by

$\alpha \xrightarrow{\sigma} \beta$ if there exists a sequence ρ_0, \dots, ρ_k of configurations such that:

$$\alpha = \rho_0 \xrightarrow{t_1} \rho_1 \cdots \xrightarrow{t_k} \rho_k = \beta$$

An additive preorder R is said to have a *finite interaction-width* if there exists $m \in \mathbb{N}$ such that for every $(\alpha, \beta) \in R$, there exists a word σ of transitions in R with an interaction-width bounded by m and such that $\alpha \xrightarrow{\sigma} \beta$. The minimal m satisfying this property is called the *interaction-width* of R , and it is denoted as $\text{width}(R)$. When R does not have a finite interaction-width, we define $\text{width}(R) \stackrel{\text{def}}{=} \omega$. The *interaction-width of a protocol* is defined as the interaction-width of its implicit additive preorder.

Finite interaction-width additive preorders are related to Petri nets as follows. A P -Petri net T (or simply a Petri net when P is clear from the context) is a finite set of P -transitions. The reachability relation of a Petri net T is the binary relation $\xrightarrow{T^*}$ over the configurations defined by $\alpha \xrightarrow{T^*} \beta$ if there exists $\sigma \in T^*$ such that $\alpha \xrightarrow{\sigma} \beta$. Observe that the reachability relation of a Petri net T is an additive preorder with an interaction-width bounded by $\max_{t \in T} |t|$. Moreover, if R is an additive preorder with a finite interaction-width then R is the reachability relation of the Petri net $\{t \in R \mid |t| \leq \text{width}(R)\}$. It follows that the class of additive preorders with finite interaction-widths is equal to the class of Petri net reachability relations.

4 STATE-COMPLEXITY OF PROTOCOLS

The state-complexity of a predicate is intuitively the minimal number of states of a protocol that stably computes it. In this section, we show that the state-complexity must involve the interaction-width and the number of leaders to discard trivial results.

Since our paper focuses on the so-called counting predicates, let us first introduce those predicates. The $(i \geq n)$ *predicate* where i is a state and n a positive natural number is the predicate $\phi_{i \geq n} : \mathbb{N}^I \rightarrow \{0, 1\}$ where $I \stackrel{\text{def}}{=} \{i\}$ satisfying $\phi_{i \geq n}(\rho) = 1$ if, and only if $\rho(i) \geq n$ for every configuration $\rho \in \mathbb{N}^I$. Such a predicate is called a *counting predicate*.

The following two examples show that counting the minimal number of states of protocols stably computing a counting predicate without taking into account the interaction-width or the number of leaders is not relevant.

Example 4.1. This example shows that the predicate $\phi_{i \geq n}$ is stably computable by a leaderless conservative protocol with a number of states bounded 2. We introduce the protocol $(P, \xrightarrow{*}, \rho_L, I, \gamma)$ where $P \stackrel{\text{def}}{=} \{i, p\}$ for some state $p \neq i$, ρ_L is the zero configuration, $I \stackrel{\text{def}}{=} \{i\}$, $\gamma^{-1}(\{0\}) = \{i\}$, $\gamma^{-1}(\{1\}) = \{p\}$, and $\xrightarrow{*}$ is the additive preorder defined by $\alpha \xrightarrow{*} \beta$ if there exists $m \in \mathbb{N}$ such that $\beta + m \cdot i = \alpha + m \cdot p$ and such that $m = 0 \vee |\alpha| \geq n$. In fact, notice that for every $\rho \in \mathbb{N}^I$ and $\alpha \in \mathbb{N}^P$ such that $\rho|_P \xrightarrow{*} \alpha$, then either $\rho(i) < n$ and in that case $\alpha = \rho|_P$ is a 0-output stable configuration, or $\rho(i) \geq n$ and in that case $\alpha \xrightarrow{*} \beta$ where β is the 1-output stable configuration defined as $\rho(i) \cdot p$. We have proved that the protocol stably computes $\phi_{i \geq n}$. The interaction-width of the previous protocol is bounded by n . In fact, just observe that the reachability relation of the Petri net $\{(\rho + i, \rho + p) \mid \rho \in \mathbb{N}^P \wedge |\rho| = n - 1\}$ is equal to $\xrightarrow{*}$. It follows that $\text{width}(\xrightarrow{*}) \leq n$. In fact, one can easily prove that $\text{width}(\xrightarrow{*}) = n$.

Example 4.2. This example shows that the predicate $\phi_{i \geq n}$ is stably computable by a conservative protocol with a number of states bounded 6 and an interaction-width bounded by 2. We introduce the protocol $(P, \xrightarrow{T^*}, n \cdot \bar{i}, I, \gamma)$ where $P \stackrel{\text{def}}{=} \{i, \bar{i}, p, \bar{p}, q, \bar{q}\}$, $I \stackrel{\text{def}}{=} \{i\}$, $\gamma^{-1}(\{1\}) = \{i, p, q\}$, $\gamma^{-1}(\{0\}) = \{\bar{i}, \bar{p}, \bar{q}\}$, and $T \stackrel{\text{def}}{=} \{t, t_p, \bar{t}_p, t_q, \bar{t}_q, \bar{t}_p, \bar{t}_q\}$ is the Petri net

defined as follows:

$$t \stackrel{\text{def}}{=} (i + \bar{i}, p + q) \quad t_p \stackrel{\text{def}}{=} (\bar{p} + i, p + i) \quad t_q \stackrel{\text{def}}{=} (\bar{q} + i, q + i) \quad t_{\bar{q}} \stackrel{\text{def}}{=} (p + \bar{q}, p + q) \\ \bar{t}_p \stackrel{\text{def}}{=} (p + \bar{i}, \bar{p} + \bar{i}) \quad \bar{t}_q \stackrel{\text{def}}{=} (q + \bar{i}, \bar{q} + \bar{i}) \quad t_{\bar{p}} \stackrel{\text{def}}{=} (q + \bar{p}, q + p)$$

Notice that each transition, except t , can only change the bar statue of an agent in states p and q . Let $\rho \in \mathbb{N}^I$ and let α be a configuration such that $n \cdot \bar{i} + \rho|_P \xrightarrow{T^*} \alpha$ and let us prove that there exists a $\phi_{i \geq n}(\rho)$ -output stable configuration β such that $\alpha \xrightarrow{T^*} \beta$. First of all, by executing $\min\{\alpha(i), \alpha(\bar{i})\}$ times the transition t from α , we can assume without loss of generality that $\alpha(i) = 0$ or $\alpha(\bar{i}) = 0$. Observe that $\alpha(\bar{i}) = 0$ if and only if $\phi_{i \geq n}(\rho) = 1$. Assume first that $\alpha(\bar{i}) > 0$. In that case, by executing the transitions \bar{t}_p and \bar{t}_q the right number of times, we get from α a configuration β such that $\beta(i) = \beta(p) = \beta(q) = 0$. Notice that β is 0-output stable. Next, assume that $\alpha(i) > 0$. In that case, by executing the transitions t_p and t_q the right number of times, we get a configuration β such that $\beta(\bar{i}) = \beta(\bar{p}) = \beta(\bar{q}) = 0$. This configuration is 1-output stable. Finally, assume that $\alpha(i) = 0$ and $\alpha(\bar{i}) = 0$. Since $n > 0$, by identifying the last time transition t is executed, we deduce that there exist configurations μ, δ such that $n \cdot \bar{i} + \rho|_P \xrightarrow{T^*} \mu \xrightarrow{t} \delta \xrightarrow{T_0^*} \alpha$ such that $\delta(i) = \delta(\bar{i}) = 0$ where $T_0 \stackrel{\text{def}}{=} \{t_{\bar{q}}, t_{\bar{p}}\}$. Notice that $\delta(p) > 0$ and $\delta(q) > 0$. Since this property is invariant by executing the transitions in T_0 , we deduce that $\alpha(p) > 0$ and $\alpha(q) > 0$. By executing $t_{\bar{p}}$ and $t_{\bar{q}}$ the right number of times, we get from α a configuration β such that $\beta(\bar{i}) = \beta(\bar{p}) = \beta(\bar{q}) = 0$. This configuration is 1-output stable. We have proved that the protocol stably computes $\phi_{i \geq n}$.

In [6], it is exhibited an infinite set of natural numbers n for which there exists a conservative population protocol stably computing the counting predicate ($i \geq n$) with an interaction-width bounded by 2, a number of leaders bounded by 2, and a number of states bounded by $O(\log \log(n))$. This paper left open the optimality of that bound. In this paper we show that such a bound is almost optimal by proving the following main theorem.

THEOREM 4.3. *For every protocol $(P, \xrightarrow{*}, \rho_L, I, F)$ with a finite interaction-width that stably computes $\phi_{i \geq n}$, we have:*

$$n \leq (4 + 4 \text{width}(\xrightarrow{*}) + 2|\rho_L|)^{|P|^{(|P|+2)^2}}$$

We deduce the following state complexity lower-bound as a corollary.

COROLLARY 4.4. *Let $h < \frac{1}{2}$ and let $m \geq 1$. The number of states of a protocol stably computing the ($i \geq n$) predicate with an interaction-width bounded by m and a number of leaders bounded by m is at least $\Omega((\log \log(n))^h)$.*

PROOF. Let us consider $\varepsilon > 0$ such that $\frac{1}{2+\varepsilon} \geq h$. Notice that for $d \in \mathbb{N}$ large enough, we have $d \leq 2^{(d+2)^\varepsilon}$. It follows that $d^{(d+2)^2} \leq 2^{(d+2)^{2+\varepsilon}}$ for d large enough. Theorem 4.3 shows that a protocol stably computing the ($i \geq n$) predicate with an interaction-width bounded by m and a number of leaders bounded by m satisfies:

$$n \leq (10m)^{|P|^{(|P|+2)^2}}$$

It follows that if n is large enough, then $|P|$ is large enough to satisfy $|P|^{(|P|+2)^2} \leq 2^{(|P|+2)^{2+\varepsilon}}$. It follows that $\log \log(n) \leq \log \log(10m) + \log(2)(|P|+2)^{2+\varepsilon}$. In particular:

$$|P| \geq \left(\frac{\log \log(n) - \log \log(10m)}{\log(2)} \right)^h - 2$$

We have proved the corollary. □

5 SMALL STABLE CONFIGURATIONS

A P -configuration ρ is said to be (T, F) -stabilized where T is a P -Petri net, and F is a subset of P if for every configuration β such that $\rho \xrightarrow{T^*} \beta$, we have $\beta(p) = 0$ for every $p \in P \setminus F$. In this section, we show that (T, F) -stabilized configurations are characterized by “small values”. This result is obtained by applying classical Rackoff’s techniques originally introduced for the Petri net coverability problem. The definition of stabilized configurations is related to the output stable configurations of protocols as shown by the following lemma.

LEMMA 5.1. *Let $(P, \xrightarrow{T^*}, \rho_L, I, \gamma)$ be a protocol where T is a Petri net and let $F \stackrel{\text{def}}{=} \gamma^{-1}(\{0\})$. A configuration is (T, F) -stabilized if, and only if, it is 0-output stable.*

PROOF. By definition. □

We first introduce some notations. Given a P -configuration ρ , we introduce $\|\rho\|_\infty \stackrel{\text{def}}{=} \max_{p \in P} \rho(p)$. Given a transition $t = (\alpha_t, \beta_t)$, we also introduce $\|t\|_\infty \stackrel{\text{def}}{=} \max\{\|\alpha_t\|_\infty, \|\beta_t\|_\infty\}$. Given a Petri net T , we define $\|T\|_\infty \stackrel{\text{def}}{=} \max_{t \in T} \|t\|_\infty$. Given a finite set Q of states, we define several restrictions related to Q as follows. Let us recall that given a P -configuration ρ , we previously defined $\rho|_Q$ as the Q -configuration defined by $\rho|_Q(q) \stackrel{\text{def}}{=} \rho(q)$ if $q \in Q$, and zero otherwise. Given a P -transition $t = (\alpha_t, \beta_t)$, we define the Q -transition $t|_Q$ as the pair $t|_Q \stackrel{\text{def}}{=} (\alpha_t|_Q, \beta_t|_Q)$. Given a P -Petri net T , we introduce the Q -Petri net $T|_Q \stackrel{\text{def}}{=} \{t|_Q \mid t \in T\}$. Given a word $\sigma = t_1 \dots t_k$ of P -transitions, we introduce the word $\sigma|_Q = t_1|_Q \dots t_k|_Q$ of Q -transitions. Notice that $\alpha \xrightarrow{\sigma} \beta$ for some P -configurations α, β implies $\alpha|_Q \xrightarrow{\sigma|_Q} \beta|_Q$. The converse property is true in some cases as shown by the following lemma.

LEMMA 5.2. *Assume that $\alpha|_Q \xrightarrow{\sigma|_Q} \rho$ for some P -configurations α, ρ , some word σ of transitions in a P -Petri net T , and some finite set Q . If $\alpha(p) \geq |\sigma| \|T\|_\infty$ for every $p \in P \setminus Q$ then there exists a configuration β such that $\alpha \xrightarrow{\sigma} \beta$, $\beta|_Q = \rho$, and $\beta(p) \geq \alpha(p) - |\sigma| \|T\|_\infty$ for every $p \in P \setminus Q$.*

PROOF. Simple induction on $|\sigma|$. □

A configuration ρ is said to be T -coverable from a configuration α where T is a Petri net if there exists a word $\sigma \in T^*$ such that $\alpha \xrightarrow{\sigma} \beta \geq \rho$ for some configuration β . The minimal length of such a word σ can be bounded using Rackoff’s techniques with respect to $\|\rho\|_\infty$, $\|T\|_\infty$, and $|Q|$ as shown by the following result introduced in [13] to prove that the T -coverability problem is decidable in exponential space.

LEMMA 5.3 ([13]). *If a configuration ρ is T -coverable from a configuration α where T is a P -Petri net, then there exists $\sigma \in T^*$ with a length bounded by $(\|\rho\|_\infty + \|T\|_\infty)^{|P|}$, and a configuration β such that $\alpha \xrightarrow{\sigma} \beta \geq \rho$.*

PROOF. This is a classical result obtained by Rackoff in [13] by induction on $|P|$. □

We deduce the following lemma that shows that (T, F) -stabilized configurations are characterized by “small values” (the values $\rho(p)$ for $p \in R$). In the statement of that lemma the relation \leq over the P -configurations is defined by $\alpha \leq \beta$ if there exists a P -configuration ρ such that $\beta = \alpha + \rho$.

LEMMA 5.4. *Let ρ be a (T, F) -stabilized P -configuration with $F \subseteq P$, let h be a positive integer satisfying $h \geq \|T\|_\infty(1 + \|T\|_\infty)^{|P|}$, and let $R \stackrel{\text{def}}{=} \{p \in P \mid \rho(p) < h\}$. Every P -configuration α such that $\alpha|_R \leq \rho|_R$ is (T, F) -stabilized.*

PROOF. Let us consider a P -configuration α such that $\alpha|_R \leq \rho|_R$. There exists a R -configuration μ such that $\rho|_R = \alpha|_R + \mu$. Assume by contradiction that α is not (T, F) -stabilized. It follows that there exists a configuration β such that $\alpha \xrightarrow{T^*} \beta$ and $\beta(p) > 0$ for some place $p \in P \setminus F$. Since $p \notin F$ and ρ is (T, F) -stabilized, we deduce that $\rho(p) = 0$. In particular $p \in R$ since $h > 0$. Since $p|_P$ is T -coverable from α , Lemma 5.3 shows there exists a word $\sigma \in T^*$ of length bounded by $(1 + \|T\|_\infty)^{|P|^{P|}}$ and a configuration η such that $\alpha \xrightarrow{\sigma} \eta \geq p|_P$. It follows that $\alpha|_R \xrightarrow{\sigma|_R} \eta|_R$. From this relation and $\rho_R = \alpha|_R + \mu$, we deduce that $\rho_R \xrightarrow{\sigma|_R} \eta|_R + \mu$. Lemma 5.2 shows that there exists a configuration δ such that $\rho \xrightarrow{\sigma} \delta$ and $\delta|_R = \eta|_R + \mu$. Since $p \in R$, we deduce that $\delta(p) = \eta(p) + \mu(p) \geq p|_P(p) = 1$. As $p \notin F$, it follows that ρ is not (T, F) -stabilized and we get a contradiction. We have proved the lemma. \square

REMARK 2. A similar result was provided in [7] in the context of conservative Petri nets with interaction-width bounded by two.

6 BOTTOM CONFIGURATIONS

Let T be a P -Petri net. The T -component of a P -configuration ρ is the set of configurations β such that $\rho \xrightarrow{T^*} \beta \xrightarrow{T^*} \rho$. A configuration ρ is said to be T -bottom if its T -component is finite and every configuration β such that $\rho \xrightarrow{T^*} \beta$ satisfies $\beta \xrightarrow{T^*} \rho$.

In this section we prove the following theorem that intuitively provides a way to reach with short words kind of bottom configurations with small size (small and short meaning doubly-exponential in that context). Other results proved in this section are only used for proving the following theorem and are no longer used in the sequel.

THEOREM 6.1. Let T be a P -Petri net, let ρ be a P -configuration, and let $b \stackrel{\text{def}}{=} (4 + 4\|T\|_\infty + 2\|\rho\|_\infty)^{d^d(1+(2+d^d)^{d+1})}$ with $d \stackrel{\text{def}}{=} |P|$. There exist two words $\sigma, w \in T^*$, a set of places $Q \subseteq P$, and two P -configurations α, β such that:

- $\rho \xrightarrow{\sigma} \alpha \xrightarrow{w} \beta$.
- $\alpha|_Q = \beta|_Q$.
- $\alpha(p) < \beta(p)$ for every state $p \in P \setminus Q$.
- $\alpha|_Q$ is $T|_Q$ -bottom.
- The cardinal of the $T|_Q$ -component of $\alpha|_Q$ is bounded by b .
- $|\sigma|, |w|, d\|\alpha\|_\infty, d\|\beta\|_\infty \leq b$.

The proof of the previous theorem is obtained by iterating the following lemma in order to obtain an increasing sequence of sets Q .

LEMMA 6.2. Let T be a P -Petri net, let ρ be a P -configuration, let Q be a set of states included in P such that $\rho|_Q$ is $T|_Q$ -bottom, let s be the cardinal of the $T|_Q$ -component of $\rho|_Q$, and let $d \stackrel{\text{def}}{=} |P \setminus Q|$.

There exist a word $\sigma \in T^*$ such that $|\sigma| \leq (1 + d(1 + s\|T\|_\infty + \|\rho\|_\infty)^{d^d})s$, and a P -configuration ρ' such that $\rho \xrightarrow{\sigma} \rho'$ and such that:

- either $\rho'|_Q = \rho|_Q$ and $\rho'(p) > \rho(p)$ for every $p \in P \setminus Q$,
- or there exists a set $Q' \subseteq P$ that strictly contains Q such that $\rho'|_{Q'}$ is $T|_{Q'}$ -bottom and the cardinal s' of the $T|_{Q'}$ -component of $\rho'|_{Q'}$ satisfies:

$$s' \leq (1 + d(1 + s\|T\|_\infty + \|\rho\|_\infty)^{d^d})s$$

PROOF. Let us introduce the sequence $\lambda_1, \dots, \lambda_d$ of natural numbers satisfying $\lambda_d \stackrel{\text{def}}{=} 1 + s\|T\|_\infty + \|\rho\|_\infty$ and satisfying $\lambda_n \stackrel{\text{def}}{=} s\lambda_{n+1}^{d-n}\|T\|_\infty + \lambda_{n+1}$ for every $n \in \{1, \dots, d-1\}$. Observe that $\lambda_1 \geq \dots \geq \lambda_d$. Moreover, $\lambda_n \leq \lambda_d \lambda_{n+1}^{d-n}$ for every $1 \leq n < d$. We deduce by induction that $\lambda_1^d \leq \lambda_d^{d^d}$.

Let $\rho_0 \stackrel{\text{def}}{=} \rho$. We are going to build by induction on n a sequence ρ_1, \dots, ρ_n of configurations, a sequence $\sigma_1, \dots, \sigma_n$ of words in T^* , and a sequence p_1, \dots, p_n of states in $P \setminus Q$ such that for every $i \in \{1, \dots, n\}$ we have:

- (i) $\rho_{i-1} \xrightarrow{\sigma_i} \rho_i$.
- (ii) $|\sigma_i| \leq \lambda_i^{d-i+1}s$.
- (iii) $\rho_i(p) \geq \lambda_i$ for every $p \in \{p_1, \dots, p_i\}$.

So, let us assume that $\rho_1, \dots, \rho_n, \sigma_1, \dots, \sigma_n$, and p_1, \dots, p_n are built for some $n \geq 0$. Since p_1, \dots, p_n are distinct elements in $P \setminus Q$, it follows that $n \leq d$.

Let us first assume that $n = d$. In that case, we have $\rho_d(p) \geq \lambda_d$ for every $p \in P \setminus Q$. As $\rho|_Q$ is a $T|_Q$ -bottom configuration and $\rho|_Q \xrightarrow{(\sigma_1 \dots \sigma_d)|_Q} \rho_d|_Q$ and since the cardinal of the $T|_Q$ -component of $\rho|_Q$ is bounded by s , we deduce that there exists a word $w \in T^*$ such that $\rho_d|_Q \xrightarrow{w|_Q} \rho|_Q$ and $|w| < s$. Since $\lambda_d \geq s\|T\|_\infty \geq |w|\|T\|_\infty$, Lemma 5.2 shows that $\rho_d \xrightarrow{w} \rho'$ for some configuration ρ' such that $\rho'|_Q = \rho|_Q$ and such that for every $p \in P \setminus Q$ we have $\rho'(p) \geq \rho(p) - |w|\|T\|_\infty \geq \lambda_d - s\|T\|_\infty > \rho(p)$ by definition of λ_d . Let us introduce $\sigma \stackrel{\text{def}}{=} \sigma_1 \dots \sigma_d w$ and notice that $|\sigma| \leq (\lambda_1^d + \dots + \lambda_d^{d-d+1} + 1)s \leq (1 + d\lambda_1^d)s$ and we have proved that the lemma holds (first case).

So we can assume that $n < d$. Let us introduce the set $R_n = (P \setminus Q) \setminus \{p_1, \dots, p_n\}$. Since $|R_n| = d - n$, the set R_n is non empty.

Assume first that for every configuration β such that $\rho_n \xrightarrow{T^*} \beta$ we have $\beta(p) < \lambda_{n+1}$ for every $p \in R_n$. In that case let $Q' \stackrel{\text{def}}{=} Q \cup R_n$. It follows that the cardinal of the set of configurations β' such that $\rho_n|_{Q'} \xrightarrow{T|_{Q'}} \beta'$ is bounded by $s\lambda_{n+1}^{d-n}$. Hence, there exists a configuration β' that is $T|_{Q'}$ -bottom and a word $w \in T^*$ such that $\rho_n|_{Q'} \xrightarrow{w|_{Q'}} \beta'$ and such that $|w| < s\lambda_{n+1}^{d-n}$. Notice that the cardinal s' of the $T|_{Q'}$ -component of β' is bounded by $s\lambda_{n+1}^{d-n} \leq s\lambda_1^d$. As $\rho_n(p) \geq \lambda_n$ for every $p \in \{p_1, \dots, p_n\}$ and $\lambda_n \geq (s\lambda_{n+1}^{d-n} - 1)\|T\|_\infty \geq |w|\|T\|_\infty$, Lemma 5.2 shows that there exists a configuration ρ' such that $\rho_n \xrightarrow{w} \rho'$, and $\rho'|_{Q'} = \beta'$. Let us consider the word $\sigma \stackrel{\text{def}}{=} \sigma_1 \dots \sigma_n w$. Notice that $|\sigma| \leq (\lambda_1^d + \dots + \lambda_n^{d-n+1} + \lambda_{n+1}^{d-n})s \leq d\lambda_1^d s$ and we have proved that the lemma holds (second case).

Finally, assume that there exists a configuration ρ_{n+1} such that $\rho_n \xrightarrow{\sigma_{n+1}} \rho_{n+1}$ for some word $\sigma_{n+1} \in T^*$ and such that $\rho_{n+1}(p_{n+1}) \geq \lambda_{n+1}$ for some state $p_{n+1} \in R_n$. We assume that $|\sigma_{n+1}|$ is minimal. Observe that every intermediate configuration β such that $\rho_n \xrightarrow{u} \beta \xrightarrow{v} \rho_{n+1}$ with $uv = \sigma_{n+1}$ and $|v| \geq 1$ satisfies $\beta(p) < \lambda_{n+1}$ for every $p \in R_n$ by minimality of $|\sigma_{n+1}|$. We deduce that there exists a word $w \in T^*$ such that $\rho_n|_{Q \cup R_n} \xrightarrow{w|_{Q \cup R_n}} \rho_{n+1}|_{Q \cup R_n}$ and such that $|w| \leq s\lambda_{n+1}^{d-n}$. As $\rho_n(p) \geq \lambda_n$ for every $p \in \{p_1, \dots, p_n\}$ and $\lambda_n \geq s\lambda_{n+1}^{d-n}\|T\|_\infty \geq |w|\|T\|_\infty$, Lemma 5.2 shows that there exists a configuration β such that $\rho_n \xrightarrow{w} \beta$ and $\beta|_{Q \cup R_n} = \rho_{n+1}|_{Q \cup R_n}$. In particular $\beta(p_{n+1}) = \rho_{n+1}(p_{n+1}) \geq \lambda_{n+1}$. By minimality of $|\sigma_{n+1}|$, we get $|\sigma_{n+1}| \leq |w| \leq s\lambda_{n+1}^{d-n}$. Now, observe that for every $p \in \{p_1, \dots, p_n\}$ we have $\rho_{n+1}(p) \geq \rho_n(p) - |\sigma_{n+1}|\|T\|_\infty \geq \lambda_n - s\lambda_{n+1}^{d-n}\|T\|_\infty \geq \lambda_{n+1}$ by definition of λ_n . We have extended our sequence in such a way (i), (ii), and (iii) are fulfilled.

We have proved the lemma. □

Now, let us prove Theorem 6.1. Observe that if $d = 0$ the theorem is trivial. So, we can assume that $d \geq 1$. Let $Q_0 \stackrel{\text{def}}{=} \emptyset$, $\rho_0 \stackrel{\text{def}}{=} \rho$, and $s_0 \stackrel{\text{def}}{=} 1$. Notice that $\rho_0|_{Q_0}$ is $T|_{Q_0}$ -bottom and the cardinal of the $T|_{Q_0}$ -component of $\rho_0|_{Q_0}$ contains s_0

elements. We build by induction on n a sequence Q_1, \dots, Q_n of subsets of P , a sequence ρ_1, \dots, ρ_n of configurations, a sequence $\sigma_1, \dots, \sigma_n$ of words in T^* such that for every $i \in \{1, \dots, n\}$:

- $\rho_{i-1} \xrightarrow{\sigma_i} \rho_i$.
- $\rho_i|_{Q_i}$ is $T|_{Q_i}$ -bottom.
- The cardinal of the $T|_{Q_i}$ -component of $\rho_i|_{Q_i}$ is equal to s_i .
- $Q_{i-1} \subset Q_i$.
- $|\sigma_i|, s_i \leq (1 + d(1 + s_{i-1}\|T\|_\infty + \|\rho_{i-1}\|_\infty)^{d^d})s_{i-1}$.

Assume the sequence built for some $n \geq 0$. Lemma 6.2 on the configuration ρ_n and the set Q_n shows that there exist a word σ_{n+1} such that $|\sigma_{n+1}| \leq (1 + d(1 + s_n\|T\|_\infty + \|\rho_n\|_\infty)^{d^d})s_n$, and a configuration ρ_{n+1} such that $\rho_n \xrightarrow{\sigma_{n+1}} \rho_{n+1}$, such that:

- either $\rho_{n+1}|_{Q_n} = \rho_n|_{Q_n}$ and $\rho_{n+1}(p) > \rho_n(p)$ for every $p \in P \setminus Q_n$,
- or there exists Q_{n+1} such that $Q_n \subset Q_{n+1} \subseteq P$ such that $\rho_{n+1}|_{Q_{n+1}}$ is $T|_{Q_{n+1}}$ -bottom and the cardinal s_{n+1} of its $T|_{Q_{n+1}}$ -component satisfies:

$$s_{n+1} \leq (1 + d(1 + s_n\|T\|_\infty + \|\rho_n\|_\infty)^{d^d})s_n$$

Observe that in the second case we have extended the sequences. In the first case, let $\alpha \stackrel{\text{def}}{=} \rho_n$, $\beta \stackrel{\text{def}}{=} \rho_{n+1}$, $\sigma \stackrel{\text{def}}{=} \sigma_1 \dots \sigma_n$, $w \stackrel{\text{def}}{=} \sigma_{n+1}$, and $Q \stackrel{\text{def}}{=} Q_n$. Since $Q_0 \subset Q_1 \dots \subset Q_n$ are subsets of P , we deduce that $n \leq d$. Let us introduce $a = (1 + d)(2 + 2\|T\|_\infty + \|\rho\|_\infty)^{d^d}$ and $h = 2 + d^d$ and let us prove by induction on i that we have $|\sigma_i|, s_i \leq a^{h^i}$ and $\|\rho_i\|_\infty \leq (1 + \|T\|_\infty)a^{h^i}$ with the convention $\sigma_0 = \varepsilon$.

The rank $i = 0$ is immediate. Assume the rank $i - 1$ proved. We have:

$$\begin{aligned} |\sigma_i|, s_i &\leq (1 + d(1 + s_{i-1}\|T\|_\infty + \|\rho_{i-1}\|_\infty)^{d^d})s_{i-1} \\ &\leq (1 + d)(2 + 2\|T\|_\infty)^{d^d} a^{h^{i-1}(d^d+1)} \\ &\leq a^{1+h^{i-1}(h-1)} \\ &\leq a^{h^i} \end{aligned}$$

Since $\rho_{i-1} \xrightarrow{\sigma_i} \rho_i$, we deduce that $\|\rho_i\|_\infty \leq \|\rho_{i-1}\|_\infty + |\sigma_i|\|T\|_\infty \leq (1 + \|T\|_\infty)a^{h^i}$. The induction is proved.

It follows that $|\sigma| \leq da^{h^d}$, $|w| \leq a^{h^{d+1}}$, and $d\|\alpha\|_\infty, d\|\beta\|_\infty \leq d(1 + \|T\|_\infty)a^{h^{d+1}} \leq a^{1+h^{d+1}}$. Since $d \geq 1$, we deduce that $(1 + d) \leq 2^{d^d}$. In particular $a \leq (4 + 4\|T\|_\infty + 2\|\rho\|_\infty)^{d^d}$. We have proved Theorem 6.1.

7 PETRI NETS WITH CONTROL-STATES

A P -Petri net with control-states (or simply a Petri net with control-states when the finite set of states P is clear from the context) is a triple (S, T, E) where S is a non empty finite set of elements called *control-states*, T is a P -Petri net, and $E \subseteq S \times T \times S$ is a set of elements called *edges*. The *Parikh image* of a word $\pi = e_1 \dots e_k$ of edges is the mapping $\#\pi \in \mathbb{N}^E$ defined by $\#\pi(e) = |\{j \in \{1, \dots, k\} \mid e_j = e\}|$. The *displacement* of a transition $t = (\alpha_t, \beta_t)$ is the function $\Delta(t) \in \mathbb{Z}^P$ defined by $\Delta(t)(p) = \beta_t(p) - \alpha_t(p)$ for every $p \in P$. The *displacement* of an edge $e = (s, t, s')$ is defined as $\Delta(e) \stackrel{\text{def}}{=} \Delta(t)$. The displacement of a word $\pi = e_1 \dots e_k$ of edges is $\Delta(\pi) \stackrel{\text{def}}{=} \sum_{1 \leq j \leq k} \Delta(e_j)$. We denote by $|\pi| \stackrel{\text{def}}{=} k$ the *length* of π . A *path* π from a control-state s to a control-state s' is a word $\pi = e_1 \dots e_k$ of edges in E such that there exists control-states s_0, \dots, s_k in S and transitions t_1, \dots, t_k in T such that $s_0 = s, s_k = s'$, and such that $e_j = (s_{j-1}, t_j, s_j)$ for every $1 \leq j \leq k$. Such a path is called a *cycle* if $s = s'$. A cycle θ of a Petri net with control-states is said to be *total*

if $\#\theta(e) > 0$ for every $e \in E$. The cycle is said to be *simple* if the control-states s_1, \dots, s_k are distinct. A *multicycle* Θ is a sequence $\theta_1, \dots, \theta_k$ of cycles. We denote by $|\Theta| \stackrel{\text{def}}{=} \sum_{j=1}^k |\theta_j|$, the *length* of a multicycle Θ . We introduce the *Parikh image* $\#\Theta \stackrel{\text{def}}{=} \sum_{j=1}^k \#\theta_j$ and the *displacement* $\Delta(\Theta) \stackrel{\text{def}}{=} \sum_{j=1}^k \Delta(\theta_j)$ of such a multicycle Θ . A multicycle Θ is said to be *total* if $\#\Theta(e) > 0$ for every $e \in E$.

A Petri net with control-states (S, T, E) is said to be *strongly connected* if for every pair (s, s') of control-states in S , there exists a path from s to s' . Let us recall the classical Euler lemma in the context of Petri nets with control-states.

LEMMA 7.1 (EULER LEMMA). *For every total multicycle Θ in a strongly connected Petri net with control-states there exists a total cycle θ such that $\#\theta = \#\Theta$.*

We deduce the following lemma.

LEMMA 7.2. *For any strongly connected Petri net with control-states (S, T, E) , there exists a total cycle θ with a length bounded by $|E||S|$.*

PROOF. Every edge $e \in E$ occurs in at least one simple cycle θ_e . It follows that the multicycle $\Theta = (\theta_e)_{e \in E}$ is total. From Lemma 7.1 we deduce that there exists a total cycle θ such that $\#\theta = \#\Theta$. Notice that $|\theta| = \sum_{e \in E} |\theta_e| \leq |E||S|$. \square

A mapping $a \in \mathbb{Z}^P$ is called a *P-action* (or simply an action if P is clear from the context). Notice that displacements of transitions, edges, paths, and multicycles are actions. We associate with an action a the value $\|a\|_1 \stackrel{\text{def}}{=} \sum_{p \in P} |a(p)|$. Given a finite set Q , we denote by $a|_Q$ the action defined for every $q \in Q$ by $a|_Q(q) \stackrel{\text{def}}{=} a(q)$ if $q \in P$, and zero otherwise.

LEMMA 7.3. *Let Θ be a multicycle of a P-Petri net with control-states (S, T, E) with $\|T\|_\infty > 0$, let $Q \subseteq P$, let $d \stackrel{\text{def}}{=} |P|$, and let $k > \|\Delta(\Theta)|_Q\|_1 (1 + 2|S|\|T\|_\infty)^{d(d+1)}$.*

There exists a multicycle Θ' such that:

- For every $p \in P$ we have:
 - $\Delta(\Theta')(p) \leq 0$ if $\Delta(\Theta)(p) \leq 0$.
 - $\Delta(\Theta')(p) < 0$ if $\Delta(\Theta)(p) \leq -k$.
 - $\Delta(\Theta')(p) \geq 0$ if $\Delta(\Theta)(p) \geq 0$.
 - $\Delta(\Theta')(p) > 0$ if $\Delta(\Theta)(p) \geq k$.
- For every $q \in Q$ we have $\Delta(\Theta')(q) = 0$.
- For every edge $e \in E$ we have $\#\Theta'(e) > 0$ if $\#\Theta(e) \geq k$.
- $|\Theta'| \leq (|E| + d)(1 + 2|S|\|T\|_\infty)^{d(d+1)}$.

PROOF. Since every cycle can be decomposed into a sequence of simple cycles without changing the Parikh image, we can assume without loss of generality that Θ is a sequence of simple cycles. We introduce the set A of actions $\Delta(\theta)$ where θ ranges over the simple cycles, and $n \stackrel{\text{def}}{=} |A|$ its cardinal. Notice that for every $a \in A$ and for every $p \in P$, we have $|a(p)| \leq |S|\|T\|_\infty$. It follows that $n \leq (1 + 2|S|\|T\|_\infty)^d$.

We denote by s the sign function of a formally defined by $s(c) \stackrel{\text{def}}{=} 1$ if $\Delta(\Theta)(c) \geq 0$, $s(c) \stackrel{\text{def}}{=} -1$ otherwise. We also introduce the P -configuration f defined by $f(p) \stackrel{\text{def}}{=} |\Delta(\Theta)(p)|$ for every $p \in P$, and the function $g : A \rightarrow \mathbb{N}$ such that $g(a)$ is the number of simple cycle θ that occurs in Θ such that $\Delta(\theta) = a$.

Notice that $s(p)f(p) = \sum_{a \in A} g(a)a(p)$ for every $p \in P$. We introduce the following linear system over the free variables $(\alpha, \beta) \in \mathbb{N}^P \times \mathbb{N}^A$:

$$\bigwedge_{p \in P} s(p)\alpha(p) = \sum_{a \in A} \beta(a)a(p) \tag{1}$$

Notice that (f, g) is a solution of that system. From [12], there exists a finite set H of solutions (α, β) of that system such that $(f, g) = \sum_{(\alpha, \beta) \in H} (\alpha, \beta)$ and such that for every $(\alpha, \beta) \in H$, we have $\|\alpha\|_1 + \|\beta\|_1 \leq (2 + \sum_{a \in A} \|a\|_\infty)^d$. As $\sum_{a \in A} \|a\|_\infty \leq (1 + 2|S| \|T\|_\infty)^d |S| \|T\|_\infty$ we deduce (by using $|S| \|T\|_\infty \geq 1$):

$$\|\alpha\|_1 + \|\beta\|_1 \leq (1 + 2|S| \|T\|_\infty)^{d(d+1)} \quad (2)$$

We introduce the set H_0 of pairs $(\alpha, \beta) \in H$ such that $\alpha(q) = 0$ for every $q \in Q$. Observe that $\sum_{q \in Q} \sum_{(\alpha, \beta) \in H} \alpha(q)$ is equal to $\sum_{q \in Q} |\Delta(\Theta)(q)| = \|\Delta(\Delta)|_Q\|_1$, and it is also equals to $\sum_{q \in Q} \sum_{(\alpha, \beta) \in H \setminus H_0} \alpha(q) \geq |H \setminus H_0|$. In particular we have:

$$|H \setminus H_0| \leq \|\Delta(\Delta)|_Q\|_1$$

We introduce the set F of edges $e \in E$ such that $\Theta(e) \geq k$. Let $e \in F$. The sum $\sum_{(\alpha, \beta) \in H} \beta(e)$ is equals to $\#\Theta(e)$ and it is also equals to $\sum_{(\alpha, \beta) \in H \setminus H_0} \beta(e) + \sum_{(\alpha, \beta) \in H_0} \beta(e)$. As $\sum_{(\alpha, \beta) \in H \setminus H_0} \beta(e) \leq |H \setminus H_0| (1 + 2|S| \|T\|_\infty)^{d(d+1)}$ we deduce that $\sum_{(\alpha, \beta) \in H_0} \beta(e) > 0$. In particular there exists $(\alpha, \beta) \in H_0$ such that $\beta(e) > 0$.

We also introduce the set R of $p \in P$ such that $|\Delta(\Theta)(p)| \geq k$. Let $p \in R$. The sum $\sum_{(\alpha, \beta) \in H} \alpha(p)$ is equals to $|\Delta(\Theta)(p)|$ and it is also equals to $\sum_{(\alpha, \beta) \in H \setminus H_0} \alpha(p) + \sum_{(\alpha, \beta) \in H_0} \alpha(p)$. As $\sum_{(\alpha, \beta) \in H \setminus H_0} \alpha(p) \leq |H \setminus H_0| (1 + 2|S| \|T\|_\infty)^{d(d+1)}$ we deduce that $\sum_{(\alpha, \beta) \in H_0} \alpha(p) > 0$. In particular there exists $(\alpha, \beta) \in H_0$ such that $\alpha(p) > 0$.

Now, let us introduce for each $e \in F$ a pair $(\alpha_e, \beta_e) \in H_0$ such that $\beta_e(e) > 0$, and let us introduce for each $p \in R$ a pair $(\alpha_p, \beta_p) \in H_0$ such that $\alpha_p(p) > 0$. Let us introduce $(\alpha', \beta') \stackrel{\text{def}}{=} \sum_{e \in F} (\alpha_e, \beta_e) + \sum_{p \in R} (\alpha_p, \beta_p)$ and observe that $\alpha'(e) > 0$ for every $e \in F$, $\beta'(p) > 0$ for every $p \in R$, and $\beta'(q) = 0$ for every $q \in Q$. Moreover, since (α', β') is a solution of (1), it follows that there exists a multicycle Θ' such that $\#\Theta' = \beta'$. In particular $\Delta(\Theta') = \Delta(\beta')$. Notice that $\Delta(\Theta') = \alpha'$, and $|\Theta'| = \|\beta'\|_1 \leq (|F| + |R|)(1 + 2|S| \|T\|_\infty)^{d(d+1)} \leq (|E| + d)(1 + 2|S| \|T\|_\infty)^{d(d+1)}$ and we have proved the lemma. \square

8 PROOF OF THEOREM 4.3

In this section we provide a proof of Theorem 4.3.

We consider a finite interaction-width protocols $(P, \xrightarrow{*}, \rho_L, I, \gamma)$ that stably computes the $(n > i)$ predicate. Notice that $I = \{i\}$. We introduce the Petri net $T \stackrel{\text{def}}{=} \{t \in \xrightarrow{*} \mid \text{width}(t) \leq \text{width}(\xrightarrow{*})\}$. Let us recall that the additive preorder $\xrightarrow{*}$ is equal to $\xrightarrow{T^*}$. Let $d \stackrel{\text{def}}{=} |Q|$ and $F = \gamma^{-1}(\{0\})$. Notice that if $d = 1$ then $n = 1$ and the proof of the theorem is done in that case. So, we can assume that $d \geq 2$. We introduce the following numbers:

$$\begin{aligned} b &\stackrel{\text{def}}{=} (4 + 4\|T\|_\infty + 2\|\rho_L\|_\infty)^{(d-1)^{d-1} (1 + (2 + (d-1)^{d-1})^d)} \\ h &\stackrel{\text{def}}{=} d(1 + \|T\|_\infty)b \\ k &\stackrel{\text{def}}{=} dh^{d^2+d+1} \\ a &\stackrel{\text{def}}{=} h^{2d+3} \\ \ell &\stackrel{\text{def}}{=} h^{5d^2} \\ r &\stackrel{\text{def}}{=} 2(d-1)^{d-1} (1 + (2 + (d-1)^{d-1})^d) (5d^2 + 2d + 4) \end{aligned}$$

We introduce $P' \stackrel{\text{def}}{=} P \setminus I$. It follows that $|P'| = d - 1$. Theorem 6.1 applied on the Petri net $T|_{P'}$ and the configuration $\rho_L|_{P'}$ shows that there exist two words $\sigma, w \in T^*$, a set $Q \subseteq P'$, and two configurations α, β such that:

- $\rho_L|_{P'} \xrightarrow{\sigma|_{P'}} \alpha \xrightarrow{w|_{P'}} \beta$.

- $\alpha|_Q = \beta|_Q$.
- $\alpha(p) < \beta(p)$ for every $p \in P' \setminus Q$
- $\alpha|_Q$ is $T|_Q$ -bottom.
- The cardinal of the $T|_Q$ -component of $\alpha|_Q$ is bounded by b .
- $|\sigma|, |w|, d\|\alpha\|_\infty, d\|\beta\|_\infty \leq b$.

Notice that $|T| \leq (1 + 2\|T\|_\infty)^{2d} \leq h^{2d}$.

We introduce the Petri net with control-states (S, T, E) where S is the $T|_Q$ -component of $\alpha|_Q$, and E is the set of edges $(s, t, s') \in S \times T \times S$ such that $s \xrightarrow{t|_Q} s'$. Observe that $|E| \leq |S||T|$ since for every (s, t, s') in E the value of s' is determined by the pair (s, t) . It follows that we have:

$$|E| \leq h^{2d+1}$$

Lemma 7.2 shows that there exists a total cycle θ_E of (S, T, E) with a length bounded by $|S||E|$. Without loss of generality we can assume that this total cycle is on the control-state $\alpha|_Q$ by considering a rotation of that cycle. We denote by σ_E the label in T^* of this total cycle. Observe that $\|T\|_\infty|\sigma_E| \leq a$.

Since $\alpha \xrightarrow{w|_{P'}} \beta$, $\alpha|_Q = \beta|_Q$, and $\alpha(p) < \beta(p)$ for every $p \in P' \setminus Q$, we deduce that there exists a configuration η such that $\eta(p) \geq a\ell$ for every $p \in P' \setminus Q$, such that $\alpha|_Q = \eta|_Q$, and such that:

$$\alpha \xrightarrow{w^{\alpha\ell}|_{P'}} \eta$$

Moreover, since σ_E is the label of a cycle on $\alpha|_Q$ we deduce that $\alpha|_Q \xrightarrow{\sigma_E|_Q} \alpha|_Q$. From $\eta|_Q = \alpha|_Q$ it follows that $\eta|_Q \xrightarrow{\sigma_E^{\ell}|_Q} \alpha|_Q$. As $\eta(p) \geq a\ell \geq \|T\|_\infty|\sigma_E^{\ell}|$ for every $p \in P' \setminus Q$, Lemma 5.2 shows that there exists a P' -configuration δ such that $\delta|_Q = \alpha|_Q$ and such that:

$$\eta \xrightarrow{\sigma_E^{\ell}|_{P'}} \delta$$

Observe that $|\sigma w^{\alpha\ell} \sigma_E^{\ell}| \|T\|_\infty \leq (b + ba\ell)\|T\|_\infty + a\ell \leq 2ba\ell(\|T\|_\infty + 1) \leq a\ell h \leq h^{2d+4}\ell$.

Assume by contradiction that $n > h^{2d+4}\ell$, and let us introduce the configuration ρ' defined by $\rho' \stackrel{\text{def}}{=} \rho_L + (n-1).i$ where i is the state such that $I = \{i\}$. Lemma 5.2 shows that there exist P -configurations α', η', δ' such that $\alpha'|_{P'} = \alpha$, $\eta'|_{P'} = \eta$, $\delta'|_{P'} = \delta$ and such that:

$$\rho' \xrightarrow{\sigma} \alpha' \xrightarrow{w^{\alpha\ell}} \eta' \xrightarrow{\sigma_E^{\ell}} \delta'$$

Since the population protocol is stably computing the $(i \geq n)$ predicate and $n-1 < n$, there exists a 0-output stable configuration μ and a word $\sigma' \in T^*$ such that $\delta' \xrightarrow{\sigma'} \mu$. Lemma 5.1 shows that μ is (T, F) -stabilized. Observe that $w^{\alpha\ell} \sigma_E^{\ell} \sigma'$ is the label of a path of (S, T, E) from $\alpha|_Q$ to $\mu|_Q$. It follows that the Parikh image of that path can be decomposed as the Parikh image of a multicycle Θ and the Parikh image of an elementary path π . Observe $\Delta(\Theta) + \Delta(\pi) = \Delta(w^{\alpha\ell} \sigma_E^{\ell} \sigma') = \mu - \alpha'$. Notice that $\#\Theta(e) \geq \ell$ for every $e \in E$ since σ_E is the label of a total cycle on $\alpha|_Q$. Since π is an elementary path, we deduce that $\|\Delta(\pi)\|_1 \leq d|S|\|T\|_\infty \leq db\|T\|_\infty \leq h - db$.

We introduce the set $R \stackrel{\text{def}}{=} \{p \in P \mid \mu(p) < h\}$. Since $h \geq \|T\|_\infty(1 + \|T\|_\infty)^{d^d}$, Lemma 5.4 shows that every configuration μ' such that $\mu|_R = \mu'|_R$ is (T, F) -stabilized. Observe that if $i \notin R$ then $\mu + i$ is (T, F) -stabilized, and by additivity, we deduce that $\rho_L + n.i \xrightarrow{T^*} \mu + i$. Since $\mu + i$ is (T, F) -stabilized, this configuration cannot reach a 1-output

stable configuration. In particular the protocol is not stably computing the $(i \geq n)$ predicate and we get a contradiction. It follows that $i \in R$.

We introduce $R' \stackrel{\text{def}}{=} R \setminus I$. Since $d\|\alpha\|_\infty \leq b$ and $\alpha'|_{P'} = \alpha|_{P'}$, we deduce that $d\|\alpha'|_{R'}\|_\infty \leq b$. From $\Delta(\Theta) = \mu - \alpha' - \Delta(\pi)$ we deduce:

$$\|\Delta(\Theta)|_{R'}\|_1 \leq (d-1)h + b + h - db \leq dh$$

As $1 + 2|S|\|T\|_\infty \leq 1 + h - 2b < h$, we deduce that $k > \|\Delta(\Theta)|_{R'}\|_1(1 + 2|S|\|T\|_\infty)^{d(d+1)}$, Lemma 7.3 shows that there exists a multicyle Θ' such that:

- For every $p \in P$ we have:
 - $\Delta(\Theta')(p) \leq 0$ if $\Delta(\Theta)(p) \leq 0$.
 - $\Delta(\Theta')(p) < 0$ if $\Delta(\Theta)(p) \leq -k$.
 - $\Delta(\Theta')(p) \geq 0$ if $\Delta(\Theta)(p) \geq 0$.
 - $\Delta(\Theta')(p) > 0$ if $\Delta(\Theta)(p) \geq k$.
- For every $p \in R'$ we have $\Delta(\Theta')(p) = 0$.
- For every $e \in E$ we have $\#\Theta'(e) > 0$ if $\#\Theta(e) \geq k$.
- $\|\Theta'\|_1 \leq (|E| + d)(1 + 2|S|\|T\|_\infty)^{d(d+1)}$

Let $m \stackrel{\text{def}}{=} -\Delta(\Theta')(i)$ and let us prove that $m > 0$. We have $\Delta(\Theta)(i) = \mu(i) - \alpha'(i) - \Delta(\pi)(i)$. Since $i \in R$, we get $\mu(i) < h$. Since $\rho_L + (n-1).i \xrightarrow{\sigma} \alpha'$, we deduce that $\alpha'(i) = \rho_L(i) + (n-1) - \Delta(\sigma)(i) \geq n - h$ since $|\sigma| \leq b$. We deduce that $\Delta(\Theta)(i) < h - n + h + h \leq 3h - n \leq -k$. Hence $\Delta(\Theta')(i) < 0$. It follows that $m > 0$.

Let $\eta \stackrel{\text{def}}{=} m.i + \Delta(\Theta')$. Notice that $\eta(i) = 0$ and $\eta(p) = 0$ for every $p \in R'$. In particular $\eta(p) = 0$ for every $p \in R$. Let us prove that η is a configuration. For every $p \in P \setminus R$ we have $\Delta(\Theta)(p) = \mu(p) - \eta'(p) - \Delta(\pi)(p) \geq db\|T\|_\infty + b - b - db\|T\|_\infty \geq 0$. It follows that $\Delta(\Theta')(p) \geq 0$. In particular $\eta(p) \geq 0$. We have proved that η is a configuration.

Finally, observe that $\#\Theta(e) \geq \ell \geq k$ for every $e \in E$. In particular $\#\Theta'(e) > 0$. Lemma 7.1 shows that $\#\Theta'$ is the Parikh image of a cycle θ on $x|_Q$. Let u be the label of that cycle. Since $|u| = \|\Theta'\|_1$, we deduce that:

$$\begin{aligned} |u|\|T\|_\infty &\leq \|T\|_\infty(|E| + d)(1 + 2b\|T\|_\infty)^{d(d+1)} \\ &\leq d(1 + \|T\|_\infty)^{2d}(1 + 2b\|T\|_\infty)^{d^2+d+1} \\ &\leq \ell \end{aligned}$$

Lemma 5.2 shows that:

$$\eta' + m.i \xrightarrow{u} \eta' + \eta$$

We have proved:

$$\rho_L + (n-1+m).i \xrightarrow{\sigma w^{\alpha\ell} u \sigma_E^\ell \sigma'} \mu + \eta$$

Since $(\mu + \eta)|_R = \mu|_R$ we deduce that $\mu + \eta$ is (T, F) -stabilized. It follows that this configuration cannot reach a 1-output stable configuration. In particular the protocol is not stably computing the $(i \geq n)$ predicate and we get a contradiction. It follows that $n \leq h^{2d+4}\ell = h^{5d^2+2d+4}$.

Notice that $d(1 + \|T\|_\infty) \leq 2^d(1 + \|T\|_\infty)^d \leq b$. Thus $h \leq b^2$. We deduce that $n \leq (4 + 4\|T\|_\infty + 2\|\rho_L\|_\infty)^r$. Since $d \geq 2$, we deduce that $d^d = ((d-1) + 1)^d \geq (d-1)^d + d(d-1)^{d-1} + 1 \geq (d-1)^{d-1} + 2 + 1$. Hence $1 + (2 + (d-1)^{d-1})^d \leq 1 + (d^d - 1)^d \leq d^{d^2}$. Moreover, $2(d-1)^{d-1} \leq d^d$. Notice that $2d \leq d^2$ and $4 \leq d^2$. Hence

$5d^2 + 2d + 4 \leq 7d^2 \leq d^5$ since $7 \leq d^3$. We deduce that r is bounded by d^{d^2+d+3} . As $d^2 + d + 3 \leq (d + 2)^2$, we get $r \leq d^{(d+2)^2}$.

We have proved Theorem 4.3 just by observing that $\|\rho_L\|_\infty \leq |\rho_L|$ and $\|T\|_\infty \leq \text{width}(\xrightarrow{*})$.

9 CONCLUSION

This paper introduces protocols that allow agents destructions/creations and leaders. Our definition of stably computing is a straight-forward extension of the one introduced by Dana Angluin, James Aspnes, and David Eisenstat in [3].

We provided in this paper state complexity lower-bounds of the form $\Omega(\log \log(n)^h)$ for any $h < \frac{1}{2}$ for protocols stably computing the counting predicates when the number of leaders and the interaction-width are bounded. This lower-bound almost matches the upper-bound $O(\log \log(n))$ introduced in [6] by Blondin, Esparza, and Jaax. We left as open the exact asymptotic state complexity.

Notice that for leaderless protocols, the state complexity is still open since there is an exponential gap between the upper-bound $O(\log(n))$ given in [6] and the lower-bound introduced in this paper.

ACKNOWLEDGMENTS

The author is supported by the grant ANR-17-CE40-0028 of the French National Research Agency ANR (project BRAVAS)

REFERENCES

- [1] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. 2004. Computation in networks of passively mobile finite-state sensors. In *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Distributed Computing, PODC 2004, St. John's, Newfoundland, Canada, July 25-28, 2004*, Soma Chaudhuri and Shay Kutten (Eds.). ACM, 290–299. <https://doi.org/10.1145/1011767.1011810>
- [2] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. 2006. Computation in networks of passively mobile finite-state sensors. *Distributed Comput.* 18, 4 (2006), 235–253. <https://doi.org/10.1007/s00446-005-0138-3>
- [3] Dana Angluin, James Aspnes, and David Eisenstat. 2006. Stably computable predicates are semilinear. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing, PODC 2006, Denver, CO, USA, July 23-26, 2006*, Eric Ruppert and Dahlia Malkhi (Eds.). ACM, 292–299. <https://doi.org/10.1145/1146381.1146425>
- [4] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. 2007. The computational power of population protocols. *Distributed Comput.* 20, 4 (2007), 279–304. <https://doi.org/10.1007/s00446-007-0040-2>
- [5] Michael Blondin, Javier Esparza, Blaise Genest, Martin Helfrich, and Stefan Jaax. 2020. Succinct Population Protocols for Presburger Arithmetic. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France (LIPIcs, Vol. 154)*, Christophe Paul and Markus Bläser (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 40:1–40:15. <https://doi.org/10.4230/LIPIcs.STACS.2020.40>
- [6] Michael Blondin, Javier Esparza, and Stefan Jaax. 2018. Large Flocks of Small Birds: on the Minimal Size of Population Protocols. In *35th Symposium on Theoretical Aspects of Computer Science, STACS 2018, February 28 to March 3, 2018, Caen, France (LIPIcs, Vol. 96)*, Rolf Niedermeier and Brigitte Vallée (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 16:1–16:14. <https://doi.org/10.4230/LIPIcs.STACS.2018.16>
- [7] Philipp Czermer and Javier Esparza. 2021. Lower Bounds on the State Complexity of Population Protocols. In *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen (Eds.). ACM, 45–54. <https://doi.org/10.1145/3465084.3467912>
- [8] Wojciech Czerwinski and Lukasz Orlikowski. 2021. Reachability in Vector Addition Systems is Ackermann-complete. *CoRR* abs/2104.13866 (2021). arXiv:2104.13866 <https://arxiv.org/abs/2104.13866>
- [9] Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. 2015. Verification of Population Protocols. In *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, September 14, 2015 (LIPIcs, Vol. 42)*, Luca Aceto and David de Frutos-Escrig (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 470–482. <https://doi.org/10.4230/LIPIcs.CONCUR.2015.470>
- [10] Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. 2017. Verification of population protocols. *Acta Informatica* 54, 2 (2017), 191–215. <https://doi.org/10.1007/s00236-016-0272-3>

- [11] Jérôme Leroux. 2021. The Reachability Problem for Petri Nets is Not Primitive Recursive. *CoRR* abs/2104.12695 (2021). arXiv:2104.12695 <https://arxiv.org/abs/2104.12695>
- [12] Loic Pottier. 1991. Minimal Solutions of Linear Diophantine Systems: Bounds and Algorithms. In *Proceedings of the 4th International Conference on Rewriting Techniques and Applications (RTA '91)*. Springer-Verlag, London, UK, UK, 162–173. <http://dl.acm.org/citation.cfm?id=647192.720494>
- [13] Charles Rackoff. 1978. The Covering and Boundedness Problems for Vector Addition Systems. *Theor. Comput. Sci.* 6 (1978), 223–231. [https://doi.org/10.1016/0304-3975\(78\)90036-1](https://doi.org/10.1016/0304-3975(78)90036-1)