



Verification of Population Protocols

Javier Esparza, Pierre Ganty, Jérôme Leroux, Rupak Majumdar

► To cite this version:

Javier Esparza, Pierre Ganty, Jérôme Leroux, Rupak Majumdar. Verification of Population Protocols. 26th International Conference on Concurrency Theory, CONCUR 2015, Sep 2015, Madrid, Spain. pp.470-482, <10.4230/LIPIcs.CONCUR.2015.470>. <hal-02391875>

HAL Id: hal-02391875

<https://hal.science/hal-02391875v1>

Submitted on 18 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 CC BY 4.0 - Attribution - International License



Verification of population protocols

Javier Esparza, Pierre Ganty, Jérôme Leroux, Rupak Majumdar

► To cite this version:

Javier Esparza, Pierre Ganty, Jérôme Leroux, Rupak Majumdar. Verification of population protocols. Acta Informatica, 2017, 54 (2), pp.191-215. 10.1007/s00236-016-0272-3 . hal-02391804

HAL Id: hal-02391804

<https://hal.science/hal-02391804v1>

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.

Verification of Population Protocols

Javier Esparza¹, Pierre Ganty², Jérôme Leroux³, and
Rupak Majumdar⁴

¹ TUM, Germany

² IMDEA Software Institute, Spain

³ LaBRI, CNRS & Université Bordeaux, France

⁴ MPI-SWS, Germany

Abstract

Population protocols (Angluin et al., PODC, 2004) are a formal model of sensor networks consisting of identical mobile devices. Two devices can interact and thereby change their states. Computations are infinite sequences of interactions satisfying a strong fairness constraint.

A population protocol is well-specified if for every initial configuration C of devices, and every computation starting at C , all devices eventually agree on a consensus value depending only on C . If a protocol is well-specified, then it is said to compute the predicate that assigns to each initial configuration its consensus value.

While the predicates computable by well-specified protocols have been extensively studied, the two basic verification problems remain open: is a given protocol well-specified? Does a protocol compute a given predicate? We prove that both problems are decidable. Our results also prove decidability of a natural question about home spaces of Petri nets.

1998 ACM Subject Classification C.2.2 Network Protocols, D.2.4 Software/Program Verification, F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases population protocols, Petri nets, parametrized verification

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2015.470

1 Introduction

Population protocols [2] are a model of distributed computation by anonymous, interacting finite-state agents. In each step, a fixed number of agents are chosen nondeterministically, and the agents interact and update their states according to a joint transition function. A population protocol is said to compute a predicate on the initial states of the agents if, in all fair executions, all agents eventually converge to the correct value of the predicate. An execution is fair if it is finite and cannot be extended, or it is infinite and every configuration of agent states that is reachable at infinitely many positions along the execution is also reached infinitely often along that execution.

The original motivation for population protocols was to model distributed computation in passively mobile sensors [2], but the model captures the essence of distributed computation in diverse areas such as trust propagation [7] and chemical reactions [15].

Much of the work on population protocols has concentrated on characterizing what predicates on the input values can be computed by *well-specified* protocols. A protocol is well-specified if, on every input, every fair execution eventually converges to configurations in which every agent agrees on a consensus value that depends only on the input. Angluin et al. [2] gave explicit well-specified protocols to compute every predicate definable in Presburger arithmetic. Later, Angluin et al. [4] showed that well-specified population protocols compute exactly the Presburger-definable predicates.



© Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar;
licensed under Creative Commons License CC-BY

26th International Conference on Concurrency Theory (CONCUR 2015).

Editors: Luca Aceto and David de Frutos Escrig; pp. 470–482

Leibniz International Proceedings in Informatics



LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Since it is easy to erroneously design protocols that are not well-specified, one can ask the natural verification question: given a population protocol, is it well-specified? In this paper, we show that the well-specification problem for population protocols is decidable. We also study the correctness problem: given a protocol and a Presburger specification, does the protocol compute the specification? Our techniques show decidability of the correctness problem as well.

The semantics of a population protocol is an infinite family of finite-state transition systems, one for each possible input. Whether the protocol reaches consensus for a given input can be decided by inspecting only one of these transition systems. However, the well-specification problem asks if consensus is reached for *all* inputs, and so it is not obviously decidable; indeed, similar questions are undecidable for many parameterized systems [5]. Moreover, the set of configurations where all agents agree on a value is not upward-closed; thus, coverability-like techniques are not immediately applicable.

Our main result is a characterization of well-specification using Presburger-definable predicates. We show that for every well-specified protocol, one can find a *witness* consisting of four Presburger-definable predicates $(\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1)$ and a bounded regular language W such that:

- each predicate is inductive (closed under taking a step of the protocol),
- each initial state is either in \mathcal{S}_0 or in \mathcal{S}_1 , but not in both,
- for $i \in \{0, 1\}$, all configurations of \mathcal{B}_i agree on the consensus value i ; moreover, \mathcal{B}_i is reachable from each configuration in \mathcal{S}_i using a string from W .

Using the decidability of Presburger arithmetic, we show that each condition above is decidable. Our proof of correctness uses recent results from the theory of Petri nets. We use the existence of Presburger-definable inductive sets that separate unreachable markings [11] to identify \mathcal{S}_0 and \mathcal{S}_1 . We use the Presburger-definability of the mutual reachability relation [12] to identify \mathcal{B}_0 and \mathcal{B}_1 . Finally, we use the theory of accelerations [14] to identify W . Along the way, we obtain an alternative proof of the theorem that well-specified protocols compute only Presburger-definable predicates.

Ultimately, our decision procedure consists of running two semi-decision procedures in parallel and does not provide a complexity upper bound. For lower bounds, we show that reachability for Petri nets can be reduced in polynomial-time to the complement of the well-specification problem.

While we focus on population protocols, our techniques also lead to new results for the theory of Petri nets. The *home space* problem asks, given a Petri net and two sets \mathcal{I} and \mathcal{H} of markings, if every marking reachable from \mathcal{I} can also reach \mathcal{H} . De Frutos and Johnen [8] showed that the home space problem is decidable if \mathcal{I} is a single marking and \mathcal{H} is a linear set. They left the case in which \mathcal{H} is a Presburger-definable set open. We make the first partial progress on this problem. Our results show that the home space problem is decidable for Presburger-definable sets \mathcal{I} and \mathcal{H} , provided the set of markings reachable from any marking in \mathcal{I} is finite.

The paper is organized as follows. Section 2 introduces population protocols. Section 3 formally defines witnesses of well-specification, shows decidability of the conditions to be met by a witness, and proves that existence of a witness implies well-specification. The proof of the converse (well-specification implies existence of a witness) is more involved. Section 4 introduces the results of Petri net theory needed for the proof, and Section 5 the proof itself. Section 6 reduces Petri net reachability to the complement of the well-specification problem. Finally, Section 7 proves the result about home spaces in Petri nets.

2 Population Protocols

A *population* on a finite set E is a mapping $P: E \rightarrow \mathbb{N}$ such that $P(e) > 0$ for some $e \in E$. Intuitively, $P(e)$ denotes the number of individuals of type $e \in E$ in the population. The set of all populations on E is denoted by $\text{Pop}(E)$. Operations on populations, like addition or maximum, are implicitly defined component wise. Given $e \in E$, we denote by \mathbf{e} the population consisting of one individual of type e , that is, the population satisfying $\mathbf{e}(e) = 1$ and $\mathbf{e}(e') = 0$ for every $e' \neq e$. The *support* of a population $P \in \mathbb{N}^E$, denoted by $\text{Sup}(P)$, is the subset of E given by $\{e \in E \mid P(e) > 0\}$. A set of populations $\mathcal{C} \subseteq \text{Pop}(E)$ is said to be *Presburger* if it can be denoted by a formula in *Presburger arithmetic*, i.e., in the first-order theory of addition $FO(\mathbb{N}, +)$.

► **Example 1.** Let $E = \{a, b\}$. The set of populations $\{P \in \text{Pop}(E) \mid P(a) \geq P(b)\}$ is Presburger, since it is denoted by the Presburger formula $F(X_a, X_b) = \exists Y: X_a = Y + X_b$. The set $\{P \in \text{Pop}(E) \mid P(a) = P(b)^2\}$ is not Presburger.

2.1 Protocol Scheme

A *protocol scheme* $\mathcal{A} = (Q, \Delta)$ consists of a finite non-empty set Q of states and a set $\Delta \subseteq Q^4$. If $(q_1, q_2, q'_1, q'_2) \in \Delta$, we write $(q_1, q_2) \mapsto (q'_1, q'_2)$ and call it a *transition*. The populations of $\text{Pop}(Q)$ are called *configurations*. Intuitively, a configuration C describes a collection of identical finite-state *agents* with Q as set of states, containing $C(q)$ agents in state q for every $q \in Q$. Pairs of agents interact using transitions from Δ .¹ Formally, given two configurations C and C' and a transition $\delta = (q_1, q_2) \mapsto (q'_1, q'_2)$, we write $C \xrightarrow{\delta} C'$ if

$$C \geq (\mathbf{q}_1 + \mathbf{q}_2) \text{ holds, and } C' = C - (\mathbf{q}_1 + \mathbf{q}_2) + (\mathbf{q}'_1 + \mathbf{q}'_2) .$$

We write $C \xrightarrow{w} C'$ for a word $w = \delta_1 \dots \delta_k$ of transitions if there exists a sequence C_0, \dots, C_k of configurations satisfying $C = C_0 \xrightarrow{\delta_1} C_1 \dots \xrightarrow{\delta_k} C_k = C'$. In this case, we say that C' is *reachable from* C . We also write $C \rightarrow C'$ if $C \xrightarrow{\delta} C'$ for some transition $\delta \in \Delta$. We have:

► **Lemma 2.** *For every configuration C , the set of configurations reachable from C is finite.*

Proof. Follows immediately from the fact that an interaction does not create or destroy agents, just changes their current states. Since Q is finite, there are only finitely many configurations C' satisfying $\sum_{q \in Q} C(q) = \sum_{q \in Q} C'(q)$. ◀

Observe that $(\text{Pop}(Q), \rightarrow)$ defines a directed graph with infinitely many vertices and edges. Consider the partition $\{\text{Pop}(Q)_i\}_{i \geq 1}$ of $\text{Pop}(Q)$, where $\text{Pop}(Q)_i = \{C \in \text{Pop}(Q) \mid \sum_{q \in Q} C(q) = i\}$. (Note that i starts at 1 because every population contains at least one agent.) Since interactions do not create or destroy agents, the set $\{\rightarrow_i\}_{i \geq 1}$, where $\rightarrow_i = \rightarrow \cap \text{Pop}(Q)_i^2$, is also a partition of \rightarrow . Therefore $(\text{Pop}(Q), \rightarrow)$ consists of the infinitely many disjoint and finite subgraphs $\{(\text{Pop}(Q)_i, \rightarrow_i)\}_{i \geq 1}$.

An *execution* of \mathcal{A} is a finite or infinite sequence of configurations C_0, C_1, \dots such that $C_i \rightarrow C_{i+1}$ for each $i \geq 0$. An execution is *fair* if it is finite and cannot be extended, or it is infinite and for every step $C \rightarrow C'$, if C occurs infinitely often along the execution, then C' also occurs infinitely often. It follows from Lemma 2 that every execution reaches a strongly

¹ While protocol schemes model pairwise interactions only, one can model k -way interactions for a fixed $k > 2$ by adding additional states.

connected component (SCC) of $(\text{Pop}(Q), \rightarrow)$ and never leaves it. We deduce the following lemma, where a bottom SCC of $(\text{Pop}(Q), \rightarrow)$ is an SCC such that every edge of \rightarrow whose source is in the SCC also belongs to the SCC. (In particular, a single vertex with no outgoing transition forms a bottom SCC.)

► **Lemma 3.** *Every fair execution eventually reaches a bottom SCC of $(\text{Pop}(Q), \rightarrow)$.*

Proof. If the execution is finite, then, since it cannot be extended, its last configuration is a bottom SCC with one single vertex and no outgoing transitions. If the execution is infinite, then the fairness condition forces it to eventually leave every non-bottom SCC it enters. ◀

2.2 Computation by Population Protocols

We define what it means for a protocol scheme to compute a predicate $\Pi: \text{Pop}(\Sigma) \rightarrow \{0, 1\}$, where Σ is a non-empty, finite set of *inputs*.

An *initial mapping* of a protocol scheme $\mathcal{A} = (Q, \Delta)$ is a function $I: \text{Pop}(\Sigma) \rightarrow \text{Pop}(Q)$ that maps each input population X to a configuration of \mathcal{A} . The set of *initial configurations* is $\mathcal{I} = \{I(X) \mid X \in \text{Pop}(\Sigma)\}$. An initial mapping I is *Presburger* if the predicate $C = I(X)$, where $C \in \text{Pop}(Q)$ and $X \in \text{Pop}(\Sigma)$, is definable in Presburger arithmetic. An initial mapping I is *simple* if there exists a sequence $(q_\sigma)_{\sigma \in \Sigma}$ of states of Q satisfying

$$I(X) = \sum_{\sigma \in \Sigma} X(\sigma) \mathbf{q}_\sigma$$

for every input population X on Σ .

An *output mapping* of a protocol scheme $\mathcal{A} = (Q, \Delta)$ is a function $O: \text{Pop}(Q) \rightarrow \{0, \perp, 1\}$ that associates to each configuration C of \mathcal{A} an output value in $\{0, \perp, 1\}$. A population C on Q such that $O(C) = b$ for some $b \in \{0, \perp, 1\}$ is called a *b-population*. An output mapping O is *Presburger* if the predicate $O(C) = b$ where $C \in \text{Pop}(Q)$ and $b \in \{0, 1\}$ is definable in Presburger arithmetic. An output mapping O is *simple* if there exists a partition (Q_0, Q_1) of Q such that

$$O(C) = \begin{cases} 0 & \text{if } \text{Sup}(C) \subseteq Q_0 \\ 1 & \text{if } \text{Sup}(C) \subseteq Q_1 \\ \perp & \text{otherwise} \end{cases}$$

for every configuration C . Notice that O is well-defined because $\text{Sup}(C) \neq \emptyset$. An execution C_0, C_1, \dots *stabilizes to b* for a given $b \in \{0, \perp, 1\}$ if there exists $n \in \mathbb{N}$ such that $O(C_m) = b$ for every $m \geq n$ (if the execution is finite, then this means for every m between n and the length of the execution). So, intuitively, an execution stabilizes to b if from some moment on all agents stay within the subset of states with output b . Notice that there may be many different executions from a given configuration C_0 , each of which may stabilize to 0, 1, or \perp , or not stabilize at all.

Most papers only consider population protocols with simple initial and output mappings. We study the more general class of Presburger initial and output mappings. In our general setting, a *population protocol* is a triple $(\mathcal{A}, \mathbf{I}, \mathbf{O})$, where \mathcal{A} is a protocol scheme, $\mathbf{I}(X, C)$ is a formula in Presburger arithmetic denoting a Presburger initial mapping $C = I(X)$, and $\mathbf{O}(C, b)$ is a formula in Presburger arithmetic denoting a Presburger output mapping $O(C) = b$. This definition encompasses population protocols with leader [3]. In these protocols the initial configuration contains one agent, called the *leader*, occupying a distinguished initial

state q_l not initially occupied by any other agent. This corresponds to the initial mapping $I(X) = \mathbf{q}_l + \sum_{\sigma \in \Sigma} X(\sigma) \mathbf{q}_\sigma$ which is obviously Presburger.

A population protocol $(\mathcal{A}, \mathbf{I}, 0)$ is *well-specified* if for every input population $X \in \text{Pop}(\Sigma)$, every fair execution of \mathcal{A} starting at $I(X)$ stabilizes to the same value, and *ill-specified* otherwise. A population protocol $(\mathcal{A}, \mathbf{I}, 0)$ *computes* a predicate Π if every fair execution of \mathcal{A} starting at $I(X)$ stabilizes to $\Pi(X)$ for every $X \in \text{Pop}(\Sigma)$.

The *well-specification problem* asks if a given protocol $(\mathcal{A}, \mathbf{I}, 0)$ is well-specified. The *correctness problem* asks if a given population protocol $(\mathcal{A}, \mathbf{I}, 0)$ computes a given Presburger predicate Π . Note that the correctness problem does not assume $(\mathcal{A}, \mathbf{I}, 0)$ to be well-specified. Consequently, if $(\mathcal{A}, \mathbf{I}, 0)$ does not compute Π then either the population protocol is ill-specified; otherwise it stabilizes to b for some input $X \in \text{Pop}(\Sigma)$ such that $\Pi(X) = 1 - b$.

3 A Decidable Criterion for Well-Specification

In this paper, the well-specification problem is shown to be decidable thanks to a decidable criterion based on Presburger arithmetic. This criterion is defined as follows. Let $\mathcal{A} = (Q, \Delta)$ be a protocol scheme. A set \mathcal{C} of configurations of \mathcal{A} is said to be *inductive* if $C \in \mathcal{C}$ and $C \rightarrow C'$ implies $C' \in \mathcal{C}$. Given a language $W \subseteq \Delta^*$ and a set \mathcal{C} of configurations, we denote by $\text{pre}_{\mathcal{A}}(\mathcal{C}, W)$ the set of configurations C such that $C \xrightarrow{w} C'$ for some word $w \in W$ and some configuration $C' \in \mathcal{C}$.

► **Definition 4.** Let $\mathcal{A} = (Q, \Delta)$ be a protocol scheme. A *witness of well-specification* of the population protocol $(\mathcal{A}, \mathbf{I}, 0)$ is a tuple $(\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1, w_1, \dots, w_k)$, where $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$ are predicates in Presburger arithmetic denoting Presburger sets of configurations $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$, and w_1, \dots, w_k are words in Δ^* denoting the language $W = w_1^* \dots w_k^*$, such that:

- (1) $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$ are inductive.
- (2) The pair $(\mathcal{I}_0, \mathcal{I}_1)$, where $\mathcal{I}_0 = \mathcal{S}_0 \cap \mathcal{I}$ and $\mathcal{I}_1 = \mathcal{S}_1 \cap \mathcal{I}$, is a partition of \mathcal{I} .
- (3) \mathcal{B}_0 is a set of 0-populations and $\mathcal{S}_0 \subseteq \text{pre}_{\mathcal{A}}(\mathcal{B}_0, W)$.
- (4) \mathcal{B}_1 is a set of 1-populations and $\mathcal{S}_1 \subseteq \text{pre}_{\mathcal{A}}(\mathcal{B}_1, W)$.

► **Lemma 5.** *The set of witnesses of well-specification is recursive.*

Proof. Let $(\mathcal{A}, \mathbf{I}, 0)$ and $(\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1, w_1, \dots, w_k)$ be as in Definition 4. We show that conditions (1)–(4) can be effectively expressed in Presburger arithmetic. For (1), a set \mathcal{M} of configurations denoted by a predicate $\mathbf{M}(C)$ in Presburger arithmetic is inductive iff the following Presburger formula is valid:

$$\forall C, C': \mathbf{M}(C) \wedge C \rightarrow C' \Rightarrow \mathbf{M}(C') .$$

So the inductiveness of $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$ is expressible. For (2), $(\mathcal{I}_0, \mathcal{I}_1)$ is a partition of \mathcal{I} iff

$$\forall C: (\exists X: \mathbf{I}(X, C)) \Leftrightarrow ((\mathbf{I}_0(C) \wedge \neg \mathbf{I}_1(C)) \vee (\neg \mathbf{I}_0(C) \wedge \mathbf{I}_1(C)))$$

is valid, where $\mathbf{I}_b(C) = (\exists X: \mathbf{I}(X, C)) \wedge \mathbf{S}_b(C)$. For (3-4), \mathcal{B}_b is a set of b -populations iff

$$\forall C: \mathbf{B}_b(C) \Rightarrow \mathbf{0}(C, b)$$

is valid. It remains to express $\mathcal{S}_b \subseteq \text{pre}_{\mathcal{A}}(\mathcal{B}_b, W)$. Observe that for every word $w \in \Delta^*$, the relation $\xrightarrow{w^*}$ defined by $C \xrightarrow{w^*} C'$ if $C \xrightarrow{w^n} C'$ for some $n \in \mathbb{N}$ is effectively definable in Presburger arithmetic. (For $w = \delta$, where $\delta = (q_1, q_2) \mapsto (q'_1, q'_2)$, this follows easily from $C' = C - (\mathbf{q}_1 + \mathbf{q}_2) + (\mathbf{q}'_1 + \mathbf{q}'_2)$. For the general case, see [14].) So the inclusion holds iff

$$\forall C_0: (\mathbf{S}_b(C_0) \Rightarrow \exists C_1, \dots, C_k: C_0 \xrightarrow{w_1^*} C_1 \cdots \xrightarrow{w_k^*} C_k \wedge \mathbf{B}_b(C_k))$$

is valid. ◀

3.1 The Criterion is Sound

We show that every population protocol satisfying the criterion is well-specified.

► **Lemma 6.** *Every population protocol $(\mathcal{A}, \mathcal{I}, 0)$ admitting a witness $(\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1, w_1, \dots, w_k)$ of well-specification is well-specified. Moreover, in this case the population protocol computes the predicate $\Pi : \text{Pop}(\Sigma) \rightarrow \{0, 1\}$ defined by:*

$$\Pi(X) = \begin{cases} 0 & \text{if } \exists C : \mathcal{I}(X, C) \wedge \mathcal{S}_0(C) \\ 1 & \text{if } \exists C : \mathcal{I}(X, C) \wedge \mathcal{S}_1(C) \end{cases} .$$

Proof. Let $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$ be the Presburger sets of configurations denoted by $\mathcal{S}_0, \mathcal{S}_1, \mathcal{B}_0, \mathcal{B}_1$, respectively. Let $W = w_1^* \dots w_k^*$. Since \mathcal{I}_0 and \mathcal{I}_1 form a partition of \mathcal{I} , it suffices to prove that every fair execution starting at \mathcal{I}_b stabilizes to b . Let $C \in \mathcal{I}_b$ and let C_0, C_1, \dots be a fair execution starting at C . Lemma 3 shows that the execution ends up in a bottom SCC. Hence, there exists $n \in \mathbb{N}$ such that C_n is in a bottom SCC. As \mathcal{S}_b is inductive, it follows that C_n is in this set. Moreover, as $\mathcal{S}_b \subseteq \text{pre}_{\mathcal{A}}(\mathcal{B}_b, W)$, there exists a word $w \in W$ and a configuration $C' \in \mathcal{B}_b$ such that $C_n \xrightarrow{w} C'$. Since C_n is in a bottom SCC, there exists a word $w' \in \Delta^*$ such that $C' \xrightarrow{w'} C_n$. Now, let $m \geq n$. Since C_m is reachable from C_n , it follows that C_m is reachable from C' . As $C' \in \mathcal{B}_b$ and \mathcal{B}_b is inductive, it follows that $C_m \in \mathcal{B}_b$. As \mathcal{B}_b is a set of b -populations, it follows that $O(C_m) = b$; thus, the execution stabilizes to b . ◀

In the rest of the paper we prove the converse of Lemma 6: every well-specified protocol admits a witness of well-specification. But before, we close the section with an example.

3.2 Example

Let $\Sigma = \{\sigma\}$, and consider the predicate $\Pi : \text{Pop}(\Sigma) \rightarrow \{0, 1\}$, where $\Pi(X)$ is the parity of $X(\sigma)$. In other words, $\Pi(X) = 0$ if $X(\sigma)$ is even, and $\Pi(X) = 1$ otherwise. This predicate is computed by a simple well-specified population protocol. The protocol scheme $\mathcal{A} = (Q, \Delta)$ has $Q = \{A_0, A_1, P_0, P_1\}$ as set of states. We call agents in A_0 and A_1 *active*, and those in P_0 and P_1 *passive*. Further, we say that agents in A_b and P_b carry the value b . The set Δ of transitions is $\{\delta_{x,y}, \delta_x \mid x, y \in \{0, 1\}\}$. Transitions $\delta_{x,y}$ allow two active agents to add their numbers modulo 2 and deactivate one of them:

$$\delta_{x,y} = (A_x, A_y) \mapsto (A_{x+y}, P_{x+y}) .$$

Transitions δ_x allow an active agent to change the value of a passive agent:

$$\delta_x = (A_x, P_{1-x}) \mapsto (A_x, P_x) .$$

The simple population protocol computing Π is given by $(\mathcal{A}, \mathcal{I}, 0)$, where the simple input mapping is defined by

$$I(X) = X(\sigma)\mathbf{A}_0$$

and the simple output mapping by $Q_0 = \{A_0, P_0\}$ and $Q_1 = \{A_1, P_1\}$.

Let us provide a witness of well-specification explaining why the protocol computes Π . We choose $\mathcal{B}_0(C) = (C(A_0) = 1 \wedge C(A_1) = 0 \wedge C(P_1) = 0)$ and $\mathcal{B}_1(C) = (C(A_1) = 1 \wedge C(A_0) = 0 \wedge C(P_0) = 0)$. Notice that the set of configurations \mathcal{B}_b denoted by \mathcal{B}_b is inductive for every $b \in \{0, 1\}$. In fact, since a configuration $C \in \mathcal{B}_b$ only has one active agent, and all agents

carry the same value b , no transition in Δ is enabled at C . Further, we define $\mathcal{S}_0(C)$ as “ $C(A_1)$ is even” and $\mathcal{S}_1(C)$ as “ $C(A_1)$ is odd”. Inspection of the transitions in Δ immediately shows that the sets \mathcal{S}_0 and \mathcal{S}_1 denoted by these two Presburger predicates are inductive. Notice that $\mathcal{I} \cap \mathcal{S}_0$ and $\mathcal{I} \cap \mathcal{S}_1$ is a partition of \mathcal{I} .

It remains to define the language W . Let us first describe a strategy to reach $\mathcal{B}_0 \cup \mathcal{B}_1$ from any configuration C . We first execute the transition $\delta_{0,0}$ as long as possible, until there is at most one active agent carrying a 0. Then we execute $\delta_{1,1}$ as long as possible, until there is at most one active agent carrying a 1. Then we execute $\delta_{1,0}$ if possible, reaching a configuration with exactly one active agent carrying a value b . Finally, we execute δ_0 as long as possible, followed by δ_1 as long as possible, leading to a configuration in which every passive agent also carries the value b . The language W models this strategy:

$$W = \delta_{0,0}^* \delta_{1,1}^* \delta_{1,0}^* \delta_0^* \delta_1^* .$$

4 Petri Net Theory for the Population Protocols Aficionados

The computation of a population protocol can be simulated by an associated Petri net. This allows us to apply results on Petri nets to population protocols.

A Petri net $N = (P, T, F)$ consists of a finite set P of *places*, a finite set T of *transitions*, and a *flow function* $F: (P \times T) \cup (T \times P) \rightarrow \mathbb{N}$. A *marking* is a mapping from P to \mathbb{N} , i.e. a mapping in \mathbb{N}^P . A transition $t \in T$ is *enabled* at marking M , written $M[t]$, if $F(p, t) \leq M(p)$ for each place $p \in P$. A transition t that is enabled at M can *fire*, yielding a marking M' such that $M'(p) = M(p) - F(p, t) + F(t, p)$ for each $p \in P$. We write this fact as $M[t] M'$. We extend enabledness and firing inductively to words of transitions as follows. Let $w = t_1 \dots t_k$ be a finite word of transitions $t_j \in T$. We define $M[w] M'$ if, and only if, there exists a sequence M_0, \dots, M_k of markings such that $M = M_0[t_1] M_1 \dots [t_k] M_k = M'$. In that case, we say that M' is *reachable* from M .

4.1 From Population Protocols To Petri Nets

Given a protocol scheme $\mathcal{A} = (Q, \Delta)$, we define the Petri net $N(\mathcal{A}) = (Q, \Delta, F)$, whose places and transitions are the states and transitions of the protocol, respectively, and where F is defined for every transition $\delta = (q_1, q_2) \mapsto (q'_1, q'_2)$ in Δ and every state $q \in Q$ by $F(q, \delta) = \mathbf{q}_1(q) + \mathbf{q}_2(q)$ and $F(\delta, q) = \mathbf{q}'_1(q) + \mathbf{q}'_2(q)$. Note that a configuration of the protocol scheme \mathcal{A} is a marking of the Petri net $N(\mathcal{A})$. Further, whenever $C \xrightarrow{\delta} C'$ for configurations C and C' , we have $C[\delta] C'$ in the Petri net, and vice versa.

The correspondence between \mathcal{A} and $N(\mathcal{A})$ allows us to transfer results from Petri nets to population protocols. Next, we briefly recall the results we need.

4.2 Acceleration Technique

Given a Petri net $N = (P, T, F)$, a set \mathcal{M} of markings, and a language $W \subseteq T^*$, we introduce the sets:

$$\begin{aligned} post_N(\mathcal{M}, W) &= \{M' \in \mathbb{N}^P \mid \exists M \in \mathcal{M} \exists w \in W : M[w] M'\} \\ pre_N(\mathcal{M}, W) &= \{M \in \mathbb{N}^P \mid \exists M' \in \mathcal{M} \exists w \in W : M[w] M'\} . \end{aligned}$$

When $W = T^*$ these sets are denoted by $post_N^*(\mathcal{M})$ and $pre_N^*(\mathcal{M})$, respectively.

The theory of acceleration (see for instance [14]) will provide a simple way for extracting the language W introduced in Definition 4. A language $W \subseteq T^*$ is said to be *bounded* [10] if

there exists a sequence w_1, \dots, w_k of words in T^* such that $W \subseteq w_1^* \dots w_k^*$. The following result will be useful for extracting the language W introduced in Definition 4.

► **Theorem 7** ([14, Corollary XI.3]). *For every Petri net $N = (P, T, F)$ and for every Presburger sets of markings \mathcal{S} and \mathcal{B} such that $\mathcal{S} \subseteq \text{pre}_N^*(\mathcal{B})$, there exists a bounded language $W \subseteq T^*$ such that $\mathcal{S} \subseteq \text{pre}_N(\mathcal{B}, W)$.*

4.3 Separators For Reachability Problems

Recently, the reachability problem for Petri nets was proved to be decidable using a very simple algorithm based on Presburger inductive sets of markings. Let us recall that a set \mathcal{M} of markings is *inductive* for a Petri net $N = (P, T, F)$, if $\text{post}_N(\mathcal{M}, T) \subseteq \mathcal{M}$. The following result will provide the sets \mathcal{S}_0 and \mathcal{S}_1 introduced in Definition 4.

► **Theorem 8** ([13, Lemma 9.1]). *For every Petri net N and for every Presburger set of markings \mathcal{M} and \mathcal{M}' such that $\text{post}_N^*(\mathcal{M}) \cap \mathcal{M}' = \emptyset$, there exists a Presburger inductive set of markings \mathcal{S} for N such that $\mathcal{M} \subseteq \mathcal{S}$ and $\mathcal{S} \cap \mathcal{M}' = \emptyset$.*

4.4 Mutual Reachability Relations

The *mutual reachability relation* of a Petri net N is the binary relation over the markings that contains the pair (M, M') if M' is reachable from M and M is reachable from M' . Intuitively, M and M' coincide and otherwise they are in the same SCC for the reachability graph. The following theorem will be useful for extracting the sets \mathcal{B}_0 and \mathcal{B}_1 introduced in Definition 4.

► **Theorem 9** ([12]). *For every Petri net N , the mutual reachability relation is effectively definable in Presburger arithmetic.*

4.5 Decomposable sets

In this section, we introduce a new result for Petri nets. This result will be used for characterizing the sets \mathcal{I}_0 and \mathcal{I}_1 introduced in Definition 4. The proof of this result is based on the geometrical characterization of the reachability sets of Petri nets based on almost semi-linear sets and decomposable sets (see [12] for definitions). It uses the technical result that if the union of two disjoint decomposable sets \mathbf{X}, \mathbf{Y} is Presburger definable, then both \mathbf{X} and \mathbf{Y} are Presburger definable as well. We defer the details to the full version of the paper.

► **Theorem 10.** *For every Petri net N , and for every Presburger sets of markings $\mathcal{B}_0, \mathcal{B}_1$ and \mathcal{I} such that $\mathcal{I}_0 = \mathcal{I} \cap \text{pre}_N^*(\mathcal{B}_0)$ and $\mathcal{I}_1 = \mathcal{I} \cap \text{pre}_N^*(\mathcal{B}_1)$ is a partition of \mathcal{I} , it follows that \mathcal{I}_0 and \mathcal{I}_1 are Presburger.*

5 The Criterion is Complete

We use the previous results to prove that every well-specified protocol admits a witness.

5.1 Characterization of Bottom Strongly Connected Components

Given a protocol scheme $\mathcal{A} = (Q, \Delta)$, a bottom SCC of the graph $(\text{Pop}(Q), \rightarrow)$ of \mathcal{A} is said to be *b-bottom* ($b \in \{0, 1\}$) if all its configurations, which are called bottom configurations, are *b-populations*. When this holds, the configurations of the SCC are called *b-bottom configurations*. We denote the sets of bottom configurations and *b-bottom configurations* by \mathcal{B} and \mathcal{B}_b , respectively.

► **Proposition 11.** *Given a protocol scheme, the sets \mathcal{B} , \mathcal{B}_0 and \mathcal{B}_1 are effectively Presburger.*

Proof. We show that the predicate $B(C)$ associated to the set of bottom configurations is definable in Presburger arithmetic. Let us introduce the predicate $MR(C, C')$ associated to the mutual reachability relation. Theorem 9 shows that $MR(C, C')$ is effectively Presburger. Now, we just observe that C is a bottom configuration iff for every configuration C' such that C and C' are mutually reachable and for every C'' such that $C' \rightarrow C''$, we have C and C'' are also mutually reachable:

$$B(C) = \forall C' \forall C'' : (MR(C, C') \wedge C' \rightarrow C'' \Rightarrow MR(C, C'')) .$$

We claim that \mathcal{B}_b is a Presburger set of configurations. To prove this, we just notice that \mathcal{B}_b is denoted by the following formula:

$$B_b(C) = B(C) \wedge \forall C' : MR(C, C') \Rightarrow 0(C', b) . \quad \blacktriangleleft$$

5.2 The final piece

In the rest of this section, we show that a population protocol is well-specified if, and only if, it admits a witness of well-specification. We deduce from this characterization that the well-specification problem, and the correctness problem are decidable.

► **Theorem 12.** *A population protocol is well-specified iff it admits a witness of well-specification.*

Proof. Lemma 6 shows that a population protocol that admits a witness of well-specification is well-specified. Conversely, let us consider a population protocol $(\mathcal{A}, I, 0)$ that is well-specified. We define \mathcal{B}_0 and \mathcal{B}_1 as the 0-bottom configurations and 1-bottom configurations, respectively. Proposition 11 shows that these sets are Presburger. Notice these two sets are also inductive.

Let us show that $(\mathcal{I}_0, \mathcal{I}_1)$ defined by $\mathcal{I}_b = I \cap pre_{\mathcal{A}}^*(\mathcal{B}_b)$ is a partition of I . Since the population protocol is well specified, it follows that $\mathcal{I}_0 \cap \mathcal{I}_1 = \emptyset$. Now let C be an initial configuration in I . Notice that there exists at least one fair execution C_0, C_1, \dots with $C_0 = C$ that stabilizes to b . Lemma 3 shows that the execution ends up in a bottom SCC. It follows that there exists $n \in \mathbb{N}$ such that C_n is a bottom configuration. Thanks to the fairness of the execution, all the configurations of the strongly connected component of C_n are b -populations. Thus $C_n \in \mathcal{B}_b$. We have proved that $C \in \mathcal{I}_b$. Thus $(\mathcal{I}_0, \mathcal{I}_1)$ is a partition of I . Following Section 4.1, define $N(\mathcal{A})$ as the Petri net associated with \mathcal{A} . From Theorem 10, we derive that \mathcal{I}_0 and \mathcal{I}_1 are Presburger.

Since the population protocol is well-specified, it follows that $post_{\mathcal{A}}^*(\mathcal{I}_0) \cap (\mathcal{B} \setminus \mathcal{B}_0)$ is empty. Hence $post_{N(\mathcal{A})}^*(\mathcal{I}_0) \cap (\mathcal{B} \setminus \mathcal{B}_0)$ is also empty and Theorem 8 shows that there exists a Presburger inductive set of markings and, by extension, configurations \mathcal{S}_0 such that $\mathcal{I}_0 \subseteq \mathcal{S}_0$ and such that $\mathcal{S}_0 \cap (\mathcal{B} \setminus \mathcal{B}_0)$ is empty. Let us prove that $\mathcal{S}_0 \subseteq pre_{\mathcal{A}}^*(\mathcal{B}_0)$. Let C be a configuration in \mathcal{S}_0 and let us consider a fair execution C_0, C_1, \dots starting from $C_0 = C$. Lemma 3 shows that the execution ends up in a bottom SCC. It follows that there exists $n \in \mathbb{N}$ such that C_n is a bottom configuration. As \mathcal{S}_0 is inductive, it holds that $C_n \in \mathcal{S}_0$. Moreover, as $\mathcal{S}_0 \cap (\mathcal{B} \setminus \mathcal{B}_0)$ is empty, we derive $C_n \in \mathcal{B}_0$. We have proved that $\mathcal{S}_0 \subseteq pre_{\mathcal{A}}^*(\mathcal{B}_0)$. Theorem 7 shows that there exists a bounded language $W_0 \subseteq \Delta^*$ such that $\mathcal{S}_0 \subseteq pre_{N(\mathcal{A})}(\mathcal{B}_0, W_0)$, hence the same holds for \mathcal{A} . Symmetrically, there exists a Presburger inductive set of configurations \mathcal{S}_1 such that $\mathcal{I}_1 \subseteq \mathcal{S}_1$ and a bounded language $W_1 \subseteq \Delta^*$ such that $\mathcal{S}_1 \subseteq pre_{\mathcal{A}}(\mathcal{B}_1, W_1)$. Since W_0 and W_1 are bounded languages, it follows that $W = W_0 \cup W_1$ is also a bounded language.

Hence, there exists a sequence of words w_1, \dots, w_k in Δ^* such that $W \subseteq w_1^* \dots w_k^*$. We have proved that $(S_0, S_1, B_0, B_1, w_1, \dots, w_k)$ is a witness of well-specification. \blacktriangleleft

That well-specified population protocols can compute Presburger predicates was shown by Angluin et al. [2] using a direct construction. Showing that well-specified population protocols can not compute anything else than Presburger predicates was harder, and first proved by Angluin, Aspnes and Eisenstat [4]. Our constructions provide an alternate proof.

► **Corollary 13.** *The well-specification problem and the correctness problem are decidable. Moreover, well-specified population protocols compute Presburger predicates, and we can effectively compute formulas in Presburger arithmetic denoting the predicates computed by well-specified population protocols.*

Proof. Notice that if a population protocol is ill-specified there exists a witness of this property given by an initial input population X in $\text{Pop}(\Sigma)$ and a configuration C satisfying $I(X, C)$ such that not all the bottom configurations reachable from C are b -populations for some $b \in \{0, 1\}$.

In particular, enumerating the finite graphs $(\text{Pop}(Q)_i, \rightarrow_i)$, and checking, for each, whether it contains a witness of ill-specification shows that the problem of deciding if a population protocol is ill-specified is recursively enumerable.

By Theorem 12, when a population protocol is well-specified, the algorithm that enumerates all the tuples (S_0, S_1, B_0, B_1) of predicates in Presburger arithmetic, and all the finite sequences w_1, \dots, w_k of words in Δ^* and checks using Lemma 5 that we have a witness of well-specification, will eventually terminate with such a witness. It follows that the well-specification problem is recursively enumerable. Moreover, in that case, from the computed witness, we derive a predicate in Presburger arithmetic denoting the computed predicate Π using Lemma 6. Together with the recursive enumerability of ill-specification above, it follows that the problem is decidable. \blacktriangleleft

Clement et al. [6] proved the decidability of the well-specification problem when the number of agents is fixed. Corollary 13 shows decidability of the same problem but for an arbitrary number of agents.

6 Lower Bounds

Finally, we show hardness for the well-specification problem by showing a polynomial-time reduction from Petri net reachability to its complement.

► **Theorem 14.** *The reachability problem for Petri nets is polynomially reducible to the complement of the well-specification problem and the complement of the correctness problem for population protocols (even with simple output mappings).*

Proof. We proceed by means of a sequence of polynomial time reductions so as to reduce the reachability problem for Petri nets to the problem of reaching, in a Petri Net $N = (P, T, F)$ with initial marking M_0 , a marking M with no tokens in $z \in P$, i.e. $M(z) = 0$. Furthermore, the reduction is such that:

- (a) $M_0(z) > 0$,
- (b) N is deadlock-free, and
- (c) every transition of N has at least one output place, and at most two input and two output places.
- (d) N contains no two transitions with the same set of input and output places

The details of the reductions are standard and omitted.

Then we construct a population protocol $(\mathcal{A}, \mathbf{I}, \mathbf{0})$ with semi-linear initial mapping. We first describe the protocol scheme $\mathcal{A} = (Q, \Delta)$. The set Q of states of the protocol contains

- a state q_p for every place $p \in P$;
- a state q_t for every transition $t \in T$; and
- two states *Source* and *Sink*.

Following (d), we write $t = (P_1, P_2)$ to denote that transition t has P_1 as set of input places and P_2 as set of output places. The set Δ of transitions contains

- (1) for every Petri net transition $t = (\{p_1, p_2\}, \{p_3, p_4\})$, two protocol transitions $(q_{p_1}, q_{p_2}) \mapsto (q_t, \text{Sink})$ and $(q_t, \text{Source}) \mapsto (q_{p_3}, q_{p_4})$;
- (2) for every Petri net transition $t = (\{p_1, p_2\}, \{p_3\})$, two protocol transitions $(q_{p_1}, q_{p_2}) \mapsto (q_t, \text{Sink})$ and $(q_t, \text{Source}) \mapsto (q_{p_3}, \text{Sink})$;
- (3) for every Petri net transition $t = (\{p_1\}, \{p_2, p_3\})$, one protocol transition $(q_{p_1}, \text{Source}) \mapsto (q_{p_2}, q_{p_3})$; and
- (4) for every Petri net transition $t = (\{p_1\}, \{p_2\})$, one protocol transition $(q_{p_1}, \text{Source}) \mapsto (q_{p_2}, \text{Sink})$;
- (5) a transition $(q_p, q_z) \mapsto (\text{Sink}, q_z)$ for each place $p \neq z$.

This completes the description of \mathcal{A} .

The output mapping O , which is simple, is given by the partition Q_0, Q_1 of Q such that $Q_0 = \{z, \text{Sink}\}$. The initial mapping $I: \text{Pop}(\Sigma) \rightarrow \text{Pop}(Q)$ is defined as follows. The set Σ is a singleton $\{\sigma\}$, and I assigns to the number n – a population of $\text{Pop}(\{\sigma\})$ – the configuration that puts

- n agents in *Source*;
- $M_0(p)$ agents in q_p for every place p ; and
- 0 agents elsewhere.

Observe that I is a semi-linear mapping.

The transitions of (1)–(4) simulate the firing of t (in the case of (1) and (2), firing t is simulated by the occurrence, one after the other, of two protocol transitions). In all cases, simulating the firing of t requires one agent to leave the *Source* state. On the other hand, no agents ever enter *Source*. Hence each execution of $(\mathcal{A}, \mathbf{I}, \mathbf{0})$ contains only finitely many occurrences of transitions of (1)–(4). Further, since every transition of (5) moves an agent to *Sink*, and no agents ever leave *Sink*, the transitions of (5) also occur only finitely often. Therefore all executions of $(\mathcal{A}, \mathbf{I}, \mathbf{0})$ are finite.

Assume that some reachable marking M of N satisfies $M(z) = 0$. Let $\tau \in T^*$ be such that $M_0[\tau] M$, and let k be the length of τ . Since $M_0(z) > 0$, we have $k > 0$. We claim that \mathcal{A} has a fair (finite) execution from $I(k\sigma)$ that does not stabilize. Consider the execution that starts by simulating τ through transitions (1)–(4). At the end of this simulation the protocol reaches a configuration C such that $C(\text{Source}) = C(q_z) = 0$ and $C(\text{Sink}) > 0$. Observe that C cannot be extended because $C(\text{Source}) = 0$ disables all transitions (1)–(4) and $C(q_z) = 0$ disables all transitions (5). Further, since every transition has at least one output place, the configuration satisfies $C(q_p) > 0$ for some $p \neq z$. Since $\text{Sink} \in Q_0$ and $\{q_p \mid p \in P\} \subseteq Q_1$, we have that $O(C) = \perp$, hence that $(\mathcal{A}, \mathbf{I}, \mathbf{0})$ is ill-specified.

Assume now that every reachable marking M of N satisfies $M(z) > 0$. Let $C_0 C_1 \dots$ be an arbitrary fair execution of $(\mathcal{A}, \mathbf{I}, \mathbf{0})$. As shown above, there is a configuration C_j such that from that moment on C_j disable all transitions (1)–(4). In particular since N is deadlock-free, we necessarily have $C_j(\text{Source}) = 0$. Because some transitions of (5) might be enabled at C_j , we extend the execution by firing them as many times as possible. This can occur only finitely many times and yield a configuration C_ℓ which cannot be extended further – all

transitions of (1)–(5) are disabled – and in which all agents are in state q_z or *Sink*. We thus find that $\text{Sup}(C_\ell) \subseteq Q_0$, hence that $O(C_\ell) = 0$ and finally that $C_0 \dots C_\ell$ is a fair execution that converges to 0. Since we picked an arbitrary fair execution we conclude that every fair execution stabilizes to 0, and therefore $(\mathcal{A}, \mathbf{I}, 0)$ is well-specified.

The same reduction shows hardness for the complement of the correctness problem for the predicate *false*. \blacktriangleleft

7 Home Spaces

As a byproduct of our main result, we present a new theorem on home spaces of Petri nets. Let N be a Petri net, and let \mathcal{I}, \mathcal{H} be two sets of markings of N . We say that \mathcal{H} is a *home space* of N with respect to \mathcal{I} if $\text{post}_N^*(\mathcal{I}) \subseteq \text{pre}_N^*(\mathcal{H})$, that is, if \mathcal{H} can be reached from any marking reachable from \mathcal{I} . The *home space problem* for a given triple $(N, \mathcal{I}, \mathcal{H})$ asks whether \mathcal{H} is a home space of N with respect to \mathcal{I} .

De Frutos and Johnen [8] have proved that the home space problem is decidable when \mathcal{I} is a singleton and \mathcal{H} is a *linear set*, that is, a set of the form $\{M_0 + n_1M_1 + \dots + n_kM_k \mid n_1, \dots, n_k \in \mathbb{N}\}$ for a given *root marking* M_0 and a given finite set $\{M_1, \dots, M_k\}$ of *periods*. They also extend the result to finite unions of linear sets *having the same periods*. While every such set is a Presburger set, the converse does not hold, and De Frutos and Johnen [8] explicitly leave the case of arbitrary Presburger sets \mathcal{H} open.

We prove decidability of the home space problem for triples $(N, \mathcal{I}, \mathcal{H})$ where \mathcal{I} and \mathcal{H} are arbitrary Presburger sets, and the net N satisfies the following condition: for every marking $M_0 \in \mathcal{I}$, the set $\text{post}_N^*(\{M_0\})$ is finite. Observe that this condition is met by Petri nets modelling parameterized systems, as in the many-process systems of German and Sistla [9, 1]. Indeed, in these systems each token of $M_0 \in \mathcal{I}$ models a finite-state process, and, since the systems have no dynamic process creation, the number of tokens does not change while the net evolves. So, while our result does not close the open problem left by De Frutos and Johnen, it provides a partial answer, and the first new result in the area since 1989.

If $\text{post}_N^*(\{M_0\})$ is finite for every $M_0 \in \mathcal{I}$, then each reachable marking can reach a bottom SCC. This is the fact we exploit. Notice that this fact no longer holds for arbitrary Petri nets. For instance, it is easy to exhibit a Petri net whose reachability graph is an infinite line, and so has no bottom SCC.

► **Lemma 15.** *Let N be a net, and let \mathcal{B} be the set of bottom markings of N , i.e., the set of markings M that are reachable from any marking reachable from M . Let \mathcal{I} be a set of markings of N such that $\text{post}_N^*(\{M_0\})$ is finite for every $M_0 \in \mathcal{I}$. A set \mathcal{H} is a home space of N with respect to \mathcal{I} iff $\mathcal{B} \setminus \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is not reachable from \mathcal{I} .*

Proof. (\Rightarrow): Assume that some marking $M \in \mathcal{B} \setminus \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is reachable from some marking of \mathcal{I} . We claim that $\text{post}_N^*(\{M\}) \cap \mathcal{H} = \emptyset$, which implies that \mathcal{H} is not a home space. Let $M' \in \text{post}_N^*(\{M\})$. By the definition of \mathcal{B} , the markings M and M' are mutually reachable, and so $M \in \text{post}_N^*(\{M'\})$. If $M' \in \mathcal{H}$, then $M' \in \mathcal{B} \cap \mathcal{H}$, and so $M \in \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$, contradicting the hypothesis. So $M' \notin \mathcal{H}$, and we are done.

(\Leftarrow): Assume \mathcal{H} is not a home space. Then there exists a marking $M \in \text{post}_N^*(\mathcal{I})$ such that $\text{post}_N^*(\{M\}) \cap \mathcal{H} = \emptyset$. Let $M_0 \in \mathcal{I}$ be a marking such that $M \in \text{post}_N^*(\{M_0\})$. By hypothesis $\text{post}_N^*(\{M_0\})$ is finite, and so, since M is reachable from M_0 , some marking $M' \in \text{post}_N^*(\mathcal{I}) \cap \mathcal{B}$ is reachable from M . We prove that $M' \in \mathcal{B} \setminus \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$. Since $M' \in \mathcal{B}$, it suffices to prove $M' \notin \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$. Assume M' is reachable from some $M'' \in \mathcal{B} \cap \mathcal{H}$, hence $M'' \neq M'$. By the definition of \mathcal{B} , the markings M' and M'' are mutually reachable, and so M'' is also reachable from M . But, since $M'' \in \mathcal{H}$, then some marking of \mathcal{H} is reachable from M , contradicting the definition of M . \blacktriangleleft

► **Theorem 16.** *The home space problem is decidable for triples $(N, \mathcal{I}, \mathcal{H})$ where*

1. \mathcal{I} and \mathcal{H} are arbitrary Presburger sets of markings, and
2. $\text{post}_N^*(\{M_0\})$ is finite for every $M_0 \in \mathcal{I}$.

Proof. By Lemma 15, it suffices to decide whether $\mathcal{B} \setminus \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is reachable from \mathcal{I} . We show that $\mathcal{B} \setminus \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is an effectively Presburger set, and then apply the decidability of the reachability problem for Presburger sets of markings (i.e., given two Presburger sets $\mathcal{P}_1, \mathcal{P}_2$, decide if some marking of \mathcal{P}_2 is reachable from some marking of \mathcal{P}_1 .)

By Theorem 9 and Proposition 11, the set \mathcal{B} of bottom markings of N is effectively Presburger. So, since Presburger sets are effectively closed under boolean operations, it suffices to show that $\text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is effectively Presburger. Observe first that, since \mathcal{H} is effectively Presburger, so is $\mathcal{B} \cap \mathcal{H}$. By the definition of \mathcal{B} , if $M' \in \text{post}_N^*(M)$ for some $M \in \mathcal{B} \cap \mathcal{H}$, then $M \in \text{post}_N^*(M')$. So $M \in \text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ iff there is a marking $M' \in \mathcal{B} \cap \mathcal{H}$ such that M and M' are mutually reachable. Since the mutual reachability relation of N is effectively Presburger, $\text{post}_N^*(\mathcal{B} \cap \mathcal{H})$ is effectively Presburger. ◀

References

- 1 Benjamin Aminof, Tomer Kotek, Sasha Rubin, Francesco Spegni, and Helmut Veith. Parameterized model checking of rendezvous systems. In *CONCUR'14*, volume 8704 of *LNCS*, pages 109–124. Springer, 2014.
- 2 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. In *PODC'04*, pages 290–299. ACM, 2004.
- 3 Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. In *DISC'06*, volume 4167 of *LNCS*, pages 61–75. Springer, 2006.
- 4 Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *PODC'06*, pages 292–299. ACM, 2006.
- 5 Krzysztof R. Apt and Dexter C. Kozen. Limits for automatic verification of finite-state concurrent systems. *Information Processing Letters*, 22(6):307–309, 1986.
- 6 J. Clement, C. Delporte-Gallet, H. Fauconnier, and M. Sighireanu. Guidelines for the verification of population protocols. In *ICDCS'11*, pages 215–224, 2011.
- 7 Z. Diamadi and Michael J. Fischer. A simple game for the study of trust in distributed systems. *Wuhan University Journal of Natural Sciences*, 6(1–2):72–82, 2001.
- 8 David Frutos-Escrig and C. Johnen. Decidability of home space property. Technical Report 503, LRI, Université de Paris-Sud. Centre d'Orsay., 1989.
- 9 Steven M. German and A. Prasad Sistla. Reasoning about systems with many processes. *Journal of ACM*, 39(3):675–735, 1992.
- 10 Seymour Ginsburg. *The Mathematical Theory of Context-Free Languages*. McGraw-Hill, Inc., New York, NY, USA, 1966.
- 11 Jérôme Leroux. The general vector addition system reachability problem by presburger inductive invariants. In *LICS'09*, pages 4–13. IEEE Computer Society, 2009.
- 12 Jérôme Leroux. Vector addition system reversible reachability problem. In *CONCUR'11*, volume 6901 of *LNCS*, pages 327–341. Springer, 2011.
- 13 Jérôme Leroux. Vector addition systems reachability problem (a simpler solution). In *Turing-100: The Alan Turing Centenary Conference*, volume 10 of *EPiC Series*, pages 214–228. EasyChair, 2012.
- 14 Jérôme Leroux. Presburger vector addition systems. In *LICS'13*, pages 23–32. IEEE Computer Society, 2013.
- 15 Saket Navlakha and Ziv Bar-Joseph. Distributed information processing in biological and computational systems. *Commun. ACM*, 58(1):94–102, December 2014.