Deciding the existence of cut-off in parameterized rendez-vous networks

Florian Horn

IRIF, CNRS, Universite de Paris, France
florian.horn@irif.fr

Arnaud Sangnier

IRIF, CNRS, Universite de Paris, France
sangnier@irif.fr

Abstract

We study networks of processes which all execute the same finite-state protocol and communicate thanks to a rendez-vous mechanism. Given a protocol, we are interested in checking whether there exists a number, called a cut-off, such that in any networks with a bigger number of participants, there is an execution where all the entities end in some final states. We provide decidability and complexity results of this problem under various assumptions, such as absence/presence of a leader or symmetric/asymmetric rendez-vous.

2012 ACM Subject Classification Theory of Computation → Models of computation

Keywords and phrases Parameterized networks, Verification, Cut-offs

Digital Object Identifier 10.4230/LIPIcs..2020.

Funding Partly supported by ANR FREDDA (ANR-17-CE40-0013).

1 Introduction

Networks with many identical processes. One of the difficulty in verifying distributed systems lies in the fact that many of them are designed for an unbounded number of participants. As a consequence, to be exhaustive in the analysis, one needs to design formal methods which takes into account this characteristic. In [21], German and Sistla introduce a model to represent networks with a fix but unbounded number of entities. In this model, each participant executes the same protocol and they communicate between each other thanks to rendez-vous (a synchronization mechanism allowing two entities to change their local state simultaneously). The number of participants can then be seen as a parameter of the model and possible verification problems ask for instance whether a property holds for all the values of this parameter or seeks for some specific value ensuring a good behavior. With the increasing presence of distributed mechanisms (mutual exclusion protocols, leader election algorithms, renaming algorithms, etc) in the core of our computing systems, there has been in the last two decades a regain of attention in the study of such parameterized networks.

Surprisingly, the verification of these parameterized systems is sometimes easier than the case where the number of participants is known. This can be explained by the following reason: in the parameterized case the procedure can adapt on demand the number of participants to build a problematic execution. It is indeed what happens with the liveness verification of asynchronous shared-memory systems. This problem is Pspace-complete for a finite number of processes and in NP when this number is a parameter [14]. It is hence worth studying the complexity of the verification of such parameterized models and many recent works have attacked these problems considering networks with different means of communication. For instance in [16, 13, 7, 6] the participants communicate thanks to broadcast of messages, in [11, 2] they use a token-passing mechanism , in [10] a message passing mechanism and in [18] the communication is performed through shared registers.

Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The relative expressiveness of some of those models has been studied in [4]. Finally in his survey [15], Esparza shows that minor changes in the setting of parameterized networks, such as the presence of a controller (or equivalently a leader), might drastically change the complexity of the verification problems.

Cut-off to ease the verification. When one has to prove the correctness of a distributed algorithm designed to work for an unbounded number of participants, one technique consists in proving that the algorithm has a cut-off, i.e. a bound on the number of processes such that if it behaves correctly for this specific number of processes then it will still be correct for any bigger networks. Such a property allows to reduce the verification procedure to the analysis of the algorithm with a finite number of entities. Unfortunately, as shown in [3], many parameterized systems do not have a cut-off even for basic properties. Instead of checking whether a general class of models admits a cut-off, we propose in this work to study the following problem: given a representation of a system and a class of properties, does it admit a cutoff ? To the best of our knowledge, looking at the existence of a cutoff as a decision problem is a subject that has not received a lot of attention although it is interesting both practically and theoretically. First, in the case where this problem is decidable, it allows to find automatically cutoffs for specific systems even though they belong to a class for which there is no general results on the existence of cutoff. The search of cutoffs has been studied in [1] where the authors propose a semi-algorithm for verification of parameterized networks with respect to safety properties. This algorithm stops when a cutoff is found. However it is not stated how to determine the existence of this cutoff, neither if this is possible or not. In [25], the authors propose a way to compute dynamically a cutoff, but they consider systems and properties for which they know that a cutoff exists. Second, from the theoretical point of view, the cutoff decision problem is interesting because it goes beyond the classical problems for parameterized systems that usually seek for the existence of a number of participants which satisfies a property or check that a property hold for all possible number of participants. Note that in the latter case, one might be in a situation that for a property to hold a minimum number of participants is necessary (and below this number the property does not hold), such a situation can be detected with the existence of a cutoff but not with the simple universal quantification.

Rendez-vous networks. We focus on networks where the communication is performed by rendez-vous. There are different reasons for this choice. First, we are not aware of any technique to decide automatically the existence of a cut-off in parameterized systems, it is hence convenient to look at this problem in a well-known setting. Another aspect which motivates the choice of this model is that the rendez-vous communication corresponds to a well-known paradigm in the design of concurrent/distributed systems (for instance rendez-vous in the programming languages C or Java can be easily implemented thanks to wait/notify mechanisms). Rendez-vous communication seems as well a natural feature for parameterized systems used to model for instance crowds or biological systems (at some point we consider symmetric rendez-vous which can be seen less common in computing systems but make sense for these other applications). Last but not least, rendez-vous networks are very close to population protocols [5] for which there has been in the last years a regain of interest in the community of formal methods [17, 8, 9]. Population protocols and rendez-vous networks are both based on rendez-vous communication, but in population protocols it is furthermore required that all the fair executions converge to some accepting set of configurations (see [17] for more details). In our case, we seek for the existence of an execution ending with all the processes in a final state. The similarities between the two models let us think that the formal techniques we use could be adapted for the analysis of some population protocols.

Our contributions. We study the Cut-off Problem (C.O.P.) for rendez-vous networks. It consists in determining whether, given a protocol labeled with rendez-vous primitives, there exists a bound B, such that in any networks of size bigger than B where the processes all run the same protocol there is an execution which brings all the processes to a final state. We assume furthermore that in our network, there could be one extra entity, called the leader, that runs its own specific protocol. We first show that C.O.P. is decidable by reducing it to a new decision problem on Petri nets. Unfortunately we show as well that it is non elementary thanks to a reduction from the reachability problem in Petri nets[12]. We then show that better complexity bounds can be obtained if we assume the rendez-vous to be symmetric (i.e. any process that requests a rendez-vous can as well from the same state accept one and vice-versa) or if we assume that there is no leader. For each of these restrictions, new algorithmic techniques for the analysis of rendez-vous networks are proposed. The following table sums up the complexity bounds we obtain.

	Asymmetric rendez-vous	Symmetric rendez-vous
Presence of a leader	Decidable and non-elementary	PSPACE
Absence of leader	EXPSPACE	NP

Table 1 Complexity results obtained for the Cut-Off Problem

Due to lack of space, omitted details and proofs can be found in [23].

2 Modeling networks with rendez-vous communication

We write $\mathbb{N}$ to denote the set of natural numbers and $[i, j]$ to represent the set $\{k \in \mathbb{N} \mid i \leq k \text{ and } k \leq j\}$ for $i, j \in \mathbb{N}$. For a finite set E, the set $\mathbb{N}^E$ represents the multisets over E. For two elements $m, m' \in \mathbb{N}^E$, we denote $m+m'$ the multiset such that $(m+m')(e) = m(e)+m'(e)$ for all $e \in E$. We say that $m \leq m'$ if and only if $m(e) \leq m'(e)$ for all $e \in E$. If $m \leq m'$, then $m' - m$ is the multiset such that $(m' - m)(e) = m'(e) - m(e)$ for all $e \in E$. The size of a multiset m is given by $|m| = \Sigma_{e \in E} m(e)$. For $e \in E$, we use sometimes the notation e for the multiset m verifying $m(e) = 1$ and $m(e') = 0$ for all $e' \in E \setminus \{e\}$ and the notation $\langle\langle e1, e1, e2, e3\rangle\rangle$ to represent the multiset with four elements $e1, e1, e2$ and $e3$.

2.1 Rendez-vous protocols

We are now ready to define our model of networks. We assume that all the entities in the network (called sometimes processes) behave similarly following the same protocol except one entity, called the leader, which might behave differently. The communication in the network is pairwise and is performed by rendez-vous through a communication alphabet Σ. Each entity can either request a rendez-vous, with the primitive $?a$, or answer to a rendez-vous, with the primitive $!a$ where a belongs to Σ. The set of actions is hence $RV(\Sigma) = \{?a, !a \mid a \in \Sigma\}$.

▶ **Definition 1** (Rendez-vous protocol). *A* rendez-vous protocol $\mathcal{P}$ *is a tuple* $\langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$ *where Q is a finite set of states partitioned into the processes states Q_P and the leader states Q_L, Σ is a finite alphabet, $q_i \in Q_P$ [resp. $q_i^L \in Q_L$] is the initial state of the processes [resp. of the leader], $q_f \in Q_P$ [resp. $q_f^L \in Q_L$] is the final state of the processes [resp. of the leader], and $E \subseteq (Q_P \times RV(\Sigma) \times Q_P) \cup (Q_L \times RV(\Sigma) \times Q_L)$ is the set of edges.*

A configuration of the rendez-vous protocol $\mathcal{P}$ is a multiset $C \in \mathbb{N}^Q$ verifying that there exists $q \in Q_L$ such that $C(q) = 1$ and $C(q') = 0$ for all $q' \in Q_L \setminus \{q\}$, in other words there is a single entity corresponding to the leader. The number of processes in a configuration C is given by $|C| - 1$. We denote by $\mathcal{C}^{(n)}$ the set of configurations C involving n processes, i.e. such that $|C| = n + 1$. The initial configuration with n processes $C_i^{(n)}$ is such that $C_i^{(n)}(q_i) = n$ and $C_i^{(n)}(q_i^L) = 1$ and $C_i^{(n)}(q) = 0$ for all $q \in Q \setminus \{q_i, q_i^L\}$. Similarly the final configuration with n processes $C_f^{(n)}$ verifies $C_f^{(n)}(q_f) = n$ and $C_f^{(n)}(q_f^L) = 1$ and $C_f^{(n)}(q) = 0$ for all $q \in Q \setminus \{q_f, q_f^L\}$. Hence in an initial configuration all the entities are in their initial state and in a final configuration they are all in their final state. The notation $\mathcal{C}$ represents the whole set of configurations equals to $\bigcup_{n \in \mathbb{N}} \mathcal{C}^{(n)}$.

We are now ready to formalize the behavior of a rendez-vous protocol. In this matter, we define the relation $\rightarrow \subseteq \bigcup_{n \geq 1} \mathcal{C}^{(n)} \times \mathcal{C}^{(n)}$ as follows : $C \rightarrow C'$ if, and only if, there is $a \in \Sigma$ and two edges $(q_1, ?a, q_2), (q_1', !a, q_2') \in E$ such that $C(q_1) > 0$ and $C(q_1') > 0$ and $C(q_1) + C(q_1') \geq 2$ and $C' = C - (q_1 + q_1') + (q_2 + q_2')$. Intuitively it means that in C there is one entity in q_1 that requests a rendez-vous and one entity in q_1' that answers to it and they both change their state to respectively q_2 and q_2'. We need the hypothesis $C(q_1) + C(q_1') \geq 2$ in case $q_1 = q_1'$. We use $\rightarrow^*$ to represent the reflexive and transitive closure of $\rightarrow$. Note that if $C \rightarrow^* C'$ then $|C| = |C'|$, in other words there is no deletion or creation of processes during an execution.

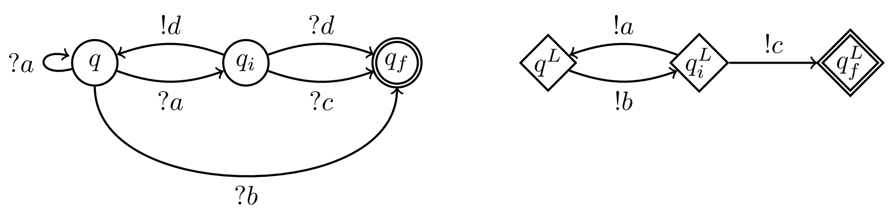

Figure 1 A rendez-vous protocol

▶ **Example 2.** Figure 1 provides an example of rendez-vous protocol where the process states are represented by circles and the leader states by diamond.

2.2 The cut-off problem

We can now describe the problem we address. It consists in determining given a protocol whether there exists a number of processes such that if we put more processes in the network it is always possible to find an execution which brings all the entities from their initial state to their final state. This **cut-off problem (C.O.P.)** can be stated formally as follows:

- **Input:** A rendez-vous protocol $\mathcal{P}$;
- **Output:** Does there exist a cut-off $B \in \mathbb{N}$ such that $C_i^{(n)} \rightarrow^* C_f^{(n)}$ for all $n \geq B$?

▶ **Example 3.** The rendez-vous network represented in Figure 1 admits a cut-off equal to 3. For $n = 3$, we have indeed an execution $C_i^{(3)} \rightarrow^* C_f^{(3)}$: $\langle\langle q_i^L, q_i, q_i, q_i \rangle\rangle \xrightarrow{d} \langle\langle q_i^L, q_i, q, q_f \rangle\rangle \xrightarrow{a} \langle\langle q^L, q_i, q, q_f \rangle\rangle \xrightarrow{b} \langle\langle q_i^L, q_i, q_f, q_f \rangle\rangle \xrightarrow{c} \langle\langle q_f^L, q_f, q_f, q_f \rangle\rangle$ (we indicate for each transition the label of the corresponding rendez-vous). For $n = 4$, the following sequence of rendez-vous leads to an execution $C_i^{(4)} \rightarrow^* C_f^{(4)}$: $\langle\langle q_i^L, q_i, q_i, q_i, q_i \rangle\rangle \xrightarrow{d} \langle\langle q_i^L, q_i, q_i, q, q_f \rangle\rangle \xrightarrow{a} \langle\langle q^L, q_i, q_i, q_i, q_f \rangle\rangle \xrightarrow{d} \langle\langle q^L, q_i, q, q_f, q_f \rangle\rangle \xrightarrow{b} \langle\langle q_i^L, q_i, q_f, q_f, q_f \rangle\rangle \xrightarrow{c} \langle\langle q_f^L, q_f, q_f, q_f, q_f \rangle\rangle$. Then for any $n > 4$, we can always come back to the case where $n = 3$ (if n is odd) or $n = 4$ (if n is even). In fact, we can always let 3 or 4 processes in q_i and move pairwise the other processes, one in q and one in q_f. Then the processes in q can be brought in q_f thanks to the rendez-vous a and b and the leader loop between q_i^L and q^L. Note that if we delete the edge $(q, ?a, q_i)$, this protocol does not admit anymore a cut-off but for all odd number $n \geq 3$, we have $C_i^{(n)} \rightarrow^* C_f^{(n)}$.

2.3 Petri nets

As we shall see there are some strong connections between rendez-vous protocols and Petri nets, this is the reason why we recall the definition of this latter model.

▶ **Definition 4** (Petri net). *A Petri net $\mathcal{N}$ is a tuple $\langle P, T, Pre, Post\rangle$ where P is a finite set of places, T is a finite set of transitions, $Pre : T \mapsto \mathbb{N}^P$ is the precondition function and $Post : T \mapsto \mathbb{N}^P$ is the postcondition function.*

A marking of a Petri net is a multiset $M \in \mathbb{N}^P$. A Petri net defines a transition relation $\Rightarrow \subseteq \mathbb{N}^P \times T \times \mathbb{N}^P$ such that $M \overset{t}{\Rightarrow} M'$ for $M, M' \in \mathbb{N}^P$ and $t \in T$ if and only if $M \geq Pre(t)$ and $M' = M - Pre(t) + Post(t)$. The intuition behind Petri nets is that marking put tokens in some places and each transition consumes with Pre some tokens and produces others thanks to $Post$ in order to create a new marking. We write $M \Rightarrow M'$ iff there exists $t \in T$ such that $M \overset{t}{\Rightarrow} M'$. Given a marking $M \in \mathbb{N}^P$, the reachability set of M is the set $Reach(M) = \{M' \in \mathbb{N}^P \mid M \Rightarrow^* M'\}$ where $\Rightarrow^*$ is the reflexive and transitive closure of $\Rightarrow$. One famous problem in Petri nets is the **reachability problem**:

- **Input:** A Petri net $\mathcal{N}$ and two markings M and M';
- **Output:** Do we have $M' \in Reach(M)$?

This problem is decidable [32, 27, 28, 29] and non elementary [12]. Another similar problem that we will refer to and which is easier to solve is the **reversible reachability problem**:

- **Input:** A Petri net $\mathcal{N}$ and two markings M and M';
- **Output:** Do we have $M' \in Reach(M)$ and $M \in Reach(M')$?

It has been shown in [31] to be EXPSPACE-complete.

3 Back and forth between rendez-vous protocols and Petri nets

3.1 From Petri nets to rendez-vous protocols

We will see here how the reachability problem for Petri nets can be reduced to the C.O.P. which gives us a non-elementary lower bound for this latter problem. We consider in the sequel a Petri net $\mathcal{N} = \langle P, T, Pre, Post\rangle$ and two markings $M, M' \in \mathbb{N}^P$. Without loss of generality we can assume that M and M' are of the following form: there exists $p_i \in P$ such that $M(p_i) = 1$ and $M(p) = 0$ for all $p \in P \setminus \{p_i\}$ and there exists $p_f \in P$ such that $M'(p_f) = 1$ and $M'(p) = 0$ for all $p \in P \setminus \{p_f\}$. Taking these restrictions on the markings does not alter the complexity of the reachability problem.

We build from $\mathcal{N}$ a rendez-vous protocol $\mathcal{P}_\mathcal{N}$ which admits a cut-off if and only if $M' \in Reach(M)$. The states of the processes in $\mathcal{P}_\mathcal{N}$ are matched to the places of $\mathcal{N}$, the number of processes in a state corresponding to the number of tokens in the associated place, and the leader is in charge to move the processes in order to simulate the changing on the number of tokens. The protocol is equipped with an extra state R, the reserve state, where the leader stores at the beginning of the simulation the number of processes which will simulate the tokens: when a transition produces a token in a place p, the leader moves a process from R to p and when it consumes a token from a place p, the leader moves a process from p to q_f. Figure 2 provides an example of a Petri net and its associated rendez-vous network. In this net, the transition letter a is used to put as many processes as necessary to simulate the number of tokens in the places in the reserve state R. The letters $pr(p_j)$

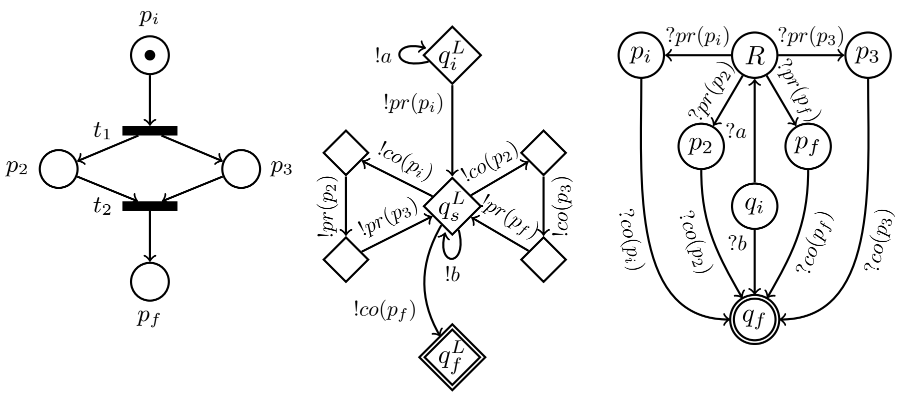

Figure 2 A Petri net $\mathcal{N}$ and its associated rendez-vous network $\mathcal{P}_{\mathcal{N}}$

are used to simulate the production of a token in the place p_j by moving a process from R to p_j and the letter $co(p_j)$ are used to simulate the consumption of a token in the place p_j by moving a process from p_j to q_f. It is then easy to see that each loop on the state q_s^L simulates a transition of the Petri net whereas the transition from q_i^L to q_s^L is used to build the initial marking and the transition from q_s^L to q_f^L is used to delete one token from the single place p_f and move the corresponding process to q_f. Finally, the letter b is used to ensure the cutoff property by moving from q_i to q_f the extra processes not needed to simulate the tokens. This construction gives us a hardness result for the C.O.P. thanks to the fact that the reachability problem in Petri nets is non-elementary [12].

▶ **Theorem 5.** *The C.O.P. is non-elementary.*

3.2 From rendez-vous protocols to Petri nets

We now show how to encode the behavior of a rendez-vous protocol into a Petri net and give a reduction from the C.O.P. to a problem on the built Petri net. We consider a rendez-vous protocol $\mathcal{P} = \langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$. From $\mathcal{P}$, we build a Petri net $\mathcal{N}_{\mathcal{P}} = \langle P, T, Pre, Post\rangle$ with $P = \{p_q \mid q \in Q\}$ and $T = \{t_i, t_f^L\} \cup \{t_{(q_1,q_2,a,q_1',q_2')} \mid q_1, q_2, q_1', q_2' \in Q$ and $a \in \Sigma$ and $(q_1, !a, q_1'), (q_2, ?a, q_2') \in E\}$. Intuitively in $\mathcal{N}_{\mathcal{P}}$, we have a place for each state of $\mathcal{P}$, the transition t_i puts tokens corresponding to new processes in the place corresponding to the initial state q_i, the transition t_f^L consumes a token in the place corresponding to the final state of the leader q_f^L and each transition $t_{(q_1,q_2,a,q_1',q_2')}$ simulates the protocol respecting the associated semantics (it checks that there is one process in q_1 another one in q_2 and that they can communicate thanks to the communication letter $a \in \Sigma$ moving to q_1' and q_2'). Figure 3 represents the Petri net $\mathcal{N}_{\mathcal{P}}$ for the protocol $\mathcal{P}$ of Figure 1 (the transitions are only labeled with the letter of the rendez-vous).

Unfortunately we did not find a way to reduce directly the C.O.P. to the reachability problem in Petri nets which would have lead directly to the decidability of C.O.P. However we will see how the C.O.P. on $\mathcal{P}$ can lead to a decision problem on $\mathcal{N}_{\mathcal{P}}$. We consider the initial marking $M_0 \in \mathbb{N}^P$ such that $M_0(p_{q_i^L}) = 1$ and $M_0(p) = 0$ for all $p \in P \setminus \{p_{q_i^L}\}$ and the family of markings $(M_f^{(n)})_{\{n \in \mathbb{N}\}}$ such that $M_f^{(n)}(p_{q_f}) = n$ and $M_f^{(n)}(p) = 0$ for all $p \in P \setminus \{p_{q_f}\}$. From the way we build the Petri net $\mathcal{N}_{\mathcal{P}}$, we deduce the following lemma:

▶ **Lemma 6.** *For all $n \in \mathbb{N}$, $C_i^{(n)} \rightarrow^* C_f^{(n)}$ in $\mathcal{P}$ iff $M_f^{(n)} \in Reach(M_0)$ in $\mathcal{N}_{\mathcal{P}}$.*

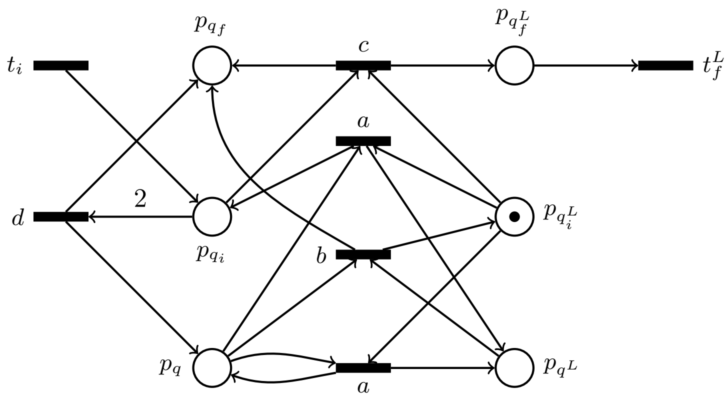

Figure 3 The Petri net $\mathcal{N}_{\mathcal{P}}$ for the protocol $\mathcal{P}$ of Figure 1

This leads us to propose a cut-off problem for Petri nets, which asks whether given an initial marking and a specific place, there exists a bound $B \in \mathbb{N}$ such that for all $n \geq B$ it is possible to reach a marking with n tokens in the specific place and none in the other. This **single place cut-off problem (single place C.O.P.)** can be stated formally as follows:

- **Input:** A Petri net $\mathcal{N}$, an initial marking M_0 and a place p_f;
- **Output:** Does there exist $B \in \mathbb{N}$ such that for all $n \geq B$, we have $M^{(n)} \in \mathit{Reach}(M_0)$ in $\mathcal{N}$ where $M^{(n)}$ is the marking verifying $M^{(n)}(p_f) = n$ and $M^{(n)}(p) = 0$ for all $p \in P \setminus \{p_f\}$?

Thanks to Lemma 6, we can then conclude the following proposition which justifies the introduction of the single place C.O.P. in our context.

▶ **Proposition 7.** *The C.O.P. reduces to the single place C.O.P.*

4 Solving C.O.P. in the general case

We show how to solve the C.O.P. by solving the single place C.O.P. To the best of our knowledge this latter problem has not yet been studied and we do not see direct connections with existing studied problems on Petri nets. It amounts to check if for some $B \in \mathbb{N}$ we have $\{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\} \text{ and } M(p_f) \geq B\} \subseteq \mathit{Reach}(M_0)$. We know from [26] that the projection of the reachability set on the single place p_f is semilinear (that can be represented by a Presburger arithmetic formula), however this does not help us since we furthermore require the other places different from p_f to be empty.

4.1 Formal tools and associated results

For $\mathbf{P}, \mathbf{P}' \subseteq \mathbb{N}^n$, we let $\mathbf{P} + \mathbf{P}' = \{p + p' \mid p \in \mathbf{P} \text{ and } p' \in \mathbf{P}'\}$ and we shall sometimes identify an element $p \in \mathbb{N}^n$ with the singleton $\{p\}$. A subset $\mathbf{P}$ of $\mathbb{N}^n$ for $n > 0$ is said to be *periodic* iff $\mathbf{0} \in \mathbf{P}$ and $\mathbf{P} + \mathbf{P} \subseteq \mathbf{P}$. Such a periodic set $\mathbf{P}$ is *finitely generated* if there exists a finite set of elements $\{\mathbf{p}_1, \ldots, \mathbf{p}_k\} \subset \mathbb{N}^n$ such that $\mathbf{P} = \{\lambda_1.\mathbf{p}_1 + \ldots + \lambda_k.\mathbf{p_k} \mid \lambda_i \in \mathbb{N} \text{ for all } i \in [1, k]\}$. A *semilinear set* of $\mathbb{N}^k$ is then a finite union of sets of the form $\mathbf{b} + \mathbf{P}$ where $\mathbf{b} \in \mathbb{N}^k$ and $\mathbf{P}$ is finitely generated. Semilinear sets are particularly useful tools because they are closed under the classical operations (union, complement and projection) and they provide a finite representation of infinite sets of vectors of naturals. Furthermore they can be represented by logical formulae expressed in Presburger arithmetic which is the decidable first-order theory of natural numbers with addition. A formula $\phi(x_1, \dots, x_k)$ of Presburger arithmetic with free variables $x_1, \dots, x_k$ defines a set $[\![\phi]\!] \subseteq \mathbb{N}^k$ given by $\{\mathbf{v} \in \mathbb{N}^k \mid \mathbf{v} \models \phi\}$ (here $\models$ is the classical satisfiability relation for Presburger arithmetic and it holds true if the formula holds when replacing each x_i by $\mathbf{v}[i]$). In [22], it was proven that a set $S \subseteq \mathbb{N}^k$ is semilinear iff there exists a Presburger formula ϕ such that $S = [\![\phi]\!]$. Note that the set $\{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\}\}$ has a single interesting component, the other being 0. We will hence need the following result to show it is indeed semilinear.

▶ **Lemma 8.** *Every periodic subset* $\mathbf{P} \subseteq \mathbb{N}$ *is semilinear.*

We now recall some connections between Petri nets and semilinear sets. Let $\mathcal{N} = \langle P, T, Pre, Post\rangle$ be a Petri net with $P = \{p_1, \dots, p_k\}$, this allows us to look at the markings as elements of $\mathbb{N}^k$ or of $\mathbb{N}^P$. Given a language of finite words of transitions $L \subseteq T^*$ and a marking M, let $Reach(M, L)$ be the reachable markings produced by L from M defined by $\{M' \subseteq \mathbb{N}^k \mid \exists w \in L \text{ such that } M \stackrel{w}{\Rightarrow} M'\}$ where we extend in the classical way the relation $\Rightarrow$ over words of transitions by saying $M \stackrel{\varepsilon}{\Rightarrow} M$ and if $w = t.w'$, we have $M \stackrel{w}{\Rightarrow} M'$ iff there exists M'' such that $M \stackrel{t}{\Rightarrow} M'' \stackrel{w'}{\Longrightarrow} M'$. A flat expression of transitions is a regular expression over T of the form $T_1T_2 \dots T_\ell$ where each T_i is either a finite word in T^* or of the form w^* with $w \in T^*$. For a flat expression FE, we denote by $L(FE)$ its associated language. In [20], the following result relating flat expressions of transitions and their produced reachability set is given (it has then been extended to more complex systems [19]).

▶ **Proposition 9.** *[20] Let* $\mathcal{N} = \langle P, T, Pre, Post\rangle$ *be a Petri net,* FE *a flat expression of transitions and* $M \in \mathbb{N}^P$ *a marking. Then* $Reach(M, L(FE))$ *is semilinear (and the corresponding Presburger formula can be computed).*

4.2 Deciding if a bound is a single-place cut-off

We prove that if one provides a bound $B \in \mathbb{N}$, we are able to decide whether it corresponds to a cut-off as defined in the single place C.O.P. Let $\mathcal{N} = \langle P, T, Pre, Post\rangle$ be a Petri net with an initial marking $M_0 \in \mathbb{N}^P$, a specific place $p_f \in P$ and a bound $B \in \mathbb{N}$. We would like to decide whether the following inclusion holds $\{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\} \text{ and } M(p_f) \geq B\} \subseteq Reach(M_0)$. An important point to decide this inclusion lies in the fact that the set $\{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\} \text{ and } M(p_f) \geq B\}$ is semilinear and this allows us to use a method similar to the one proposed in [24] to check whether the reachability set of a Petri net equipped with a semilinear set of initial markings is universal. One key point is the following result which is a reformulation of a Lemma in [30]. This result was originally stated for Vector Addition System with States (VASS), but it is well known that a Petri net can be translated into a VASS with an equivalent reachability set.

▶ **Proposition 10.** *[24, Theorem 1] Let* $\mathcal{N} = \langle P, T, Pre, Post\rangle$ *be a Petri net,* $M \in \mathbb{N}^P$ *a marking and* $S \subseteq \mathbb{N}^P$ *a semilinear set of markings. If* $S \subseteq Reach(M)$ *then there is a flat expression* FE *of transitions such that* $S \subseteq Reach(M, L(FE))$*.*

Following the technique used in [24], this proposition provides us a tool to solve our inclusion problem. We use two semi-procedures, one searches for a $M' \in \{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\} \text{ and } M(p_f) \geq B\}$ but not in $Reach(M_0)$ and the other one searches a flat expression of transitions FE such that $\{M \in \mathbb{N}^P \mid M(p) = 0 \text{ for all } p \in P \setminus \{p_f\} \text{ and } M(p_f) \geq B\} \subseteq Reach(M_0, L(FE))$.

▶ **Proposition 11.** *For a Petri net $\mathcal{N} = \langle P, T, Pre, Post\rangle$, a marking $M_0 \in \mathbb{N}^P$, a place $p_F \in P$ and a bound $B \in \mathbb{N}$, testing whether $\{M \in \mathbb{N}^P \mid M(p) = 0$ for all $p \in P \setminus \{p_f\}$ and $M(p_f) \geq B\} \subseteq Reach(M_0)$ is decidable.*

4.3 Finding the bound

We now show why the single-place C.O.P. is decidable. Let $\mathcal{N} = \langle P, T, Pre, Post\rangle$ be a Petri net with a marking $M_0 \in \mathbb{N}^P$ and a place $p_f \in P$. One key aspect is that the set of markings reachable from M_0 with no token in the other places except p_f is semilinear. This is a consequence of the following proposition.

▶ **Proposition 12.** *[30, Lemma IX.1] Let $S \subseteq \mathbb{N}^P$ be a semilinear set of markings. Then the set $Reach(M_0) \cap S$ is a finite union of sets $\mathbf{b} + \mathbf{P}$ where $\mathbf{b} \in \mathbb{N}^P$ and $\mathbf{P} \subseteq \mathbb{N}^P$ is periodic.*

From this proposition and Lemma 8, we can deduce the following result.

▶ **Proposition 13.** *$Reach(M_0) \cap \{M \in \mathbb{N}^P \mid M(p) = 0$ for all $p \in P \setminus \{p_f\}\}$ is semilinear.*

Another key point for the decidability of the single-place C.O.P. is the ability to test whether the intersection of the reachability set of a Petri net with a linear set is empty. In fact, it reduces to the reachability problem.

▶ **Lemma 14.** *If $S \subseteq \mathbb{N}^P$ is a linear set of the form $\mathbf{b} + \mathbf{P}$ where $\mathbf{P}$ is finitely generated, then testing whether $Reach(M_0) \cap S = \emptyset$ is decidable.*

The previous results allow us to design two semi-procedures to decide the single place C.O.P. The first one enumerates the $B \in \mathbb{N}$ and uses the result of Proposition 11 to check if one is a cut-off. The other one uses the fact that if there does not exist a cut-off then the set $\{M \notin Reach(M_0) \mid M(p) = 0$ for all $p \in P \setminus \{p_f\}\}$ is semi-linear (by Proposition 13) and infinite and it includes a semi-linear set of the form $\{\mathbf{b} + \lambda.\mathbf{p} \mid \lambda \in \mathbb{N}\}$ with $\mathbf{b}, \mathbf{p} \in \mathbb{N}^P$ and $\mathbf{0} < \mathbf{p}$. In this latter case we have $Reach(M_0) \cap \{\mathbf{b} + \lambda.\mathbf{p} \mid \lambda \in \mathbb{N}\} = \emptyset$ and we use the result of Lemma 14 to enumerate the $\mathbf{b}, \mathbf{p}$ and find a pair satisfying this property.

▶ **Theorem 15.** *The single place C.O.P. is decidable.*

Thanks to Proposition 7, we obtain the result which concludes this section.

▶ **Corollary 16.** *The C.O.P. is decidable.*

5 The specific case of symmetric rendez-vous

Even though the C.O.P. is decidable, the lower bound is quite bad as mentioned in Theorem 5 and the decision procedure presented in the proof of Theorem 15 is quite technical. We show here that for a specific family of rendez-vous protocols, solving C.O.P. is easier.

5.1 Definition and basic properties

A rendez-vous protocol $\mathcal{P} = \langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$ is *symmetric* if it respects the following property: for all $q, q' \in Q$ and $a \in \Sigma$, we have $(q, !a, q') \in E$ iff $(q, ?a, q') \in E$. In this context we denote such transitions by (q, a, q'). We furthermore assume w.l.o.g. that in the underlying graph of $\mathcal{P}$ for every states q in Q_P there is a path from q_i to q and a path from q to q_f (otherwise an initial configuration can never reach a configuration with a process in q or from a configuration with a process in q a final configuration can never been reached). We now work under these hypotheses.

In symmetric rendez-vous protocols, it is always possible to bring in any state as many pairs of processes one desires from the initial state q_i and to remove as many pairs of processes (and bring them to the final state q_f). To perform such actions, it is enough to move pairs of processes following the same path (as the rendez-vous are symmetric, this is allowed by the semantics of rendez-vous protocols). We now state these properties formally. Let $\mathcal{P} = \langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$ be a symmetric rendez-vous protocol.

▶ **Lemma 17.** *Let $C \in \mathcal{C}$ verifying $C_i^{(|C|-1)} \to^* C$. Then:*

1. *for all $C' \in \mathcal{C}$ such that $C(q) \leq C'(q)$ and $(C(q) = C'(q)) \mod 2$ for all $q \in Q$, we have $C_i^{(|C'|-1)} \to^* C'$, and,*
2. *for all $C' \in \mathcal{C}$ such that $|C'| = |C|$ and $C'(q) \leq C(q)$ for all $q \in Q \setminus \{q_f\}$ and $(C(q) = C'(q)) \mod 2$ for all $q \in Q$, we have $C_i^{(|C'|-1)} \to^* C'$.*

As a consequence, we show that there is a cut-off in $\mathcal{P}$ iff a final configuration with an even number and another one with an odd number of processes are reachable in $\mathcal{P}$.

▶ **Lemma 18.** *There exists $B \in \mathbb{N}$ such that $C_i^{(n)} \to^* C_f^{(n)}$ for all $n \geq B$ iff there exists an even $n_{\mathsf{E}} \in \mathbb{N}$ and an odd $n_{\mathsf{O}} \in \mathbb{N}$ such that $C_i^{(n_{\mathsf{E}})} \to^* C_f^{(n_{\mathsf{E}})}$ and $C_i^{(n_{\mathsf{O}})} \to^* C_f^{(n_{\mathsf{O}})}$.*

5.2 The even-odd abstraction

We now present our tool to decide C.O.P. for a symmetric rendez-vous protocol $\mathcal{P} = \langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$. We build an abstraction of the transition system $(\mathcal{C}, \to)$ where we only remember the state of the leader and whether the number of processes in each state is even (denoted by E) or odd (O). Let $\widehat{\mathsf{E}} = \mathsf{O}$ and $\widehat{\widehat{\mathsf{E}}} = \mathsf{E}$. The set of even-odd configurations is $\Gamma_{\mathsf{EO}} = Q_L \times \{\mathsf{E}, \mathsf{O}\}^{Q_P}$. To an even-odd configuration $(q^L, \gamma) \in \Gamma_{\mathsf{EO}}$, we associate the set of configurations $[\![(q^L, \gamma)]\!] \subseteq \mathcal{C}$ such that $[\![(q^L, \gamma)]\!] = \{C \in \mathcal{C} \mid C(q^L) = 1 \text{ and } C(q) = 0 \mod 2 \text{ iff } \gamma(q) = \mathsf{E}\}$. We now define the even-odd transition relation $\dashrightarrow \subseteq \Gamma_{\mathsf{EO}} \times E \times E \times \Gamma_{\mathsf{EO}}$. We have $(q_1^L, \gamma_1) \overset{e,e'}{\dashrightarrow} (q_2^L, \gamma_2)$ iff one the following conditions holds:

1. $e = (q_1^L, a, q_2^L)$ and $e' = (q_1, a, q_2)$ belongs to $Q_P \times RV(\Sigma) \times Q_P$ and if $q_1 = q_2$ then $\gamma_2 = \gamma_1$ else $\gamma_2(q_1) = \widehat{\gamma_1(q_1)}$, $\gamma_2(q_2) = \widehat{\gamma_1(q_2)}$ and $\gamma_2(q) = \gamma_1(q)$ for all $q \in Q_P \setminus \{q_1, q_2\}$.
2. $e, e' \in Q_P \times RV(\Sigma) \times Q_P$ and $q_1^L = q_2^L$ and $e = (q_1, a, q_2)$ and $e' = (q_3, a, q_4)$ and there exists $\gamma' \in \{\mathsf{E}, \mathsf{O}\}^{Q_P}$ such that:
   - if $q_1 = q_2$ then $\gamma' = \gamma_1$ else $\gamma'(q_1) = \widehat{\gamma_1(q_1)}$, $\gamma'(q_2) = \widehat{\gamma_1(q_2)}$ and $\gamma'(q) = \gamma_1(q)$ for all $q \in Q_P \setminus \{q_1, q_2\}$, and,
   - if $q_3 = q_4$ then $\gamma_2 = \gamma'$ else $\gamma_2(q_3) = \widehat{\gamma'(q_3)}$, $\gamma_2(q_4) = \widehat{\gamma'(q_4)}$ and $\gamma_2(q) = \gamma'(q)$ for all $q \in Q_P \setminus \{q_3, q_4\}$.

The relation $\overset{e,e'}{\dashrightarrow}$ reflects how the parity of the number of processes changes when performing a rendez-vous involving edges e and e'. For instance, the first case illustrates a rendez-vous between the leader and a process, hence the parity of the number of states in q_1 and in q_2 changes except when these two control states are equal. The second case deals with a rendez-vous between two processes and it is cut in two steps to take care of the cases like for instance $q_1 \neq q_2$ and $q_3 \neq q_4$ and $q_1 \neq q_4$ and $q_2 = q_3$; in fact here the parity of the number of processes in q_2 should not change, since the first transition adds one process to q_2 and the second one removes one from it. We write $(q_1^L, \gamma_1) \dashrightarrow (q_2^L, \gamma_2)$ iff there exists $e, e' \in E$ such that $(q_1^L, \gamma_1) \overset{e,e'}{\dashrightarrow} (q_2^L, \gamma_2)$ and $\dashrightarrow^*$ denotes the reflexive and transitive closure of $\dashrightarrow$.

As said earlier, $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$ is an abstraction of $(\mathcal{C}, \rightarrow)$. We will prove that this abstraction is enough to solve the C.O.P. For this, we define the following abstract configurations in Γ_{EO}:

- $(q_i^L, \gamma_i^{\mathsf{E}})$ and $(q_f^L, \gamma_f^{\mathsf{E}})$ are such that $\gamma_i^{\mathsf{E}}(q) = \gamma_f^{\mathsf{E}}(q) = \mathsf{E}$ for all $q \in Q_P$;
- $(q_i^L, \gamma_i^{\mathsf{O}})$ and $(q_f^L, \gamma_f^{\mathsf{O}})$ are such that $\gamma_i^{\mathsf{O}}(q) = \gamma_f^{\mathsf{O}}(q) = \mathsf{E}$ for all $q \in Q_P \setminus \{q_i, q_f\}$ and $\gamma_i^{\mathsf{O}}(q_f) = \gamma_f^{\mathsf{O}}(q_i) = \mathsf{E}$ and $\gamma_i^{\mathsf{O}}(q_i) = \gamma_f^{\mathsf{O}}(q_f) = \mathsf{O}$.

Note that we have then $\{C_i^{(n)} \mid n \text{ is even}\} \subseteq [\![(q_i^L, \gamma_i^{\mathsf{E}})]\!]$ and $\{C_i^{(n)} \mid n \text{ is odd}\} \subseteq [\![(q_i^L, \gamma_i^{\mathsf{O}})]\!]$ and $\{C_f^{(n)} \mid n \text{ is even}\} \subseteq [\![(q_f^L, \gamma_f^{\mathsf{E}})]\!]$ and $\{C_f^{(n)} \mid n \text{ is odd}\} \subseteq [\![(q_f^L, \gamma_f^{\mathsf{O}})]\!]$. According to the definitions of the relations $\rightarrow$ and $\dashrightarrow$, we can easily deduce this first result.

▶ **Lemma 19** (Completeness). *Let $n \in \mathbb{N}$. If $C_i^{(n)} \rightarrow^* C_f^{(n)}$ and n is even [resp. n is odd] then $(q_i^L, \gamma_i^{\mathsf{E}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{E}})$ [resp. $(q_f^L, \gamma_i^{\mathsf{O}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{O}})$].*

The two next lemmas show that our abstraction is sound for C.O.P. The first one can be proved by induction on the length of the path in $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$ using Point 1. of Lemma 17.

▶ **Lemma 20.** *If $(q_i^L, \gamma_i^{\mathsf{E}}) \dashrightarrow^* (q^L, \gamma)$ [resp. $(q_i^L, \gamma_i^{\mathsf{O}}) \dashrightarrow^* (q^L, \gamma)$] then there exists $n \in \mathbb{N} \setminus \{0\}$ such that n is even [resp. n is odd] and $C_i^{(n)} \rightarrow^* C$ with $C \in [\![(q^L, \gamma)]\!]$.*

Using Point 2. of Lemma 17 we obtain the soundness of our abstraction.

▶ **Lemma 21** (Soundness). *If $(q_i^L, \gamma_i^{\mathsf{E}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{E}})$ [resp. $(q_i^L, \gamma_i^{\mathsf{O}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{O}})$] then there exists $n \in \mathbb{N}$ such that n is even [resp. n is odd] and $C_i^{(n)} \rightarrow^* C_f^{(n)}$.*

Thanks to the Lemmas 18, 19 and 21 to solve the C.O.P. when the considered rendez-vous protocol is symmetric it is enough to check whether $(q_i^L, \gamma_i^{\mathsf{E}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{E}})$ and $(q_i^L, \gamma_i^{\mathsf{O}}) \dashrightarrow^* (q_f^L, \gamma_f^{\mathsf{O}})$. But since the transition system $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$ has a finite number of vertices whose number is bounded by $|Q_L| \cdot 2^{|Q_P|}$, these two reachability questions can be solved in NPSPACE in $|Q|$. By Savitch's theorem, we obtain the following result.

▶ **Theorem 22.** *C.O.P. restricted to symmetric rendez-vous protocols is in* PSPACE.

6 Supressing the leader

6.1 Definition and properties

A rendez-vous protocol $\mathcal{P} = \langle Q, Q_P, Q_L, \Sigma, q_i, q_f, q_i^L, q_f^L, E\rangle$ has *no leader* when $Q_L = \{q_f^L\}$ and $q_i^L = q_f^L$ and the transition relation does not refer to the state in Q_L, i.e. $E \subseteq Q_P \times RV(\Sigma) \times Q_P$. We can then assume that $\mathcal{P} = \langle Q_P, \Sigma, q_i, q_f, E\rangle$ and delete any reference to the leader state. We suppose again w.l.o.g. that in the considered rendez-vous protocols without leader there is a path from q_i to q and a path from q to q_f for all q in Q_P. Rendez-vous protocols with no leader enjoy some properties easing the resolution of the C.O.P.

▶ **Lemma 23.** *Let $\mathcal{P} = \langle Q_P, \Sigma, q_i, q_f, E\rangle$ be a rendez-vous protocol with no leader. Then the following properties hold:*

1. *If $C_i^{(n)} \rightarrow^* C_f^{(n)}$ and $C_i^{(m)} \rightarrow^* C_f^{(m)}$ for $m, n \in \mathbb{N}$, then $C_i^{(n+m)} \rightarrow^* C_f^{(n+m)}$.*
2. *There exists $B \in \mathbb{N}$ such that $C_i^{(n)} \rightarrow^* C_f^{(n)}$ for all $n \geq B$ iff there exists $N \in \mathbb{N}$ such that $C_i^{(N)} \rightarrow^* C_f^{(N)}$ and $C_i^{(N+1)} \rightarrow^* C_f^{(N+1)}$.*

Proof. **1.** This point is a direct consequence of the semantics of rendez-vous protocols associated with the fact that there is no leader. In fact assume $C_i^{(n)} \rightarrow^* C_f^{(n)}$ and $C_i^{(m)} \rightarrow^* C_f^{(m)}$. And consider the configuration C such that $C(q_i) = m$, $C(q_f) = n$ and

$C(q) = 0$ for all $q \in Q_P \setminus \{q_i, q_f\}$. Then it is clear that we have $C_i^{(n+m)} \to^* C \to^* C_f^{(n+m)}$, the first part of this execution mimicking the execution $C_i^{(n)} \to^* C_f^{(n)}$ and the last part mimics the execution $C_i^{(m)} \to^* C_f^{(m)}$ on the m processes left in q_i in C.

2. If there exists $B \in \mathbb{N}$ such that $C_i^{(n)} \to^* C_f^{(n)}$ for all $n \geq B$, then we have $C_i^{(B)} \to^* C_f^{(B)}$ and $C_i^{(B+1)} \to^* C_f^{(B+1)}$. Assume now that there exists $N \in \mathbb{N}$ such that $C_i^{(N)} \to^* C_f^{(N)}$ and $C_i^{(N+1)} \to^* C_f^{(N+1)}$. We show that for all $n \geq N^2$, we have $C_i^{(n)} \to^* C_f^{(n)}$. Let $n \geq N^2$ and let $R \in [0, N-1]$ be such that $(n = R) \mod N$. By definition of the modulo, there exists $A \geq 0$ such that $n = A \cdot N + R$. Since $n \geq N^2$, we have necessarily $A \geq N$. As a consequence we can rewrite n as: $n = R \cdot (N+1) + (A - R) \cdot N$. But then since $C_i^{(N)} \to^* C_f^{(N)}$, by 1. we have $C_i^{((A-R)\cdot N)} \to^* C_f^{((A-R)\cdot N)}$ and since $C_i^{(N+1)} \to^* C_f^{(N+1)}$, by 1. we have $C_i^{(R\cdot(N+1))} \to^* C_f^{(R\cdot(N+1))}$. By a last application of 1. we get $C_i^{(n)} \to^* C_f^{(n)}$.

◀

6.2 The symmetric case

We will now see how the procedure proposed in the proof of Theorem 22 to solve in polynomial space the C.O.P. for symmetric rendez-vous protocols can be simplified when there is no leader. Let $\mathcal{P} = \langle Q_P, \Sigma, q_i, q_f, E\rangle$ be a symmetric rendez-vous protocol with no leader and let $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$ be the abstract transition system of $(\mathcal{C}, \to)$ as defined in Section 5.2. If we adapt the results of Lemmas 18, 19 and 21 to the no leader case, we deduce that to solve the C.O.P. it is enough to check whether $\gamma_i^{\mathsf{E}} \dashrightarrow^* \gamma_f^{\mathsf{E}}$ and $\gamma_i^{\mathsf{O}} \dashrightarrow^* \gamma_f^{\mathsf{O}}$ (we have deleted the leader states from these results). Note that by definition $\gamma_i^{\mathsf{E}} = \gamma_f^{\mathsf{E}}$, hence the only thing to verify is if $\gamma_i^{\mathsf{O}} \dashrightarrow^* \gamma_f^{\mathsf{O}}$ holds. This check can be made efficiently using the fact that there is no leader, because any reodering of a path is still a path in $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$ (since we do not need to worry anymore about the leader state) and we can delete the pairs of edges that consecutively repeat since they have the same action on the parity.

▶ **Lemma 24.** *If $\gamma \dashrightarrow^* \gamma'$ then there exists $k \leq |E|^2$ and $e_1, e_1', e_2, e_2', \ldots, e_k, e_k' \in E$ such that $\gamma \overset{e_1,e_1'}{\dashrightarrow} \gamma_1 \overset{e_2,e_2'}{\dashrightarrow} \ldots \overset{e_k,e_k'}{\dashrightarrow} \gamma'$.*

It means that if $\gamma_i^{\mathsf{O}} \dashrightarrow^* \gamma_f^{\mathsf{O}}$ then there is a path of polynomial length (in the size of $\mathcal{P}$) between these two abstract configurations. It is hence enough to guess such a sequence of polynomial length and to check that it effectively corresponds to a path in $(\Gamma_{\mathsf{EO}}, \dashrightarrow)$.

▶ **Theorem 25.** *C.O.P. for symmetric rendez-vous protocols with no leader is in* NP.

6.3 Upper bound for the C.O.P. with no leader

We now prove that the C.O.P. for rendez-vous protocols with no leader reduces to the reversible reachability problem in Petri nets. Let $\mathcal{P} = \langle Q_P, \Sigma, q_i, q_f, E\rangle$ be a rendez-vous protocol with no leader and such that w.l.o.g. there is no edge going out of q_f[1].

Let $\mathcal{N}_{\mathcal{P}} = \langle P, T, Pre, Post\rangle$ be the Petri net whose construction is provided in Section 3.2 (where we have removed all the places corresponding to leader states as well as the transition t_f^L). From $\mathcal{N}_{\mathcal{P}}$, we build the reverse Petri net $\mathcal{N}_{\mathcal{P}}^R$ obtained by keeping the same set of places and reversing all the transitions. Formally $\mathcal{N}_{\mathcal{P}}^R = \langle P^R, T^R, Pre^R, Post^R\rangle$, where

[1] To achieve this, we can simply duplicate q_f adding a new final state q_f' and for each edge going into q_f we add an edge from the same state to q_f'

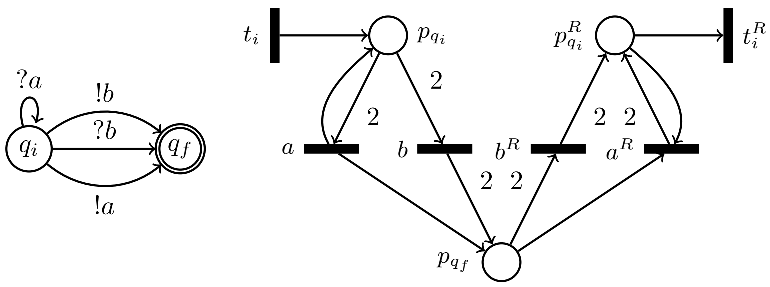

Figure 4 A rendez-vous protocol with no leader $\mathcal{P}$ and the associated Petri net $\mathcal{N}'_{\mathcal{P}}$

$P^R = \{p^R \mid p \in P\}$, $T^R = \{t^R \mid t \in T\}$ and for all $p^R \in P^R$ and $t^R \in T^R$, we have $Pre^R(t^R)(p^R) = Post(t)(p)$ and $Post^R(t^R)(p^R) = Pre(t)(p)$. Let M_0^R be the marking such that $M_0^R(p^R) = 0$ for all $p^R \in P^R$ and $(M_f^{R,(n)})_{\{n \in \mathbb{N}\}}$ be the family of markings verifying $M_f^{R,(n)}(p^R_{q_f}) = n$ and $M_f^{R,(n)}(p) = 0$ for all $p \in P^R \setminus \{p^R_{q_f}\}$. A direct consequence of Lemma 6 and of the definition of $\mathcal{N}^R_{\mathcal{P}}$ is that $C_i^{(n)} \rightarrow^* C_f^{(n)}$ iff $M_0^R \in Reach(M_f^{R,(n)})$ for all $n \in \mathbb{N}$.

From $\mathcal{N}_{\mathcal{P}}$ and $\mathcal{N}^R_{\mathcal{P}}$, we build the Petri net $\mathcal{N}'_{\mathcal{P}}$ obtained by taking the disjoint unions of places and transitions of the two nets except for the place p_{q_f} and $p^R_{q_f}$ which are merged in a single place p_{q_f}. Formally, $\mathcal{N}'_{\mathcal{P}} = \langle P', T', Pre', Post' \rangle$ where $P' = (P \cup P^R) \setminus \{p^R_{q_f}\}$, $T' = T \cup T^R$, $Pre'(t)(p) = Pre(t)(p)$ and $Post'(t)(p) = Post(t)(p)$ and $Pre'(t)(p^R) = Post'(t)(p^R) = 0$ for all $p \in P$, $p^R \in P^R$ and $t \in T$, $Pre'(t^R)(p^R) = Pre^R(t^R)(p^R)$ and $Post'(t^R)(p^R) = Post^R(t^R)(p^R)$ and $Pre'(t^R)(p) = Post'(t^R)(p) = 0$ for all $p^R \in P^R$, $p \in P \setminus \{p_{q_f}\}$ and $t \in T$, and $Pre'(t^R)(p_{q_f}) = Pre^R(t^R)(p^R_{q_f}))$ and $Post'(t^R)(p_{q_f}) = Post^R(t^R)(p^R_{q_f}))$ (this last case corresponds to the merging of p_{q_f} and $p^R_{q_f}$). Figure 4 provides an example of this latter Petri net.

We now explain why this new net is useful to solve the C.O.P. when there is no leader. First remember that thanks to Point 2. of Lemma 23 it is enough to check whether there exists $N \in \mathbb{N}$ such that $C_i^{(N)} \rightarrow^* C_f^{(N)}$ and $C_i^{(N+1)} \rightarrow^* C_f^{(N+1)}$. Intuitively, in $\mathcal{N}'_{\mathcal{P}}$ this property will be witnessed by the fact that we can bring $N+1$ tokens in p_{q_f} using transitions in T and remove N tokens from p_{q_f} thanks to the transitions in T^R letting hence one token in p_{q_f} and similarly if there is already a token in p_{q_f} we can bring N others and remove afterwards $N+1$. As for $\mathcal{N}_{\mathcal{P}}$, we let M_0 be the marking with no token, and $(M^{(n)})_{\{n \in \mathbb{N}\}}$ be the family of markings such that $M^{(n)}(p_{q_f}) = n$ and $M^{(n)}(p) = 0$ for all $p \in P' \setminus \{p_{q_f}\}$. Note that since there is no leader, we have here $M_0 = M^{(0)}$. The next lemma states the correctness of our reduction to the reversible reachability problem.

▶ **Lemma 26.** *There exists $N \in \mathbb{N}$ such that $C_i^{(N)} \rightarrow^* C_f^{(N)}$ and $C_i^{(N+1)} \rightarrow^* C_f^{(N+1)}$ iff $M^{(1)} \in Reach(M_0)$ and $M_0 \in Reach(M^{(1)})$ in the Petri net $\mathcal{N}'_{\mathcal{P}}$.*

Since we know that the reversible reachability problem for Petri net is EXPSPACE-complete [31], we obtain the following complexity result.

▶ **Theorem 27.** *C.O.P. restricted to rendez-vous protocols with no leader is in* EXPSPACE.

We were not able to propose a lower bound for the C.O.P. apart for the general case, but when there is no leader, we know that there is a protocol which admits a cut-off whose value is exponential in the size of a protocol. This protocol is shown on Figure 5. To bring a process in q_1, we need in fact two processes, to bring a process in q_2 and empty q_1, we need four processes and so on. The letter a is then used to ensure that as soon as we have

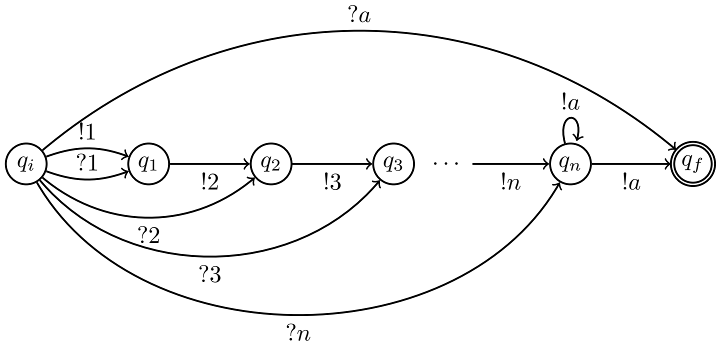

Figure 5 A rendez-vous protocol with no leader and an exponential cut-off

processes only in q_n and in q_i (and at least one of them in each of these states), there is a way to bring all of them in q_f.

7 Conclusion

We have shown here that the C.O.P. is decidable for rendez-vous networks. Furthermore we have provided complexity upper bounds when considering restrictions on the networks such as symmetric rendez-vous or absence of leader. Unfortunately, we did not succeed in finding matching lower bounds. Reducing other problems to the C.O.P. is in fact tedious without leader or when allowing only symmetric rendez-vous, because it is then quite hard to enforce that a specific number of processes are in some states which is a property that is in general needed to design reductions. However we have some hope to either improve our upper bounds or find matching lower bounds. We wish as well to understand in which matters the techniques we used could be adapted to other parameterized systems and more specifically to population protocols. Finally, one of the justification to consider the cutoff problem is that in some distributed systems it could be the case that a correctness property does not hold for any number of processes, but that a minimal number of participants is needed to reach a goal. It could be interesting to study a variant of our cutoff problem where we do not require all the processes to reach a final state but we want to know given a number of processes how many among them can be brought in such a state. An interesting property could be to check whether there exists a bound b such that for any number of processes, the minimal number that can not be brought to a final state by any execution is always lower than b. In such networks, it would mean that at most b entities have to be sacrificed to let the others reach the final state.

References

1. Parosh Aziz Abdulla, Frédéric Haziza, and Lukás Holík. Parameterized verification through view abstraction. *STTT*, 18(5):495–516, 2016.
2. Benjamin Aminof, Swen Jacobs, Ayrat Khalimov, and Sasha Rubin. Parametrized model checking of token-passing systems. In *VMCAI'14*, volume 8318 of *LNCS*, pages 262–281. Springer-Verlag, 2014.
3. Benjamin Aminof, Tomer Kotek, Sasha Rubin, Francesco Spegni, and Helmut Veith. Parameterized model checking of rendezvous systems. *Distributed Computing*, 31(3):187–222, 2018.

4 Benjamin Aminof, Sasha Rubin, and Florian Zuleger. On the expressive power of communication primitives in parameterised systems. In *LPAR'15*, volume 9450 of *LNCS*, pages 313–328. Springer-Verlag, 2015.

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

6 Nathalie Bertrand, Patricia Bouyer, and Anirban Majumdar. Reconfiguration and message losses in parameterized broadcast networks. In *CONCUR'19*, volume 140 of *LIPIcs*, pages 32:1–32:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

7 Nathalie Bertrand, Miheer Dewaskar, Blaise Genest, Hugo Gimbert, and Adwait Amit Godbole. Controlling a population. *Logical Methods in Computer Science*, 15(3), 2019.

8 Michael Blondin, Javier Esparza, and Stefan Jaax. Peregrine: A tool for the analysis of population protocols. In *CAV'18*, volume 10981 of *LNCS*, pages 604–611. Springer, 2018.

9 Michael Blondin, Javier Esparza, and Stefan Jaax. Expressive power of broadcast consensus protocols. In *CONCUR'19*, volume 140 of *LIPIcs*, pages 31:1–31:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

10 Benedikt Bollig, Paul Gastin, and Len Schubert. Parameterized verification of communicating automata under context bounds. In *RP'14*, volume 8762 of *LNCS*, pages 45–57. Springer-Verlag, 2014.

11 Edmund M. Clarke, Muralidhar Talupur, Tayssir Touili, and Helmut Veith. Verification by network decomposition. In *CONCUR'04*, volume 3170 of *LNCS*, pages 276–291. Springer-Verlag, 2004.

12 Wojciech Czerwinski, Slawomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. The reachability problem for petri nets is not elementary. In *STOC'19*, pages 24–33. ACM, 2019.

13 Giorgio Delzanno, Arnaud Sangnier, and Gianluigi Zavattaro. Parameterized verification of ad hoc networks. In *CONCUR'10*, volume 6269 of *LNCS*, pages 313–327. Springer-Verlag, 2010.

14 Antoine Durand-Gasselin, Javier Esparza, Pierre Ganty, and Rupak Majumdar. Model checking parameterized asynchronous shared-memory systems. *Formal Methods in System Design*, 50(2-3):140–167, 2017.

15 Javier Esparza. Keeping a crowd safe: On the complexity of parameterized verification (invited talk). In *STACS'14*, volume 25 of *LIPIcs*, pages 1–10. Leibniz-Zentrum für Informatik, 2014.

16 Javier Esparza, Alain Finkel, and Richard Mayr. On the verification of broadcast protocols. In *LICS'99*, pages 352–359. IEEE Comp. Soc. Press, July 1999.

17 Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of population protocols. *Acta Inf.*, 54(2):191–215, 2017.

18 Javier Esparza, Pierre Ganty, and Rupak Majumdar. Parameterized verification of asynchronous shared-memory systems. In *CAV'13*, volume 8044 of *LNCS*, pages 124–140. Springer-Verlag, 2013.

19 Alain Finkel and Jérôme Leroux. How to compose presburger-accelerations: Applications to broadcast protocols. In *FST TCS'02*, volume 2556 of *LNCS*, pages 145–156. Springer, 2002.

20 Laurent Fribourg. Petri nets, flat languages and linear arithmetic. In *WFLP'00*, pages 344–365, 2000.

21 Steven M. German and A. Prasad Sistla. Reasoning about systems with many processes. *J. ACM*, 39(3):675–735, 1992.

22 Seymour Ginsburg and Edwin H. Spanier. Semigroups, presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.

23 Florian Horn and Arnaud Sangnier. Deciding the existence of cut-off in parameterized rendez-vous networks. *CoRR*, abs/2007.05789, 2020. URL: `http://arxiv.org/abs/2007.05789`.

24 Petr Jancar, Jérôme Leroux, and Grégoire Sutre. Co-finiteness and co-emptiness of reachability sets in vector addition systems with states. In *PETRI NETS'18*, volume 10877 of *LNCS*, pages 184–203. Springer, 2018.

25 Alexander Kaiser, Daniel Kroening, and Thomas Wahl. Dynamic cutoff detection in parameterized concurrent programs. In *CAV'10*, volume 6174 of *LNCS*, pages 645–659. Springer, 2010.

26 Hans Kleine Büning, Theodor Lettmann, and Ernst W. Mayr. Projections of vector addition system reachability sets are semilinear. *Theor. Comput. Sci.*, 64(3):343–350, 1989.

27 S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *STOC'82*, pages 267–281. ACM, 1982.

28 Jean-Luc Lambert. A structure to decide reachability in petri nets. *Theor. Comput. Sci.*, 99(1):79–104, 1992.

29 Jérôme Leroux. Vector addition system reachability problem: a short self-contained proof. In *POPL'11*, pages 307–316. ACM, 2011.

30 Jérôme Leroux. Presburger vector addition systems. In *LICS'13*, pages 23–32. IEEE Computer Society, 2013.

31 Jérôme Leroux. Vector addition system reversible reachability problem. *Logical Methods in Computer Science*, 9(1), 2013.

32 Ernst W. Mayr. An algorithm for the general petri net reachability problem. *SIAM J. Comput.*, 13(3):441–460, 1984.