


# Verification of Population Protocols with Unordered Data

**Steffen van Bergerem** ✉ 

Humboldt-Universität zu Berlin, Germany

**Roland Guttenberg** ✉ 

Technische Universität München, Germany

**Sandra Kiefer** ✉ 

University of Oxford, United Kingdom

**Corto Mascle** ✉

LaBRI, Université de Bordeaux, France

**Nicolas Waldburger** ✉ 

IRISA, Université de Rennes, France

**Chana Weil-Kennedy** ✉ 

IMDEA Software Institute, Spain

---

## Abstract

Population protocols are a well-studied model of distributed computation in which a group of anonymous finite-state agents communicates via pairwise interactions. Together they decide whether their initial configuration, i. e., the initial distribution of agents in the states, satisfies a property. As an extension in order to express properties of multisets over an infinite data domain, Blondin and Ladouceur (ICALP'23) introduced population protocols with unordered data (PPUD). In PPUD, each agent carries a fixed data value, and the interactions between agents depend on whether their data are equal or not. Blondin and Ladouceur also identified the interesting subclass of immediate observation PPUD (IOPPUD), where in every transition one of the two agents remains passive and does not move, and they characterised its expressive power.

We study the decidability and complexity of formally verifying these protocols. The main verification problem for population protocols is well-specification, that is, checking whether the given PPUD computes some function. We show that well-specification is undecidable in general. By contrast, for IOPPUD, we exhibit a large yet natural class of problems, which includes well-specification among other classic problems, and establish that these problems are in  $\text{EXPSPACE}$ . We also provide a lower complexity bound, namely  $\text{CONEXPTIME}$ -hardness.

**2012 ACM Subject Classification** Theory of computation → Verification by model checking; Theory of computation → Distributed computing models

**Keywords and phrases** Population protocols, Parameterized verification, Distributed computing, Well-specification

**Funding** *Steffen van Bergerem*: This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) — project number 431183758 (gefördert durch die Deutsche Forschungsgemeinschaft (DFG) — Projektnummer 431183758).

*Sandra Kiefer*: This research was supported by the Glasstone Benefaction, University of Oxford [Violette and Samuel Glasstone Research Fellowships in Science 2022] as well as Jesus College in Oxford, UK.

*Chana Weil-Kennedy*: This work was supported by the grant PID2022-138072OB-I00, funded by MCIN, FEDER, UE and partially supported by PRODIGY Project (TED2021-132464B-I00) funded by MCIN and the European Union NextGeneration.

**Acknowledgements** This project started at and has benefitted substantially from the research camp Autóbóz 2023 in Kassel, Germany. We would like to thank the host, sponsors, and organisers of the research camp for bringing us together.

## 1 Introduction

Population protocols (PP) model distributed computation and have received a lot of attention [1, 2, 8, 11, 13, 17] since their introduction in 2004 [3]. In a PP, a collection of indistinguishable mobile agents with constant-size memory communicate via pairwise interactions. When two agents meet, they exchange information about their states and update their states accordingly. The agents collectively compute whether their input configuration, i. e., the initial distribution of agents in each state, satisfies a certain predicate. For a PP to compute a predicate, the protocol must be *well-specified*, i. e., for every initial configuration, all fair runs starting in this configuration must converge to the same answer. It was shown that PP compute exactly the predicates of Presburger arithmetic [4]. Moreover, well-specification is known to be decidable but as hard as the reachability problem for Petri nets [16]. Note that deciding well-specification is a problem that concerns *parameterised verification* in the sense of [7, 14], i. e., one must decide that something holds with respect to every value of the parameter. Here the parameter is the number of agents that are present – the PP must converge to one answer for every initial configuration, no matter the number of agents.

*Population protocols with unordered data* (PPUD) were introduced by Blondin and Ladouceur as a means to compute predicates over arbitrarily large domains [9]. In this setting, each agent holds a read-only datum from an infinite set  $\mathbb{D}$ . When interacting, agents may check (dis)equality of their data. While PP can compute properties like “there are more than 5 agents in state  $q_1$ ”, PPUD can express, e. g., “there are more than 2 data with 5 agents each in state  $q_1$ ”. In [9], the authors construct a PPUD computing the absolute majority predicate, i. e., whether a datum is held by more than half of the agents. They also characterise the expressive power of *immediate observation PPUD* (IOPPUD), a subclass of interest in which interactions are restricted to observations. That is, in every interaction, one of the two agents is passive and does not change its state. The decidability and complexity of the main verification question for PPUD, namely well-specification, is left open in Blondin’s and Ladouceur’s article [9]. It is the subject of this paper.

**Contributions** We start by showing that well-specification is *undecidable* for PPUD. This follows from a reduction from 2-counter machines; in fact, the presence of data allows us to encode zero-tests. Contrasting this, we show that deciding well-specification is in EXPSpace for IOPPUD. To this end, we define *generalised reachability expressions* (GRE) and establish that, for IOPPUD, deciding whether the set of configurations that satisfy a given GRE is empty is in EXPSpace. This decidability result is powerful; indeed, this emptiness problem subsumes classic verification problems like reachability and coverability, as well as parameterised verification problems such as well-specification and correctness, where the latter asks whether a given protocol computes a given predicate. Lastly, we exhibit a CONEXPTIME lower bound for deciding emptiness of GRE for IOPPUD.

**Related work** For a recent survey of the research on verification of PP (without data), see [15]. In particular, the well-specification problem for PP is known to be decidable, but as hard as Petri net reachability [16] and therefore Ackermann-complete [12, 23, 24]. In their seminal paper on the computational power of PP, Angluin, Aspnes, Eisenstat, and Ruppert also introduced five subclasses of PP that model one-way communication [4]. One of these is immediate observation population protocols (IOPP), which correspond to IOPPUD without data. The complexity of well-specification for all five subclasses is determined in [17]. In particular, the paper shows that well-specification for IOPP is PSPACE-complete. IOPP

were modelled by immediate observation Petri nets, where classic parameterised problems can be decided in polynomial space. The notion of generalised reachability expression was first phrased in this setting, and one of the consequences is that the emptiness problem of GRE for IOPP is PSPACE-complete [28]. Our result shows that adding data to the model as in [8] (and extending GRE naturally) pushes the emptiness problem between CONEXPTIME and EXPSpace.

While the introduction of data in the PP model happened recently [9], a similar approach has been studied in the related model of Petri nets, under the name of *data nets*. In this setting, the classic problem of coverability (or control-state reachability) is decidable but non-primitive recursive [22] and in fact  $\mathbf{F}_{\omega\omega}$ -complete [26]. While PPUD can be encoded into data nets, our results show that the problems that we study cannot be reduced to coverability. Another related model is formed by broadcast networks of register automata (BNRA) [19], an extension of reconfigurable broadcast networks (RBN) with data. RBN subsume IOPP [6], and consequently BNRA subsume IOPPUD. However, the complexity of coverability in BNRA is known to be  $\mathbf{F}_{\omega\omega}$ -complete, hence non-primitive recursive, and more complex problems quickly become undecidable [19]. These hardness results contrast with the EXPSpace membership.

**Organisation** In Section 2, we introduce the models of PPUD and IOPPUD, the notion of GRE, and we state our main results. We prove undecidability of well-specification for PPUD in Section 3. The next sections are dedicated to the study of IOPPUD. In Section 4, we establish bounds on the number of observed agents. In Section 5, we introduce the technical notions of boxes and containers and use the bounds from the previous section to translate GRE into containers. We present the complexity bounds for emptiness of GRE in Section 6.

## 2 Population Protocols and Main Results

We use the notation  $[m, n] := \{\ell \in \mathbb{N} \mid m \leq \ell \leq n\}$  for  $m, n \in \mathbb{N}$  and  $[m, +\infty) := \{\ell \in \mathbb{N} \mid m \leq \ell\}$ .

### 2.1 Population Protocols with Unordered Data

We fix an infinite *data* domain  $\mathbb{D}$ , an infinite set of *agents*  $\mathbb{A}$ , and a function  $\mathbf{dat}: \mathbb{A} \rightarrow \mathbb{D}$  such that  $\mathbf{dat}^{-1}(d)$  is infinite for all  $d \in \mathbb{D}$ . For  $d \in \mathbb{D}$ , a *d-agent* is an agent  $a \in \mathbb{A}$  with  $\mathbf{dat}(a) = d$ .

► **Definition 1.** A *population protocol with unordered data (PPUD)* is a tuple  $(Q, \Delta, I, O)$  where  $Q$  is a finite set of states,  $\Delta \subseteq Q^2 \times \{=, \neq\} \times Q^2$  the set of transitions,  $I \subseteq Q$  the set of initial states, and  $O: Q \rightarrow \{\top, \perp\}$  the output function.

The size of a PPUD  $\mathcal{P}$ , written  $|\mathcal{P}|$ , is its number of states. We fix a PPUD  $(Q, \Delta, I, O)$ .

A *configuration* is a function  $\gamma: \mathbb{A} \rightarrow Q \cup \{*\}$  such that  $\gamma(a) \in Q$  only holds for finitely many agents  $a \in \mathbb{A}$  (the agents *appearing* in  $\gamma$ ). We denote by  $\Gamma$  the set of all configurations, and by  $\Gamma_{\text{init}} := \{\gamma \in \Gamma \mid \forall a \in \mathbb{A}, \gamma(a) \notin Q \setminus I\}$  the set of *initial configurations*. Given  $\gamma \in \Gamma$  and  $d \in \mathbb{D}$ , we let  $\gamma_d^\# : Q \rightarrow \mathbb{N}$  be the function that maps each state  $q$  to the number of  $d$ -agents in  $q$  in  $\gamma$ .

Given  $\gamma, \gamma' \in \Gamma$ , we write  $\gamma \rightarrow \gamma'$ , and call it a *step* from  $\gamma$  to  $\gamma'$ , when there are states  $q_1, q_2, q_3, q_4 \in Q$  and two distinct agents  $a_1, a_2 \in \mathbb{A}$  such that  $((q_1, q_2), \bowtie, (q_3, q_4)) \in \Delta$ ,  $\gamma(a_1) = q_1$ ,  $\gamma(a_2) = q_2$ ,  $\gamma'(a_1) = q_3$ ,  $\gamma'(a_2) = q_4$ ,  $\gamma(a) = \gamma'(a)$  for all  $a \in \mathbb{A} \setminus \{a_1, a_2\}$ , and, additionally, if  $\bowtie$  is an equality (resp. disequality), then  $\mathbf{dat}(a_1) = \mathbf{dat}(a_2)$  (resp.  $\mathbf{dat}(a_1) \neq$

$\text{dat}(a_2)$ ). A *run*  $\rho$  is a (finite or infinite) sequence of consecutive steps  $\rho: \gamma_1 \rightarrow \gamma_2 \rightarrow \gamma_3 \rightarrow \dots$ . We write  $\rho: \gamma \xrightarrow{*} \gamma'$  to denote that  $\rho$  is a finite run from  $\gamma$  to  $\gamma'$ , and simply  $\gamma \xrightarrow{*} \gamma'$  to denote the existence of such a run. For every  $\gamma \in \Gamma$ , let  $\text{Post}^*(\gamma) := \{\gamma' \in \Gamma \mid \gamma \xrightarrow{*} \gamma'\}$  and  $\text{Pre}^*(\gamma) := \{\gamma' \in \Gamma \mid \gamma' \xrightarrow{*} \gamma\}$ . A run  $\rho$  *covers* a state  $q \in Q$  if there is a configuration  $\gamma$  in  $\rho$  such that  $\gamma(a) = q$  for some agent  $a$ .

In accordance with [3] and [9], we consider a run  $\gamma_1 \rightarrow \gamma_2 \rightarrow \dots$  *fair* if it is infinite<sup>1</sup> and for every configuration  $\gamma$  with  $|\{i \in \mathbb{N} \mid \gamma_i \xrightarrow{*} \gamma\}| = \infty$ , it holds that  $|\{i \in \mathbb{N} \mid \gamma_i = \gamma\}| = \infty$ . That is, every infinitely often reachable configuration also occurs infinitely often along the run. For  $b \in \{\top, \perp\}$ , a *b-consensus* is a configuration  $\gamma$  in which, for all agents  $a \in \mathbb{A}$  appearing in  $\gamma$ , it holds that  $O(\gamma(a)) = b$ . A fair run  $\rho: \gamma_1 \rightarrow \gamma_2 \rightarrow \dots$  *stabilises* to  $b \in \{\top, \perp\}$  if there is an  $n \in \mathbb{N}$  such that for every  $i \geq n$ ,  $\gamma_i$  is a *b-consensus*. A protocol is *well-specified* if, for every initial configuration  $\gamma_0 \in \Gamma_{\text{init}}$ , there is  $b \in \{\top, \perp\}$  such that all fair runs starting in  $\gamma_0$  stabilise to  $b$ . The *well-specification problem* for PPUD asks, given a PPUD  $\mathcal{P}$ , whether  $\mathcal{P}$  is well-specified. Given a PPUD  $\mathcal{P}$  and a function  $\Pi: \Gamma_{\text{init}} \rightarrow \{\top, \perp\}$ ,  $\mathcal{P}$  *computes*  $\Pi$  if, for every  $\gamma_0 \in \Gamma_{\text{init}}$ , every fair run of  $\mathcal{P}$  starting in  $\gamma_0$  stabilises to  $\Pi(\gamma_0)$ .

► **Example 2.** Consider the following PPUD  $\mathcal{P}$ . Its set of states is  $Q = \{\ell_0, \ell_1, f_0, f_1, \text{dead}\}$ , with  $I = \{\ell_1\}$ ,  $O(\ell_0) = O(f_0) = \top$ ,  $O(\ell_1) = O(f_1) = O(\text{dead}) = \perp$  and its transitions are:

$$\begin{aligned} \forall b, b' \in \{0, 1\}, (\ell_b, \ell_{b'}) &\mapsto (\ell_{b \oplus b'}, f_{b \oplus b'}) & \forall b, b' \in \{0, 1\}, (\ell_b, f_{b'}) &\mapsto (\ell_b, f_b) \\ \forall q, q' \in Q, (q, q') &\mapsto_{=} (\text{dead}, \text{dead}) & \forall q \in Q, (q, \text{dead}) &\mapsto (\text{dead}, \text{dead}) \end{aligned}$$

where  $\mapsto_{=}$  denotes that the data of the agents must be equal,  $\mapsto$  without subscript means no condition on data (or equivalently, the transition exists both for equality and disequality), and  $\oplus$  denotes the XOR operator.  $\mathcal{P}$  is well-specified and computes the function  $\Pi$  that is equal to  $\top$  whenever there is an even number of appearing data and they all have exactly one corresponding agent. To see this, if there are two agents of equal datum, then all fair runs eventually have all agents on *dead* and stabilise to  $\perp$ . Otherwise, there will eventually be a single agent in  $\{\ell_0, \ell_1\}$ , and it will be on  $\ell_b$  if and only if the number of agents has parity  $b$ , in which case all other agents will eventually go to  $f_b$  and the run stabilises to  $\perp$  if  $b = 1$  (odd number of agents) and to  $\top$  if  $b = 0$  (even number of agents).

A more interesting but also more complex example is the majority protocol described in [9, Section 3]; it computes whether a datum has the absolute majority, i. e., strictly more agents than all other data combined.

Well-specification is the fundamental verification problem for population protocols. However, as we will see in Section 3, this problem is undecidable for PPUD.

► **Theorem 3.** *The well-specification problem for PPUD is undecidable.*

This motivates the study of the restricted class of *immediate observation* PPUD.

## 2.2 Immediate Observation Protocols

Immediate observation protocols [4] are a restriction of population protocols where, when two agents interact, one of the two agents does not change its state. The restriction of the model with data to immediate observation was first considered in [9].

<sup>1</sup> One often considers that a finite run  $\gamma_f \xrightarrow{*} \gamma_\ell$  is fair when there is no  $\gamma$  such that  $\gamma_\ell \rightarrow \gamma$ . In the following, we rule out this possibility by implicitly assuming that, for all  $q_1, q_2 \in Q$  and  $\bowtie \in \{=, \neq\}$ , it holds that  $((q_1, q_2), \bowtie, (q_1, q_2)) \in \Delta$ , and ignoring the trivial cases of runs with at most one agent.

► **Definition 4.** An *immediate observation population protocol with unordered data* (or *IOPPPUD*) is a PPUD  $\mathcal{P} = (Q, \Delta, I, O)$  where every transition  $\delta \in \Delta$  is of the form  $(q_1, q_2, \bowtie, q_1, q_3)$ , with  $q_1, q_2, q_3 \in Q$  and  $\bowtie \in \{=, \neq\}$ , i. e., the first agent does not change its state.

For IOPPPUD, we denote a transition  $(q_1, q_2, \bowtie, q_1, q_3)$  by  $q_2 \xrightarrow{\bowtie q_1} q_3$ . If we have a step  $\gamma \rightarrow \gamma'$  with transition  $q_2 \xrightarrow{\bowtie q_1} q_3$  that involves agents  $a, a_o \in \mathbb{A}$  where  $a$  is the agent moving from  $q_2$  to  $q_3$  and  $a_o$  is the agent in  $q_1$ , we denote it by  $\gamma \xrightarrow{\bowtie a_o} a \gamma'$ . We say that agent  $a$  *observes* agent  $a_o$ , and call  $a_o$  the *observed* agent. Intuitively,  $a$  “observes”  $a_o$  and reacts, whereas  $a_o$  may not even know it has been observed.

► **Example 5.** Consider the following IOPPPUD  $\mathcal{P} = (Q, \Delta, I, O)$ , with  $Q := \{q_0, q_1, q_2, q_3\}$ ,  $I := \{q_0, q_1\}$ ,  $O(q_3) = \top$ ,  $O(q) = \perp$  for all  $q \neq q_3$ , and transitions in  $\Delta$  as follows:

$$q_0 \xrightarrow{=q_1} q_2 \quad q_1 \xrightarrow{=q_0} q_2 \quad q_2 \xrightarrow{\neq q_2} q_3 \quad \forall q \in \{q_1, q_2\}, \forall \bowtie \in \{=, \neq\}, q \xrightarrow{\bowtie q_3} q_3$$

This protocol is well-specified: from  $\gamma_0 \in \Gamma_{\text{init}}$ , all fair runs stabilise to  $\top$  if two data have agents on both  $q_0$  and  $q_1$ , and all fair runs stabilise to  $\perp$  otherwise. Indeed, if there is a datum with agents on both  $q_0$  and  $q_1$ , by fairness eventually an agent with this datum is sent to  $q_2$ ; if there are two such data, then eventually some agent covers  $q_3$ , and then all agents are sent to  $q_3$  and the run stabilises to  $\top$ . Conversely, if it is not the case, then  $q_3$  cannot be covered and all fair runs stabilise to  $\perp$ .

Let  $\rho: \gamma_{\text{start}} \xrightarrow{*} \gamma_{\text{end}}$  be a run. Agent  $a_o$  is *internally observed* (resp. *externally observed*) in  $\rho$  if  $\rho$  contains a step of the form  $\gamma_1 \xrightarrow{=a_o} a \gamma_2$  (resp.  $\gamma_1 \xrightarrow{\neq a_o} a \gamma_2$ ); it is *observed* if one of the two cases holds. Similarly, a datum  $d$  is *observed* in  $\rho$  if an agent  $a$  with  $\text{dat}(a) = d$  is observed in  $\rho$ ; we define similarly a datum being internally or externally observed.

While the set of functions that can be computed by PPUD remains an open question, it is known that IOPPPUD exactly compute interval predicates [9], defined as follows. Let  $S$  be a finite set. A *simple interval predicate* over  $S$  is a formula  $\psi$  of the form  $\exists d_1, \dots, d_m, \bigwedge_{q \in S} \bigwedge_{j=1}^m \#(q, d_j) \in [A_{q,j}, B_{q,j}]$  where, for all  $q \in S$  and  $j \in [1, m]$ , we have  $A_{q,j} \in \mathbb{N}$  and  $B_{q,j} \in \mathbb{N} \cup \{+\infty\}$ . The dotted quantifiers quantify over *pairwise distinct* data. Formally, given a protocol  $\mathcal{P}$  with set of states  $Q$  such that  $S \subseteq Q$  and given  $\gamma \in \Gamma$ , the predicate  $\psi$  is *satisfied* by  $\gamma$  if there exist pairwise distinct data  $d_1, \dots, d_m \in \mathbb{D}$  such that for all  $q \in S$  and  $j \in [1, m]$ , it holds that  $\gamma_{d_j}^{\#}(q) \in [A_{q,j}, B_{q,j}]$  (resp.  $\gamma_{d_j}^{\#}(q) \in [A_{q,j}, B_{q,j}]$  in the case that  $B_{q,j} = +\infty$ ). An *interval predicate* over  $S$  is a Boolean combination  $\varphi$  of simple interval predicates over  $S$ ; we define that  $\varphi$  is *satisfied* by a configuration  $\gamma$  if the simple interval predicates satisfied by  $\gamma$  satisfy the Boolean combination.

► **Theorem 6** ([9], Theorem 18 and Corollary 29). *Given a finite set  $I$ , the functions computed by IOPPPUD with set of initial states<sup>2</sup>  $I$  are exactly the interval predicates over  $I$ .*

► **Example 7.** The protocol described in Example 2 and the majority protocol of [9, Section 3] cannot be turned into immediate observation protocols, as they compute functions that cannot be expressed as interval predicates. The immediate observation protocol from Example 5 computes the following interval predicate, which is actually a simple interval predicate:

$$\exists d_1, d_2, (\#(q_0, d_1) \geq 1) \wedge (\#(q_1, d_1) \geq 1) \wedge (\#(q_0, d_2) \geq 1) \wedge (\#(q_1, d_2) \geq 1).$$

<sup>2</sup> This does not limit the number of states of said protocols, as their set of states  $Q$  may be larger than  $I$ .

Given a simple interval predicate  $\psi = \exists d_1, \dots, d_m, \bigwedge_{q \in I} \bigwedge_{j=1}^m \#(q, d_j) \in [A_{q,j}, B_{q,j}]$ , we define its *width* as  $m$ , its *height*  $h$  as the maximum of all finite  $A_{q,j}$  and  $B_{q,j}$ , and its *size* as  $|I| \cdot m \cdot \log(h)$ . We also define the *width* (resp. *height*) of an interval predicate as the maximum of the widths (resp. heights) of its simple interval predicates, and its *size*, measuring the space taken by its encoding, as the sum of their sizes plus its number of Boolean operators.

► **Remark 8.** In [9], predicates refer to an input alphabet  $\Sigma$ , which is converted into initial states using an input mapping. For convenience, we have not included the input alphabet in our model, which is why we arbitrarily fix a set of initial states  $I$  in Theorem 6.

### 2.3 Generalised Reachability Expressions

We define a general class of specifications, called generalised reachability expressions, which are formulas constructed using interval predicates as atoms and using union, complement,  $\text{Post}^*$ , and  $\text{Pre}^*$  as operators. This concept is inspired by [28, Section 2.4], although our choice of atoms is more general and adapted to the data setting.

► **Definition 9.** Let  $\mathcal{P} = (Q, \Delta, I, O)$  be a protocol.

**Generalised Reachability Expressions (GRE)** over  $\mathcal{P}$  are produced by the grammar

$$E ::= \varphi \mid E \cup E \mid \overline{E} \mid \text{Post}^*(E) \mid \text{Pre}^*(E),$$

where  $\varphi$  ranges over interval predicates over  $Q$ .

Given a GRE  $E$ , we define the set of configurations defined by  $E$ , denoted  $\llbracket E \rrbracket_{\mathcal{P}}$ , as the set containing all configurations of  $\mathcal{P}$  that satisfy the formula, where the predicates are interpreted as above and the other operators are interpreted naturally (the overline denotes set complementation). This set is denoted  $\llbracket E \rrbracket$  when  $\mathcal{P}$  is clear from context.

The *length*  $|E|$  of a GRE  $E$  is its number of operators. Letting  $\varphi_1, \dots, \varphi_k$  be the interval predicates used as atoms in  $E$ , the *norm*  $\|E\|$  of  $E$  is the maximum of the heights and widths of the  $\varphi_i$ . Its *size* is the sum of the sizes of the  $\varphi_i$  plus  $|E|$ . The *emptiness problem for GRE* asks, given as input a protocol  $\mathcal{P}$  and a GRE  $E$  over  $\mathcal{P}$ , whether  $\llbracket E \rrbracket_{\mathcal{P}} = \emptyset$ . We will show in Section 6 that, for IOPPUD, this problem is decidable.

► **Theorem 10.** The emptiness problem for GRE over IOPPUD is in  $\text{EXPSPACE}$ .

We now argue that this decidability result is powerful, as it implies decidability of many classic problems on IOPPUD. We start with well-specification. We use the notation  $\forall d, \varphi$  as a short form for  $\neg \exists d, \neg \varphi$ . Given a PPUD  $\mathcal{P}$  and  $b \in \{\top, \perp\}$ , let  $\text{Out}_b := \forall d, \bigwedge_{q \in O^{-1}(\{b\})} \#(q, d) = 0$  be the GRE for  $b$ -consensus configurations; moreover, let  $\text{Stable}_b := \overline{\text{Pre}^*(\text{Out}_b)}$  be the GRE for *stable*  $b$ -consensus, i. e., configurations from which all runs lead to a  $b$ -consensus.

► **Proposition 11.** Let  $\mathcal{P}$  be a PPUD,  $E_{ws} := \Gamma_{\text{init}} \cap \text{Pre}^*(\overline{\text{Pre}^*(\text{Stable}_{\top})}) \cap \text{Pre}^*(\overline{\text{Pre}^*(\text{Stable}_{\perp})})$ .  $\mathcal{P}$  is well-specified if and only if  $\llbracket E_{ws} \rrbracket_{\mathcal{P}} = \emptyset$ .

**Proof.** First,  $\Gamma_{\text{init}} = \llbracket \forall d, \bigwedge_{q \notin I} \#(q, d) = 0 \rrbracket$  and  $E_{ws}$  is indeed a GRE over  $\mathcal{P}$ . For every  $\gamma \in \Gamma$ ,  $\text{Post}^*(\gamma)$  is finite as all configurations reachable from  $\gamma$  have the same number of agents. Therefore, a fair run  $\rho$  that visits  $\text{Pre}^*(\mathcal{S})$  infinitely often for  $\mathcal{S} \subseteq \Gamma$  must visit  $\mathcal{S}$  infinitely often. Let  $\gamma_0 \in \Gamma_{\text{init}}$  and  $b \in \{\top, \perp\}$ ; it suffices to prove that there is a fair run from  $\gamma_0$  that does *not* stabilise to  $b$  if and only if  $\gamma_0 \in \llbracket \text{Pre}^*(\overline{\text{Pre}^*(\text{Stable}_b)}) \rrbracket$ . If, from  $\gamma_0$ , one can reach  $\gamma \notin \llbracket \text{Pre}^*(\text{Stable}_b) \rrbracket$ , then one can build a fair run from  $\gamma_0$  that first goes to  $\gamma$ , and



then forever performs arbitrary steps in a fair way; since  $\gamma \notin \llbracket \text{Pre}^*(\text{Stable}_b) \rrbracket$ , it will stay in  $\llbracket \text{Pre}^*(\text{Out}_b) \rrbracket$ , so by fairness it visits  $\llbracket \text{Out}_b \rrbracket$  infinitely often and does not stabilise to  $b$ . Conversely, if there is a fair run that does not stabilise to  $b$ , then it never visits  $\llbracket \text{Stable}_b \rrbracket$ , hence by fairness it eventually stops visiting  $\llbracket \text{Pre}^*(\text{Stable}_b) \rrbracket$ ; this proves that it visits a configuration  $\gamma \in \llbracket \text{Pre}^*(\text{Stable}_b) \rrbracket$ , and  $\gamma$  is reachable from  $\gamma_0$  hence  $\gamma_0 \in \llbracket \text{Pre}^*(\text{Pre}^*(\text{Stable}_b)) \rrbracket$ . ◀

Many other problems can be expressed as emptiness problems for GRE; we list a few.

- The *correctness* problem for IOPPUD asks, given an IOPPUD  $\mathcal{P} = (Q, I, O, \Delta)$  and an interval predicate  $\varphi$  over  $I$ , whether  $\mathcal{P}$  computes  $\varphi$ . This can be equivalently phrased as  $\llbracket \varphi^{-1}(b) \cap \text{Pre}^*(\overline{\text{Pre}^*(\text{Stable}_b)}) \rrbracket = \emptyset$  for all  $b \in \{\top, \perp\}$ , where  $\varphi^{-1}(b)$  is the set of initial configurations that  $\varphi$  maps to  $b$ . Note that the previous expression is a GRE because, in Definition 9, we chose as atoms interval predicates such as  $\varphi$ .
- The *set-reachability* problem (called cube-reachability in [5]) asks, given two sets of configurations  $\mathcal{S}_1, \mathcal{S}_2$ , whether  $\mathcal{S}_2$  is reachable from  $\mathcal{S}_1$ ; this typically expresses safety problems where  $\mathcal{S}_2$  represents “bad configurations” that must not be reached. If  $\mathcal{S}_1 = \llbracket E_1 \rrbracket$  and  $\mathcal{S}_2 = \llbracket E_2 \rrbracket$ , then this amounts to checking whether  $\llbracket E_1 \cap \text{Pre}^*(E_2) \rrbracket$  is empty.
- The *home-space problem* asks, given a protocol  $\mathcal{P}$  and a set of configurations  $\mathcal{H}$ , whether  $\mathcal{H}$  can be reached from every configuration reachable from an initial configuration. If  $\mathcal{H}$  can be expressed as a GRE  $E$ , then it suffices to check whether  $\llbracket \text{Post}^*(\Gamma_{\text{init}}) \rrbracket \subseteq \llbracket \text{Pre}^*(E) \rrbracket$ . This problem has been studied in Petri nets [21], but also in probabilistic settings, for example in [10] for asynchronous shared-memory systems; indeed, in systems with uniform probabilistic schedulers where  $\text{Post}^*(\gamma_0)$  is finite for every initial configuration  $\gamma_0$ , this problem is equivalent to asking whether the probability of reaching  $\mathcal{H}$  is equal to 1.

Theorem 10 entails that, for IOPPUD, all these problems are decidable and in EXPSPACE.

### 3 Undecidability of Verification of Population Protocols with Unordered Data


In this section, we establish that the most fundamental verification problem for PPUD, i. e., the well-specification problem, is already undecidable.

► **Theorem 3.** *The well-specification problem for PPUD is undecidable.*

We proceed by reduction from the halting problem for 2-counter machines with zero-tests, a famously undecidable problem [25]. Here, we give a proof sketch. The detailed reduction can be found in Appendix A.

We fix a 2-counter machine and build a protocol  $\mathcal{P}$  which is *not* well-specified if and only if the counter machine halts. A *2-counter machine* performs increments, decrements and zero-tests on two counters. The main difficulty are the zero-tests. Let us first recall how increments and decrements are simulated in many prior undecidability results for population protocols and Petri nets [20, 26]. The protocol has a control part  $Q_{CM} := \{q_i \mid i \in [1, n]\}$  where a single *instruction agent* evolves; this part has one state per instruction of the machine. Increments and decrements are simulated as follows: The instruction agent interacts with states of  $\{R\} \cup \{x, y\}$ , where  $R$  is a *reservoir state* and  $x$  and  $y$  are states in which the number of agents represents the value of the counters. For example an increment on counter  $x$  moves one agent from the reservoir  $R$  to  $x$  and advances the instruction agent to the next instruction. The reservoir is hence implicitly assumed to start with arbitrarily many agents.

The main difficulty is that one does not want to take the  $= 0$  branch of a zero-test when the value of the counter is non-zero. Actually, similar to [20, 26], we will not prevent the



■ **Figure 1** For simplicity, we use Petri net notation: circles are states, rectangles are Petri net transitions. To encode this into our protocols, we split each transition into pairwise interactions.

existence of such runs. Instead, our protocol will have “violating” runs which take the wrong branch of a zero-test, but our well-specification check will consider only *violation-free* runs. The correctness of the reduction is then established in two steps: The CM halts if and only if some *violation-free* run to the halting state exists, and this is true if and only if our protocol is not well-specified. We establish the connection between non-well-specification and the existence of a *violation-free* run in our protocol.

In the first place, we guarantee that every initial configuration of our protocol has a fair run stabilising to  $\perp$ , so that  $\mathcal{P}$  is not well-specified if and only if there exists a fair run which does not stabilise to  $\perp$ <sup>3</sup>. Second, we introduce *violation detection*, a mechanism which guarantees that *fair* runs which contain a violation stabilise to  $\perp$ , hence preserving well-specification. To do so, we add a sink state  $q_{\perp}$ , which has output  $\perp$  and is attracting, i. e., all other states have a transition to  $q_{\perp}$  available when observing that  $q_{\perp}$  is non-empty. Violation detection then entails adding transitions into  $q_{\perp}$  that will be available infinitely often if the run (or its initial configuration) contained a violation. By fairness, any run containing a violation will then eventually put an agent into  $q_{\perp}$ , and hence, because  $q_{\perp}$  is attracting, the run ends in a deadlock with all agents in  $q_{\perp}$ . In particular, any fair run containing a violation will output  $\perp$  as claimed. There are two types of violation detection.

First, we want to only mark those runs as violation-free that start in initial configurations where  $U \in Q$  has at most one agent of each datum. To do so, we make agents remember whether their initial state was  $U$  or not (by encoding it into the state space), and, from every state, we add a transition to  $q_{\perp}$  such that this transition is enabled when an agent who started in  $U$  observes another agent of same datum that also started in  $U$ .

Second, we want to detect violations which consist in falsely simulating a zero-test, as discussed above. Here our technique shares some similarities with [26]. Let  $c \in \{x, y\}$  be a counter; for every zero-test of the counter machine, we add two types of transitions to the protocol. The first type simulates the  $c \neq 0$  branch and can be taken by the instruction agent upon interacting with some agent on state  $c$ ; by contrast, the  $c = 0$  branch can always be taken. However, if it was taken with  $c \neq 0$ , then *violation detection* will eventually detect this. For this mechanism, we introduce a counter control state  $\bar{c} \in Q$ . At any point in time,  $\bar{c}$  contains one agent, similar to the instruction agent. The crux of our violation detection is that only agents which share the datum with the agent in  $\bar{c}$  will be allowed to move in and out of state  $c$ , as illustrated in Figure 1a.

The  $= 0$  branch of a zero-test is depicted in Figure 1b. It replaces the agent on  $\bar{c}$  with an

<sup>3</sup> This can be done with the addition of a fresh state that is the only initial state and that has internal transitions to all former initial states and an internal transition to a sink state that has output  $\perp$ .



agent with fresh datum from state  $U$ . Thus, when the  $c = 0$  branch is taken, any remaining agent in  $c$  is stuck in  $c$  as it will never again share datum with the agent in  $\bar{c}$ . Violation detection then sends an agent in  $c$  to  $q_\perp$  upon observing an agent in  $\bar{c}$  with different datum.

Now that we have violation detection in place, it only remains to explain the connection to halting. The halting instruction  $q_n$  in  $Q_{CM}$  is the only state with output  $\top$ . Hence, any run not outputting  $\perp$  must contain an agent in the halting instruction at some point, and be violation-free by the above. That is, the counter machine reached the halting state without violations. Conversely, if the machine halts, one can build a finite run that puts an agent into the halting state without any violation occurring. The corresponding configuration is then a deadlock, and hence the extension to an infinite run (by staying there forever) is a fair run not outputting  $\perp$ . This proves that well-specification is undecidable for PPUD, which motivates restricting ourselves to immediate observation PPUD.

## 4 An Analysis of Immediate Observation Protocols with Data

To obtain our complexity bounds on the emptiness problem for GRE, we first show some transformations on runs that allow us to bound the number of observed agents. All runs that we consider in this section are finite, and we therefore write them as  $\gamma_1 \rightarrow \dots \rightarrow \gamma_m$  or  $\gamma_{start} \xrightarrow{*} \gamma_{end}$ . In the rest of this section, we fix an IOPUD  $\mathcal{P} = (Q, \Delta, I, O)$ .

We introduce some notation for agents in runs. Let  $\rho: \gamma_1 \xrightarrow{*} \gamma_m$  and  $d \in \mathbb{D}$ . We let  $\mathbb{A}_\rho$  be the set of agents appearing in  $\rho$ , and set  $\mathbb{A}_\rho^d := \{a \in \mathbb{A}_\rho \mid \mathbf{dat}(a) = d\}$ . We let  $\mathbb{A}_{\rho, o}^d$  be the set of agents with datum  $d$  that are observed in  $\rho$ , i. e., the  $a_o \in \mathbb{A}_\rho^d$  such that there exists a step  $\gamma_i \xrightarrow{\times a_o} \gamma_{i+1}$  in  $\rho$ . For all  $q_1, q_m \in Q$ , we let  $\mathbb{A}_{\rho, q_1, q_m}^d$  be the set of agents with datum  $d$  that start in  $q_1$  and end in  $q_m$ , i. e., the  $a \in \mathbb{A}_\rho^d$  such that  $\gamma_1(a) = q_1$  and  $\gamma_m(a) = q_m$ . Moreover, we let  $\mathbb{D}_\rho := \{d \in \mathbb{D} \mid \mathbb{A}_\rho^d \neq \emptyset\}$  be the set of data appearing in  $\rho$ . We may omit  $\rho$  in the subscript if the run is clear from the context.

### 4.1 Bounds on the Number of Observed Agents per Datum


Let  $\rho: \gamma_1 \rightarrow \dots \rightarrow \gamma_m$  be a run. For  $i \in [1, m]$ , we call  $\gamma_i \rightarrow \gamma_{i+1}$  the  $i$ -th step in  $\rho$ . Let  $\rho[\rightarrow i]$  (resp.  $\rho[i \rightarrow]$ ) denote the prefix of  $\rho$  ending on its  $i$ -th configuration (resp. the suffix of  $\rho$  starting on its  $i$ -th configuration). Let  $a, b \in \mathbb{A}_\rho$ . Agent  $a$  is *active* in the  $i$ -th step if  $\gamma_i \xrightarrow{\times a_o} \gamma_{i+1}$  for some agent  $a_o$ . Otherwise,  $a$  is *idle* in that step. We say  $b$  *copies*  $a$  in  $\rho$  if after every step  $\gamma_i \xrightarrow{\times a_o} \gamma_{i+1}$  in  $\rho$  via some transition  $t$ , there is a step  $\gamma_{i+1} \xrightarrow{\times a_o} \gamma_{i+2}$  via  $t$  and, additionally,  $b$  is idle in every step not immediately following an active step of  $a$ .

The following lemma allows us to add agents to a run that copy an agent of the same datum.

► **Lemma 12** (Agents copycat). *Let  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run. Let  $a \in \mathbb{A}_\rho$  and  $\tilde{a} \in \mathbb{A} \setminus \mathbb{A}_\rho$  with  $\mathbf{dat}(a) = \mathbf{dat}(\tilde{a})$ . Then there exist configurations  $\tilde{\gamma}_{start}, \tilde{\gamma}_{end}$  and a run  $\tilde{\rho}: \tilde{\gamma}_{start} \xrightarrow{*} \tilde{\gamma}_{end}$  such that:*

- (i)  $\mathbb{A}_{\tilde{\rho}} = \mathbb{A}_\rho \uplus \{\tilde{a}\}$ , and for all  $a' \in \mathbb{A}_\rho$ ,  $\tilde{\gamma}_{start}(a') = \gamma_{start}(a')$  and  $\tilde{\gamma}_{end}(a') = \gamma_{end}(a')$ ;
- (ii)  $\tilde{\gamma}_{start}(\tilde{a}) = \gamma_{start}(a)$  and  $\tilde{\gamma}_{end}(\tilde{a}) = \gamma_{end}(a)$ ;
- (iii)  $\tilde{a}$  is not observed in  $\tilde{\rho}$ .

**Proof.** We let  $\tilde{\gamma}_{start}$  be such that  $\tilde{\gamma}_{start}(\tilde{a}) = \gamma_{start}(a)$  and  $\tilde{\gamma}_{start}(a') = \gamma_{start}(a')$  for all  $a' \neq \tilde{a}$ . We construct  $\tilde{\rho}$  by going through  $\rho$  step by step, making  $\tilde{a}$  copy  $a$ : whenever  $\rho$  takes a step  $\xrightarrow{\times a_o} \gamma_a$ , then we take this step followed by step  $\xrightarrow{\times a_o} \tilde{\gamma}_{\tilde{a}}$  to  $\tilde{\rho}$ . We can do so because  $\mathbf{dat}(\tilde{a}) = \mathbf{dat}(a)$  and because agent  $a_o \neq a$  has not moved and thus can be observed again. These are the only steps where  $\tilde{a}$  is involved, hence it is never observed. ◀



■ **Figure 2** An example of a run with six steps on a protocol with two states  $q_1, q_2$ .  $a, b, c, d, e$  denote agents;  $a, b, c$  have the same datum and  $d, e$  have the same datum. Dashed lines are observations.

The following result shows that, given a run  $\rho$ , we can construct a new run with a small subset of the agents of  $\mathbb{A}_\rho$  such that, for all  $d \in \mathbb{D}$  and all states  $q_1$  and  $q_2$ , if there is a  $d$ -agent starting in  $q_1$  and ending in  $q_2$  in  $\rho$ , then this is also true in the new run. The full proof can be found in Appendix B.

► **Lemma 13 (Agents core).** *Let  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run. Then there exist configurations  $\gamma'_{start}, \gamma'_{end}$  and a run  $\rho': \gamma'_{start} \xrightarrow{*} \gamma'_{end}$  with  $\mathbb{A}_{\rho'} \subseteq \mathbb{A}_\rho$  such that:*

- (i) for all  $a \in \mathbb{A}_{\rho'}$ ,  $\gamma'_{start}(a) = \gamma_{start}(a)$  and  $\gamma'_{end}(a) = \gamma_{end}(a)$ ;
- (ii) for all  $d \in \mathbb{D}$  and  $q_s, q_e \in Q$ , if  $\mathbb{A}_{\rho, q_s, q_e}^d \neq \emptyset$ , then  $\mathbb{A}_{\rho', q_s, q_e}^d \neq \emptyset$ ;
- (iii) for all  $d \in \mathbb{D}$ , we have  $|\mathbb{A}_{\rho'}^d| \leq |Q|^3$ .

**Proof sketch.** We adapt the bunch argument from the case of IO protocols without data [18]. Suppose there is  $d \in \mathbb{D}$  and  $q_s, q_e \in Q$  such that  $|\mathbb{A}_{\rho, q_s, q_e}^d| > |Q|$ . Let  $\mathcal{R}$  be the set of states visited by agents of  $\mathbb{A}_{\rho, q_s, q_e}^d$  in  $\rho$ . Notice that  $|\mathcal{R}| \leq |Q|$ . We define a family  $(a_q)_{q \in \mathcal{R}}$  of pairwise distinct agents such that reducing  $\mathbb{A}_{\rho, q_s, q_e}^d$  in  $\rho$  to  $(a_q)_{q \in \mathcal{R}}$  still yields a valid run.

We iterate through  $\mathcal{R}$  as follows. Let  $q \in \mathcal{R}$  and let  $f$  be the first moment  $q$  is reached in  $\rho$ , i. e., the minimal index such that there exists an  $a \in \mathbb{A}_{\rho, q_s, q_e}^d$  with  $\gamma_f(a) = q$ . Let  $\ell$  be the last moment  $q$  is occupied in  $\rho$ , i. e., the maximal index such that there exists an  $a \in \mathbb{A}_{\rho, q_s, q_e}^d$  with  $\gamma_\ell(a) = q$ . Let  $\alpha_q$  be the agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$  that reaches  $q$  first, i. e.,  $\gamma_f(\alpha_q) = q$ , and let  $\beta_q$  be the agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$  that leaves  $q$  last, i. e.,  $\gamma_\ell(\beta_q) = q$ . Note that these agents do not have to be distinct. We pick a fresh agent  $a_q \notin \mathbb{A}_\rho$  with  $\mathbf{dat}(a_q) = d$  and modify  $\rho$  as follows. We let  $a_q$  copy  $\alpha_q$  in  $\rho[\rightarrow f]$ , then  $a_q$  stays idle until  $\beta_q$  leaves  $q$  (for the last time) and then  $a_q$  copies  $\beta_q$  in  $\rho[\ell \rightarrow]$ . We do this for every  $q \in \mathcal{R}$ .

Then, for every step in which an  $a_o$  in  $\mathbb{A}_{\rho, q_s, q_e}^d$  is observed in state  $q$ , let  $a_q$  be observed instead, i. e., replace steps  $\xrightarrow{a_o} a$  with  $\xrightarrow{a_q} a$ . Finally, remove all the agents of  $\mathbb{A}_{\rho, q_s, q_e}^d$  from the run, and identify (or substitute) each  $a_q$  with a distinct agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$ , so that  $(a_q)_{q \in \mathcal{R}} \subseteq \mathbb{A}_{\rho, q_s, q_e}^d$ . We do this for every  $d \in \mathbb{D}$  and  $q_s, q_e \in Q$  such that  $|\mathbb{A}_{\rho, q_s, q_e}^d| > |Q|$ . ◀

► **Example 14.** Consider the run  $\rho$  depicted in Figure 2. Applying Lemma 13 on  $\rho$  yields a new run  $\rho'$  with 4 agents instead of 5. Indeed, let  $d$  denote the datum of  $a, b$  and  $c$ ; we have  $|\mathbb{A}_{\rho, q_1, q_1}^d| = |\{a, b, c\}| = 3$  whereas  $|Q| = 2$ . In  $\rho$ , agents  $a$  and  $b$  successively go from  $q_1$  to  $q_2$  and back to  $q_1$ . In  $\rho'$ , these two agents are replaced by a single agent (named  $b$  again) who goes to  $q_2$  on the first step and only leaves  $q_2$  on the last step. In  $\rho'$ , the new agent  $b$  is observed by  $d$  in the second step, and by  $e$  in the penultimate step.

## 4.2 Bounds on the Number of Observed Data

Given a run and a datum  $d$  appearing in it, we define the *trace* of  $d$  in  $\rho$  as the function  $\mathbf{tr}_\rho^d: Q^2 \rightarrow \mathbb{N}$  such that for all  $q_1, q_2 \in Q$ , it holds that  $\mathbf{tr}_\rho^d(q_1, q_2) = |\mathbb{A}_{\rho, q_1, q_2}^d|$ . For each pair

of states  $q_1, q_2$ , the trace counts the number of  $d$ -agents starting in  $q_1$  and ending in  $q_2$ . For example, the trace of the run  $\rho$  of Example 14, with  $d$  the datum of agents  $a, b$  and  $c$ , is such that  $\text{tr}_\rho^d(q_1, q_1) = 3$  and  $\text{tr}_\rho^d(q, q') = 0$  for all  $(q, q') \neq (q_1, q_1)$ . The trace is the information we need to copy data: if there is a datum  $d$  with trace  $tr$  in a run, then we can add data to the run that mimic  $d$  and have the same trace. The following lemma echoes Lemma 12.

► **Lemma 15** (Data copycat). *Let  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run. Let  $d \in \mathbb{D}_\rho$  and  $\tilde{d} \in \mathbb{D} \setminus \mathbb{D}_\rho$ . Then there exist configurations  $\tilde{\gamma}_{start}, \tilde{\gamma}_{end}$  and a run  $\tilde{\rho}: \tilde{\gamma}_{start} \xrightarrow{*} \tilde{\gamma}_{end}$  such that:*

- (i)  $\mathbb{A}_{\tilde{\rho}} = \mathbb{A}_\rho \uplus \mathbb{A}_{\tilde{\rho}}^{\tilde{d}}$ , and for all  $a \in \mathbb{A}_\rho$ ,  $\tilde{\gamma}_{init}(a) = \gamma_{init}(a)$  and  $\tilde{\gamma}_{end}(a) = \gamma_{end}(a)$ ,
- (ii)  $\text{tr}_{\tilde{\rho}}^{\tilde{d}} = \text{tr}_\rho^d$  and  $\text{tr}_{\tilde{\rho}}^{d'} = \text{tr}_\rho^{d'}$  for all  $d' \neq \tilde{d}$ ,
- (iii)  $\tilde{d}$  is not externally observed in  $\tilde{\rho}$ .

**Proof.** For all  $q_s, q_e \in Q$  and all  $a \in \mathbb{A}_{q_s, q_e}^d$ , we add an agent  $\tilde{a}$  with datum  $\tilde{d}$  in  $q_s$  at the start. We do this in a way similar to Lemma 12: after every step  $\xrightarrow{\neq a_o} a$  in  $\rho$ , we insert a step  $\xrightarrow{\neq \tilde{a}_o} \tilde{a}$ , and after every step  $\xrightarrow{= a_o} a$  in  $\rho$ , we insert a step  $\xrightarrow{= \tilde{a}_o} \tilde{a}$ . We thus maintain the fact that each added agent  $\tilde{a}$  is in the same state as its counterpart  $a$ . In particular, they are in the same state at the end of the run. This yields a run  $\tilde{\rho}$  with  $\text{tr}_{\tilde{\rho}}^{\tilde{d}} = \text{tr}_\rho^d$ , and such that for all  $d' \neq \tilde{d}$ ,  $\text{tr}_{\tilde{\rho}}^{d'} = \text{tr}_\rho^{d'}$ . Since  $\tilde{d} \notin \mathbb{D}_\rho$ , it is not externally observed in  $\tilde{\rho}$ . ◀

Like we showed for the agents, we show that we can reduce the number of data in a run. We lift the proof strategy of Lemma 13 from agents to data, exploiting the sets of data with equal traces. The full proof is in Appendix B.

► **Lemma 16** (Data core). *Let  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run and let  $K$  be a number such that there are at most  $K$  agents of each datum in  $\rho$ . Then there exist configurations  $\gamma'_{start}, \gamma'_{end}$ , a run  $\rho': \gamma'_{start} \xrightarrow{*} \gamma'_{end}$ , and a subset of data  $\mathbb{D}_{\rho'} \subseteq \mathbb{D}_\rho$  such that:*

- (i) for all  $d \in \mathbb{D}_{\rho'}$  and all agents  $a$  of datum  $d$ ,  $\gamma_{start}(a) = \gamma'_{start}(a)$  and  $\gamma_{end}(a) = \gamma'_{end}(a)$ ,
- (ii) for all  $d \in \mathbb{D}_\rho$ , there exists  $d' \in \mathbb{D}_{\rho'}$  such that  $\text{tr}_{\rho'}^{d'} = \text{tr}_\rho^d$ ,
- (iii)  $|\mathbb{D}_{\rho'}| \leq (K + 1)^{|Q|^3 + |Q|^2}$ .

**Proof sketch.** We define the notion of split trace. The split trace of a datum  $d$  at the  $i$ -th configuration of a run  $\rho$  maps every triple of states  $(q_1, q_2, q_3)$  to the number of  $d$ -agents that are in  $q_1$  at the start of  $\rho$ , then in  $q_2$  in the  $i$ -th configuration, and finally in  $q_3$  at the end. Since there are at most  $K$  agents per datum, there are at most  $(K + 1)^{|Q|^2}$  possible traces and  $M = (K + 1)^{|Q|^3}$  possible split traces.

For every trace  $tr$ , if there are more than  $M$  data that have trace  $tr$  in  $\rho$ , we apply a similar argument to Lemma 13: we select one datum for each possible split trace, and use it to cover all external observations of agents whose datum matches that split trace. We remove the other data, and show that this is still a valid run. The bound on the total number of data comes from the number of traces and split traces. ◀

► **Corollary 17.** *For every run  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$ , there exists a run  $\tilde{\rho}: \gamma_{start} \xrightarrow{*} \gamma_{end}$  such that for all  $d \in \mathbb{D}$ , it holds that  $|\mathbb{A}_{\tilde{\rho}, o}^d| \leq |Q|^3$  and that agents of at most  $(|Q|^3 + 1)^{|Q|^3 + |Q|^2}$  different data are externally observed.*

**Proof.** We first apply Lemma 13 to  $\rho$  to obtain  $\rho^{(1)}: \gamma_{start}^{(1)} \xrightarrow{*} \gamma_{end}^{(1)}$  over the same data such that for all  $d \in \mathbb{D}$ , it holds that  $|\mathbb{A}_{\rho^{(1)}}^d| \leq |Q|^3$ . Then we apply Lemma 16 to obtain  $\rho^{(2)}: \gamma_{start}^{(2)} \xrightarrow{*} \gamma_{end}^{(2)}$  with at most  $(|Q|^3 + 1)^{|Q|^3 + |Q|^2}$  data. By Lemma 13-(i) and Lemma 16-(i), the remaining agents have the same initial and final states in  $\rho$  and  $\rho^{(2)}$ . It remains to put

back the agents and data we removed, without increasing the number of externally observed data or observed agents per datum.

By Lemma 16-(ii), every trace of a datum in  $\rho^{(1)}$  appears as the trace of a datum in  $\rho^{(2)}$ . Thus, it is possible to re-add data of  $\mathbb{D}_{\rho^{(1)}}$  to  $\mathbb{D}_{\rho^{(2)}}$  using repeated applications of Lemma 15. By Lemma 15-(iii), this does not add any external observation. So we obtain a run  $\tilde{\rho}^{(1)}$  from  $\gamma_{start}^{(1)}$  to  $\gamma_{end}^{(1)}$  such that at most  $(|Q|^3 + 1)^{|Q|^3 + |Q|^2}$  data are externally observed by Lemma 15-(iii). Recall that there are at most  $|Q|^3$  agents per datum in  $\gamma_{start}^{(1)}$  by Lemma 13-(iii); in particular there are at most  $|Q|^3$  observed agents per datum in  $\tilde{\rho}^{(1)}$ .

By Lemma 13-(ii), for each datum  $d$  and states  $q_s, q_e$ , if there is a  $d$ -agent  $a$  such that  $\gamma_{start}(a) = q_s$  and  $\gamma_{end}(a) = q_e$  then there is an agent  $a'$  such that  $\gamma_{start}^{(1)}(a') = q_s$  and  $\gamma_{end}^{(1)}(a') = q_e$  in  $\rho$ . Therefore, due to Lemma 13-(ii), we can apply Lemma 12 repeatedly to add back the missing agents in  $\tilde{\rho}^{(1)}$  and obtain a run  $\check{\rho}$  from  $\gamma_{start}$  to  $\gamma_{end}$ . By Lemma 12-(iii), this does not add any observation. As a result, we obtain a run from  $\gamma_{start}$  to  $\gamma_{end}$  in which at most  $(|Q|^3 + 1)^{|Q|^3 + |Q|^2}$  data are externally observed and for all datum  $d$ , at most  $|Q|^3$   $d$ -agents are observed. ◀

## 5 From Expressions to Containers

In this section, we define the technical notions of boxes and containers, which are meant to represent sets of configurations defined by counting agents and data up to some thresholds. In Proposition 21, we will prove that the set of configurations defined by a generalised reachability expression  $E$  can be described as a union of containers whose thresholds are exponential in the length of  $E$  and polynomial in its norm. To do so, we will leverage the bounds on the number of observed agents from Section 4 to bound the description of the GRE  $\text{Post}^*(F)$  with respect to the one of GRE  $F$ . The key result of Proposition 21 will be used in Section 6 to obtain the decidability of the emptiness problem for GRE.


### 5.1 Equivalence of Predicates and Containers

In this subsection, we fix an IOPPUD  $\mathcal{P} = (Q, \Delta, I, O)$ .

Let  $n, M \in \mathbb{N}$ . An  $n$ -box is a vector  $\mathbf{b}: Q \rightarrow [0, n]$ . Given a configuration  $\gamma$  and a datum  $d \in \mathbb{D}$ , we define the  $n$ -box of  $d$  in  $\gamma$  as  $\lceil \gamma, d \rceil^n: Q \rightarrow [0, n]$  such that for all  $q \in Q$ ,  $\lceil \gamma, d \rceil^n(q) = \min\{n, \gamma_d^\#(q)\}$ ; in words, the  $n$ -box of  $d$  truncates the number of agents of  $d$  if it exceeds  $n$ . We write  $\mathbf{Boxes}_n$  for the set of all  $n$ -boxes. We define the equivalence relation  $\equiv_n$  over  $\Gamma \times \mathbb{D}$  by  $(\gamma_1, d_1) \equiv_n (\gamma_2, d_2)$  whenever  $\lceil \gamma_1, d_1 \rceil^n = \lceil \gamma_2, d_2 \rceil^n$ . An equivalence class of  $\equiv_n$  is a set of the form  $\{(\gamma, d) \in \Gamma \times \mathbb{D} \mid \lceil \gamma, d \rceil^n = \mathbf{b}\}$  for  $\mathbf{b} \in \mathbf{Boxes}_n$ ; we represent such an equivalence class for  $\equiv_n$  by the associated  $n$ -box  $\mathbf{b}$ .

To lift this concept to data, we count the number of data with the same  $n$ -box up to bound  $M$ . The  $(n, M)$ -container of a configuration  $\gamma$  is the function  $\lceil \gamma \rceil^{n, M}: \mathbf{Boxes}_n \rightarrow [0, M]$  such that  $\lceil \gamma \rceil^{n, M}(\mathbf{b}) = \min\{M, |\{d \in \mathbb{D} \mid \lceil \gamma, d \rceil^n = \mathbf{b}\}|\}$  for all  $\mathbf{b} \in \mathbf{Boxes}_n$ . We define the equivalence relation  $\equiv_{n, M}$  over  $\Gamma$  by  $\gamma_1 \equiv_{n, M} \gamma_2$  whenever  $\lceil \gamma_1 \rceil^{n, M} = \lceil \gamma_2 \rceil^{n, M}$ . An equivalence relation for  $\equiv_{n, M}$  is the preimage of some  $(n, M)$ -container by the previously described function; we represent such an equivalence class by the associated  $(n, M)$ -container. Figure 3 illustrates the function mapping a given configuration to its container.

In all the following, we use the terms  $n$ -boxes and  $(n, M)$ -containers to designate both the vectors and the equivalence classes of  $\equiv_n$  and  $\equiv_{n, M}$  that they represent. For instance, we write *union of  $n$ -boxes* for the union of the corresponding equivalence classes of  $\equiv_n$ .



■ **Figure 3** How a configuration is mapped to a  $(4, 2)$ -container. Here, the protocol has five states  $q_0, \dots, q_4$ . Five distinct data appear in the configuration, and they are represented using symbols.

The partition of  $\Gamma$  into  $(n, M)$ -containers becomes finer as  $n$  and  $M$  grow.

► **Lemma 18.** *Let  $n_1, n_2, M_1, M_2 \in \mathbb{N}$ . If  $n_1 \leq n_2$  and  $M_1 \leq M_2$ , then every  $(n_1, M_1)$ -container is a union of  $(n_2, M_2)$ -containers.*

Algorithmically, we represent an  $n$ -box as a list of appearing states with associated numbers from  $[1, n]$  encoded in binary. Similarly, we represent an  $(n, M)$ -container as a list of appearing  $n$ -boxes with associated numbers from  $[1, M]$  encoded in binary.

In fact, interval predicates exactly describe finite unions of containers.

► **Proposition 19.** *The sets of configurations defined by interval predicates of height at most  $n$  and width at most  $M$  are exactly the sets formed by unions of  $(n, M)$ -containers.*

**Proof sketch.** For the translation from predicates to containers, consider a simple interval predicate  $\exists d_1, \dots, d_M, \bigwedge_{q \in Q} \bigwedge_{j=1}^M \#(q, d_j) \in [A_{q,j}, B_{q,j}]$  of height  $n$ . This predicate cannot distinguish data mapped to the same  $n$ -box, hence cannot distinguish configurations in the same equivalence class for  $\equiv_{n,M}$ , i. e.,  $(n, M)$ -containers. The same directly extends to interval predicates.

For the other direction, we prove that a given  $(n, M)$ -container can be expressed as an interval predicate of height at most  $n$  and width at most  $M$ . To do so, given a box  $\mathbf{b} \in \mathbf{Boxes}_n$  and  $m \leq M$ , we define the simple interval predicate  $\varphi_{\mathbf{b}, \geq m}$  expressing that at least  $m$  data are mapped to box  $\mathbf{b}$ . Formally,  $\varphi_{\mathbf{b}, \geq m} := \exists d_1, \dots, d_m, \bigwedge_{q \in Q} \bigwedge_{j=1}^m \#(q, d_j) \in [A_q, B_q]$ , where, for all  $q \in Q$ ,  $A_q := \mathbf{b}(q)$ ,  $B_q := \mathbf{b}(q)$  if  $\mathbf{b}(q) < n$  and  $B_q := +\infty$  if  $\mathbf{b}(q) = n$ . This predicate has height at most  $n$  and width at most  $M$ . A Boolean combination of such predicates allows us to express an  $(n, M)$ -container. The full proof can be found in Appendix C.1. ◀

We therefore have two equivalent representations. Both are useful: interval predicates allow us to express properties more naturally, but containers are more convenient for the

proofs in the remainder of this section. While they are equally expressive, each can be much more succinct than the other, as stated below. We refer to Appendix C.1 for details.

► **Remark 20.** Containers can be exponentially more succinct than interval predicates, while interval predicates can be doubly exponentially more succinct than unions of containers.

## 5.2 A Translation from Expressions to Containers

Based on the translation from interval predicates to containers from Proposition 19, we can now show that for all generalised reachability expressions  $E$  over an IOPPUD  $\mathcal{P}$ , the set  $\llbracket E \rrbracket_{\mathcal{P}}$  is a union of  $(n, M)$ -containers with  $n$  and  $M$  bounded in terms of  $E$  and  $\mathcal{P}$ .

► **Proposition 21.** *There is a polynomial function  $\text{poly}: \mathbb{N} \rightarrow \mathbb{N}$  such that for all IOPPUD  $\mathcal{P}$  and GRE  $E$ , the set  $\llbracket E \rrbracket_{\mathcal{P}}$  is a union of  $(\|E\| \cdot (\text{poly}(|\mathcal{P}|))^{|E|}, \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |E|^2})$ -containers.*

The detailed proof of Proposition 21 can be found in Appendix C.3. We prove the result by structural induction on  $E$ . The base case, when  $E$  is an interval predicate, is provided by Proposition 19. For the induction step, handling Boolean operators is straightforward; the difficulty lies in operators  $\text{Pre}^*$  and  $\text{Post}^*$ . This is handled by the following lemma, which relies on the bounds from Section 4.

Equivalence classes for fixed values of  $n$  and  $M$  do not behave well with respect to the reachability relation, in the sense that it can happen that  $\gamma_{start} \xrightarrow{*} \gamma_{end}$  and  $\gamma_{start} \equiv_{n,M} \chi_{start}$ , but there is no  $\chi_{end} \equiv_{n,M} \gamma_{end}$  such that  $\chi_{start} \xrightarrow{*} \chi_{end}$ . However, this will hold if we take some margin on the equivalence relation of configurations at the start; the following two functions express this margin. For all  $n, M \in \mathbb{N}$ , let  $f(n) := (n + |\mathcal{P}|^3) \cdot |\mathcal{P}|$  and  $g(n, M) := (M + (|\mathcal{P}|^3 + 1)^{|\mathcal{P}|^3 + |\mathcal{P}|^2})(n + 1)^{|\mathcal{P}|}$ .

The following lemma states that, if a set of configurations  $C$  cannot distinguish  $\equiv_{n,M}$ -equivalent configurations, then  $\text{Pre}^*(C)$  cannot distinguish  $\equiv_{f(n), g(n, M)}$ -equivalent configurations. In other words, if  $C$  is a union of  $\equiv_{n,M}$ -equivalence classes, (i. e., of  $(n, M)$ -containers), then  $\text{Pre}^*(C)$  is a union of  $\equiv_{f(n), g(n, M)}$ -equivalence classes.

► **Lemma 22.** *For all  $n, M \in \mathbb{N}$  and all configurations  $\gamma_{start}, \gamma_{end}, \chi_{start} \in \Gamma$ , if there is a run  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  and  $\gamma_{start} \equiv_{f(n), g(n, M)} \chi_{start}$ , then there is a configuration  $\chi_{end} \in \Gamma$  with  $\gamma_{end} \equiv_{n, M} \chi_{end}$  and a run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$ .*

**Proof sketch.** We first apply Corollary 17 to  $\rho$ , so that we can assume that  $\rho$  has a limited number of externally observed data and of observed agents per datum.

In this proof sketch, we first handle the case with only one datum. Then, we explain how to generalise this. The full proof can be found in Appendix C.2.

Suppose that all agents in  $\gamma_{start}$  and  $\chi_{start}$  share a single datum  $d$ , and suppose  $(\gamma_{start}, d) \equiv_{f(n)} (\chi_{start}, d)$ . Let  $\mathbb{A}_{\gamma}$  and  $\mathbb{A}_{\chi}$  be the agents in  $\gamma_{start}$  and  $\chi_{start}$ , respectively. For all  $q, q' \in Q$ , we set  $\mathbb{A}_{\gamma}^{q \rightarrow} := \{a \in \mathbb{A}_{\gamma} \mid \gamma_{start}(a) = q\}$ ,  $\mathbb{A}_{\chi}^{q \rightarrow} := \{a \in \mathbb{A}_{\chi} \mid \chi_{start}(a) = q\}$ ,  $\mathbb{A}_{\gamma}^{\rightarrow q'} := \{a \in \mathbb{A}_{\gamma} \mid \gamma_{end}(a) = q'\}$ , and  $\mathbb{A}_{\gamma}^{q \rightarrow q'} := \mathbb{A}_{\gamma}^{q \rightarrow} \cap \mathbb{A}_{\gamma}^{\rightarrow q'}$ .

Our aim is to assign to each agent in  $\chi_{start}$  an agent in  $\gamma_{start}$  to mimic. To do so, we construct a mapping  $\nu: \mathbb{A}_{\chi} \rightarrow \mathbb{A}_{\gamma}$  such that

- (A) for all  $a \in \mathbb{A}_{\chi}$ , we have  $\chi_{start}(a) = \gamma_{start}(\nu(a))$ ,
- (B) for all  $a' \in \mathbb{A}_{\gamma}$  observed in  $\rho$ , we have  $\nu^{-1}(a') \neq \emptyset$ , and
- (C) for all  $q' \in Q$ , we have  $|\nu^{-1}(\mathbb{A}_{\gamma, d}^{\rightarrow q'})| = |\mathbb{A}_{\chi, d}^{\rightarrow q'}|$ , or  $|\nu^{-1}(\mathbb{A}_{\gamma, d}^{\rightarrow q'})| \geq n$  and  $|\mathbb{A}_{\gamma, d}^{\rightarrow q'}| \geq n$ .



We build  $\nu$  separately on each set  $\mathbb{A}_\chi^{q \rightarrow}$  by defining, for each  $q \in Q$ , a mapping  $\nu_q: \mathbb{A}_\chi^{q \rightarrow} \rightarrow \mathbb{A}_\gamma^{q \rightarrow}$ . Let  $q \in Q$ . As  $(\chi_{start}, d) \equiv_{f(n)} (\gamma_{start}, d)$ , either  $|\mathbb{A}_\gamma^{q \rightarrow}| = |\mathbb{A}_\chi^{q \rightarrow}|$ , or both  $|\mathbb{A}_\gamma^{q \rightarrow}|$  and  $|\mathbb{A}_\chi^{q \rightarrow}|$  are at least  $f(n)$ . If  $|\mathbb{A}_\gamma^{q \rightarrow}| = |\mathbb{A}_\chi^{q \rightarrow}|$ , we let  $\nu_q$  form a bijection between  $\mathbb{A}_\chi^{q \rightarrow}$  and  $\mathbb{A}_\gamma^{q \rightarrow}$ . Consider now the second case, where  $|\mathbb{A}_\gamma^{q \rightarrow}|$  and  $|\mathbb{A}_\chi^{q \rightarrow}|$  are at least  $f(n)$ . We aim at selecting, for every  $q' \in Q$ , a set  $A_{q \rightarrow q'} \subseteq \mathbb{A}_\gamma^{q \rightarrow q'}$  of agents that must be copied in  $\pi$ . If  $|\mathbb{A}_\gamma^{q \rightarrow q'}| \leq n$ , then we let  $A_{q \rightarrow q'} := \mathbb{A}_\gamma^{q \rightarrow q'}$ . Otherwise, we first put in  $A_{q \rightarrow q'}$  all agents in  $\mathbb{A}_\gamma^{q \rightarrow q'}$  that are observed in  $\rho$ , at most  $|\mathcal{P}|^3$  in total by Corollary 17. If  $|A_{q \rightarrow q'}| < n$ , we add arbitrary agents from  $\mathbb{A}_\gamma^{q \rightarrow q'}$  to  $A_{q \rightarrow q'}$  until  $|A_{q \rightarrow q'}| \geq n$ . Either way, we have selected  $A_{q \rightarrow q'}$  of size at most  $|\mathcal{P}|^3 + n$  for each  $q'$ , hence at most  $f(n)$  agents in total. For every  $q'$ , we have  $|A_{q \rightarrow q'}| \leq |\mathbb{A}_\gamma^{q \rightarrow q'}|$ , and either  $|A_{q \rightarrow q'}| = |\mathbb{A}_\gamma^{q \rightarrow q'}|$  or the two sets have size more than  $n$ .

We now build  $\nu_q$  such that its image over  $\mathbb{A}_\chi^{q \rightarrow}$  is  $\bigcup_{q' \in Q} A_{q \rightarrow q'}$ . We build this in two steps. First, we assign to each  $\bigcup_{q' \in Q} A_{q \rightarrow q'}$  one antecedent by  $\nu_q$  in  $\mathbb{A}_\chi^{q \rightarrow}$ . This is possible because  $|\bigcup_{q' \in Q} A_{q \rightarrow q'}| \leq f(n) \leq |\mathbb{A}_\chi^{q \rightarrow}|$ . We then identify some  $q''$  such that  $|A_{q \rightarrow q''}| > n$  and map all remaining agents of  $\mathbb{A}_\chi^{q \rightarrow}$  to an arbitrary agent in  $A_{q \rightarrow q''}$ . Such a  $q''$  exists because  $|\mathbb{A}_\gamma^{q \rightarrow}| \geq f(n) \geq n \cdot |\mathcal{P}|$ , so there is a  $q'$  such that  $|\mathbb{A}_\gamma^{q \rightarrow q'}| \geq n$ , and hence  $|A_{q \rightarrow q'}| \geq n$  by construction.

This concludes the construction of  $\nu$ . It remains to prove that  $\nu$  fulfils Items (A)–(C). Items (A) and (B) are immediate from the definition. We prove Item (C). Let  $q' \in Q$ . We distinguish two cases:

- if  $|\mathbb{A}_\gamma^{q \rightarrow q'}| < n$  for all  $q \in Q$ , then for all  $q$ , we have  $|\nu^{-1}(\mathbb{A}_\gamma^{q \rightarrow q'})| = |\nu^{-1}(A_{q \rightarrow q'})| = |A_{q \rightarrow q'}| = |\mathbb{A}_\gamma^{q \rightarrow q'}|$ , so  $|\nu^{-1}(\mathbb{A}_\gamma^{q \rightarrow q'})| = |\mathbb{A}_\gamma^{q \rightarrow q'}|$ ;
- if  $|\mathbb{A}_\gamma^{q \rightarrow q'}| \geq n$  for some  $q \in Q$ , then  $|A_{q \rightarrow q'}| \geq n$ , so  $|\nu^{-1}(\mathbb{A}_\gamma^{q \rightarrow q'})| \geq n$ , and thus both  $|\nu^{-1}(\mathbb{A}_\gamma^{q \rightarrow q'})|$  and  $|\mathbb{A}_\gamma^{q \rightarrow q'}|$  are at least  $n$ .

We construct a run  $\pi$  from  $\chi_{start}$  by copying  $\rho$  as follows. For each step of  $\rho$  where an agent  $a$  performs some transition  $t$ , we make  $|\nu^{-1}(a)|$  steps in  $\pi$  so that all agents in  $\nu^{-1}(a)$  perform transition  $t$  one by one. If  $a$  observed some agent  $a'$ , there is  $a''$  in  $\pi$  that can be observed because  $\nu^{-1}(a') \neq \emptyset$ : we made sure to map an agent to each observed agent in  $\rho$ .

For the general case with more data, we similarly construct two mappings  $\mu$  and  $\nu$ . First we define  $\mu$ , which maps each datum  $d$  of  $\chi_{start}$  to one of  $\gamma_{start}$  such that  $(\gamma_{start}, \mu(d)) \equiv_{f(n)} (\chi_{start}, d)$ . Then, for each datum  $d$ ,  $\nu$  maps each agent  $a$  with datum  $d$  of  $\chi_{start}$  to one with datum  $\mu(d)$  of  $\gamma_{start}$ .

Once  $\mu$  and  $\nu$  are defined, we build a run from  $\chi_{start}$  to a configuration  $\chi_{end}$  in which each agent  $a$  mimics the behaviour of  $\nu(a)$  in  $\rho$ . We make sure that agents (resp. data) observed in  $\rho$  have agents (resp. data) mapped to them, so that we can take the same transitions in  $\rho$  and  $\pi$ . The construction of  $\nu$  ensures that, for all data  $d$ , we have  $(\gamma_{end}, \mu(d)) \equiv_{f(n)} (\chi_{end}, d)$ . The construction of  $\mu$  ensures that  $\gamma_{end} \equiv_{f(n), g(n, M)} \chi_{end}$ . ◀

## 6 Decidability and Complexity Bounds

### 6.1 Decidability in Exponential Space

In this section, we use the results on GRE from Section 5 to provide an EXPSPACE upper bound for the emptiness problem for GRE. In the following, we assume that the representation of a GRE  $E$  takes  $|E| + \log(|E|)$  space.

We first prove that we can decide membership of a configuration (encoded in a naive way) in a GRE in PSPACE. A configuration is represented *data-explicitly* if it is represented

as a list of vectors of  $\mathbb{N}^Q$ , one vector for each datum. The *size* of this representation is  $k \cdot |\mathcal{P}| \cdot \log(m)$  where  $k$  is the number of data and  $m$  is the number of agents appearing in  $\gamma$ .

► **Proposition 23.** *The following problem is decidable in PSPACE: given a PPUD  $\mathcal{P}$ , a GRE  $E$ , and a configuration  $\gamma$  described data-explicitly, decide if  $\gamma \in \llbracket E \rrbracket_{\mathcal{P}}$ .*

The proof is given in Appendix D. It uses a relatively straightforward induction on  $E$  to show that this problem can be decided in polynomial space using a recursive algorithm (with a polynomial whose degree does not depend on  $E$ ). For the case where  $E = \text{Post}^*(F)$ , we rely on the fact that the numbers of agents and data remain the same throughout a run; we therefore can guess the configuration  $\gamma'$  such that  $\gamma' \in \llbracket F \rrbracket_{\mathcal{P}}$  (which can be checked with a recursive call) and  $\gamma \xrightarrow{*} \gamma'$  (which can be checked by exploration of the graph containing configurations with as many agents and data as  $\gamma$ ). The case  $E = \text{Pre}^*(F)$  is similar.

Proposition 23 allows us to check if a given configuration of a PPUD is in the set described by a GRE<sup>4</sup>. In the case of IOPPUD, Proposition 21 allows us to search for a witness configuration within some bounded set, yielding decidability.

► **Theorem 10.** *The emptiness problem for GRE over IOPPUD is in EXPSpace.*

**Proof.** Suppose  $\llbracket E \rrbracket_{\mathcal{P}}$  is not empty. By Proposition 21, it contains an  $(n, M)$ -container  $\text{cont}$  with  $n := \lceil |E| \rceil \cdot \text{poly}(|\mathcal{P}|)^{|E|}$  and  $M := \lceil |E| \rceil^{\text{poly}(|\mathcal{P}|) \cdot |E|^2}$ . We construct a configuration  $\gamma \in \text{cont}$  as follows. For each  $n$ -box  $\mathbf{b}$ , we select  $\text{cont}(\mathbf{b})$  many data such that over all  $n$ -boxes, the selected data are pairwise distinct. Then, for each  $n$ -box  $\mathbf{b}$ , each state  $q \in Q$  of  $\mathcal{P}$ , and each datum  $d_{\mathbf{b}}$  selected for  $\mathbf{b}$ , we put  $\mathbf{b}(q)$  many agents with datum  $d_{\mathbf{b}}$  in  $q$ . Note that the configuration  $\gamma$  is in  $\text{cont}$ , and the number of agents it contains it at most  $n \cdot |\mathcal{P}| \cdot |\mathbf{Boxes}_n| \cdot M$ . We have  $|\mathbf{Boxes}_n| = (n+1)^{|\mathcal{P}|} = \lceil |E| \rceil \cdot \text{poly}(|\mathcal{P}|)^{|E| \cdot |\mathcal{P}|}$ . We assumed at the beginning of Section 6 that the encoding of  $E$  uses memory  $|E| + \log(|E|)$ . As a result,  $n$ ,  $M$ ,  $|\mathcal{P}|$  and  $|\mathbf{Boxes}_n|$  are all at most exponential in the size of the input. Therefore, if  $\llbracket E \rrbracket_{\mathcal{P}}$  is not empty, then it contains a configuration with at most exponentially many agents. We can guess the data-explicit description of such a configuration in non-deterministic exponential space, and then check that the guessed configuration is in  $\llbracket E \rrbracket_{\mathcal{P}}$  in exponential space by Proposition 23 (we apply the PSPACE algorithm on an exponential input). As a result, deciding emptiness of  $\llbracket E \rrbracket_{\mathcal{P}}$  is in NEXPSpace, which is identical with EXPSpace. ◀

## 6.2 A Lower Complexity Bound


We now provide the following lower complexity bound.

► **Theorem 24.** *The emptiness problem for GRE over IOPPUD is CONEXPTIME-hard.*

**Proof sketch.** We proceed by reduction from the problem of tiling an exponentially large grid, a NEXPTIME-complete problem [27], to the complement of the emptiness problem for GRE. The full proof can be found in Appendix E.

A *tiling instance* is a tuple  $(2^n, \mathcal{C}, \mathcal{T})$ , with  $n \geq 1$ ,  $\mathcal{C}$  a finite set of *colours* with special colour white, and  $\mathcal{T} = \{t_1, \dots, t_m\} \subseteq \mathcal{C}^4$  a finite set of *tiles*. We can view a tile as a square whose four edges are coloured. The *tiling problem* asks whether there is a *tiling*, that is, a mapping  $\tau: [0, 2^n - 1] \times [0, 2^n - 1] \rightarrow \mathcal{T}$  such that the colours of neighbouring tiles match and the borders of the grid are white.

<sup>4</sup> This implies that the emptiness problem for GRE over PPUD, while undecidable due to Theorem 3, is semi-decidable: one can simply enumerate all configurations and test membership for each of them.



■ **Figure 4** Partial depiction of the protocol constructed in Theorem 24.

Given a tiling instance  $(2^n, \mathcal{C}, \mathcal{T})$ , we build an instance  $(\mathcal{P}, E)$  of the emptiness problem for GRE. In  $\mathcal{P}$ , witness tilings can be encoded in the configurations, and we construct  $(\mathcal{P}, E)$  such that  $\llbracket E \rrbracket$  contains exactly the configurations that correspond to a correctly encoded witness tiling. More precisely,  $\gamma \in \llbracket E \rrbracket$  when:

- (Cond1) for all  $(i, j) \in [0, 2^n - 1]^2$ , some datum encodes coordinates  $(i, j)$  and a tile type;
- (Cond2) for all  $(i, j) \in [0, 2^n - 1]^2$ , there is at most one datum encoding  $(i, j)$ ;
- (Cond3) the mapping  $[0, 2^n - 1]^2 \rightarrow \mathcal{C}$  defined by the data is a tiling.

The GRE  $E$  will be of the form of a conjunction, i. e., a list of *constraints* that the configuration must satisfy. Our first constraint is  $\overline{\text{Pre}^*(\text{Pres}(q_{\perp}))}$  where  $q_{\perp}$  is a special error state and  $\text{Pres}(q)$  is the GRE expressing that some agent is in  $q$ . This forbids, in  $\llbracket E \rrbracket$ , configurations from which  $q_{\perp}$  can be covered.

(Cond1) is obtained using the *tiling gadget* in Figure 4. States  $t_1, \dots, t_m$  represent the available tiles of  $\mathcal{T}$ , and coordinate states allow for a binary representation of the horizontal and vertical coordinates of a square in the grid. For a datum  $d$ , the agents of datum  $d$  in the coordinate states encode the position of the square corresponding to  $d$ , and an agent of datum  $d$  in state  $t_i$  indicates that the square in the grid corresponding to  $d$  should be coloured according to tile  $t_i$ . Configurations in  $\llbracket E \rrbracket$  are not allowed to have two agents of same datum playing the same role; otherwise, one of them may observe the other and go to  $q_{\perp}$ . In particular, each datum has at most  $2n + 1$  agents in the tiling gadget.

To obtain (Cond2), we use a *duplication gadget*, partially represented in Figure 4. We enforce that any configuration in  $\llbracket E \rrbracket$  has one agent of each datum in  $D_a$ , one in  $D_b$  and none in the rest of the duplication gadget. The blue part implements a test (depicted in Figure 5b in Appendix E) where two agents of distinct data, one from  $D_a$  and one from  $D_b$ , may test that their data encode the same coordinates; if this is the case, they may go to  $q_{\text{dup}}$ . If there are more than two agents in the blue part, this test is not reliable but  $q_{\geq 3}$  can be

covered. Item (Cond2) can therefore be achieved by enforcing that configurations in  $\llbracket E \rrbracket$  are *not* in  $\text{Pre}^*(\text{Pres}(q_{\text{dup}}) \cap \overline{\text{Pre}^*(\text{Pres}(q_{\geq 3}))})$ .

Finally, we explain how (Cond3) is achieved; we describe only how the horizontal (left-right) borders are verified. We use a gadget, named *horizontal verifier* in Figure 4. In this gadget, a single agent, called *verifier*, is in charge of verifying that colours of left-right borders match. The verifier uses  $2n$  auxiliary agents to encode two variables  $x, y \in [0, 2^n - 1]$  in binary. Again, transitions to  $q_{\perp}$  detect when two agents play the same role, so that there is only one verifier and so that variables  $x$  and  $y$  can be implemented faithfully. The initialisation  $x = y = 0$  is enforced as a constraint in  $E$ . We now sketch how the verifier reads the encoded tiling; to do that, it must synchronise with the datum encoding  $(x, y)$ .

This is done using the synchronisation gadget of Figure 4. In  $\llbracket E \rrbracket$ , all agents in the synchronisation gadget are in  $S^{(h)}$ . Moreover, we add a constraint in  $E$  so that  $\gamma \in \llbracket E \rrbracket$  requires that there is a run from  $\gamma$  where all agents in the synchronisation gadget end in  $F^{(h)}$  and where the verifier ends in  $q_f^{(h)}$ . The synchronisation tests guarantee that, whenever there is an agent in  $T_i^{(h)}$ , this agent's datum encodes square  $(i, j)$  where  $i$  is equal to the current value of  $x$  and  $j$  is equal to the current value of  $y$ . The synchronisation is challenging to design because the values of  $x$  and  $y$  may change throughout a run and only one bit can be tested at a time. However, as proved in Lemma 37 of Appendix E, this can be achieved by having a first synchronisation test that checks equality of bits from most to least significant, and a second test that checks equality from least to most significant. ◀

### 6.3 Discussion on Complexity Gaps

We now discuss some complexity gaps left open by this paper. First, there remains a complexity gap for the emptiness problem for GRE, which is known to be between  $\text{CONEXPTIME}$  (Theorem 24) and  $\text{EXPSpace}$  (Theorem 10). Closing the gap appears challenging. On one hand, if the problem is below  $\text{EXPSpace}$ , then this probably requires developing new techniques. On the other hand, proving  $\text{EXPSpace}$ -hardness does not seem easy. In particular, the synchronisation techniques from Theorem 24 assumes that each datum synchronises only once with the verifier. This synchronisation technique would not be suitable for, e.g., multiple interactions between the head and the cells of a Turing machine.

Another, arguably more important open question is the exact complexity of well-specification, which is only known to be between  $\text{PSPACE}$  (model without data, [18]) and  $\text{EXPSpace}$  (Theorem 10). On the one hand, it is unclear whether relevant configurations can be stored in polynomial space:

▷ **Claim 25.** The number of data to consider for well-specification may have to be exponential.

The claim is formalised and proven in Appendix E.5. As a consequence, proving that the problem is in  $\text{PSPACE}$  cannot be achieved with a procedure that explicitly stores configurations. On the other hand, in order to build a reduction from the tiling problem as in Theorem 24, we need a new idea to enforce that *at most* one datum encodes each tile. In Theorem 24, we had states  $q_{\text{dup}}$  and  $q_{\geq 3}$  and duplication meant being able to cover  $q_{\text{dup}}$  and, at the same, forbid that  $q_{\geq 3}$  can ever be covered in the future. We do not know how to encode this constraint when working with an instance of well-specification.

## 7 Conclusion

We have studied the verification of population protocols with unordered data [9], an extension of population protocols where agents carry data from an infinite unordered set. We first

proved that the well-specification problem is undecidable (Theorem 3), which then led us to consider the restriction to protocols with immediate observation. This subclass was defined in [9], where the authors proved that these protocols compute exactly the interval predicates. We defined a general class of problems on this model, which consists in deciding the existence of a configuration satisfying a so-called generalised reachability expression; this class of problems subsumes many classic problems, one of which is well-specification. Despite its generality, we showed the problem to be decidable in exponential space (Theorem 10); we also provided a CONEXPTIME lower bound. A remaining open question is the exact complexity of well-specification for immediate observation population protocols with unordered data, which is located between PSPACE (model without data, [18]) and EXPSPACE (Theorem 10).

---

## References

---

- 1 Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. Time-space trade-offs in population protocols. In *28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 2560–2579, 2017. doi:10.1137/1.9781611974782.169.
- 2 Dan Alistarh and Rati Gelashvili. Recent algorithmic advances in population protocols. *SIGACT News*, 49(3):63–73, 2018. doi:10.1145/3289137.3289150.
- 3 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. In *33rd Annual ACM Symposium on Principles of Distributed Computing, PODC 2004*, pages 290–299. ACM, 2004. doi:10.1145/1011767.1011810.
- 4 Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Comput.*, 20(4):279–304, 2007. doi:10.1007/S00446-007-0040-2.
- 5 A. R. Balasubramanian, Lucie Guillou, and Chana Weil-Kennedy. Parameterized analysis of reconfigurable broadcast networks. In *Foundations of Software Science and Computation Structures - 25th International Conference, FoSSaCS 2022*, pages 61–80, 2022. doi:10.1007/978-3-030-99253-8\_4.
- 6 A. R. Balasubramanian and Chana Weil-Kennedy. Reconfigurable broadcast networks and asynchronous shared-memory systems are equivalent. In *12th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2021*, pages 18–34, 2021. doi:10.4204/EPTCS.346.2.
- 7 Roderick Bloem, Swen Jacobs, Ayrat Khalimov, Igor Konnov, Sasha Rubin, Helmut Veith, and Josef Widder. *Decidability of Parameterized Verification*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers, 2015. doi:10.2200/S00658ED1V01Y201508DCT013.
- 8 Michael Blondin, Javier Esparza, Stefan Jaax, and Philipp J. Meyer. Towards efficient verification of population protocols. *Formal Methods Syst. Des.*, 57(3):305–342, 2021. doi:10.1007/S10703-021-00367-3.
- 9 Michael Blondin and François Ladouceur. Population protocols with unordered data. In *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023*, pages 115:1–115:20, 2023. doi:10.4230/LIPICS.ICALP.2023.115.
- 10 Patricia Bouyer, Nicolas Markey, Mickael Randour, Arnaud Sangnier, and Daniel Stan. Reachability in networks of register protocols under stochastic schedulers. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016*, pages 106:1–106:14, 2016. doi:10.4230/LIPICS.ICALP.2016.106.
- 11 Philipp Czerner, Roland Guttenberg, Martin Helfrich, and Javier Esparza. Fast and succinct population protocols for Presburger arithmetic. *J. Comput. Syst. Sci.*, 140:103481, 2024. doi:10.1016/J.JCSS.2023.103481.

- 12 Wojciech Czerwinski and Lukasz Orlikowski. Reachability in vector addition systems is Ackermann-complete. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*, pages 1229–1240, 2021. doi:10.1109/FOCS52979.2021.00120.
- 13 Robert Elsässer and Tomasz Radzik. Recent results in population protocols for exact majority and leader election. *Bull. EATCS*, 126, 2018. URL: <http://bulletin.eatcs.org/index.php/beatcs/article/view/549/546>.
- 14 Javier Esparza. Keeping a crowd safe: On the complexity of parameterized verification (invited talk). In *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5-8, 2014, Lyon, France*, pages 1–10, 2014. doi:10.4230/LIPICS.STACS.2014.1.
- 15 Javier Esparza. Population protocols: Beyond runtime analysis. In *Reachability Problems - 15th International Conference, RP 2021*, pages 28–51, 2021. doi:10.1007/978-3-030-89716-1\_3.
- 16 Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of population protocols. *Acta Informatica*, 54(2):191–215, 2017. doi:10.1007/S00236-016-0272-3.
- 17 Javier Esparza, Stefan Jaax, Mikhail A. Raskin, and Chana Weil-Kennedy. The complexity of verifying population protocols. *Distributed Comput.*, 34(2):133–177, 2021. doi:10.1007/S00446-021-00390-X.
- 18 Javier Esparza, Mikhail A. Raskin, and Chana Weil-Kennedy. Parameterized analysis of immediate observation Petri nets. In *Application and Theory of Petri Nets and Concurrency - 40th International Conference, PETRI NETS 2019*, pages 365–385, 2019. doi:10.1007/978-3-030-21571-2\_20.
- 19 Lucie Guillou, Corto Mascle, and Nicolas Waldburger. Parameterized broadcast networks with registers: from NP to the frontiers of decidability. In *Foundations of Software Science and Computation Structures - 27th International Conference, FoSSaCS 2024*, pages 250–270, 2024. doi:10.1007/978-3-031-57231-9\_12.
- 20 Petr Jancar. Undecidability of bisimilarity for Petri nets and some related problems. *Theor. Comput. Sci.*, 148(2):281–301, 1995. doi:10.1016/0304-3975(95)00037-W.
- 21 Petr Jancar and Jérôme Leroux. The semilinear home-space problem is Ackermann-complete for Petri nets. In *34th International Conference on Concurrency Theory, CONCUR 2023*, pages 36:1–36:17, 2023. doi:10.4230/LIPICS.CONCUR.2023.36.
- 22 Ranko Lazic, Thomas Christopher Newcomb, Joël Ouaknine, A. W. Roscoe, and James Worrell. Nets with tokens which carry data. In *28th International Conference on Applications and Theory of Petri Nets and Other Models of Concurrency, ICATPN 2007*, pages 301–320, 2007. doi:10.1007/978-3-540-73094-1\_19.
- 23 Jérôme Leroux. The reachability problem for Petri nets is not primitive recursive. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*, pages 1241–1252, 2021. doi:10.1109/FOCS52979.2021.00121.
- 24 Jérôme Leroux and Sylvain Schmitz. Reachability in vector addition systems is primitive-recursive in fixed dimension. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–13, 2019. doi:10.1109/LICS.2019.8785796.
- 25 Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., 1967.
- 26 Fernando Rosa-Velardo and David de Frutos-Escrig. Decidability and complexity of Petri nets with unordered data. *Theor. Comput. Sci.*, 412(34):4439–4451, 2011. doi:10.1016/J.TCS.2011.05.007.
- 27 François Schwarzentruber. The complexity of tiling problems. *CoRR*, abs/1907.00102, 2019. arXiv:1907.00102.
- 28 Chana Weil-Kennedy. *Observation Petri Nets*. PhD thesis, Technical University of Munich, Germany, 2023. URL: <https://nbn-resolving.org/urn:nbn:de:bvb:91-diss-20230320-1691161-1-3>.



## A Undecidability of Verification of Population Protocols with Unordered Data: Proof for Section 3

► **Theorem 3.** *The well-specification problem for PPUD is undecidable.*

A 2-counter machine is a transition system using two counters  $x$  and  $y$ . It consists of a list of instructions of the form:  $\text{inc}(c)$ , increment on counter  $c$ ;  $\text{dec}(c)$ , decrement on counter  $c$ ;  $\text{Test}_0(c, k)$ , zero-test that moves to instruction  $k$  if  $c = 0$  (and to the next instruction otherwise); and  $\text{Halt}$ , halting instruction. The machine starts at the first instruction with both counters equal to 0. It is well known that deciding whether a 2-counter machine eventually reaches the  $\text{Halt}$  instruction is an undecidable problem.

Let  $\text{CM}$  be a 2-counter machine with  $n$  instructions  $i_1, \dots, i_n$ .

Let  $Q_{\text{CM}} := \{i_1, \dots, i_n, i'_1, \dots, i'_n\}$ . We define a population protocol with unordered data  $\mathcal{P}$  over infinite domain  $\mathbb{D}$  as follows:

1.  $Q_{\text{op}} := \{\text{idle}, \text{inc}, \text{dec}, \text{done}, \text{Test}_0, = 0, > 0\}$
2.  $Q_{\text{main}} := Q_{\text{CM}} \cup \{\mathbf{x}, \mathbf{y}\} \cup \{\bar{\mathbf{x}}, \bar{\mathbf{y}}\} \times Q_{\text{op}} \cup \{R, U, q_\perp\}$ .
3.  $Q := Q_{\text{main}} \times \{R, \text{other}\}$ .
4.  $\delta$  will be defined step by step later.
5.  $I = \{(R, R), (U, \text{other}), ((\bar{\mathbf{x}}, \text{idle}), \text{other}), ((\bar{\mathbf{y}}, \text{idle}), \text{other}), (i_1, \text{other})\}$ .
6.  $O(i_m, \text{other}) = O(i_m, R) = \top$  if  $i_m = \text{halt}$ , else  $O(q) = \perp$  for all other  $q \in Q$ .

The actual states of the protocol are  $Q$ , but the second component is never updated; it simply remembers whether the initial state was  $R$  or not. We will hence mainly refer to  $Q_{\text{main}}$  and specify all except one transition in terms of  $Q_{\text{main}}$  only.

Recall the role of each state:

- $R$  is a reservoir, which can contain many agents of all data.
- $U$  is a state in which all agents should have different data
- $Q_{\text{CM}}$  is the set of control states in which a single agent should evolve.
- for both  $c \in \{x, y\}$ ,  $\bar{c} \times Q_{\text{op}}$  must contain a single agent at all times, with a datum that is not in  $U$ , which is in charge of synchronising with the agent in  $Q_{\text{CM}}$  and apply the transitions of the machine to counter  $c$ .
- $\mathbf{x}, \mathbf{y}$  represent the counters: all agents in each  $\mathbf{c}$  should have the same datum as the current agent in  $\bar{c} \times Q_{\text{op}}$ , and the number of agents in  $\mathbf{c}$  represents the current counter value.

We will ensure that those conditions are met using the violation detection mechanisms below.

We use the following notation: For  $p, p', q, q' \in Q$  we denote  $((p, p'), \bowtie, (q, q')) \in \delta$  by  $p, p' \mapsto_{\bowtie} q, q'$ , or leave away  $\bowtie$  if this transition is enabled in either case. The first transition we present allows us to guarantee that at most one agent per datum is not in  $R$  initially.

$$(p, \text{other}), (p', \text{other}) \mapsto_{=} (q_\perp, \text{other}), (q_\perp, \text{other}) \quad p, p' \in Q_{\text{main}} \quad (\text{Input Violation})$$

This transition detects that a violation occurred; multiple agents of the same datum started in  $U$  or in the control states. It overwrites whatever other transition would be defined between the states in  $Q_{\text{main}}$ . It is important that it is also impossible for one of the counter control agents in  $\{\bar{\mathbf{x}}, \bar{\mathbf{y}}\} \times Q_{\text{op}}$  to initially have a datum which is still in the pool  $U$ , otherwise one cheat, i. e. an incorrectly performed zero-test, would be undetectable at the start.

We continue by specifying the other violation transitions; namely if two agents are in  $Q_{CM}$ , if two agents are in  $\bar{x}$ , or if an agent in  $\bar{x}$  meets an agent in  $x$  of different datum:

$$\begin{array}{lll}
(\bar{c}, b), c \mapsto_{\neq} q_{\perp}, q_{\perp} & b \in Q_{op}, c \in \{x, y\} & \text{(Counter Colour Violation)} \\
(\bar{c}, b), (\bar{c}, b') \mapsto q_{\perp}, q_{\perp} & b, b' \in Q_{op}, c \in \{x, y\} & \text{(Counter Control Violation)} \\
q, q' \mapsto q_{\perp}, q_{\perp} & q, q' \in Q_{CM} & \text{(Control State Violation)} \\
q_{\perp}, q \mapsto q_{\perp}, q_{\perp} & q \in Q_{main} & \text{(Convert To Sink)}
\end{array}$$

The (Convert To Sink) transition informs other agents about violations.

Next we explain the actual counter machine simulation. Increments and Decrements are performed via a sequence of three transitions each as follows.

$$\begin{array}{lll}
i_m, (\bar{c}, \text{idle}) \mapsto i'_m, (\bar{c}, \text{inc}) & c \in \{x, y\}, i_m = \text{inc}(c) & \text{(Start Increment } c) \\
(\bar{c}, \text{inc}), R \mapsto_{=} (\bar{c}, \text{done}), c & c \in \{x, y\} & \text{(Increment } c) \\
(\bar{c}, \text{done}), i'_m \mapsto (\bar{c}, \text{idle}), i_{m+1} & c \in \{x, y\}, m \in \{1, \dots, n\} & \text{(End Operation on } c) \\
i_m, (\bar{c}, \text{idle}) \mapsto i'_m, (\bar{c}, \text{dec}) & c \in \{x, y\}, i_m = \text{dec}(c) & \text{(Start Decrement } c) \\
(\bar{c}, \text{dec}), c \mapsto_{=} (\bar{c}, \text{done}), R & c \in \{x, y\} & \text{(Decrement } c)
\end{array}$$

To simulate zero test instructions, the counter control agent in  $\bar{c}$  is informed about the instruction; afterwards they either meet an agent in  $c$  of their own datum, in which case they decide  $> 0$ , or an agent in  $U$ , in which case this new agent takes their place and the counter is assumed to be 0.

$$\begin{array}{lll}
i_m, (\bar{c}, \text{idle}) \mapsto i'_m, (\bar{c}, \text{Test}_0) & i_m = \text{Test}_0(c, k), c \in \{x, y\}, k \in \mathbb{N} & \text{(Zerotest Start)} \\
(\bar{c}, \text{Test}_0), U \mapsto R, (\bar{c}, = 0) & c \in \{x, y\} & (c=0) \\
(\bar{c}, \text{Test}_0), c \mapsto (\bar{c}, > 0), c & c \in \{x, y\} & (c>0) \\
i'_m, (\bar{c}, = 0) \mapsto i_k, (\bar{c}, \text{idle}) & i_m = \text{Test}_0(c, k), c \in \{x, y\}, k \in \mathbb{N} & \text{(End Zerotest } =0) \\
i'_m, (\bar{c}, > 0) \mapsto i_{m+1}, (\bar{c}, \text{idle}) & i_m = \text{Test}_0(c, k), c \in \{x, y\}, k \in \mathbb{N} & \text{(End Zerotest } >0)
\end{array}$$

**Correctness:** First assume that the counter machine CM does halt.

Let  $k$  be the number of steps CM requires to halt. Let  $\gamma_0$  be an initial configuration with  $\gamma_0(d_i, R) \geq k$  and  $\gamma_0(d_i, U) = 1$  for at least  $k$  different data  $d_1, \dots, d_k$ . Moreover, assume that in  $\gamma_0$  there is no input violation, i. e. all transitions with “Violation” as part of the name are disabled. In particular,  $\gamma_0(\mathbb{D}, i_1) = \gamma_0(\mathbb{D}, \bar{x}) = \gamma_0(\mathbb{D}, \bar{y}) = 1$ , i. e. we do not have a violation in the counter or the control states.

We perform the following transition sequence  $\sigma$ : Using the  $k$  data above, we correctly simulate the counter machine until it halts. Observe that any simulating transition takes at most 1 agent out of  $R$  and at most one new datum from  $U$ , hence we do not run out of agents in  $R, U$ . Call the reached configuration  $\gamma$ . This configuration is clearly a deadlock, since we do not have any violations to detect, and the control agent in  $Q_{CM}$  does not start any new operation. However, some agent is in the halt instruction, i. e. we indeed have a fair (in fact terminal) run which does not output  $\perp$ .

Now assume that CM does *not* halt. We start with three important observations, which will be used to prove that violation transitions can only be disabled by occurring:

1.  $\gamma(Q_{CM}, \mathbb{D})$ ,  $\gamma(\{\bar{x}\} \times Q_{op}, \mathbb{D})$  and  $\gamma(\{\bar{y}\} \times Q_{op}, \mathbb{D})$  are preserved by all transitions except transitions moving an agent to the sink state.
2. Agents who enter  $\{\bar{x}, \bar{y}\} \times Q_{op}$  are always from  $U$ , a previously unused datum.
3. Agents who enter or leave  $x, y$  are always the same datum as the corresponding  $\bar{x}, \bar{y}$ .

We have to prove that every initial configuration has a unique output. Let  $\gamma_0$  be any initial configuration. We claim that every fair run  $\pi = (\gamma_0, \gamma_1, \dots)$  stabilises to output  $\perp$ . Assume that in  $\pi$  no violation occurs, i. e.  $\gamma_m(q_\perp) = 0$  for all  $m \in \mathbb{N}$ , otherwise  $\pi$  has output  $\perp$  via transition (Convert To Sink), since agents cannot leave the sink.

Initially all agents not starting in  $R$  have a different datum, otherwise (Input Violation) would eventually occur. Furthermore, by observation 1,  $Q_{CM}, (\bar{x}, \text{idle})$  and  $(\bar{y}, \text{idle})$  all start with one agent each, otherwise the corresponding violation transition would eventually occur by fairness. We claim that in  $\pi$  the counter machine is simulated faithfully. Assume the opposite. The only way to not simulate the counter machine correctly is by performing a zero-test wrong. That is, at some  $\gamma_m$  the goto  $k$  part of a  $\text{Test}_0(c, k)$  instruction was applied with  $\gamma_m(c, d) > 0$  for some  $d \in \mathbb{D}$ .

However, by observation 2, the unique agent with a state in the set  $\{\bar{c}\} \times Q_{op}$  will always have a data  $d' \neq d$  for the rest of the run. By observation 3, the agent in state  $c$  with data  $d$  can hence never leave, i. e. we have  $\gamma_l(c, d) > 0$  for all  $l \geq m$ . This implies that (Counter Colour Violation) is always enabled, and would eventually occur, contradiction.

Hence in  $\pi$  the counter machine is faithfully simulated. Since CM does not halt, we have  $\gamma_l(i_k, \mathbb{D}) = 0$  for all  $l \in \mathbb{N}$  and halt instructions  $i_k$ . Since this is the only state with output  $\top$ , every agent always has output  $\perp$ , and the run  $\pi$  hence has output  $\perp$  as required, concluding the proof.

As a final remark, this reduction in fact also establishes that the following problem is undecidable: given two interval predicates  $\varphi_1, \varphi_2$  and a PPUD  $\mathcal{P}$ , decide whether the GRE  $\text{Post}^*(\varphi_1) \cap \varphi_2$  is empty. To see this, instead of using violation detection, use  $\varphi_1$  to encode the restrictions on initial configurations and  $\varphi_2$  to encode that at the end, no agent in  $x$  or  $y$  is supposed to have a different datum than the corresponding agent in  $\bar{x}$  or  $\bar{y}$ .

## B An Analysis of Immediate Observation Protocols with Data: Proofs for Section 4

We recall a few notations: Let  $\rho: \gamma_1 \rightarrow \gamma_2 \rightarrow \dots \rightarrow \gamma_m$  be a run. For  $i \in [1, m]$ , let  $\rho[\rightarrow i]$  (resp.  $\rho[i \rightarrow]$ ) denote the prefix of  $\rho$  ending on its  $i$ -th configuration (resp. the suffix of  $\rho$  starting on its  $i$ -th configuration).

We let  $\mathbb{A}_\rho$  be the set of agents appearing in  $\rho$ , and set  $\mathbb{A}_\rho^d := \{a \in \mathbb{A}_\rho \mid \text{dat}(a) = d\}$ . We let  $\mathbb{A}_{\rho, o}^d$  be the set of agents with datum  $d$  that are observed in  $\rho$ , i. e., the  $a_o \in \mathbb{A}_\rho^d$  such that there exists a step  $\gamma \xrightarrow{\text{data}_o} \gamma'$  in  $\rho$ . For all  $q_1, q_m \in Q$ , we let  $\mathbb{A}_{\rho, q_1, q_m}^d$  be the set of agents with datum  $d$  that start in  $q_1$  and end in  $q_m$ , i. e., the  $a \in \mathbb{A}_\rho^d$  such that  $\gamma_1(a) = q_1$  and  $\gamma_m(a) = q_m$ .

► **Lemma 13 (Agents core).** *Let  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run. Then there exist configurations  $\gamma'_{start}, \gamma'_{end}$  and a run  $\rho': \gamma'_{start} \xrightarrow{*} \gamma'_{end}$  with  $\mathbb{A}_{\rho'} \subseteq \mathbb{A}_\rho$  such that:*

- (i) for all  $a \in \mathbb{A}_{\rho'}$ ,  $\gamma'_{start}(a) = \gamma_{start}(a)$  and  $\gamma'_{end}(a) = \gamma_{end}(a)$ ;
- (ii) for all  $d \in \mathbb{D}$  and  $q_s, q_e \in Q$ , if  $\mathbb{A}_{\rho, q_s, q_e}^d \neq \emptyset$ , then  $\mathbb{A}_{\rho', q_s, q_e}^d \neq \emptyset$ ;
- (iii) for all  $d \in \mathbb{D}$ , we have  $|\mathbb{A}_{\rho'}^d| \leq |Q|^3$ .

**Proof.** Let  $\rho : \gamma_1 \rightarrow \gamma_2 \rightarrow \dots \rightarrow \gamma_m$ . Suppose there are  $d \in \mathbb{D}$  and  $q_s, q_e \in Q$  such that  $|\mathbb{A}_{\rho, q_s, q_e}^d| > |Q|$  (otherwise we can set  $\rho' := \rho$ ). Let  $\mathcal{R}$  be the set of states visited by agents of  $\mathbb{A}_{\rho, q_s, q_e}^d$  during  $\rho$ . Notice that  $|\mathcal{R}| \leq |Q|$ . We are going to define a family  $(a_q)_{q \in \mathcal{R}}$  of pairwise distinct agents and corresponding trajectories and identify them with agents in  $\mathbb{A}_{\rho, q_s, q_e}^d$  such that reducing  $\mathbb{A}_{\rho, q_s, q_e}^d$  in  $\rho$  to  $(a_q)_{q \in \mathcal{R}}$  still yields a valid run. Repeating the operation for every  $\mathbb{A}_{\rho, q_s, q_e}^d$  will yield a run fulfilling the conditions of the lemma.

We iterate through  $\mathcal{R}$  as follows. Let  $q$  be a state in  $\mathcal{R}$  and let  $f$  be the first moment  $q$  is reached in  $\rho$ , i. e., the minimal index such that there exists an  $a \in \mathbb{A}_{\rho, q_s, q_e}^d$  with  $\gamma_f(a) = q$ . Let  $\ell$  be the last moment  $q$  is occupied in  $\rho$ , i. e., the maximal index such that there exists an  $a \in \mathbb{A}_{\rho, q_s, q_e}^d$  with  $\gamma_\ell(a) = q$ . Let  $\alpha_q$  be an agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$  that reaches  $q$  first, i. e.,  $\gamma_f(\alpha_q) = q$ , and let  $\beta_q$  be an agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$  that leaves  $q$  last, i. e.,  $\gamma_\ell(\beta_q) = q$ . Note that these agents do not have to be distinct.

We pick a fresh agent  $a_q \notin \mathbb{A}_\rho$  with  $\mathbf{dat}(a_q) = d$  and modify  $\rho$  as follows. We let  $a_q$  copy  $\alpha_q$  in  $\rho[\rightarrow f]$ , then  $a_q$  stays idle until  $\beta_q$  leaves  $q$  (for the last time) and then  $a_q$  copies  $\beta_q$  in  $\rho[\ell \rightarrow]$ . We do this for every  $q \in \mathcal{R}$ . To see that this still yields a valid run  $\rho''$ , we can apply Lemma 12 stepwise.

In  $\rho''$ , for every step in which an  $a_o$  in  $\mathbb{A}_{\rho, q_s, q_e}^d$  is observed in state  $q$ , let  $a_q$  be observed instead, i. e., replace steps  $\xrightarrow{\times a_o} a$  with  $\xrightarrow{\times a_q} a$ . Now remove all steps involving agents in  $\mathbb{A}_{\rho'', q_s, q_e}^d \setminus \{a_q \mid q \in \mathcal{R}\}$ . To see that this yields a valid run, note that we only need to verify that every step is valid. This is the case, since, whenever an agent in  $\mathbb{A}_{\rho'', q_s, q_e}^d$  is observed in state  $q$  in  $\rho''$ , by construction of the trajectory for  $a_q$ , the agent  $a_q$  is also in state  $q$ . To render  $\mathbb{A}_{\rho'', q_s, q_e}^d$  a subset of  $\mathbb{A}_{\rho, q_s, q_e}^d$ , we identify (or substitute) each  $a_q$  with a distinct agent in  $\mathbb{A}_{\rho, q_s, q_e}^d$ .

By applying this transformation to all data in  $\mathbb{D}_\rho$  and all pairs of states  $q_s, q_e$  for which  $|\mathbb{A}_{\rho, q_s, q_e}^d| > |Q|$ , we obtain a run  $\rho'$  in which, for all  $d, q_s, q_e$ , at most  $|Q|$  agents go from  $q_s$  to  $q_e$  with datum  $d$ . In total, for each  $d$ , there are most  $|Q|^3$   $d$ -agents. Lemma 13-(i) is guaranteed by the fact that all agents of  $\mathbb{A}_{\rho', q_s, q_e}^d$  are from  $\mathbb{A}_{\rho, q_s, q_e}^d$ . Lemma 13-(ii) follows from the fact that for all  $d, q_s, q_e$ , either  $|\mathbb{A}_{\rho', q_s, q_e}^d| \leq |Q|$  and  $\mathbb{A}_{\rho', q_s, q_e}^d = \mathbb{A}_{\rho, q_s, q_e}^d$  or  $|\mathbb{A}_{\rho', q_s, q_e}^d| > |Q|$  and  $|\mathbb{A}_{\rho', q_s, q_e}^d| = |Q|$  as we have replaced those agents by the agents  $a_q$ .  $\blacktriangleleft$

**► Lemma 16 (Data core).** *Let  $\rho : \gamma_{start} \xrightarrow{*} \gamma_{end}$  be a run and let  $K$  be a number such that there are at most  $K$  agents of each datum in  $\rho$ . Then there exist configurations  $\gamma'_{start}, \gamma'_{end}$ , a run  $\rho' : \gamma'_{start} \xrightarrow{*} \gamma'_{end}$ , and a subset of data  $\mathbb{D}_{\rho'} \subseteq \mathbb{D}_\rho$  such that:*

- (i) for all  $d \in \mathbb{D}_{\rho'}$  and all agents  $a$  of datum  $d$ ,  $\gamma_{start}(a) = \gamma'_{start}(a)$  and  $\gamma_{end}(a) = \gamma'_{end}(a)$ ,
- (ii) for all  $d \in \mathbb{D}_\rho$ , there exists  $d' \in \mathbb{D}_{\rho'}$  such that  $\mathbf{tr}_{\rho'}^{d'} = \mathbf{tr}_\rho^d$ ,
- (iii)  $|\mathbb{D}_{\rho'}| \leq (K + 1)^{|Q|^3 + |Q|^2}$ .

**Proof.** Let  $\rho : \gamma_1 \rightarrow \gamma_2 \rightarrow \dots \rightarrow \gamma_m$ . We proceed similarly as in the proof of Lemma 13, lifting the proof from agents to data.

For a datum  $d$  and  $i \in [1, m]$ , we let  $\mathbf{str}_\rho^d(i) : Q^3 \rightarrow \mathbb{N}$ , called the *split trace of  $d$  in  $\rho$  at  $i$* , be such that for all  $q_1, q_2, q_3 \in Q$ , we have

$$\mathbf{str}_\rho^d(i)(q_1, q_2, q_3) := |\mathbb{A}_{\rho[\rightarrow i], q_1, q_2}^d \cap \mathbb{A}_{\rho[i \rightarrow], q_2, q_3}^d|.$$

Thus, for each triple of states  $(q_1, q_2, q_3)$ , the value of  $\mathbf{str}_\rho^d(i)(q_1, q_2, q_3)$  is the number of  $d$ -agents that were in  $q_1$  at the start of  $\rho$ , in  $q_2$  at the  $i$ th configuration, and in  $q_3$  at the end.

Let  $\mathcal{S} := \{\mathbf{str}_\rho^d(i) \mid d \in \mathbb{D}_\rho, i \in [1, m]\}$ . As there are no more than  $K$  agents of each datum in  $\rho$ , it holds that  $\mathbf{str}_\rho^d(i)(q_1, q_2, q_3) \in [0, K]$  for all data  $d$ , for  $i \in [1, m]$ , and states  $q_1, q_2, q_3 \in Q$ . Hence,  $|\mathcal{S}| \leq M := (K + 1)^{|Q|^3}$ .

For  $tr: Q^2 \rightarrow [0, K]$ , let  $\mathbb{D}_\rho^{tr} := \{d \in \mathbb{D}_\rho \mid \text{tr}_\rho^d = tr\}$  be the data in  $\rho$  with trace  $tr$ . If  $|\mathbb{D}_\rho^{tr}| \leq M$  for all  $tr: Q^2 \rightarrow [0, K]$ , then  $|\mathbb{D}_\rho| \leq M \cdot (K+1)^{|Q|^2} = (K+1)^{|Q|^3+|Q|^2}$ , so  $\rho' := \rho$  already fulfils the requirements of the lemma.

Hence, in the following, let  $tr: Q^2 \rightarrow [0, K]$  be a trace with  $|\mathbb{D}_\rho^{tr}| > M$ . We define  $\mathcal{S}_{tr} = \{\text{str}_\rho^d(i) \mid d \in \mathbb{D}_\rho^{tr}, i \in [1, m]\}$  the set of split traces corresponding to some datum  $d$  of run  $\rho$  with trace  $tr$  at some point  $i$ . Note that  $\mathcal{S}_{tr} \subseteq \mathcal{S}$ , and thus,  $|\mathcal{S}_{tr}| \leq M$ .

For each split trace  $str \in \mathcal{S}_{tr}$ , we let  $f_{str}$  be the minimal index such that there exists a datum  $\delta_{str}$  in  $\mathbb{D}_\rho^{tr}$  such that  $\text{str}_\rho^{\delta_{str}}(f_{str}) = str$ . Similarly, for each  $str \in \mathcal{S}_{tr}$ , we let  $\ell_{str}$  be the maximal index such that there exists a datum  $\varepsilon_{str}$  in  $\mathbb{D}_\rho^{tr}$  such that  $\text{str}_\rho^{\varepsilon_{str}}(\ell_{str}) = str$ . Note that those data do not have to be distinct. Also note that, by definition, we have  $f_{str} \leq \ell_{str}$  for all  $str \in \mathcal{S}_{tr}$ . As  $\text{str}_\rho^{\delta_{str}}(f_{str}) = str = \text{str}_\rho^{\varepsilon_{str}}(\ell_{str})$ , we can define a bijection  $\mathbf{g}_{str}: \mathbb{A}_\rho^{\delta_{str}} \rightarrow \mathbb{A}_\rho^{\varepsilon_{str}}$  between  $\delta_{str}$ -agents and  $\varepsilon_{str}$ -agents appearing in  $\rho$  such that  $\gamma_1(a) = \gamma_1(\mathbf{g}_{str}(a))$ ,  $\gamma_m(a) = \gamma_m(\mathbf{g}_{str}(a))$ , and  $\gamma_{f_{str}}(a) = \gamma_{\ell_{str}}(\mathbf{g}_{str}(a))$  for all  $\delta_{str}$ -agents  $a \in \mathbb{A}_\rho^{\delta_{str}}$ .

We pick  $|\mathcal{S}_{tr}| \leq M$  arbitrary, but pairwise distinct, data  $(\eta_{str})_{str \in \mathcal{S}_{tr}}$  in  $\mathbb{D}_\rho^{tr}$ , one for each split trace reached by a datum of  $\mathbb{D}_\rho^{tr}$  in  $\rho$ . For each  $str \in \mathcal{S}_{tr}$ , since  $\delta_{str}, \eta_{str} \in \mathbb{D}_\rho^{tr}$  have the same trace  $tr$  in  $\rho$ , we can define a bijection  $\mathbf{f}_{str}: \mathbb{A}_\rho^{\eta_{str}} \rightarrow \mathbb{A}_\rho^{\delta_{str}}$  between  $\eta_{str}$ -agents and  $\delta_{str}$ -agents appearing in  $\rho$  so that  $\gamma_1(a) = \gamma_1(\mathbf{f}_{str}(a))$  and  $\gamma_m(a) = \gamma_m(\mathbf{f}_{str}(a))$  for all  $\eta_{str}$ -agents  $a \in \mathbb{A}_\rho^{\eta_{str}}$ .

The function  $\mathbf{l}_{str} := \mathbf{g}_{str} \circ \mathbf{f}_{str}: \mathbb{A}_\rho^{\eta_{str}} \rightarrow \mathbb{A}_\rho^{\varepsilon_{str}}$  is then a bijection between  $\eta_{str}$ -agents and  $\varepsilon_{str}$ -agents such that  $\gamma_1(a) = \gamma_1(\mathbf{l}_{str}(a))$  and  $\gamma_m(a) = \gamma_m(\mathbf{l}_{str}(a))$  for all  $\eta_{str}$ -agents  $a \in \mathbb{A}_\rho^{\eta_{str}}$ .

Intuitively, we will replace  $\mathbb{D}_\rho^{tr}$  with a subset  $\{\eta_{str} \mid str \in \mathcal{S}_{tr}\}$ . Agents of each  $\eta_{str}$  will be in charge of allowing all external observations on agents of data  $d \in \mathbb{D}_\rho^{tr}$  in carried when that datum matches split trace  $str$ .

To do so, agents of  $\eta_{str}$  start by mimicking the moves of  $\delta_{str}$  (the first datum of  $\mathbb{D}_\rho^{tr}$  to reach  $str$ ) until  $f_{str}$ . This is done by making each agent copy its image by  $\mathbf{f}_{str}$ . Then, agents of  $\eta_{str}$  remain idle and allow all aforementioned external observations to be carried out. Finally, when point  $\ell_{str}$  is reached in  $\rho$ , the agents of  $\eta_{str}$  copy their images by  $\mathbf{l}_{str}$  until the end.

This is formalised by the following claim.

▷ **Claim 26.** There is a run  $\tilde{\rho}: \tilde{\gamma}_1 \xrightarrow{*} \tilde{\gamma}_2 \xrightarrow{*} \dots \xrightarrow{*} \tilde{\gamma}_m$  with  $\mathbb{D}_{\tilde{\rho}} = (\mathbb{D}_\rho \setminus \mathbb{D}_\rho^{tr}) \cup \{\eta_{str} \mid str \in \mathcal{S}_{tr}\}$  such that for all  $i \in [1, m]$ , the following holds.

1. For all  $d \in \mathbb{D}_{\tilde{\rho}}$ , we have  $\mathbb{A}_{\tilde{\rho}}^d = \mathbb{A}_\rho^d$ .
2. For all  $d \in \mathbb{D}_\rho \setminus \mathbb{D}_\rho^{tr}$  and  $a \in \mathbb{A}_\rho^d$ , we have  $\tilde{\gamma}_i(a) = \gamma_i(a)$ .
3. For all  $str \in \mathcal{S}_{tr}$  and all  $a \in \mathbb{A}_\rho^{\eta_{str}}$ ,
  - if  $i \leq f_{str}$ , then  $\tilde{\gamma}_i(a) = \gamma_i(\mathbf{f}_{str}(a))$ ,
  - if  $f_{str} \leq i \leq \ell_{str}$ , then  $\tilde{\gamma}_i(a) = \gamma_{f_{str}}(\mathbf{f}_{str}(a)) = \gamma_{\ell_{str}}(\mathbf{l}_{str}(a))$ , and
  - if  $i \geq \ell_{str}$  then  $\tilde{\gamma}_i(a) = \gamma_i(\mathbf{l}_{str}(a))$ .

*Proof.* We inductively define configurations  $\tilde{\gamma}_i$  and runs  $\tilde{\gamma}_1 \xrightarrow{*} \dots \xrightarrow{*} \tilde{\gamma}_i$  for all  $i \in [1, m]$  that satisfy Items 1–3. For all  $a \in \mathbb{A}$ , we set  $\tilde{\gamma}_1(a) := \gamma_1(a)$  if  $\mathbf{dat}(a) \notin \mathbb{D}_\rho^{tr}$  or if  $\mathbf{dat}(a) = \delta_{str}$  for some  $str \in \mathcal{S}_{tr}$ , and we set  $\tilde{\gamma}_1(a) := *$  otherwise. Now assume we have defined  $\tilde{\rho}$  up to  $\tilde{\gamma}_i$  with  $i < m$ . We continue with the run up to  $\tilde{\gamma}_{i+1}$ . Let  $\gamma_i \xrightarrow{\triangleleft a_o} \gamma_{i+1}$  be the step at hand, and let  $q \xrightarrow{\triangleleft q_o} p$  be the used transition.

- If  $\mathbf{dat}(a_o) \notin \mathbb{D}_\rho^{tr}$ , by Items 1 and 2, we have  $\tilde{\gamma}_i(a_o) = \gamma_i(a_o) = q_o$ , and we set  $\tilde{a}_o := a_o$ .

- If  $d_o := \mathbf{dat}(a_o) \in \mathbb{D}_\rho^{tr}$ , then let  $str_o := \mathbf{str}_\rho^{d_o}(i)$ . Since  $f_{str_o}$  and  $\ell_{str_o}$  were defined as the minimal and maximal indices where  $str_o$  occurs, we have  $f_{str_o} \leq i \leq \ell_{str_o}$ . Moreover, since  $\gamma_i(a_o) = q_o$ , we must have  $str_o(q_{1,o}, q_o, q_{m,o}) > 0$  for some  $q_{1,o}, q_{m,o} \in Q$ . As a result, there exists a  $\delta_{str_o}$ -agent  $a'_o$  such that  $\gamma_{f_{str_o}}(a'_o) = q_o$ . We let  $\tilde{a}_o := \mathbf{f}_{str_o}^{-1}(a'_o)$ . Then, by Item 3, we have  $\tilde{\gamma}_i(\tilde{a}_o) = \gamma_{f_{str_o}}(a'_o) = q_o$ .

In both cases, we have  $\tilde{\gamma}_i(\tilde{a}_o) = q_o$ . We now define the steps  $\tilde{\gamma}_i \xrightarrow{*} \tilde{\gamma}_{i+1}$  based on another case distinction.

- If  $\mathbf{dat}(a) \notin \mathbb{D}_\rho^{tr}$ , then we use the step  $\tilde{\gamma}_i \xrightarrow{\bowtie \tilde{a}_o} \tilde{\gamma}_{i+1}$  based on the transition  $q \xrightarrow{\bowtie q_o} p$ . If  $\mathbf{dat}(a_o) \notin \mathbb{D}_\rho^{tr}$ , then  $\tilde{a}_o = a_o$ , and this transition can be taken because  $\tilde{\gamma}_i(a) = \gamma_i(a)$  and  $\tilde{\gamma}_i(a_o) = \gamma_i(a_o)$ . Otherwise, if  $\mathbf{dat}(a_o) \in \mathbb{D}_\rho^{tr}$ , then  $\mathbf{dat}(a_o) \neq \mathbf{dat}(a)$ , so  $\bowtie$  is  $\neq$ . Further,  $\mathbf{dat}(\tilde{a}_o) \in \{\eta_{str} \mid str \in \mathcal{S}_{tr}\}$ , so  $\mathbf{dat}(\tilde{a}_o) \neq \mathbf{dat}(a)$ . Also, we defined  $\tilde{a}_o$  so that  $\tilde{\gamma}_i(\tilde{a}_o) = \gamma_i(a_o)$ . Then the transition can be taken because  $\tilde{\gamma}_i(a) = \gamma_i(a)$  and  $\tilde{\gamma}_i(\tilde{a}_o) = \gamma_i(a_o)$ .
- If  $\mathbf{dat}(a) \in \mathbb{D}_\rho^{tr} \setminus \{\delta_{str}, \varepsilon_{str} \mid str \in \mathcal{S}_{tr}\}$ , then we ignore the step and set  $\tilde{\gamma}_{i+1} := \tilde{\gamma}_i$ .
- If  $\mathbf{dat}(a) \in \{\delta_{str}, \varepsilon_{str} \mid str \in \mathcal{S}_{tr}\}$ , then, for all  $str \in \mathcal{S}_{tr}$  we sequentially apply the following.
  - If  $i + 1 \leq f_{str}$  and  $\mathbf{dat}(a) = \delta_{str}$ , then we move the agent  $\tilde{a} := \mathbf{f}_{str}^{-1}(a)$ . By Item 3, we know that  $\tilde{\gamma}_i(\tilde{a}) = \gamma_i(\mathbf{f}_{str}(\tilde{a})) = \gamma_i(a) = q$ . Furthermore, we have  $\tilde{\gamma}_i(\tilde{a}_o) = \gamma_i(a_o) = q_o$ . Finally, if  $\bowtie$  is  $=$ , then it holds that  $\mathbf{dat}(a_o) = \delta_{str}$ , and thus  $\mathbf{dat}(\tilde{a}) = \eta_{str} = \mathbf{dat}(\tilde{a}_o)$ . If  $\bowtie$  is  $\neq$ , then by definition of  $f_{str}$ , as  $i < f_{str}$ , at this point, no datum of trace  $tr$  matches split trace  $str$ . That is, for all data  $d \in \mathbb{D}_\rho^{tr}$ , we have  $\mathbf{str}_\rho^d(i) \neq str$ . Thus, for  $d_o := \mathbf{dat}(a_o)$ , we have  $d_o \notin \mathbb{D}_\rho^{tr}$  or  $str_o := \mathbf{str}_\rho^{d_o}(i) \neq str$ . In the first case, we have  $\tilde{a}_o = a_o$ , so  $\mathbf{dat}(\tilde{a}_o) \notin \mathbb{D}_\rho^{tr}$ , but  $\mathbf{dat}(\tilde{a}) = \eta_{str} \in \mathbb{D}_\rho^{tr}$ . In the second case, we have  $\mathbf{dat}(\tilde{a}_o) = \eta_{str_o} \neq \eta_{str} = \mathbf{dat}(\tilde{a})$ , since the data  $(\eta_{str'})_{str' \in \mathcal{S}_{tr}}$  are pairwise distinct. In both cases, we have  $\mathbf{dat}(\tilde{a}_o) \neq \eta_{str}$ . Hence, in all cases, we can let the agent  $\tilde{a}$  take transition  $q \xrightarrow{\bowtie q_o} p$ .
  - Similarly, if  $i \geq \ell_{str}$  and  $\mathbf{dat}(a) = \varepsilon_{str}$ , then we move the agent  $\tilde{a} := \mathbf{l}_{str}^{-1}(a)$ . By Item 3, we know that  $\tilde{\gamma}_i(\tilde{a}) = \tilde{\gamma}_i(\mathbf{l}_{str}^{-1}(a)) = \gamma_i(a) = q$ . Furthermore, we have  $\tilde{\gamma}_i(\tilde{a}_o) = \gamma_i(a_o) = q_o$ . Finally, if  $\bowtie$  is  $=$ , then it holds that  $\mathbf{dat}(a_o) = \varepsilon_{str}$ , and thus  $\mathbf{dat}(\tilde{a}) = \eta_{str} = \mathbf{dat}(\tilde{a}_o)$ . If  $\bowtie$  is  $\neq$  then by definition of  $\ell_{str}$ , as  $i \geq \ell_{str}$ ,  $\mathbf{dat}(a_o)$  cannot have trace  $tr$  and match split trace  $str$ , as otherwise it would still match it at step  $i + 1 > \ell_{str}$ . Therefore, analogously to the case  $i + 1 \leq f_{str}$  and  $\mathbf{dat}(a) = \delta_{str}$ , for  $d_o := \mathbf{dat}(a_o)$ , we have  $d_o \notin \mathbb{D}_\rho^{tr}$  or  $str_o := \mathbf{str}_\rho^{d_o}(i) \neq str$ . Again, this implies that  $\mathbf{dat}(\tilde{a}_o) \neq \eta_{str}$ . In all cases, we can let the agent  $\tilde{a}$  take transition  $q \xrightarrow{\bowtie q_o} p$ .

Note that for each  $str$ , at most one of the two cases applies. Furthermore, all the moving agents  $\tilde{a}$  are distinct, as they have different data (since the  $\eta_{str}$  are distinct). Moreover,  $\tilde{a}_o$  does not move at any point. Therefore, as all those steps are enabled in  $\tilde{\gamma}_i$ , they can all be taken sequentially to get to  $\tilde{\gamma}_{i+1}$ .

It is straightforward to show that the induction hypothesis is maintained in all these cases. Since the moves of agents  $a$  with  $\mathbf{dat}(a) \notin \mathbb{D}_\rho^{tr}$  do not change, clearly,  $\tilde{\gamma}_{i+1}(a) = \gamma_{i+1}(a)$ . For agents  $a$  with  $\mathbf{dat}(a) = \eta_{str}$  for some  $str \in \mathcal{S}_{tr}$ , note that we make  $a$  follow the same steps as  $\mathbf{f}(a)$  until point  $f_{str}$ , then stay idle until  $\ell_{str}$ , and then follow the same steps as  $\mathbf{l}(a)$ .

This concludes our induction.  $\triangleleft$

The following claim is a direct consequence of the previous one, as all remaining data have preserved their initial and final configuration, thus their traces. Moreover, we have only



deleted data with trace  $tr$ . Since there is a datum with trace  $tr$  in  $\rho$ , the set  $\mathcal{S}_{tr}$  is not empty, and thus there are data in  $\rho'$  with trace  $tr$ , i. e., the datum  $\eta_{str}$  for every  $str \in \mathcal{S}_{tr}$ . Hence, as all traces appearing in  $\rho$  are represented in  $\rho'$ .

▷ **Claim 27.** For all  $tr: Q^2 \rightarrow [0, K]$ , there is a run  $\tilde{\rho}: \tilde{\gamma}_1 \xrightarrow{*} \tilde{\gamma}_m$  over the set of agents appearing in  $\rho$  with datum in  $(\mathbb{D}_\rho \setminus \mathbb{D}_\rho^{tr}) \cup \{\eta_{str} \mid str \in \mathcal{S}_{tr}\}$  such that

- for all  $d \in \mathbb{D}_{\rho'}$  and all  $d$ -agent  $a$ , we have  $\gamma_1(a) = \tilde{\gamma}_1(a)$  and  $\gamma_m(a) = \tilde{\gamma}_m(a)$ ,
- for all  $d \in \mathbb{D}_\rho$ , there exists  $d'$  such that  $\mathbf{tr}_{\rho'}^{d'} = \mathbf{tr}_\rho^d$ , and
- there are at most  $M$  data  $d'$  such that  $\mathbf{tr}_{d'}^{\rho'} = tr$ .

The statement of Lemma 16 follows by iteratively applying the claim for each trace  $tr: Q^2 \rightarrow [0, K]$  with  $|\mathbb{D}_\rho^{tr}| > M$ ; as argued above, once we have  $|\mathbb{D}_\rho^{tr}| \leq M$  for all  $tr: Q^2 \rightarrow [0, K]$ , then  $|\mathbb{D}_\rho| \leq M \cdot (K + 1)^{|Q|^2} = (K + 1)^{|Q|^3 + |Q|^2}$ , so  $\rho' := \rho$  fulfils the requirements of the lemma. ◀

## C From Expressions to Containers: Proofs for Section 5

► **Lemma 18.** *Let  $n_1, n_2, M_1, M_2 \in \mathbb{N}$ . If  $n_1 \leq n_2$  and  $M_1 \leq M_2$ , then every  $(n_1, M_1)$ -container is a union of  $(n_2, M_2)$ -containers.*

**Proof.** First, we show that every  $n_1$ -box is a union of  $n_2$ -boxes. Let  $\gamma, \chi \in \Gamma$  and  $d, d' \in \mathbb{D}$  with  $\lceil \gamma, d \rceil^{n_2} = \lceil \chi, d' \rceil^{n_2}$ . For every state  $q$ , there are either at least  $n_2$  agents with datum  $d$  in  $\gamma$  and at least  $n_2$  agents with datum  $d'$  in  $\chi$ , in which case there are at least  $n_1$  agents in both, or the two numbers are the same.

As a result,  $\lceil \gamma, d \rceil^{n_1} = \lceil \chi, d' \rceil^{n_1}$ . Hence, the partition of  $\Gamma \times \mathbb{D}$  induced by  $n_2$ -boxes is at least as fine as the one induced by  $n_1$ -boxes.

Now let  $\gamma, \chi \in \Gamma$  such that  $\lceil \gamma \rceil^{n_2, M_2} = \lceil \chi \rceil^{n_2, M_2}$ . We show that  $\lceil \gamma \rceil^{n_1, M_1} = \lceil \chi \rceil^{n_1, M_1}$ . Let  $\mathbf{b} \in \mathbf{Boxes}_{n_1}$ . Then  $\mathbf{b}$  is a union of  $n_2$ -boxes  $\mathbf{b}_1, \dots, \mathbf{b}_k$ . If one of the two configurations has less than  $M_1$  data mapped to  $\mathbf{b}$ , then, as  $M_2 \geq M_1$ , it has less than  $M_2$  data mapped to each  $\mathbf{b}_i$ . As a consequence, the other configuration has the same number of data mapped to each  $\mathbf{b}_i$ , and thus the same number of data mapped to  $\mathbf{b}$ .

This shows that  $(n_2, M_2)$ -containers form a partition of  $\Gamma$  that is at least as fine as  $(n_1, M_1)$ -containers, concluding our proof. ◀

### C.1 Proof of Proposition 19 and Comparison Between Sizes of Representations

► **Proposition 19.** *The sets of configurations defined by interval predicates of height at most  $n$  and width at most  $M$  are exactly the sets formed by unions of  $(n, M)$ -containers.*

We start with the translation from interval predicates to containers. First, in Lemma 28, we show that a simple interval predicate of height  $n$  and width  $M$  cannot distinguish configurations that are equivalent with respect to  $\equiv_{n, M}$ . Then, in Lemma 29, we use this fact to prove that an interval predicate of height  $n$  and width  $M$  can be translated into a union of  $(n, M)$ -containers.

► **Lemma 28.** *Let  $n, M \in \mathbb{N}$ , and let  $\gamma, \chi$  be configurations such that  $\gamma \equiv_{n, M} \chi$ . Furthermore, let  $\psi$  be a simple interval predicate of height at most  $n$  and width at most  $M$ . Then  $\gamma$  satisfies  $\psi$  if and only if  $\chi$  does.*

**Proof.** Let  $\psi = \exists d_1, \dots, d_M, \bigwedge_{q \in Q} \bigwedge_{j=1}^M \#(q, d_j) \in [A_{q,j}, B_{q,j}]$  be a simple interval predicate of height  $n$  and width  $M$ .

Suppose  $\gamma$  satisfies  $\psi$ . Then there are pairwise distinct data  $d_1, \dots, d_M$  such that for all  $q \in Q$  and  $j \in [1, M]$ , it holds that  $\gamma_{d_j}^\#(q) \in [A_{q,j}, B_{q,j}]$ .

Let  $\mathbf{b}$  be an  $n$ -box, and let  $d_{\mathbf{b},1}, \dots, d_{\mathbf{b},k_{\mathbf{b}}}$  be the pairwise distinct data among  $(d_i)_{1 \leq i \leq M}$  with  $[\gamma, d_i]^n = \mathbf{b}$ . As  $k_{\mathbf{b}} \leq M$  and  $\gamma \equiv_{n,M} \chi$ , there are pairwise distinct data  $d'_{\mathbf{b},1}, \dots, d'_{\mathbf{b},k_{\mathbf{b}}}$  such that  $[\chi, d'_{\mathbf{b},j}]^n = \mathbf{b}$  for all  $j \in [1, k_{\mathbf{b}}]$ . By doing this for every  $n$ -box, we obtain pairwise distinct data  $d'_1, \dots, d'_M$  such that  $(\gamma, d_i) \equiv_n (\chi, d'_i)$  for all  $i \in [1, M]$ . Moreover, for all  $i \in [1, M]$ , since  $A_{q,i}, B_{q,i} \in [0, n] \cup \{+\infty\}$ , and since  $\gamma_{d_i}^\#(q) \in [A_{q,i}, B_{q,i}]$  and  $(\gamma, d_i) \equiv_n (\chi, d'_i)$ , we have  $\chi_{d'_i}^\#(q) \in [A_{q,i}, B_{q,i}]$ .

This shows that  $\chi$  satisfies  $\psi$ . The other direction follows by symmetry.  $\blacktriangleleft$

**► Lemma 29.** *Let  $\varphi$  be an interval predicate of height at most  $n$  and width at most  $M$ . Then  $\llbracket \varphi \rrbracket_{\mathcal{P}}$  is a union of  $(n, M)$ -containers.*

**Proof.** Let  $\gamma, \chi$  be configurations such that  $\gamma \equiv_{n,M} \chi$ . By definition,  $\varphi$  is a Boolean combination of simple interval predicates  $\psi_1, \dots, \psi_p$  for some  $p \in \mathbb{N}$ . Furthermore, by Lemma 28, for every  $i \in [1, p]$ , the predicate  $\psi_i$  is satisfied by  $\gamma$  if and only if it is satisfied by  $\chi$ . Thus,  $\varphi$  is satisfied by  $\gamma$  if and only if it is satisfied by  $\chi$ .

As a result, we obtain that each equivalence class of  $\equiv_{n,M}$  (i. e., each  $(n, M)$ -container) is either fully contained in  $\llbracket \varphi \rrbracket_{\mathcal{P}}$  or disjoint from it. Since  $(n, M)$ -containers form a partition of the set of configurations, this implies the statement of Lemma 29.  $\blacktriangleleft$

The next result takes care of the other direction of the proof of Proposition 19. That is, we show that any finite union of  $(n, M)$ -containers can be expressed as an interval predicate of height  $n$  and width  $M$ . Combined with the first direction, this shows that the two formalisms are equally expressive.

**► Lemma 30.** *Let  $\text{cont}_1, \dots, \text{cont}_k$  be sets of configurations such that for all  $i$ ,  $\text{cont}_i$  is an  $(n_i, M_i)$ -container. Then there is an interval predicate of width  $\max_i n_i$  and height  $\max_i M_i$  that defines the set  $\bigcup_{i=1}^k \text{cont}_i$ .*

**Proof.** As interval predicates are closed under disjunction, we only need to show that every  $(n, M)$ -container can be expressed as an interval predicate of width  $n$  and height  $M$ .

Let  $\text{cont}$  be an  $(n, M)$ -container. We construct an interval predicate expressing the same set of configurations. Let  $\mathbf{b}$  be an  $n$ -box, and let  $m \in \mathbb{N}$ . We set

$$\varphi_{\mathbf{b}, \geq m} := \exists d_1, \dots, d_m, \bigwedge_{q \in Q} \bigwedge_{j=1}^m \#(q, d_j) \in [A_{q,j}, B_{q,j}]$$

where for all  $q \in Q$  and  $j \in [1, m]$ , we set  $A_{q,j} := \mathbf{b}(q)$ , and we set  $B_{q,j} := \mathbf{b}(q)$  if  $\mathbf{b}(q) < n$  and  $B_{q,j} := +\infty$  otherwise. This simple interval predicate expresses that there are at least  $m$  pairwise distinct data that match the  $n$ -box  $\mathbf{b}$ . Let

$$\psi_{\text{cont}} := \bigwedge_{\substack{\mathbf{b} \in \text{Boxes}_n \\ \text{cont}(\mathbf{b}) < M}} (\varphi_{\mathbf{b}, \geq \text{cont}(\mathbf{b})} \wedge \neg \varphi_{\mathbf{b}, \geq \text{cont}(\mathbf{b})+1}) \wedge \bigwedge_{\substack{\mathbf{b} \in \text{Boxes}_n \\ \text{cont}(\mathbf{b}) = M}} \varphi_{\mathbf{b}, \geq M}.$$

This interval predicate expresses that for every  $n$ -box  $\mathbf{b}$ , the number of data mapped to  $\mathbf{b}$  matches the corresponding number in  $\text{cont}$ , or that the number of data mapped to  $\mathbf{b}$  is at least  $M$  if the corresponding number in  $\text{cont}$  is  $M$ . Hence, the interval predicate  $\psi_{\text{cont}}$  is satisfied by a configuration  $\gamma$  if and only if  $\gamma \in \text{cont}$ .  $\blacktriangleleft$

► **Remark 20.** Containers can be exponentially more succinct than interval predicates, while interval predicates can be doubly exponentially more succinct than unions of containers.

**Proof.** To see this, suppose we want to express that there are exactly  $2^k - 1$  data that all have one agent in  $q_1$  and no agents in  $q_0$ . This can be expressed with a  $(2, 2^k)$ -container (whose binary encoding uses  $\mathcal{O}(k)$  bits for each number). Meanwhile, this cannot be expressed by a union of  $(n, M)$ -containers with  $M < 2^k$ , as they cannot distinguish a configuration with  $2^k - 1$  such data from one with more than  $2^k - 1$ . By the proposition above, this means that an interval predicate for this set requires width  $2^k$ , and thus its encoding must be of size  $\Omega(2^k)$  (as we must have at least this many data variables).

Conversely, consider the set of configurations containing at least one agent in state  $q$ . This is expressible by a trivial interval predicate, but it corresponds to the union of all  $(1, 1)$ -containers  $\text{cont}$  such that  $\text{cont}(\mathbf{b}) > 0$  for some box  $\mathbf{b}$  with  $\mathbf{b}(q) > 0$ . One can show that there are  $2^{2^{|\mathcal{Q}|}} - 2^{2^{|\mathcal{Q}|-1}}$  such containers. ◀

## C.2 Proof of Lemma 22

Recall that we chose  $f: \mathbb{N} \rightarrow \mathbb{N}$  and  $g: \mathbb{N}^2 \rightarrow \mathbb{N}$  to be the functions with  $f(n) := (n + |\mathcal{P}|^3) \cdot |\mathcal{P}|$  and  $g(n, M) := (M + (|\mathcal{P}|^3 + 1)^{|\mathcal{P}|^3 + |\mathcal{P}|^2})(n + 1)^{|\mathcal{P}|}$ . for all  $n, M \in \mathbb{N}$

► **Lemma 22.** *For all  $n, M \in \mathbb{N}$  and all configurations  $\gamma_{start}, \gamma_{end}, \chi_{start} \in \Gamma$ , if there is a run  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  and  $\gamma_{start} \equiv_{f(n), g(n, M)} \chi_{start}$ , then there is a configuration  $\chi_{end} \in \Gamma$  with  $\gamma_{end} \equiv_{n, M} \chi_{end}$  and a run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$ .*

**Proof.** By Corollary 17, we can assume that  $\rho$  has at most  $|\mathcal{P}|^3$  observed agents per datum and at most  $(|\mathcal{P}|^3 + 1)^{|\mathcal{P}|^3 + |\mathcal{P}|^2}$  externally observed data. Let  $\mathbb{A}_\gamma$  and  $\mathbb{A}_\chi$  be the sets of agents appearing in  $\gamma_{start}$  and  $\chi_{start}$ , respectively. Let  $\mathbb{D}_\gamma$  and  $\mathbb{D}_\chi$  be their sets of data.

We are going to define maps  $\mu: \mathbb{D}_\chi \rightarrow \mathbb{D}_\gamma$  and  $\nu: \mathbb{A}_\chi \rightarrow \mathbb{A}_\gamma$ . Intuitively,  $\mu$  (resp.  $\nu$ ) will map each datum in  $\mathbb{D}_\chi$  (resp. agent in  $\mathbb{A}_\chi$ ) to a datum in  $\gamma_{start}$  (resp. agent in  $\mathbb{A}_\gamma$ ) that it should mimic. We will ensure that data and agents are mapped to a counterpart in  $\gamma_{start}$  that is “compatible” in the sense that they can mimic that counterpart while keeping the same initial and final configurations (up to  $n$ -approximation for the data). We will also ensure that  $\mu$  maps a datum to each externally observed datum in  $\rho$ , so that the former can fulfil the same roles in the run from  $\chi_{start}$ . For the same reason, for each datum  $d \in \mathbb{D}_\chi$ , all observed agents in  $\rho$  with datum  $\mu(d)$  must have an agent of  $\chi_{start}$  with datum  $d$  mapped to them. We will be able to satisfy these conditions as  $\gamma_{start}$  and  $\chi_{start}$  are equivalent up to sufficient bounds.

The proof is split in three parts. We first prove two claims establishing the existence of mappings  $\mu$  and  $\nu$  with the desired properties. The rest of the proof is dedicated to the construction of the run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$ .

For all  $\mathbf{b} \in \mathbf{Boxes}_{f(n)}$  and  $\mathbf{b}' \in \mathbf{Boxes}_n$ , we set  $\mathbb{D}_\gamma^{\mathbf{b} \rightarrow} := \{d \in \mathbb{D}_\gamma \mid \lceil \gamma_{start}, d \rceil^{f(n)} = \mathbf{b}\}$ ,  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow} := \{d \in \mathbb{D}_\chi \mid \lceil \chi_{start}, d \rceil^{f(n)} = \mathbf{b}\}$ ,  $\mathbb{D}_\gamma^{\rightarrow \mathbf{b}'} := \{d \in \mathbb{D}_\gamma \mid \lceil \gamma_{end}, d \rceil^n = \mathbf{b}'\}$ , and  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow \mathbf{b}'} := \mathbb{D}_\gamma^{\mathbf{b} \rightarrow} \cap \mathbb{D}_\chi^{\rightarrow \mathbf{b}'}$ .

► **Claim 31.** There exists a mapping  $\mu: \mathbb{D}_\chi \rightarrow \mathbb{D}_\gamma$  such that

1. for all  $d \in \mathbb{D}_\chi$ , we have  $(\chi_{start}, d) \equiv_{f(n)} (\gamma_{start}, \mu(d))$ ,
2. for all data  $d' \in \mathbb{D}_\gamma$  externally observed in  $\rho$ ,  $\mu^{-1}(d') \neq \emptyset$ , and
3. for all  $\mathbf{b}' \in \mathbf{Boxes}_n$ , we have  $|\mu^{-1}(\mathbb{D}_\chi^{\rightarrow \mathbf{b}'})| = |\mathbb{D}_\gamma^{\rightarrow \mathbf{b}'}|$ , or  $|\mathbb{D}_\chi^{\rightarrow \mathbf{b}'}| \geq M$  and  $|\mu^{-1}(\mathbb{D}_\chi^{\rightarrow \mathbf{b}'})| \geq M$ .

Proof. As  $\gamma_{start} \equiv_{f(n), g(n, M)} \chi_{start}$ , we know that for all  $\mathbf{b} \in \mathbf{Boxes}_{f(n)}$ , it holds that  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow}| = |\mathbb{D}_\chi^{\mathbf{b} \rightarrow}|$ , or both  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow}|$  and  $|\mathbb{D}_\chi^{\mathbf{b} \rightarrow}|$  are at least  $g(n, M)$ .

In the first case, we let  $\mu$  map a datum of  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow}$  to each datum  $\mathbb{D}_\gamma^{\mathbf{b} \rightarrow}$  to form a bijection. Otherwise, in the second case, for each  $\mathbf{b}' \in \mathbf{Boxes}_n$ , if  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| \leq M$ , then we map a datum of  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow}$  to each datum of  $\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}$ . If  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| > M$ , then we select  $M$  data from it, we add all those which are externally observed in  $\rho$ , and we map a datum of  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow}$  to each one of them. As there are at most  $(|\mathcal{P}|^3 + 1)^{|\mathcal{P}|^3 + |\mathcal{P}|^2}$  externally observed data in  $\rho$ , we have selected at most  $M + ((|\mathcal{P}|^3 + 1)^{|\mathcal{P}|^3 + |\mathcal{P}|^2})$  data per  $n$ -box  $\mathbf{b}'$ . Further, as  $|\mathbf{Boxes}_n| \leq (n + 1)^{|\mathcal{P}|}$ , we have selected at most  $g(n, M)$  data in total. Hence, we can indeed map a datum of  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow}$  to each one of them. As  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow}| \geq g(n, M) \geq M \cdot |\mathbf{Boxes}_n|$ , there exists  $\mathbf{b}' \in \mathbf{Boxes}_n$  such that  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| \geq M$ . We pick any  $d \in \mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}$  and map all remaining data of  $\mathbb{D}_\chi^{\mathbf{b} \rightarrow}$  to  $d$ .

This concludes the construction of  $\mu$ . It remains to prove that  $\mu$  fulfils the requirements of the claim. The first two items follow directly from the definition of  $\mu$ . Now let  $\mathbf{b}' \in \mathbf{Boxes}_n$ .

- If  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| \leq M$  for all  $\mathbf{b} \in \mathbf{Boxes}_{f(n)}$ , then, by definition of  $\mu$ ,  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| = |\mu^{-1}(\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'})|$  for all  $\mathbf{b} \in \mathbf{Boxes}_{f(n)}$ . As  $\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'} = \bigsqcup_{\mathbf{b} \in \mathbf{Boxes}_{f(n)}} \mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}$ , we obtain  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| = |\mu^{-1}(\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'})|$ .
- If  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}| \geq M$  for some  $\mathbf{b} \in \mathbf{Boxes}_{f(n)}$ , then  $|\mu^{-1}(\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'})| \geq M$  by definition of  $\mu$ . As  $\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'} \subseteq \mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}$ , we obtain that both  $|\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'}|$  and  $|\mu^{-1}(\mathbb{D}_\gamma^{\mathbf{b} \rightarrow \mathbf{b}'})|$  are at least  $M$ .

This concludes the proof of the claim.  $\triangleleft$

We now use similar arguments as above to define the mapping  $\nu$ . The proof is almost identical, but we provide it in full to avoid any confusion. For every  $d \in \mathbb{D}_\chi$ , we let  $\mathbb{A}_{\gamma, d}$  (resp.  $\mathbb{A}_{\chi, d}$ ) be the set of agents in  $\mathbb{A}_\gamma$  (resp.  $\mathbb{A}_\chi$ ) with datum  $d$ . Moreover, for all  $q, q' \in Q$ , we set  $\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow} := \{a \in \mathbb{A}_{\gamma, \mu(d)} \mid \gamma_{start}(a) = q\}$ ,  $\mathbb{A}_{\chi, d}^{q \rightarrow} := \{a \in \mathbb{A}_{\chi, d} \mid \chi_{start}(a) = q\}$ ,  $\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'} := \{a \in \mathbb{A}_{\gamma, \mu(d)} \mid \gamma_{end}(a) = q'\}$ , and  $\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'} := \mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow} \cap \mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'}$ .

$\triangleright$  Claim 32. Let  $\mu$  be the mapping constructed in the previous claim. There exists a mapping  $\nu: \mathbb{A}_\chi \rightarrow \mathbb{A}_\gamma$  such that, for all  $d \in \mathbb{D}_\chi$ , it holds that

- (A) for all  $a \in \mathbb{A}_{\chi, d}$ , we have  $\chi_{start}(a) = \gamma_{start}(\nu(a))$ ,
- (B) for all observed  $a' \in \mathbb{A}_{\gamma, \mu(d)}$ , we have  $\nu^{-1}(a') \neq \emptyset$ , and
- (C) for all  $q' \in Q$ , we have  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'})| = |\mathbb{A}_{\chi, d}^{\rightarrow q'}|$ , or  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'})| \geq n$  and  $|\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'}| \geq n$ .

Proof. We proceed datum by datum. That is, for each  $d \in \mathbb{D}_\chi$ , we define the map  $\nu$  over agents in  $\chi_{start}$  with datum  $d$ .

Let  $d \in \mathbb{D}_\chi$ . By Item 1 of Claim 31, we have  $(\chi_{start}, d) \equiv_{f(n)} (\gamma_{start}, \mu(d))$ . Thus, we know that for all  $q \in Q$ , it holds that  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow}| = |\mathbb{A}_{\chi, d}^{q \rightarrow}|$ , or both  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow}|$  and  $|\mathbb{A}_{\chi, d}^{q \rightarrow}|$  are at least  $f(n)$ . In the first case, we let  $\nu$  form a bijection between  $\mathbb{A}_{\chi, d}^{q \rightarrow}$  and  $\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow}$ . Otherwise, in the second case, for each  $q' \in Q$ , if  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}| \leq n$ , then we select  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}|$  agents from  $\mathbb{A}_{\chi, d}^{q \rightarrow}$  and let  $\nu$  form a bijection between them and  $\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}$ . Otherwise, if  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}| > n$ , then we select  $n$  agents from it, add all those which are observed in  $\rho$ , and we let  $\nu$  be surjective on this set by picking an agent from  $\mathbb{A}_{\chi, d}^{q \rightarrow}$  for every one of them.

As there are at most  $|\mathcal{P}|^3$  observed agents with datum  $d$  in  $\rho$  (by Corollary 17, as mentioned at the beginning of the proof), we have selected at most  $n + |\mathcal{P}|^3$  agents per state  $q'$ . Hence, we have selected at most  $f(n)$  agents in total. Thus, we can indeed map an agent of  $\mathbb{A}_{\chi, d}^{q \rightarrow}$  to each one of them.

As  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow}| \geq f(n) \geq n \cdot |\mathcal{P}|$ , there exists  $q' \in Q$  such that  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}| \geq n$ . We pick any  $a \in \mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}$  and map all remaining agents of  $\mathbb{A}_{\chi, d}^{q \rightarrow}$  to  $a$ .

This concludes the construction of  $\nu$ . It remains to prove that  $\nu$  fulfils the requirements of the claim. The two first items are immediate from the definition. The third item is obtained by observing that

- if  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}| \leq n$  for all  $q \in Q$ , then  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'})| = |\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}|$  for all  $q$ , and thus  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'})| = |\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'}|$ , and
- if  $|\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'}| \geq n$  for some  $q \in Q$ , then  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{q \rightarrow q'})| \geq n$ , and thus both  $|\nu^{-1}(\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'})|$  and  $|\mathbb{A}_{\gamma, \mu(d)}^{\rightarrow q'}|$  are at least  $n$ .

This concludes the proof of the claim.  $\triangleleft$

In the remainder of the proof, we construct a run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$  with  $\chi_{end} \equiv_{n, M} \gamma_{end}$  as required in the statement of the lemma. For that, we decompose  $\rho$  into steps  $\rho = \gamma_0 \rightarrow \gamma_1 \rightarrow \dots \rightarrow \gamma_m$ , with  $\gamma_0 = \gamma_{start}$  and  $\gamma_m = \gamma_{end}$ . By induction, we construct a sequence of configurations  $\chi_0, \dots, \chi_m$  and runs  $\pi_1, \dots, \pi_m$  such that, for all  $i \in [0, m]$ , we have  $\gamma_i(\nu(a)) = \chi_i(a)$  for all  $a \in \mathbb{A}_X$  and, for all  $j \in [1, m]$ , it holds that  $\pi_j: \chi_{j-1} \xrightarrow{*} \chi_j$ .

We start by setting  $\chi_0 := \chi_{start}$ . By Item (A) of Claim 32, we have  $\gamma_0(\nu(a)) = \chi_0(a)$  for all  $a$ . Now assume we constructed  $\chi_i$  and  $\pi_i$  for all  $i \leq k$  for some  $k < m$ . We construct  $\chi_{k+1}$  and  $\pi_{k+1}$  as follows. Consider the step  $\gamma_k \rightarrow \gamma_{k+1}$ , and let  $\delta = q \xrightarrow{\bowtie q_0} q'$  be the used transition, let  $a^\gamma$  be the observing agent, and let  $a_o^\gamma$  be the observed agent in that step. Let  $\{a_1^X, \dots, a_p^X\} := \nu^{-1}(a^\gamma)$ . Note that  $\mu(\mathbf{dat}(a_j^X)) = \mathbf{dat}(a^\gamma)$  for all those agents.

- If  $\bowtie$  is  $=$ , then  $\mathbf{dat}(a_o^\gamma) = \mathbf{dat}(a^\gamma)$ . Let  $\{d_1, \dots, d_r\} := \mu^{-1}(\mathbf{dat}(a_o^\gamma))$ . By Item (B) of Claim 32, for all  $\ell \in [1, r]$ , since  $a_o^\gamma$  is observed in  $\rho$ , there exists  $a_{o, \ell}^X$  such that  $\nu(a_{o, \ell}^X) = a_o^\gamma$  and  $\mathbf{dat}(a_{o, \ell}^X) = d_\ell$ . By the induction hypothesis, for all  $\ell$ , we have  $\chi_k(a_{o, \ell}^X) = \gamma_k(a_o^\gamma)$ , and for all  $j$ , we have  $\chi_k(a_j^X) = \gamma_k(a^\gamma)$ . Thus, for every  $\ell \in [1, r]$  and sequentially for every  $j \in [1, d]$ , we can execute a step  $\xrightarrow{\bowtie a_{o, \ell}^X} a_j^X$  using transition  $\delta$ . We call this sequence of steps  $\pi_{k+1}$ , and we call the reached configuration  $\chi_{k+1}$ .
- If  $\bowtie$  is  $\neq$ , then  $\mathbf{dat}(a_o^\gamma)$  is externally observed in  $\rho$ , and by Item 2 of Claim 31, there exists some datum  $d$  such that  $\mu(d) = \mathbf{dat}(a_o^\gamma)$ . Furthermore, as  $a_o^\gamma$  is observed in  $\rho$ , by Item (B) of Claim 32, there exists  $a_o^X$  such that  $\nu(a_o^X) = a_o^\gamma$  and  $\mathbf{dat}(a_o^X) = \mu(\mathbf{dat}(a_o^\gamma)) = d$ . Moreover, for all  $j$ , it holds that  $\mu(\mathbf{dat}(a_j^X)) = \mathbf{dat}(a^\gamma) \neq \mathbf{dat}(a_o^\gamma) = \mu(\mathbf{dat}(a_o^X))$ . By the induction hypothesis, we have  $\chi_k(a_o^X) = \gamma_k(a_o^\gamma)$  and, for all  $j$ , it holds that  $\chi_k(a_j^X) = \gamma_k(a^\gamma)$ .

Thus, sequentially for every  $j \in [1, p]$ , we can execute a step  $\xrightarrow{\bowtie a_o^X} a_j^X$  using transition  $\delta$ . We call this sequence of steps  $\pi_{k+1}$ , and we call the reached configuration  $\chi_{k+1}$ .

It is easy to check that  $\gamma_{k+1}(\nu(a)) = \chi_{k+1}(a)$  for all  $a$  in both cases. This concludes our induction. As a result, we obtain a configuration  $\chi_{end} = \chi_m$  and a run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$  such that  $\gamma_i(\nu(a)) = \chi_i(a)$  for all  $a$  appearing in  $\chi_{init}$  and  $i \in [1, m]$ .

All that is left to do is establish that  $\chi_{end} \equiv_{n, M} \gamma_{end}$ . As a direct consequence of Item (C) of Claim 32, it holds that for all  $d$ , we have  $(\gamma_{end}, \mu(d)) \equiv_n (\chi_{end}, d)$ . Using Item 3 of Claim 31, we obtain that  $\gamma_{end} \equiv_{n, M} \chi_{end}$ , proving the lemma.  $\triangleleft$

We obtain the following analogous statement for Post\*.

► **Corollary 33.** *For all  $n, M \in \mathbb{N}$  and all configurations  $\gamma_{start}, \gamma_{end}, \chi_{end} \in \Gamma$ , if there is a run  $\rho: \gamma_{start} \xrightarrow{*} \gamma_{end}$  and  $\gamma_{end} \equiv_{f(n), g(n, M)} \chi_{end}$ , then there is a configuration  $\chi_{start} \in \Gamma$  with  $\chi_{start} \equiv_{n, M} \gamma_{start}$  and a run  $\pi: \chi_{start} \xrightarrow{*} \chi_{end}$ .*

**Proof.** This result is obtained by reversing every transition in the IOPPUD, i. e., replacing every  $q_2 \xrightarrow{\bowtie q_1} q_3$  with  $q_3 \xrightarrow{\bowtie q_1} q_2$ . We can then apply Lemma 22 on the reversed system.  $\triangleleft$

### C.3 Proof of Proposition 21

► **Proposition 21.** *There is a polynomial function  $\text{poly}: \mathbb{N} \rightarrow \mathbb{N}$  such that for all IOPPU  $\mathcal{P}$  and GRE  $E$ , the set  $\llbracket E \rrbracket_{\mathcal{P}}$  is a union of  $\left( \|E\| \cdot (\text{poly}(|\mathcal{P}|))^{|E|}, \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |E|^2} \right)$ -containers.*

**Proof.** In order to improve the readability of this proof, instead of directly using the bounds in terms of  $f(n)$  and  $g(n, M)$ , we use two polynomial functions  $\text{poly}_1, \text{poly}_2$ . Although the degree of the polynomials will be larger than necessary, the bounds will suffice to prove EXPSPACE membership.

We let  $\text{poly}_1, \text{poly}_2: \mathbb{N} \rightarrow \mathbb{N}$  be polynomial functions such that, for all  $n, M \in \mathbb{N}$  with  $n > 0$ , it holds that  $f(n) \leq n \cdot \text{poly}_1(|\mathcal{P}|)$  and  $g(n, M) \leq M \cdot n^{\text{poly}_2(|\mathcal{P}|)}$ . Note that such functions exist by the definition of  $f$  and  $g$ . Let  $\text{poly}: \mathbb{N} \rightarrow \mathbb{N}, n \mapsto \text{poly}_1(n) \cdot \text{poly}_2(n)$ . Moreover, let  $\alpha$  and  $\beta$  be functions with  $\alpha(\mathcal{P}, E, F) := \|E\| \cdot \text{poly}_1(|\mathcal{P}|)^{|F|}$  and  $\beta(\mathcal{P}, E, F) := \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |F|^2}$  for all protocols  $\mathcal{P}$  and all GRE  $E, F$ .

For every sub-expression  $F$  of  $E$ , we prove by induction on  $|F|$  that  $\llbracket F \rrbracket_{\mathcal{P}}$  can be described as a union of  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers.

- If  $F$  is an interval predicate then, by Lemma 29,  $\llbracket F \rrbracket_{\mathcal{P}}$  is a union of  $(\|F\|, \|F\|)$ -containers, and thus, by Lemma 18, it is also a union of  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers.
- Suppose  $F = \overline{G}$  for some GRE  $G$ . By the induction hypothesis,  $\llbracket G \rrbracket_{\mathcal{P}}$  is a union of  $(\alpha(\mathcal{P}, E, G), \beta(\mathcal{P}, E, G))$ -containers, and thus also of  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers by Lemma 18. Since  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers form a partition of the set of configurations, this shows that  $\llbracket E \rrbracket_{\mathcal{P}}$  is also a union of  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers.
- If  $F = G_1 \cup G_2$ , then again, by the induction hypothesis and Lemma 18, both  $\llbracket G_1 \rrbracket_{\mathcal{P}}$  and  $\llbracket G_2 \rrbracket_{\mathcal{P}}$  are unions of  $(\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F))$ -containers, and so is  $\llbracket F \rrbracket_{\mathcal{P}}$ .
- Suppose  $F = \text{Pre}^*(G)$  for some GRE  $G$ . By the induction hypothesis,  $\llbracket G \rrbracket_{\mathcal{P}}$  is a union of  $(\alpha(\mathcal{P}, E, G), \beta(\mathcal{P}, E, G))$ -containers. We show that for all configurations  $\gamma_1, \gamma_2$ , if we have  $\gamma_1 \equiv_{\alpha(\mathcal{P}, E, F), \beta(\mathcal{P}, E, F)} \gamma_2$  and  $\gamma_1 \in \llbracket F \rrbracket_{\mathcal{P}}$ , then  $\gamma_2 \in \llbracket F \rrbracket_{\mathcal{P}}$ . By Lemma 22, it suffices to show that  $\gamma_1 \equiv_{f(\alpha(\mathcal{P}, E, G)), g(\alpha(\mathcal{P}, E, G), \beta(\mathcal{P}, E, G))} \gamma_2$ . In the following, we assume that the constant term of  $\text{poly}_1$  is positive and that  $\text{poly}_1$  does not have any negative coefficients. Hence, for all  $\mathcal{P}, E, G$ , we have  $\alpha(\mathcal{P}, E, G) > 0$ . Then it holds that

$$\begin{aligned}
 f(\alpha(\mathcal{P}, E, G)) &\leq \alpha(\mathcal{P}, E, G) \cdot \text{poly}_1(|\mathcal{P}|) \\
 &= \|E\| \cdot \text{poly}_1(|\mathcal{P}|)^{|G|} \cdot \text{poly}_1(|\mathcal{P}|) \\
 &= \|E\| \cdot \text{poly}_1(|\mathcal{P}|)^{|G|+1} \\
 &= \|E\| \cdot \text{poly}_1(|\mathcal{P}|)^{|F|} \\
 &= \alpha(\mathcal{P}, E, F)
 \end{aligned}$$

and

$$\begin{aligned}
 g(\alpha(\mathcal{P}, E, G), \beta(\mathcal{P}, E, G)) &\leq \beta(\mathcal{P}, E, G) \cdot \alpha(\mathcal{P}, E, G)^{\text{poly}_2(|\mathcal{P}|)} \\
 &= \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |G|^2} \cdot (\|E\| \text{poly}_1(|\mathcal{P}|))^{|G| \cdot \text{poly}_2(|\mathcal{P}|)} \\
 &\leq \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |G|^2} \|E\|^{|G| \cdot \text{poly}(|\mathcal{P}|)} \\
 &\leq \|E\|^{\text{poly}(|\mathcal{P}|) \cdot |F|^2} \\
 &= \beta(\mathcal{P}, E, F).
 \end{aligned}$$

Using Lemma 18, we conclude that  $\gamma_1 \equiv_{f(\alpha(\mathcal{P}, E, G)), g(\alpha(\mathcal{P}, E, G), \beta(\mathcal{P}, E, G))} \gamma_2$ .

- If  $F = \text{Post}^*(G)$  for some GRE  $G$ , we proceed the same way as in the previous case, using Corollary 33 instead of Lemma 22.



All in all, by induction, this proves that  $\llbracket E \rrbracket_{\mathcal{P}}$  is a union of  $(\alpha(\mathcal{P}, E, E), \beta(\mathcal{P}, E, E))$ -containers, which concludes the proof of Proposition 21.  $\blacktriangleleft$

## D Decidability and Upper Complexity Bounds: Proofs for Section 6

► **Proposition 23.** *The following problem is decidable in PSPACE: given a PPUD  $\mathcal{P}$ , a GRE  $E$ , and a configuration  $\gamma$  described data-explicitly, decide if  $\gamma \in \llbracket E \rrbracket_{\mathcal{P}}$ .*

We prove the following auxiliary result:

► **Lemma 34.** *The following problem is decidable in PSPACE: given a PPUD  $\mathcal{P}$  and two configurations  $\gamma_1, \gamma_2$  described data-explicitly, decide if  $\gamma_1 \xrightarrow{*} \gamma_2$ .*

**Proof.** Let  $k$  be the number of data appearing in  $\gamma_1$ , and  $m$  the number of agents in  $\gamma_1$ . All steps preserve the number of data and the number of agents. As a result, the graph of configurations reachable from  $\gamma_1$  contains only configurations with  $k$  data and  $m$  agents; we can look for a path between  $\gamma_1$  and  $\gamma_2$  in this graph in PSPACE.  $\blacktriangleleft$

We now prove Proposition 23.

**Proof of Proposition 23.** By Lemma 34, there exists a polynomial function  $\text{poly}_{reach}$  such that we can check reachability between two configurations represented data-explicitly with  $k$  data and  $m$  agents in deterministic space  $\text{poly}_{reach}(|\mathcal{P}|, k, \log(m))$ . We denote the size of the encoding of a GRE  $E$  by  $|\langle E \rangle|$ .

Let  $\psi = \exists d_1, \dots, d_M, \bigwedge_{q \in Q} \bigwedge_{j=1}^M \#(q, d_j) \in [A_{q,j}, B_{q,j}]$  be a simple interval predicate of height  $n$  and width  $M$ , and let  $\gamma$  a configuration described data-explicitly with  $k$  data and  $m$  agents. One can check whether  $\gamma$  satisfies  $\psi$  by enumerating all injective mappings from  $\{d_i \mid 1 \leq i \leq M\}$  to the data appearing in  $\gamma$ , and checking whether one of them satisfies the interval conditions on the number of agents in each state. This requires polynomial space in  $k + \log(m) + |\langle \psi \rangle|$ .

Given an interval predicate  $\varphi$  and a configuration  $\gamma$ , it is then straightforward to check which simple interval predicates appearing in  $\varphi$  are satisfied by  $\gamma$  and infer the satisfaction of the interval predicate from this, all in polynomial space. Let  $\text{poly}_{sat}$  be a polynomial such that  $\text{poly}_{sat}(k, \log(m), |\langle \psi \rangle|)$  bounds the required space. We define  $\text{poly}(x, y, z, t) := \text{poly}_{sat}(y, z, t) + \text{poly}_{reach}(x, y, z) + xyz$ .

We proceed by induction on  $E$  to show that checking if a configuration  $\gamma$  is in  $\llbracket E \rrbracket_{\mathcal{P}}$  can be done in space  $\mathcal{O}(|E| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle E \rangle|))$ , where  $k$  and  $m$  are the numbers of data and agents appearing in  $\gamma$ , respectively.

- If  $E$  is an interval predicate, then we simply evaluate whether  $\gamma$  satisfies it. It is straightforward to get the truth value of each simple interval predicate over a configuration in space  $\text{poly}_{sat}(k, \log(m), |\langle E \rangle|)$ , and then infer the satisfaction of the interval predicate from it.
- If  $E = \overline{F}$ , then it suffices to check if  $\gamma$  is in  $\llbracket F \rrbracket_{\mathcal{P}}$  and reverse the answer, which requires no extra space. Hence, by the induction hypothesis, we need space at most  $|F| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle E \rangle|) \leq |E| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle E \rangle|)$ .
- If  $E = F_1 \cup F_2$ , then it suffices to check if  $\gamma$  is in  $\llbracket F_1 \rrbracket_{\mathcal{P}}$ , accept if it is the case, and erase the previous computation and check if  $\gamma$  is in  $\llbracket F_2 \rrbracket_{\mathcal{P}}$  otherwise. By the induction hypothesis, the required space is then the maximum of the space required by the two sub-computations, which is at most  $|E| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle E \rangle|)$ .

- If  $E = \text{Post}^*(F)$ , then we use the fact that a step of a run never changes the number of data or the number of agents in a configuration. Hence,  $\gamma \in \llbracket E \rrbracket_{\mathcal{P}}$  if and only if there exists  $\gamma'$  with  $k$  data and  $m$  agents such that  $\gamma' \in \llbracket F \rrbracket_{\mathcal{P}}$  and  $\gamma' \xrightarrow{*} \gamma$ . Thus, we can enumerate all configurations with  $k$  data and  $m$  agents, which requires space  $k \cdot |\mathcal{P}| \cdot \log(m)$ , and check for each one of them whether it is in  $\llbracket F \rrbracket_{\mathcal{P}}$ , which requires space  $|F| \cdot \text{poly}(k, \log(m), |\langle F \rangle|)$  by the induction hypothesis. For every such configuration  $\gamma'$ , we check whether  $\gamma$  is reachable from it, which requires space  $\text{poly}_{\text{reach}}(|\mathcal{P}|, k, \log(m))$ . As a result, we can check whether  $\gamma$  is in  $\llbracket E \rrbracket_{\mathcal{P}}$  in space  $|F| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle F \rangle|) + k \cdot |\mathcal{P}| \cdot \log(m) + \text{poly}_{\text{reach}}(|\mathcal{P}|, k, \log(m)) \leq |E| \cdot \text{poly}(|\mathcal{P}|, k, \log(m), |\langle E \rangle|)$ .
- The case  $E = \text{Pre}^*(F)$  is analogous to the previous one.

Our induction shows that we can check if  $\gamma \in \llbracket E \rrbracket_{\mathcal{P}}$  in deterministic space at most  $|E| \cdot \text{poly}(k, |Q|, \log(m), |\langle E \rangle|)$ . As a result, the problem at hand is in PSPACE. ◀

## E Details for Theorem 24

We proceed by reduction from the problem of tiling an exponentially large grid. Recall the following definitions: A *tiling instance* is a tuple  $(2^n, \mathcal{C}, \mathcal{T})$  where  $n \geq 1$ ,  $\mathcal{C}$  is a finite set of *colours* with special colour *white* and  $\mathcal{T} = \{t_1, \dots, t_m\} \subseteq \mathcal{C}^4$  is a finite set of *tiles*. A tile represents a square whose 4 edges have a colour; we write  $t \in \mathcal{T}$  as  $t =: (\text{top}(t), \text{bottom}(t), \text{left}(t), \text{right}(t))$ . The *size* of the tiling problem is  $n + |\mathcal{C}| + m$  (implicitly,  $2^n$  is encoded in binary). Informally, the tiling problem asks for the existence of a tiling of the  $2^n \times 2^n$  grid, i.e., a mapping  $\tau : [0, 2^n - 1] \times [0, 2^n - 1] \rightarrow \mathcal{T}$  where the colours of neighbouring tiles match and the borders of the grid are white:

► **Definition 35** ([27]). *The following tiling problem is NEXPTIME-complete.*

**Input:** A tiling instance  $(2^n, \mathcal{C}, \mathcal{T})$ ,

**Question:** Does there exist a *tiling*, i.e., a mapping  $\tau : [0, 2^n - 1] \times [0, 2^n - 1] \rightarrow \mathcal{T}$  such that:

1. for all  $i \in [0, 2^n - 1]$ ,  $\text{left}(\tau(i, 0)) = \text{right}(\tau(i, 2^n - 1)) = \text{white}$ ,
2. for all  $i \in [0, 2^n - 1]$ ,  $j \in [0, 2^n - 2]$ ,  $\text{right}(\tau(i, j)) = \text{left}(\tau(i, j + 1))$ ,
3. for all  $j \in [0, 2^n - 1]$ ,  $\text{top}(\tau(0, j)) = \text{bottom}(\tau(2^n - 1, j)) = \text{white}$ ,
4. for all  $i \in [0, 2^n - 2]$ ,  $j \in [0, 2^n - 1]$ ,  $\text{bottom}(\tau(i, j)) = \text{top}(\tau(i + 1, j))$ ?

This problem is NEXPTIME-complete [27]. We will therefore provide a polynomial-time reduction from this problem to the negation of the emptiness problem for GRE.

We use the following notations for interval predicates:

- given a state  $q$ ,  $\text{Pres}(q) := \exists d, \#(q, d) > 0$  is the simple interval predicate indicating that  $q$  is populated,
- given a set of states  $S \subseteq Q$ ,  $\text{Abs}(S) := \bigwedge_{q \in S} (\forall d, \#(q, d) = 0) = \bigwedge_{q \in S} \neg(\exists d, \#(q, d) > 0)$  is the interval predicate indicating that all states in  $S$  are empty,
- given a state  $q$ ,  $\text{Mandatory}(q) := \bigwedge_{p \in Q} \forall d, (\#(p, d) = 0 \vee \#(q, d) > 0)$  is the interval predicate indicating that all data appearing in the configuration has some agent on  $q$ .

### E.1 Encoding the Tiling and Checking for Duplicates

Figure 5 displays how one encodes the tiling into a configuration and how one checks for duplicates. A *duplicate* designates that two values encode the same tile coordinates. Figure 5a corresponds to where  $\tau$  is encoded. We only want  $\llbracket \mathcal{E} \rrbracket$  to contain configurations where no data type has two agents in the tile type part. To do so, we make it so that, if this condition

is violated, one will be able to cover  $q_{\perp}$ . For example, from  $h_1(0)$ , there is a transition to  $q_{\perp}$  labelled “ $h_1(1), =$ ”, so that a configuration where a datum has agents on both  $h_1(0)$  and  $h_1(1)$  will be in  $\text{Pre}^*(\text{Pres}(q_{\perp}))$ . Similarly, we ensure that a datum does not have agents on states  $\mathfrak{t}_i$  and  $\mathfrak{t}_j$  for  $i \neq j$ .

The subprotocol of Figure 5b allows us to enforce that no two data encode the same tile. We will impose that every datum has at least one agent on  $D_a$  and one on  $D_b$ .  $q_{\geq 3}$  can be covered if at least three agents go to the blue part and not to the sink state; therefore, in order to reach a configuration covering  $q_{\text{dup}}$  but from which one cannot cover  $q_{\geq 3}$ , one needs to send at most two agents to the blue part. Because the left and the right track must observe each other with disequality tests, this is in fact only possible with agents of different data, one in the left track and one in the right track. But then, the two agents can make it all the way to  $q_{\text{dup}}$  if and only if their data have an agent in common on every bit of horizontal and vertical coordinate which, assuming that they have only one agent per bit, means that the two data encode the same coordinates. Overall, the predicate that we add to  $\mathcal{E}$  for the tiling encoding is  $\mathcal{E}_{\text{encoding}}$ , defined as follows:

$$\begin{aligned} \mathcal{E}_{\text{encoding}} := & \overline{\text{Pre}^*(\text{Pres}(q_{\perp}))} \cap \text{Mandatory}(D_a) \cap \text{Mandatory}(D_b) \\ & \cap \overline{\text{Abs}(Q_{\text{blue}} \cup Q_{\text{red}})} \cap \overline{\text{Pre}^*(\text{Pres}(q_{\text{dup}}) \cap \overline{\text{Pre}^*(\text{Pres}(q_{\geq 3}))})} \end{aligned}$$

where  $Q_{\text{blue}}$  and  $Q_{\text{red}}$  are the states in the blue and red parts, respectively.

For a configuration  $\gamma$ , let  $\tau(\gamma) : [0, 2^n - 1]^2 \rightarrow 2^{\mathcal{T}}$  be the function that, to  $x, y \in [0, 2^n - 1]$ , maps the set of every  $t_i \in \mathcal{T}$  such that there is a datum  $d$  in  $\gamma$  with an agent on  $\mathfrak{t}_i$  and:

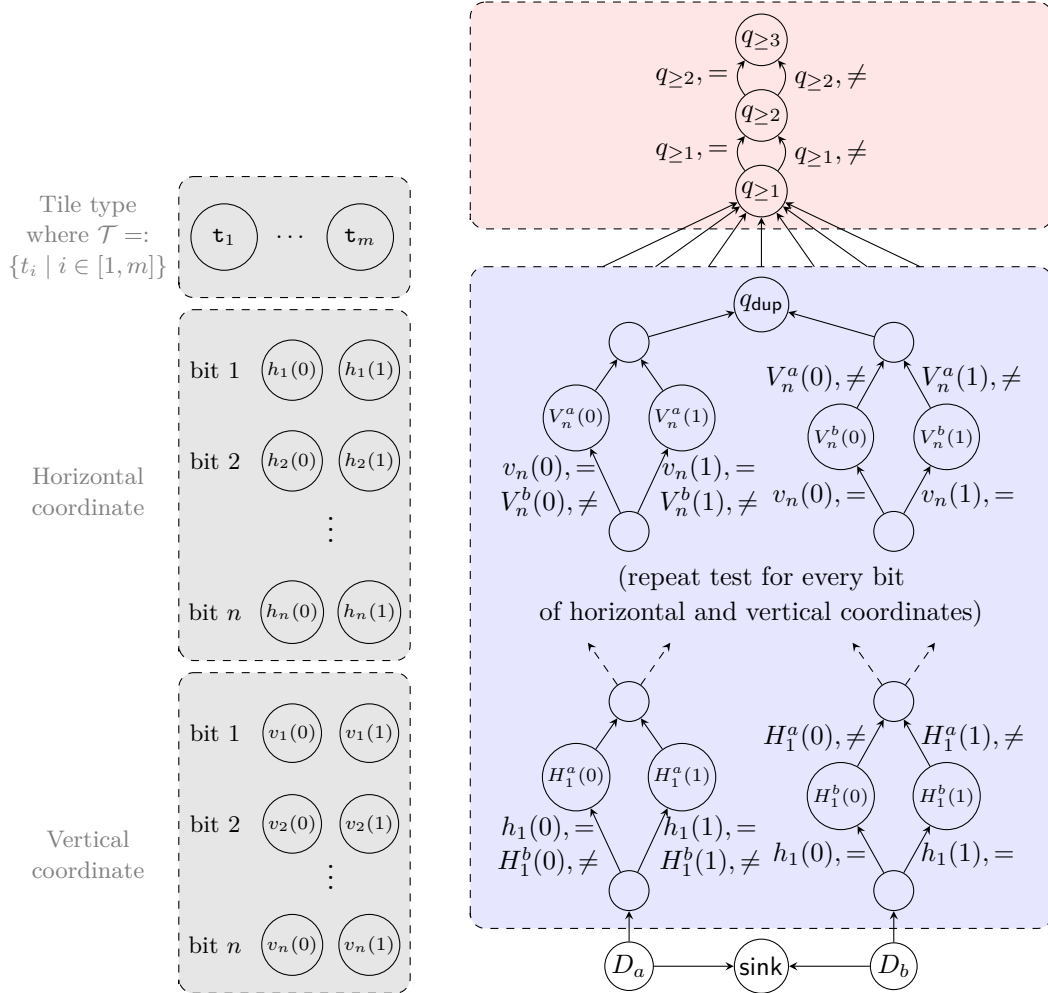
- if the  $i$ -th bit of  $x$  is  $b$  then  $d$  has an agent in  $h_i(b)$  in  $\gamma$ ,
- if the  $i$ -th bit of  $y$  is  $b$  then  $d$  has an agent in  $v_i(b)$  in  $\gamma$ .

► **Lemma 36.** *For every  $\gamma \in \llbracket \mathcal{E}_{\text{encoding}} \rrbracket$ , for every  $i, j \in [0, 2^n - 1]$ ,  $|\tau(\gamma)(i, j)| \leq 1$ .*

**Proof.** Let  $\gamma \in \llbracket \mathcal{E}_{\text{encoding}} \rrbracket$ . Assume by contradiction that we have  $i, j$  such that  $\{t_i, t_j\} \subseteq \tau(\gamma)(i, j)$  with  $i \neq j$ . Let  $d$  (resp.  $d'$ ) the data witnessing that  $t \in \tau(\gamma)(i, j)$ . If  $d = d'$  then there is an agent in  $\mathfrak{t}_i$  and one in  $\mathfrak{t}_j$  with same datum, but then one can cover  $q_{\perp}$  from  $\gamma$  which contradicts that  $\gamma \in \llbracket \overline{\text{Pre}^*(\text{Pres}(q_{\perp}))} \rrbracket$ . If  $d \neq d'$ , because  $\gamma \in \llbracket \text{Mandatory}(D_a) \cap \text{Mandatory}(D_b) \rrbracket$ ,  $d$  has an agent  $a$  in  $D_a$  and  $d'$  has an agent  $a'$  in  $D_b$ . By making  $a$  and  $a'$  evolve in the blue part, taking the values of the bits of  $x$  and then  $y$  and observing each other, they can cover  $q_{\text{dup}}$ ; by sending all other agents in  $D_a$  and  $D_b$  to sink, we reach a configuration  $\gamma' \in \llbracket \text{Pres}(q_{\text{dup}}) \cap \overline{\text{Pre}^*(\text{Pres}(q_{\geq 3}))} \rrbracket$ , a contradiction. ◀

## E.2 Verifying that the Encoded Tiling is Valid


We here describe the horizontal verifier, i. e., the part of  $\mathcal{P}$  responsible for verifying left-right frontiers of the tiling. The data of agents in this part are irrelevant. Two variables,  $x, y \in [0, 2^n - 1]$ , are encoded in binary in similar fashion to Figure 5a. These two variables will be private to the horizontal verifier. There will be exactly one agent encoding each bit; the novelty with respect to Figure 5a is that these agents will be able to move between values 0 and 1, so that the values of  $x$  and  $y$  can be changed. The rest of the verifier will contain a single agent, called the *main agent*. Under the guarantee that there is only one agent per bit, the main agent can directly read and modify the values of  $x$  and  $y$ . For example, to change bit  $i$  of  $x$  from 0 to 1, the main agent goes to state  $x_i(0 \rightarrow 1)$ , the agent of the bit observes it and moves to value 1, the main agent observes this change and continues. We represent the part of the protocol encoding the behaviour of the main agent in Figure 6



**(a)** The part of  $\mathcal{P}$  encoding the value of  $\tau$ . Not depicted are the following transition to  $q_{\perp}$ . For every  $i \neq j \in [1, m]$ ,  $\tau_i$  has a transition to  $q_{\perp}$  labelled  $\tau_j, =$ . For every  $i \in [1, n]$  and  $b \in \{0, 1\}$ ,  $h_i(b)$  has a transition to  $q_{\perp}$  labelled  $h_i(1-b), =$ ; and similarly for the vertical counter.

**(b)** The test that no tile is encoded by two data. Transitions labelled by two observations are a shortcut for two chained transitions, one with each observation. If only two agents with distinct data come in the blue part (one from  $D_a$  and one from  $D_b$ ), they will be able to cover  $q_{dup}$  if and only if their data encode the same tile. However, if at least three agents are in the blue part, then  $q_{\geq 3}$  can be covered.

■ **Figure 5** The part of  $\mathcal{P}$  encoding the tiling and checking for duplicates.



■ **Figure 6** The part of the protocol encoding the horizontal verifier. The parts encoding the values of  $x$  and  $y$  are not represented. Variables  $T$  and  $\text{colour}$  are encoded directly in the state space. The operation “ $T \leftarrow T_i$  where  $T_i^{(h)}$  populated” corresponds to  $m = |\mathcal{T}|$  transitions in parallel, one for each  $t_i \in \mathcal{T}$ ; every such transition can be taken if an agent on state  $T_i^{(h)}$  is observed. Not depicted is a gadget that allows to go to  $q_\perp$  if at least two agents play the role of horizontal verifier or two agents play the same bit of  $x$  or of  $y$ .

where, for simplicity, we abstract bit encoding of  $x$  and  $y$  into direct access and modification of the values of  $x$  and  $y$  by the main agent. The main agent will also manipulate a variable  $\text{colour} \in \mathcal{C}$ , which can be directly encoded in the state space.

### E.3 Synchronisation with the Verifier


Finally, in the gadget represented in Figure 7, each datum will be represented by exactly one agent, called the *reader*, which will initially be in  $S^{(h)}$ . The reader synchronises with the horizontal verifier in order to let the main agent know what tile type its datum encodes. The goal of synchronisation is to ensure, in relevant runs, that an agent in the middle part of Figure 7 has a datum value equal to the current value of  $z := x + 2^n y$ . The reader wants to test equality of its own number in  $[0, 2^{2n} - 1]$  with  $z$ , this number being equal to its horizontal coordinate plus  $2^n$  times the vertical coordinate. However, because the value of  $z$  changes throughout the run, it would not suffice that the reader simply tests equality once per bit. However, as we will now see, it suffices that the first synchronisation test tests for equality all bits from most to least significant, and that the second synchronisation test tests for equality all bits from least to most significant.

► **Lemma 37.** *Consider a run where an agent  $a$  starts on  $S^{(h)}$  and ends on  $F^{(h)}$ ; assume that its datum faithfully encodes value  $k \in [0, 2^N - 1]$ . Then:*

- *the value of  $z$  after the first synchronisation test, which tests bits from most to least significant, is such that  $z \geq k$ ,*
- *if we have  $z > k$  at the beginning of the second synchronisation test, which tests bits from least to most significant, then the agent may not cover  $S_f^{(h)}$ .*

**Proof.** Let  $N := 2n$ . In the following, a binary number  $\ell \in [0, 2^N - 1]$  is denoted  $\overline{\ell_1 \ell_2 \dots \ell_N}^2$  where  $\ell_i \in \{0, 1\}$  is the  $i$ th most significant bit.

We first prove the first claim. Let  $z^{(m)}$  the value of  $z$  after the  $m$ -th most significant bit is tested in the first synchronisation test. We prove by induction that  $z^{(m)} \geq \overline{k_1 \dots k_m 0 \dots 0}^2$ . It is trivially true for  $m = 0$ . Assume that  $z^{(m)} \geq \overline{k_1 \dots k_m 0 \dots 0}^2$ : if we have  $k_{m+1} = 0$  then



■ **Figure 7** The synchronisation gadget.

we have the property for  $m + 1$ . If  $k_{m+1} = 1$ , then the first value greater than  $z^{(m)}$  whose  $(m + 1)$ -th bit is at least 1 is greater than  $\overline{k_1 \dots k_m k_{m+1} 0 \dots 0}^2$ , concluding the induction.

We now prove the second claim. Let  $z^{(m)}$  now denote the value of  $z$  after the test of bit  $N - m + 1$  in the second synchronisation test; we use convention that  $z^{(0)}$  is the value right before the second synchronisation test, so that by hypothesis  $z^{(0)} > k$ . We prove by induction that there is a bit  $j \leq N - m$  such that bit  $j$  has value 1 in  $z^{(m)}$  and 0 in  $k$  whereas all bits  $i < j$  are equal for  $z^{(m)}$  and  $k$ . This is true for  $m = 0$  by hypothesis that  $z^{(0)} > k$ . Assume that it is true for  $m$ ; if  $j < N - m - 1$  then it is true for  $m + 1$ . Otherwise, we have  $j = N - m$ , but the test of bit  $j$  needs  $z^{(m)}$  to have value 0 at bit  $j$ ; this means that there is a carry between  $z^{(m)}$  and  $z^{(m+1)}$  that propagates to bit  $j - 1$  (and possibly further), proving the result. For  $m = N$ , this yields a contradiction. ◀

## E.4 Summary of the Reduction

The protocol built is a combination of all the parts built previously. In Appendix E.1, we have defined a GRE  $\mathcal{E}_{\text{encoding}}$  that, due to Lemma 36, guarantees that, in considered configurations, a given coordinate will have at most one tile type. We now build a larger GRE that expresses that the horizontal and vertical verifiers are able to verifying the entire grid without anyone cheating. Overall, the constructed GRE  $\mathcal{E}$  is the following:

$$\mathcal{E} := \mathcal{E}_{\text{encoding}} \cap \overline{\text{Pre}^*(\text{Pres}(q_{\perp}))} \cap \mathcal{E}_{\text{init}}^{(h)} \cap \mathcal{E}_{\text{init}}^{(v)} \cap \text{Pre}^*(\mathcal{E}_{\text{final}}^{(h)}) \cap \text{Pre}^*(\mathcal{E}_{\text{final}}^{(v)})$$

where:

- $\mathcal{E}_{\text{init}}^{(h)}$  expresses that all states in the horizontal synchronisation gadget except  $S^{(h)}$  are empty and that all states in the horizontal verifier gadget are empty except  $q_i^{(h)}$  and the states of bits of  $x$  and  $y$  of value 0,
- $\mathcal{E}_{\text{init}}^{(v)}$  is the counterpart of  $\mathcal{E}_{\text{init}}^{(h)}$  but for the vertical verifier and vertical synchronisation,
- $\mathcal{E}_{\text{final}}^{(h)}$  expresses that  $q_f^{(h)}$  is not empty and that all states in the horizontal synchronisation gadget are empty except  $F^{(h)}$ ,
- $\mathcal{E}_{\text{final}}^{(v)}$  is the counterpart of  $\mathcal{E}_{\text{final}}^{(h)}$  but for the vertical verifier and vertical synchronisation.

► **Lemma 38.** *If  $\llbracket \mathcal{E} \rrbracket$  is not empty, then the instance of the tiling problem is positive.*

**Proof.** Let  $\gamma \in \llbracket \mathcal{E} \rrbracket$ . Due to Lemma 36, we have that  $|\tau(\gamma)(x, y)| \leq 1$  for every  $x$  and  $y$ ; let  $\tau : [0, 2^n - 1]^2 \rightarrow \mathcal{T} \cup \{*\}$  be the function that to  $(x, y)$  maps the only element in  $\tau(\gamma)(x, y)$  if it exists.



We will prove that  $\tau$  is never equal to  $*$  and that it is in fact a tiling. We will prove that the horizontal colours match, i. e., that  $\tau$  satisfies Item 3 and Item 4 of Definition 35; the vertical case is very similar.

Let  $\gamma'$  such that  $\gamma \xrightarrow{*} \gamma'$  and  $\gamma' \in \llbracket \mathcal{E}_{\text{final}}^{(h)} \rrbracket$ . By definition of  $\mathcal{E}_{\text{final}}^{(h)}$ , there is in  $\gamma'$  one agent  $a$  in  $q_f^{(h)}$ . In fact,  $a$  is the only agent in the main part of the horizontal verifier in both  $\gamma$  and  $\gamma'$ , as otherwise  $\gamma$  would be in  $\llbracket \text{Pre}^*(\text{Pres}(q_{\perp})) \rrbracket$ . We know that  $a$  is in state  $q_i^{(h)}$  in  $\gamma$  as all other states are forbidden by  $\mathcal{E}_{\text{init}}^{(h)}$ . There is at least one agent encoding each bit of  $x$  and  $y$ , otherwise  $a$  would not be able to go from  $q_i^{(h)}$  to  $q_f^{(h)}$ ; there is also at most one agent per bit, as otherwise one would be able to cover  $q_{\perp}$  from  $\gamma$ . Therefore, one can consider that  $x$  and  $y$  are encoded reliably; to enforce that, initially,  $x = y = 0$ , we can either encode it into  $\mathcal{E}_{\text{init}}^{(h)}$  or make the main agent perform an initial sequence of transitions setting the variables to the right values.

In the run from  $\gamma$  to  $\gamma'$ ,  $z := x + 2^n y$  has to go incrementally from 0 to  $2^{2n} - 1$ . Moreover, whenever the transition updated the value of  $T$  is taken, we claim that the value  $T_i$  taken is such that  $\tau(x, y) = t_i$  (with  $x$  and  $y$  the values when the transition is taken). Indeed, let  $x, y \in [0, 2^n - 1]$  and assume that some agent  $a_{\text{syn}}$  is observed on  $T_i^{(h)}$  by  $a$  with these values of  $x$  and  $y$ . Because of  $\mathcal{E}_{\text{init}}^{(h)}$  and  $\mathcal{E}_{\text{final}}^{(h)}$ , we know that  $a_{\text{syn}}$  is in  $S^{(h)}$  in  $\gamma$  and in  $F^{(h)}$  in  $\gamma'$ . Let  $z := x + 2^n y$ . Let  $d$  denote the datum of  $a_{\text{syn}}$  in  $\gamma$ . Because one cannot cover  $q_{\perp}$  from  $\gamma$ ,  $d$  has at most one agent in every bit of the horizontal and vertical coordinates. Moreover, since  $a_{\text{syn}}$  manages to pass the synchronisation tests, there is exactly agent with datum  $d$  one per bit. Let  $x_d$  and  $y_d$  the values encoded by agents of  $d$ ; let  $z_d := x_d + y_d$ . Because  $a_{\text{syn}}$  passes the two synchronisation test, by Lemma 37 we must have  $z = z_d$  hence  $\tau(x, y) = t_i$ .

We have proven that, when the verifier takes the transition observing state  $T_i^{(h)}$ , the value observed indeed corresponds to  $\tau(x, y) \in \mathcal{T}$ . We conclude that, because the main agent of the horizontal verifier ends on  $q_f^{(h)}$ ,  $\tau$  satisfies Item 3 and Item 4 of Definition 35. We can similarly prove that  $\tau$  also satisfies Item 1 and Item 2, proving that  $\tau$  is a witness that our instance of the tiling problem is positive.  $\blacktriangleleft$

► **Lemma 39.** *If the instance of the tiling problem is positive, then  $\llbracket \mathcal{E} \rrbracket \neq \emptyset$ .*

**Proof.** Let  $\tau$  be a witness that the instance is positive. Let  $\gamma$  be a configuration whose appearing data are in the set  $\{d_{i,j} \mid i, j \in [0, 2^n - 1]\}$  and, for every  $i, j \in [0, 2^n - 1]$ :

- in the states of Figure 5a,  $d_{i,j}$  has exactly  $2n + 1$  agents:  $n$  agents encoding the value of  $i$  in the “horizontal coordinate” part,  $n$  agents encoding the value of  $j$  in the “vertical coordinate” part and one agent in the “tile type” part, on the state corresponding to  $\tau(i, j)$ ;
- in the states of Figure 5b,  $d_{i,j}$  has two agents: one on  $D_a$  and one on  $D_b$ ;
- $d_{i,j}$  has one agent on  $S^{(h)}$  and one agent on  $S^{(v)}$ , and none in the rest of the horizontal and vertical synchronisation gadgets;
- $d_{i,j}$  has agents in the horizontal and vertical verifier if and only if  $i = j = 0$ ;  $d_{0,0}$  has one agent on  $q_i^{(h)}$ , one agent on  $q_i^{(v)}$  and  $4n$  agents on bits of value 0 of the four variables of the verifiers (variables  $x$  and  $y$  of the horizontal verifier and variables  $x$  and  $y$  of the vertical verifier, which are distinct although, for ease of notation, we never distinguished them above).

Because  $\gamma$  does not have two data with the same coordinates, one cannot cover  $q_{\text{dup}}$  from  $\gamma$  without putting at least three agents in the blue part, which would in turn allow reaching  $q_{\geq 3}$ . This proves that  $\gamma \in \text{Pre}^*(\text{Pres}(q_{\text{dup}}) \cap \text{Pre}^*(\text{Pres}(q_{\geq 3})))$ . It is quite easy to prove that

$\gamma \in \mathcal{E}_{\text{encoding}}$  and also that  $\gamma \in \overline{\text{Pre}^*(\text{Pres}(q_{\perp}))} \cap \mathcal{E}_{\text{init}}^{(h)} \cap \mathcal{E}_{\text{init}}^{(v)}$ . It finally remains to show that  $\gamma \in \text{Pre}^*(\mathcal{E}_{\text{final}}^{(h)} \cap \mathcal{E}_{\text{final}}^{(v)})$ . To do so, we consider the execution where the horizontal and vertical verifiers go through all tiles one by one, synchronising with the right agents every time; this execution exists because  $\tau$  satisfies all four conditions from Definition 35. Overall, we have  $\gamma \in \llbracket \mathcal{E} \rrbracket$  which proves  $\llbracket \mathcal{E} \rrbracket \neq \emptyset$ . ◀

With Lemmas 38 and 39, we have built a polynomial-time reduction from the tiling problem to the negation of the emptiness problem for GRE, which shows that the emptiness problem is CONEXPTIME-hard.

## E.5 Proof of Claim 25

We formalise Claim 25 into the following proposition:

► **Proposition 40.** *For every  $n \in \mathbb{N}$ , there is a protocol  $\mathcal{P}_n$  of size polynomial in  $n$  that is not well-specified but in which all fair runs from configurations with less than  $2^n$  data stabilise to 0.*

**Proof.** For every  $n$ , we build a protocol  $\mathcal{P}_n$  with states  $q_{\perp}$  and  $q_f$  such that:

- from every initial configuration,  $\mathcal{P}_n$  has a run stabilising to  $\perp$ ;
- $Q \setminus \{q_f\}$  has output  $\perp$ ,  $q_f$  has output  $\top$ ,
- an agent on  $q_{\perp}$  attracts agents on any other state,
- an agent on  $q_f$  attracts agents on any other state except  $q_{\perp}$ .

Here, we say that an agent on state  $q$  *attracts* agents in state  $q'$  when there is a transition from  $q'$  to  $q$  that observes the agent on  $q$ . This implies that, in a fair execution, if  $q$  is populated infinitely often then eventually all agents move from  $q'$  to  $q$ . The protocol  $\mathcal{P}_n$  is a simplified version of the protocol from Theorem 24. There are  $2n$  *encoding states*, allowing to encode in binary a number in  $[0, 2^n - 1]$ . Each datum must have at most one agent per bit, otherwise by fairness  $q_{\perp}$  is eventually covered. In a separate gadget, a verifier is able to cover  $q_f$  only after checking that, for every value in  $[0, 2^n - 1]$ , there is a datum encoding this value. This uses the same two ideas as in Theorem 24. First, the verifier manipulates a variable  $x \in [0, 2^n - 1]$  implemented using another gadget and  $n$  agents. Second, there is a synchronisation gadget that makes sure that the verifier only interacts with a datum encoding the current value of  $x$ . For this second point, there is a non-guarded transition from  $q$  to  $q_{\perp}$  for every state  $q$  of the synchronisation gadget except its last state. This guarantees that, eventually, all agents in the synchronisation gadget have successfully passed both synchronisation tests. Again, if two agents play the same role,  $q_{\perp}$  is eventually covered. In order to stabilise to  $\top$ , one needs one datum for each value in  $[0, 2^n - 1]$ , hence  $2^n$  data. ◀