

1 Population Protocols over Ordered Agents

2 Anonymous author

3 Anonymous affiliation

4 Anonymous author

5 Anonymous affiliation

6 Anonymous author

7 Anonymous affiliation

8 Anonymous author

9 Anonymous affiliation

10 Anonymous author

11 Anonymous affiliation

12 Anonymous author

13 Anonymous affiliation

14 — Abstract —

15 Population protocols are a distributed computation model in which a collection of anonymous,
16 finite-state agents interact in randomly chosen pairs and update their states according to a fixed
17 transition function. The computation is defined by the eventual stabilization of the population to a
18 consensus that represents the output. In practice, it is natural to allow each agent to carry a unique
19 identifier and compare it with that of another agent before interacting. We model this extension by
20 having agents be totally ordered and interactions between two agents to be fireable only if their pair
21 of identifiers falls in some condition set. For instance, $\text{PP}[\lt]$ allows for two agents to interact only if
22 the first one appears before the second one.

23 We study population protocols over ordered agents $\text{PP}[\mathcal{N}]$ where \mathcal{N} is a set of predicates available
24 to restrict transition firing. We also study $\text{IO-PP}[\mathcal{N}]$, the *immediate observation* fragment of $\text{PP}[\mathcal{N}]$
25 where only one agent changes state per interaction. Our main result is that $\text{IO-PP}[\lt]$ recognizes
26 exactly the unambiguous star-free languages, which admits many other characterizations, such as
27 two-variable first-order logic or two-way deterministic partially-ordered automata. We also provide
28 a logic and an automaton model that fits in $\text{PP}[\lt]$. We further show that if the successor predicate
29 appears in a set \mathcal{N} of $\text{NSPACE}(n)$ -computable predicates, then $\text{IO-PP}[\mathcal{N}] = \text{PP}[\mathcal{N}] = \text{NSPACE}(n)$.
30 Finally, we investigate the problem of deciding whether a given population protocol always stabilizes
31 to a consensus. While this problem is decidable for unordered population protocols, we show that
32 this is undecidable already for $\text{PP}[\lt]$ and $\text{IO-PP}[+1]$, but conditionally decidable for $\text{IO-PP}[\lt]$.

33 **2012 ACM Subject Classification** Theory of computation \rightarrow Formal languages and automata theory;
34 Theory of computation \rightarrow Distributed computing models; Theory of computation \rightarrow Logic and
35 verification

36 **Keywords and phrases** Population protocols, first-order logic, partially-ordered automata, unam-
37 biguous star-free languages

38 **Digital Object Identifier** 10.4230/LIPIcs.ICALP.2026.



XX:0 Population Protocols over Ordered Agents

39 To facilitate the reviewing process, we provide a table of contents below. We further provide
40 hyperlinks for navigating between statements in the main text and proofs in the appendix.
41 To go to the proof of a statement, click “ \downarrow ” in the left margin, and then click “ \uparrow ” to go
42 back to the statement in the main text.

43 Contents

44	1 Introduction	2
45	2 Preliminaries	4
46	3 Semantic restrictions of population protocols: a toolbox	6
47	3.1 Semi-deciders	7
48	3.2 Protocols with stabilizing inputs	7
49	4 Expressiveness of IO-PP[$<$]	9
50	4.1 IO-PP[$<$] \subseteq DA	9
51	4.2 DA \subseteq IO-PP[$<$]	10
52	5 Expressiveness of PP[$<$]	10
53	5.1 First-order logic over word intervals	11
54	5.2 Partially-ordered Parikh automata	12
55	6 Expressiveness of IO-PP[\mathcal{N}] and PP[\mathcal{N}] when successor is available	13
56	6.1 IO-PP[+1] = PP[+1]	13
57	6.2 NSPACE(n) \subseteq PP[+1]	14
58	7 Decidability of checking whether a population protocol is a decider	14
59	7.1 The syntax of PP[$<$], IO-PP[+1] and PP[+1] deciders are undecidable	14
60	7.2 The syntax of IO-PP[$<$] deciders is decidable, conditionally	15
61	8 Open questions	16
62	A Appendix	19
63	A.1 Example of a PP[$<$] for the median language	19
64	A.2 Missing proofs from Section 3.1	20
65	A.3 Missing proofs from Section 3.2	21
66	A.4 Missing proofs from Section 4.1	23
67	A.5 Missing proofs from Section 4.2	25
68	A.6 Missing proofs from Section 5.1	25
69	A.7 Missing proofs from Section 5.2	26
70	A.8 Missing proofs from Section 6.1	28
71	A.9 Missing proofs from Section 6.2	29
72	A.10 Missing proofs from Section 7.1	35
73	A.11 Missing proofs from Section 7.2	36
74	B Supplementary constructions and explorations	39
75	B.1 Alternative construction for $\Delta_2[<] \subseteq$ IO-PP[$<$]	39
76	B.2 Alternative construction for $\Delta_1^{\text{int}}[<, +, \equiv] \subseteq$ PP[$<$]	40
77	B.3 Alternative proof of $\Delta_2[<] \subseteq \Delta_1^{\text{int}}[<, +, \equiv]$	42

78	B.4 Partially-ordered two-way Parikh automata	42
79	B.5 Alternative proof for $\text{NSPACE}(n) \subseteq \text{IO-PP}[+1]$	46

1 Introduction

Population protocols form a well-established model of distributed computing where anonymous agents, with very limited individual computational power, work collectively to achieve a common task [2]. In this model, an input is scattered among agents that interact pairwise and must take a decision by reaching a consensus that is stable, that is, agents must eventually all agree on the output (“consensus”) and stop changing their mind (“stable”). Population protocols provide a theoretical framework for reasoning about a wide range of distributed systems, including networks of mobile sensors, chemical reaction networks, and social networks. [4, 9, 10].

To familiarize the reader with population protocols, we present a classical protocol for the task of majority voting. A population consists of n agents (who are not aware of n), each carrying a state from a finite set. Here, agents start with either a or b as their state. The population aims at collectively determining whether there are initially more a 's than b 's. At each discrete moment, a pair of agents is chosen arbitrarily and their respective state, from $Q = \{a, b, \underline{a}, \underline{b}\}$, is updated according to these rules:

<i>active to passive</i>	<i>propagation of winning side</i>	<i>tiebreaker</i>
$a, b \rightarrow \underline{a}, \underline{b}$	$a, \underline{b} \rightarrow a, \underline{a}$ $b, \underline{a} \rightarrow b, \underline{b}$	$\underline{a}, \underline{b} \rightarrow \underline{b}, \underline{b}$

Since agents are unordered, each rule $p, q \rightarrow p', q'$ also stands for $q, p \rightarrow q', p'$. Here are three possible executions of the protocol starting from three different “inputs,” that is, assignments of an initial state a or b to each agent:

$aaabb \rightarrow aa\underline{abb} \rightarrow \underline{aaabb} \rightarrow \underline{aaaab} \rightarrow \underline{aaaaa},$
 $aabbb \rightarrow a\underline{abbb} \rightarrow \underline{aabbb} \rightarrow \underline{abbbb} \rightarrow \underline{bbbbb},$
 $aabb \rightarrow \underline{aabb} \rightarrow \underline{aabb} \rightarrow \underline{abbb} \rightarrow \underline{bbbb}.$

Assuming fair scheduling (e.g., choosing agents uniformly at random), one can show that the population stabilizes (almost surely) to the correct outcome: if there is a majority of a 's initially, then agents eventually remain in $\{a, \underline{a}\}$, otherwise they eventually remain in $\{b, \underline{b}\}$.

As agents and rules are unordered, each configuration can be seen as a multiset $\mathbf{c}: Q \rightarrow \mathbb{N}$ where $\mathbf{c}(q)$ indicates the number of agents in state q . It is known that population protocols compute precisely the subsets of \mathbb{N}^Q that are semilinear [3], or, equivalently, that are definable in Presburger arithmetic (the first-order theory of the naturals with order and addition). In particular, majority voting amounts to computing the predicate $\varphi(\mathbf{c}) = \mathbf{c}(a) > \mathbf{c}(b)$.

From a modeling perspective, the fact that agents are unordered is meant to correspond to a situation in which agents are replicated and anonymous entities, and hence have no identifiers and are indistinguishable. Yet, it is natural to allow replicated agents to be totally ordered, e.g., they could be devices with a unique identifier, such as a serial number, stored in read-only memory. In this context, interactions may depend on the relationships between these identifiers.

From the perspective of automata theory, this corresponds to considering population protocols where configurations are *words* rather than multisets. Standard population protocols can be seen as computing commutative properties of words, such as $|w|_a > |w|_b$, while the word setting additionally allows for noncommutative properties, such as “the middle agent has an a ” or “agents strictly alternate between a and b .”

121 Contribution

122 Motivated by the above, we propose to study population protocols with totally-ordered
 123 agents. The class of population protocols so defined, $\text{PP}[\mathcal{N}]$, is parameterized by a set \mathcal{N}
 124 of predicates over positions that can be used to restrict transition firings. Central to our
 125 study is the class $\text{PP}[\prec]$, in which a transition $p, q \xrightarrow{\cdot} r, s$ can only be applied to two agents
 126 in respective states p and q if the first agent appears *before* the second agent. We consider
 127 a well-studied restriction called *immediate observation* [3] where an interaction can only
 128 update a single agent, called the “observer.”

129 Our main result, Theorem 10, establishes that $\text{IO-PP}[\prec]$ has the same expressive power
 130 as unambiguous star-free languages, an important subclass of regular languages that admits
 131 a trove of characterizations (see [32] for a lovely survey on the pervasiveness of this class in
 132 automata theory). Hence, $\text{IO-PP}[\prec]$ is the class of languages captured by these formalisms
 133 over finite words:

- 134 ■ partially-ordered unambiguous automata [24];
- 135 ■ partially-ordered two-way deterministic automata [30];
- 136 ■ $\text{LTL}[\mathbf{F}^{-1}, \mathbf{F}]$: linear temporal logic with past and future operators [14];
- 137 ■ $\text{FO}^2[\prec]$: the two-variable fragment of first-order logic with order [33];
- 138 ■ $\Delta_2[\prec]$: the intersection of the $\exists^*\forall^*$ and $\forall^*\exists^*$ fragments of first-order logic with order [26];
- 139 ■ languages recognized by finite monoids from the variety DA [29].

140 This provides the first characterization of this class in terms of distributed computing,
 141 rather than automata, logic, or algebra.

142 In addition, we explore systematically the classes induced by our definitions:

- 143 ■ In Section 3, we provide a toolbox to study population protocols over ordered agents
 144 $\text{PP}[\mathcal{N}]$, regardless of \mathcal{N} . We study protocols that need to stabilize only if they reach
 145 a positive consensus, which we call *semi-deciders*, as opposed to protocols that always
 146 stabilize to a consensus, dubbed *deciders*. We also refine the technology of protocols with
 147 *stabilizing inputs*, studied in [28], which enables a form of composition between protocols.
- 148 ■ In Section 4, we prove the aforementioned characterization of $\text{IO-PP}[\prec]$, and in Section 5,
 149 we provide a natural logic and an automaton model expressible in $\text{PP}[\prec]$. In this latter
 150 section, we fall short of showing exact characterizations, but provide conjectures based
 151 on our new models.
- 152 ■ In Section 6, we explore the expressiveness of $\text{IO-PP}[\mathcal{N}]$ and $\text{PP}[\mathcal{N}]$ when the successor
 153 predicate is available, that is, when transitions can be restricted to fire only if they act on
 154 two adjacent agents. We show that if all the predicates of \mathcal{N} are $\text{NSPACE}(n)$ computable,
 155 then $\text{IO-PP}[\mathcal{N}] = \text{PP}[\mathcal{N}] = \text{NSPACE}(n)$ — this is arguably less surprising than our main
 156 result on $\text{IO-PP}[\prec]$, as the successor allows for the left-to-right propagation of information.
- 157 ■ Finally, since protocols are only well-behaved when they are deciders, and since this
 158 property is semantic, we explore in Section 7 whether we can check if a given population
 159 protocol is a decider. We thus ask if the *syntax* of deciders is decidable — this is
 160 sometimes called the *well-specification* problem. We show that it is undecidable for
 161 $\text{IO-PP}[+1]$, $\text{PP}[+1]$, and $\text{PP}[\prec]$, and conditionally decidable for $\text{IO-PP}[\prec]$.

162 Related work

163 Our model is closely related to the *community protocols* of Guerraoui and Ruppert [17].
 164 These are population protocols where each agent has a unique identifier; each agent can
 165 *store* a constant number of identifiers; and interactions depend on these identifiers but only

XX:4 Population Protocols over Ordered Agents

166 with respect to their relative order. The motivation of Guerraoui and Ruppert was to devise
167 an extension of population protocols, as mild as possible, which would be fault-tolerant.
168 They proved that community protocols can decide languages from $\text{NSPACE}(n \log n)$ while
169 tolerating Byzantine failures of a constant number of agents [17]. The unique identifier of
170 each agent is considered to be stored in read-only memory, as in real-world low-cost chips,
171 and so exempt from Byzantine failures. Our model corresponds to community protocols
172 where each agent has a *single immutable* register initialized to its unique identifier.

173 Bournez, Cohen, and Rabie introduce *homonym protocols*, a parameterized restriction of
174 community protocols, where the n agents have $f(n)$ identifiers [7]. The cases of $f(n) = 1$
175 and $f(n) = n$ are respectively population protocols (everyone has the same identifier) and
176 community protocols (there are as many identifiers as agents). Identifiers are from $[0..f(n)-1]$
177 and agents can compare them with respect to $x < y$, $x = y$, $x = y + 1$ and $x = 0$.

178 In [16], Gańczorz et al. introduce *selective population protocols* as an extension of popula-
179 tion protocols where the state space is partitioned into finitely many “groups,” and where an
180 interaction picks at random an initiator in state s , and then a responder in state s' from the
181 group of s , provided it is nonempty. This is a powerful model that allows for zero-tests, i.e.,
182 checking whether no agent holds a certain state. Thus, selective population protocols are
183 orthogonal to our model. However, the authors dedicate a section to the median problem:
184 $\bigcup_{n \geq 0} \Sigma^n a \Sigma^n$. They show that if selective population protocols are extended with the possi-
185 bility of comparing keys, then they can solve the median problem in time $\mathcal{O}(\log^4 n)$. The
186 authors further provide a short proof that, without leveraging the “selective” aspect of their
187 model (i.e., groups and zero-tests), any population protocol for the median problem must
188 work in expected time $\Omega(n)$. This latter setting, only briefly discussed in [16], corresponds to
189 our model.

190 Further extensions include *mediated population protocols* [23], where communication edges
191 have an internal state; *population protocols with unordered data* [5], where the input alphabet
192 is infinite; and *population protocols for graph class identification problems* [37, 1], where
193 agents aim at determining whether the communication topology satisfies some property. For
194 other work on immediate-observation protocols, see [13, 35, 11, 34].

2 Preliminaries

196 **Automata and logic.** We assume some familiarity with formal languages, automata theory
197 and logic over finite words (e.g., see the textbook [11]). For a word $w \in \Sigma^*$, we write $w[i]$ for
198 the i -th letter of w , starting at 1, and $w[i..j]$ for the infix $w[i]w[i+1] \cdots w[j]$. For $\sigma \in \Sigma$, we
199 write $|w|_\sigma$ for the number of occurrences of σ in w .

200 We write FO to denote first-order logic over words, where quantifiers range over positions,
201 that is, the set $\{1, \dots, |w|\}$ for a given word w , and where $a(x)$ holds with respect to w iff
202 $w[x] = a$ (see, e.g., [11, Chap. 8] for formal definitions). We write FO[<] for the extension
203 of FO with the numerical predicate $<$, that allows to test whether $x < y$ for two positions
204 x and y . For example, the sentence $\varphi = (\exists x)(\forall y)[a(x) \wedge ((x < y) \rightarrow b(y))]$ describes the
205 language Σ^*ab^* . By abuse of notation, FO[<] stands for both the set of syntactic sentences
206 and for the class of languages described by these sentences. It is well known that FO[<] is
207 the class of star-free (regular) languages.

208 We write $\Sigma_i[<]$ (resp. $\Pi_i[<]$) for the fragment of FO[<] of sentences in prenex normal
209 form with i blocks of alternating quantifiers starting with \exists (resp. \forall). For example, Σ^*ab^*
210 belongs to $\Sigma_2[<]$ due to the form of our example φ . We let $\Delta_i[<] = \Sigma_i[<] \cap \Pi_i[<]$.

Population protocols over ordered agents. A population protocol (PP) describes how a totally-ordered set of finite-state agents interact and reach a decision about their overall initial states. Interactions can happen between any pair of agents, and predicates are used to restrict how transitions can be taken, based on the position of the agents in the order.

(*Syntax.*) We extend classical PPs to allow for transitions to carry a test on the *positions* of the totally-ordered agents. Let $\mathcal{N} \subseteq 2^{\mathbb{N} \times \mathbb{N}}$ be any set, whose elements we call *numerical predicates*: these will be used as the allowed tests on a transition.¹ We let $\mathbf{true} = \mathbb{N}^2$ be the always-true predicate. The set $\text{PP}[\mathcal{N}]$ of PPs over \mathcal{N} is the set of transition systems (Q, Σ, O, Δ) where Q is a finite set of *states*, $\Sigma \subseteq Q$ is a distinguished subset of *initial* states, $O: Q \rightarrow \{\top, \perp\}$ maps each state to an *opinion*, and $\Delta \subseteq Q^2 \times (\mathcal{N} \cup \{\mathbf{true}\}) \times Q^2$ is a set of *transitions*. An element of Δ is denoted $q_1, q_2 \xrightarrow{P} q_3, q_4$, expressing, intuitively, that if two distinct agents meet, the first being in position i and state q_1 , the second in position j and state q_2 , such that $(i, j) \in P$, then the first agent changes its state to q_3 and the second to q_4 .

If $\mathcal{N} = \{P_1, P_2, \dots\}$, we write $\text{PP}[P_1, P_2, \dots]$ for $\text{PP}[\mathcal{N}]$. Classical PPs can be seen as the class $\text{PP}[\emptyset]$, recalling that we assume that \mathbf{true} is always available as a numerical predicate. Our main interest is in the class $\text{PP}[\prec]$, where \prec is seen as the set of pairs (i, j) with $i < j$, but we will study more expressive predicates in Sections 3 and 6.

A PP is *immediate-observation* if at most one agent changes state in each interaction, i.e., every transition is of the form $a, b \xrightarrow{P} a, c$ or $a, b \xrightarrow{P} c, b$. We write $\text{IO-PP}[\mathcal{N}]$ for the class of protocols in $\text{PP}[\mathcal{N}]$ that are immediate-observation.

(*Semantics.*) Since our agents are totally ordered, we define the *configuration* of a system as a word of Q^+ . *Initial configurations* are words of Σ^+ . Let u and v be two configurations of the same length, we say that u *leads to* v , denoted $u \rightarrow v$, if the two configurations are equal except at potentially two distinct positions i and j , and there is a transition $u[i], u[j] \xrightarrow{P} v[i], v[j] \in \Delta$ with $(i, j) \in P$. We let \rightarrow^* be the reflexive transitive closure of \rightarrow .

Consensus and stability. With w a configuration, let $O(w) \in \{\perp, \top\}$ be the common opinion of all states appearing in w , if there is one; otherwise $O(w)$ is undefined. A configuration w is a *b-consensus* if $O(w) = b$. It is further *b-stable* if $u \rightarrow^* v$ implies that v is a *b-consensus*.

► **Example 1.** We give an example of a protocol in $\text{PP}[\prec]$. Consider the transition system $\mathcal{P} = (\{a, b, q_\perp\}, \{a, b\}, O, \Delta)$, with $O(a) = O(b) = \top$, $O(q_\perp) = \perp$, and transitions $b, a \xrightarrow{\prec} q_\perp, q_\perp$ and $q_\perp, _ \xrightarrow{\mathbf{true}} q_\perp, q_\perp$ with “ $_$ ” standing for any state. Every input configuration belongs to $(a + b)^+$ and is therefore a \top -consensus (all agents output \top initially). However, such configurations need not be \top -stable. For instance, starting from ba we can apply the rule $b, a \xrightarrow{\prec} q_\perp, q_\perp$ and obtain $q_\perp q_\perp$, which is \perp -stable. In contrast, every configuration in $a^* b^*$ is \top -stable: it is a \top -consensus and no transition is enabled.

Language of a PP. Consider an infinite sequence $w_0 \rightarrow w_1 \rightarrow w_2 \rightarrow \dots$ of configurations, which we call a *run*. We say that it is *fair* if for every w_i that appears infinitely often, each configuration of $\{v \mid w_i \rightarrow v\}$ appears infinitely often as well. By induction, fairness guarantees the same for each configuration of $\{v \mid w_i \rightarrow^* v\}$. Intuitively, fairness ensures that reachable configurations cannot be avoided forever (in a probabilistic setting, the fair runs would be the runs occurring almost surely). We will assume that any configuration can be extended into a run, and so into a fair run, by implicitly adding “no operation” transitions.

¹ These are sometimes called *uniform* numerical predicates in the literature, to emphasize the fact that they do not depend on the *total* number of agents. A natural predicate that is not uniform is $\mathbf{max}(x)$ which is true if x is the position of the last agent. This technical difference will not impact our results.

XX:6 Population Protocols over Ordered Agents

253 The *language* of a PP \mathcal{P} is the set $L(\mathcal{P})$ of initial configurations from which there is a
254 fair run that contains a \top -stable configuration. Naturally, an initial configuration can be the
255 origin of a fair run containing a \top -stable configuration, another run containing a \perp -stable
256 configuration, or even a run that contains no stable configurations. We single out PPs that
257 have more crisp behaviors. We say that a PP is a *decider* when for all $w \in \Sigma^*$, there is a
258 $b \in \{\top, \perp\}$ such that *all* fair runs from w contain a b -stable configuration. This is usually
259 called *well-specified* in the literature and we justify our nomenclature in Section 3.1. We
260 say that $L(\mathcal{P})$ is PP[\mathcal{N}]-decidable or IO-PP[\mathcal{N}]-decidable with the obvious meaning. We will
261 identify PP[\mathcal{N}] and IO-PP[\mathcal{N}] with the class of languages decidable by these protocols.

262 Note that since population protocols are ill-defined when no agents are present, we adopt
263 the convention, when working with their languages, to disregard the empty word.

264 ► **Example 2.** Consider the protocol of PP[$<$] from Example 1. We show that it decides the
265 language a^*b^* . Every input in a^*b^* is already \top -stable. Conversely, if an input word is not in
266 a^*b^* , then it contains the factor ba . Hence, in any fair run, the transition $(b, a) \rightarrow (q_\perp, q_\perp)$ is
267 eventually executed. From that point on, every remaining agent eventually interacts with a
268 q_\perp -agent and is converted to q_\perp , so the run reaches a configuration in q_\perp^* , which is \perp -stable.

269 ► **Example 3.** Appendix A.1 provides a more intricate and thorough example of a PP[$<$].

270 We provide a generic upper bound on the complexity of languages decided by population
271 protocols; we will exhibit a matching lower bound in Section 6 for IO-PP[+1]. Recall that
272 NSPACE(n) is the class of languages recognized by linear-bounded nondeterministic Turing
273 machines, i.e., nondeterministic machines that require space $O(n)$ over inputs of size n .
274 Languages of this class are exactly the context-sensitive languages [20].

275 ► **Theorem 4.** Let \mathcal{N} be a set of numerical predicates, all of which decidable in NSPACE(n).
276 We have PP[\mathcal{N}] \subseteq NSPACE(n).

277 **Proof.** Consider a protocol in PP[\mathcal{N}]. The set of configurations that are *not* \top -stable is
278 decidable in NSPACE(n). Indeed, it is sufficient to nondeterministically guess a partial run
279 (i.e., a finite sequence of configurations $w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_n$) that leads to a configuration
280 that is not a \top -consensus. Each transition can be guessed, its condition checked, and its
281 effect applied in NSPACE(n).

282 By the Immerman–Szelepcsényi theorem, NSPACE(n) is closed under complement, and so
283 the set of \top -stable configurations is in NSPACE(n). Since the protocol is a decider, to check
284 that a word w is accepted, it is sufficient to nondeterministically guess a partial run from w ,
285 and check that it ends in a \top -stable configuration. These are all tasks in NSPACE(n). ◀

3 Semantic restrictions of population protocols: a toolbox

287 In this section, we define two restrictions of population protocols that will be used to simplify
288 our constructions.

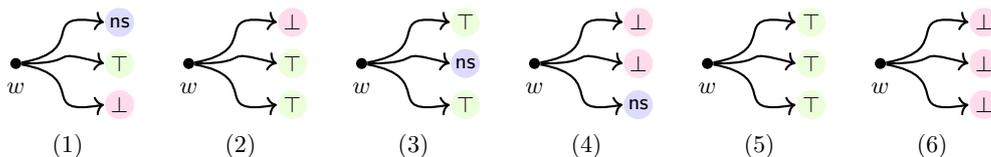
289 We first define *semi-deciders*, which do not require the full behavior of deciders with
290 respect to stable configurations. Throughout the next sections, we will see that semi-deciders
291 are much easier to define for some languages, and we will rely on the forthcoming Lemma 5
292 to combine semi-deciders into deciders.

293 We then define *protocols with stabilizing inputs*, which are protocols where agents can
294 change their mind about their input. Such protocols are harder to design, since they are
295 more robust to change, but we show, in Lemma 9, that they exhibit a strong closure property:
296 they are closed under alphabet rewriting.

3.1 Semi-deciders

► **Definition 1.** We say that a PP *semi-decides* $L \subseteq \Sigma^+$ if for all $u \in \Sigma^+$ and every fair run ρ starting from u , we have $u \in L$ iff ρ visits a \top -stable configuration. We use the terms PP[\mathcal{N}]-*semi-decidable* and IO-PP[\mathcal{N}]-*semi-decidable* with the obvious meaning.

Note that a decider is a semi-decider in which we additionally require that $u \notin L$ iff ρ visits a \perp -stable configuration; that is, all fair runs contain a stable configuration. Illustrating the differences, Figure 1 shows the situations that can occur for runs from an input w .



■ **Figure 1** Configurations labeled \top , \perp and ns are \top -stable, \perp -stable, and configurations from which no stable configurations are reachable. Arrows depict partial, finite runs. Situations (1-6) can happen in PP. Only situations (4-6) can happen in semi-deciders, while only situations (5-6) can in deciders.

Our naming convention is justified by the following property:

► **Lemma 5.** A language L is PP[\mathcal{N}]-*decidable* iff L and its complement are PP[\mathcal{N}]-*semi-decidable*. The same holds for IO-PP[\mathcal{N}].

Proof sketch. We combine the two semi-deciders for L and for $\Sigma^+ \setminus L$ by running them in parallel. Each agent stores a pair of states, one for each semi-decider, together with a *belief* indicating which semi-decider it currently trusts. Transitions either simulate one step of one of the semi-deciders, or *flip* the belief so that beliefs can align; in particular, an agent may flip when it meets an agent with the opposite belief, or when its currently trusted component produces a \perp -witness. For any input u , exactly one of the two simulations eventually provides such a witness (since exactly one of $u \in L$ or $u \notin L$ holds), which forces all agents to converge to the correct belief and yields a stable consensus, hence a decider. ◀

We note these elementary closure properties:

► **Lemma 6.** If $L_1, L_2 \subseteq \Sigma^+$ are PP[\mathcal{N}]-*semi-decidable*, then it is also the case for $L_1 \cap L_2$ and $L_1 \cup L_2$. This further holds for IO-PP[\mathcal{N}] and deciders.

3.2 Protocols with stabilizing inputs

Agents of a protocol are generally not aware that they have reached a stable consensus and hence “terminated.” To carry out a task A and use its output in a subsequent task B, a protocol has to perform *both* tasks concurrently. The protocol for task B thus guesses what the output of task A is going to be, but ought to be able to self-correct if it becomes clear that the guess was wrong. In this subsection, we introduce a formal notion for this “self-correction” which will simplify the design of composable protocols; this is inspired by a recent presentation of [28] for PP[\emptyset].

Let us consider protocols where each agent keeps a copy of its input. Formally, a protocol $\mathcal{P} = (Q, \Sigma, O, \Delta)$ is said to be *input-saving* if

- $Q = \Sigma \times R$ for some finite set R ;

XX:8 Population Protocols over Ordered Agents

- 329 ■ Each $\sigma \in \Sigma$ is identified with (σ, r_σ) for some $r_\sigma \in R$; and
- 330 ■ The first component σ of any state $(\sigma, r) \in Q$ is left unchanged by all transitions of Δ .

331 For all $w \in Q^+$, let $\iota(w) \in \Sigma^+$ be the projection of w onto its first component. Given
 332 $u, v \in Q^n$, we write $u \rightsquigarrow v$ if either $u \rightarrow v$, or v equals u except at a single position i where
 333 $u[i] = (\sigma, r)$ and $v[i] = (\sigma', r)$. This second type of transitions models a “change of mind” of
 334 agent i on its input. We write \rightsquigarrow^* for the reflexive transitive closure of \rightsquigarrow . Note what we do
 335 not change the definition of run, which still relies on \rightarrow only.

336 ► **Definition 2.** We say that \mathcal{P} semi-decides $L \subseteq \Sigma^+$ with stabilizing inputs if (a) \mathcal{P} is
 337 input-saving, and (b) for all $u \in \Sigma^+$, all $u \rightsquigarrow^* v$ and every fair run ρ starting from v , it is
 338 the case that $\iota(v) \in L$ iff ρ visits a \top -stable configuration. We use the term “decides” if (b)
 339 is strengthened with the condition that $\iota(v) \notin L$ iff ρ visits a \perp -stable configuration.

340 Note that this is more robust than simply semi-deciding: if \mathcal{P} semi-decides L with
 341 stabilizing inputs, then it semi-decides L . Intuitively, after $u \rightsquigarrow^* v$, the configuration of the
 342 population (i.e., the projection of v onto the second component) may be *incompatible* with
 343 $\iota(v)$ since agents have possibly changed their mind several times on their input. Computing
 344 with stabilizing inputs means that the protocol is able to fix its configuration so that it
 345 reflects the expected output on $\iota(v)$.

346 It is known that any language $L \in \text{PP}[\emptyset]$ can be decided, and hence semi-decided, with
 347 stabilizing inputs² [28]. Let us turn to an example, which will be useful later:

348 ► **Proposition 7.** *The language $0^*1^* \cdots k^*$ is IO-PP[<]-semi-decidable with stabilizing inputs.*

349 **Proof.** Let $\Sigma = \{0, 1, \dots, k\}$ and $L = 0^*1^* \cdots k^*$. The protocol $\mathcal{P} = (Q, \Sigma, \Delta, O)$ for L is
 350 defined by $Q = \Sigma \times \{\top, \perp\}$, $O((\sigma, o)) = o$, each $\sigma \in \Sigma$ identified with (σ, \top) , and these rules:

$$\begin{aligned}
 351 \quad & (x, \top), (y, o) \xrightarrow{<} (x, \perp), (y, o) && \text{for } x > y, \\
 352 \quad & (x, \perp), (y, o) \xrightarrow{\text{true}} (x, \top), (y, o).
 \end{aligned}$$

353 The purpose of the first rule is to detect a misordering. The second rule allows any agent to
 354 nondeterministically reset itself. Correctness is shown in the appendix. ◀

355 We now turn to closure properties of protocols with stabilizing inputs. We first cover
 356 union and intersection, then move on to (nondeterministic) alphabet rewriting.

357 ► **Lemma 8.** *If $L_1, L_2 \subseteq \Sigma^+$ are PP[\mathcal{N}]-semi-decidable with stabilizing inputs, then it is also
 358 the case for $L_1 \cap L_2$ and $L_1 \cup L_2$. This further holds for IO-PP[\mathcal{N}] and deciders.*

359 Given $f: \Sigma \rightarrow 2^\Gamma$ and $w \in \Sigma^n$, let $f(w) = \{w'_1 \cdots w'_n \mid w'_i \in f(w_i) \text{ for each } i \in [1..n]\}$.
 360 For example, if $f(0) = \{a, b\}$ and $f(1) = \{b, c\}$, then $f(01) = \{ab, ac, bb, bc\}$. We extend this
 361 notion to languages: $f(L) = \bigcup_{w \in L} f(w)$.

362 ► **Lemma 9.** *Let $f: \Sigma \rightarrow 2^\Gamma$. If $L \subseteq \Sigma^+$ is PP[\mathcal{N}]-semi-decidable with stabilizing inputs,
 363 then $f(L)$ is PP[\mathcal{N}]-semi-decidable with stabilizing inputs. This further holds for IO-PP[\mathcal{N}].*

364 **Proof sketch.** The construction builds a protocol that semi-decides $f(L)$ by simulating, in its
 365 second component, the input-saving semi-decider for L on a *guessed* word $v \in \Sigma^+$ compatible
 366 with the real input $u \in \Gamma^+$ (i.e., $u \in f(v)$). Agents may revise this guess: whenever a \perp

² It was claimed earlier by [1], but definitions and proofs were deferred to a full paper that never appeared.

opinion is observed in the simulated component, an agent is allowed to change its guessed letter to any σ consistent with its input letter (preserving $u \in f(v)$ letterwise). If $u \in f(L)$, fairness ensures that the population can eventually rewrite the guessed word into some $v' \in L$, after which the simulation of the semi-decider for L reaches a \top -stable configuration and changes become disabled; if $u \notin f(L)$, reaching a \top -stable configuration would force the simulated input to lie in L , which is not possible. ◀

4 Expressiveness of IO-PP[<]

In this section, we provide a precise characterization of the languages decided by IO-PP[<]. For any language $L \subseteq \Sigma^*$, let us write \equiv_L for the syntactic congruence of L , i.e., $w \equiv_L w'$ iff $uwv \in L \Leftrightarrow uw'v \in L$ for all $u, v \in \Sigma^*$. A language is in the class DA if it is regular, and satisfies, writing $\alpha(w) \subseteq \Sigma$ for the set of letters appearing in a word w :

$$(\forall w \in \Sigma^*) [w \equiv_L w^2 \rightarrow (\forall h \in \alpha(w)^*) [w \equiv_L w \cdot h \cdot w]], \quad (1)$$

It is a fascinating result of Pin and Weil [26] that $DA = \Delta_2[<]$. We will leverage both characterizations to show the theorem below. The left-to-right inclusion will rely on (1), and the converse on $DA = \Delta_2[<]$, semi-decidability, and protocols with stabilizing inputs.

► **Theorem 10.** $IO\text{-}PP[<] = DA$.

4.1 IO-PP[<] \subseteq DA

We first show that the set of stable configurations in a PP[<], and thus in an IO-PP[<], admits a simple description. Given a finite alphabet Σ and two words $u, v \in \Sigma^*$, we say that u is a *subword* of v , written $u \preceq v$, if u can be obtained from v by removing letters. It is well known that \preceq is a *well-quasi-order* over Σ^* [19]. A consequence is that every strictly decreasing sequence of subword-closed sets must be finite.

↓ ► **Lemma 11.** *Any protocol in PP[<] is a well-structured transition system w.r.t. \preceq , that is, for all $u, u', w \in Q^*$ with $u \preceq w$ and $u \rightarrow u'$, there exists w' such that $u' \preceq w'$ and $w \rightarrow w'$. Further, the set of b -stable configurations is subword-closed and computable.*

As expected from the definition of DA, the core of the argument showing the inclusion of IO-PP[<] in DA will rely on a pumping argument. As a first observation, we show that, in an IO-PP[<], an infix of the form wzw with $\alpha(z) \subseteq \alpha(w)$ can, in some sense, mimic the behavior of the protocol on the simpler infix w .

↓ ► **Lemma 12.** *In an IO-PP[<], if $uwv \rightarrow^* u'w'v'$ with $|u| = |u'|$ and $|v| = |v'|$, then for all z with $\alpha(z) \subseteq \alpha(w)$, there exists z' such that $uwzvw \rightarrow^* u'w'z'w'v'$ and $\alpha(z') \subseteq \alpha(w')$.*

We use this first result to show a pumping lemma on IO-PP[<]. It results from the fact that the sets of \top -stable and \perp -stable configurations are downward-closed and Lemma 12.

↓ ► **Lemma 13.** *Let $L \in IO\text{-}PP[<]$. There exists $m \geq 1$ such that, for all $w_1, \dots, w_m \in \Sigma^+$ and $z \in \Sigma^*$ with $\alpha(w_1) = \dots = \alpha(w_m) \supseteq \alpha(z)$, we have $w_1 \cdots w_m \equiv_L (w_1 \cdots w_m)z(w_1 \cdots w_m)$.*

We can now show that IO-PP[<]-decidable languages satisfy the definition of DA. In Lemma 14, we show that Equation 1 is satisfied, and in Lemma 15, that the languages are regular, concluding the proof.

► **Lemma 14.** *Let $L \in IO\text{-}PP[<]$ and let $w \in \Sigma^*$ be such that $w \equiv_L w^2$. It is the case that $w \equiv_L w \cdot h \cdot w$ for all $h \in \alpha(w)^*$.*

XX:10 Population Protocols over Ordered Agents

407 **Proof.** With m obtained from Lemma 13, the latter implies that $w^m h w^m \equiv_L w^m$, and since
408 $w \equiv_L w^2$, we obtain $w \equiv_L w \cdot h \cdot w$ as claimed. ◀

↓ 409 ▶ **Lemma 15.** *Every IO-PP[<]-decidable language is regular.*

410 **Proof sketch.** We show that beyond some length, every word must contain a pattern of the
411 form $(w_1 \cdots w_m)z(w_1 \cdots w_m)$ with $\alpha(w_1) = \dots = \alpha(w_m) \supseteq \alpha(z)$. By Lemma 13, this means
412 that every sufficiently long word is \equiv_L -equivalent to a shorter one, meaning that \equiv_L has
413 finitely many equivalence classes. ◀

4.2 DA \subseteq IO-PP[<]

414 Recall that $L \in \Pi_2[<]$ iff $\Sigma^* \setminus L \in \Sigma_2[<]$. Since DA is equal to $\Delta_2[<] = \Sigma_2[<] \cap \Pi_2[<]$, it is
415 enough to prove that $\Sigma_2[<]$ languages are IO-PP[<]-semi-decidable, appealing to Lemma 5 to
416 conclude. Note that $\Sigma_2[<]$ is the set of languages expressible as finite unions of languages of
417 the form $L = A_0^* a_1 A_2^* \cdots a_{m-1} A_m^*$ with the A_i 's being subalphabets (see, e.g., [25, Thm. 8.8]).
418 Since semi-decidable languages are closed under union by Lemma 6, we need only show that L is IO-PP[<]-
419 semi-decidable to conclude. We start with a technical proposition that extends a result of [28]
420 to immediate-observation protocols, then provide a semi-decider for L in Proposition 17.
421

↓ 422 ▶ **Proposition 16.** *The language $\{w \in \Sigma^+ \mid |w|_a = 1\}$ is IO-PP[\emptyset]-semi-decidable with
423 stabilizing inputs.*

424 **Proof sketch.** Each agent stores its current input, its last input, and a belief on the output.
425 If the current input of an agent mismatches its former one, then it resets itself to the current
426 one. Each agent eventually stabilizes to a fixed input (σ, σ, \cdot) . The belief component then
427 controls the consensus: an (a, a, \top) agent can turn any (σ, σ, \cdot) with $\sigma \neq a$ to \top , while the
428 presence of any (σ, σ, \perp) can spread \perp to other agents. Finally, if an a -agent observes another
429 a , it switches to \perp ; combined with the rule that lets an a -agent reset itself to \top , this makes
430 a -agents alternate forever between \top and \perp whenever there are at least two a 's, preventing
431 stabilization in that case. As a result, if the input contains no a then every fair run reaches
432 a \perp -stable consensus; if it contains exactly one a then every fair run reaches a \top -stable
433 consensus; and if it contains at least two a 's then some agent flips its belief forever. ◀

434 ▶ **Proposition 17.** *The language $L = A_0^* a_1 A_2^* \cdots a_{m-1} A_m^*$ is IO-PP[<]-semi-decidable (with
435 stabilizing inputs).*

436 **Proof.** Let $\Gamma = \{0, 1, \dots, m\}$, $K = 0^* 1^* \cdots m^*$ and $K' = \bigcap_{a \in \Gamma, a \text{ odd}} \{w \in \Gamma^+ \mid |w|_a = 1\}$.
437 By Proposition 7, K is IO-PP[<]-semi-decidable with stabilizing inputs. Furthermore, by
438 Proposition 16 and Lemma 8, K' is IO-PP[\emptyset]-semi-decidable with stabilizing inputs. By
439 Lemma 8, $K \cap K'$ is IO-PP[<]-semi-decidable with stabilizing inputs. Let $f(i) = A_i$ for even
440 i , and $f(i) = \{a_i\}$ otherwise. We are done by Lemma 9 since $L = f(K \cap K')$. ◀

5 Expressiveness of PP[<]

441 The crisp characterization of the previous section ties IO-PP[<] to a wealth of computational
442 models with strikingly different flavors. Chief among them, DA is characterized by a logic,
443 $\Delta_2[<]$, and by partially-ordered unambiguous automata. A natural question is thus whether
444 PP[<] also admits such a diverse array of characterizations. We fall short of providing exact
445 characterizations, but offer, in this section, two large classes of languages, one logically-defined
446 and one based on partially-ordered automata, that are PP[<]-decidable.
447

Our formalisms will involve Presburger arithmetic, the first-order theory of the naturals with order and addition; e.g., $\phi(x) = (\exists y \in \mathbb{N})[y \geq 1 \wedge x = 2y]$ holds iff x is a positive even number. Write \equiv_c for the equivalence of naturals modulo c . It is well known that Presburger arithmetic together with \equiv_c with any $c \geq 2$ admits quantifier elimination. For our purposes, a *Presburger formula* is a (quantifier-free) Boolean combination of predicates of the form $\sum_{i=1}^n a_i x_i < b$ or $\sum_{i=1}^n a_i x_i \equiv_c b$, where $a_i, b \in \mathbb{Z}$, $c \geq 2$ and variables x_i are over \mathbb{N} .

5.1 First-order logic over word intervals

Definition 3. For $w \in \Sigma^n$ and $\sigma \in \Sigma$, let $\#_\sigma: \{1, \dots, n\} \times \{1, \dots, n\} \rightarrow \mathbb{N}$ be the function that counts the number of occurrences of σ between two positions of w , i.e., $\#_\sigma(x, y) = |w[x..y]|_\sigma$. We define FO^{int} as first-order logic over word *intervals*, that is, with access to numerical values $\#_a(x, y)$ where x and y are either first-order variables, the first position (denoted “1”) or the last position (denoted “**max**”). We will study $\text{FO}^{\text{int}}[<, +, \equiv]$, where we allow numerical values to be compared, added, and tested modulo c for any constant c . Note that a variable x can be expressed as $\sum_{\sigma \in \Sigma} \#_\sigma(1, x)$, we will thus assume that the atomic formulas only have terms of the form $\#_a(x, y)$, though we write 1 for $\sum_{\sigma \in \Sigma} \#_\sigma(1, 1)$, and **max** for $\sum_{\sigma \in \Sigma} \#_\sigma(1, \text{max})$. The logics Σ_k^{int} , Π_k^{int} , and Δ_k^{int} are naturally defined.

Example 18. Let us show that the *median language* $\bigcup_{n \in \mathbb{N}} \Sigma^n a \Sigma^n$ belongs to $\Delta_1^{\text{int}}[<, +, \equiv]$. The following predicate asserts that x is the middle position: $\psi(x) = \sum_{\sigma \in \Sigma} \#_\sigma(1, x) = \sum_{\sigma \in \Sigma} \#_\sigma(x, \text{max})$. The median language can either be expressed by $(\exists x)[\psi(x) \wedge a(x)]$, or by $(\forall x)[(\psi(x) \rightarrow a(x)) \wedge \text{max} \equiv_2 1]$. Note that $a(x)$ is the same as $\#_a(x, x) = 1$.

We now provide a convenient characterization of languages in $\Sigma_1^{\text{int}}[<, +, \equiv]$.

Lemma 19. Any language from $\Sigma_1^{\text{int}}[<, +, \equiv]$ is a finite union of languages of the form $K = \{a_0 w_1 a_1 \dots w_m a_m \mid w_i \in \Sigma^*, a_i \in \Sigma, \varphi(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto |a_i|_\sigma)\}$ where $m \geq 0$ and φ is a Presburger formula over variables $\{x_{i,\sigma} \mid i \in [1..m], \sigma \in \Sigma\} \cup \{y_{i,\sigma} \mid i \in [0..m], \sigma \in \Sigma\}$.

Thanks to the previous lemma, we are able to build a protocol $\text{PP}[<]$ semi-deciding any language in $\Sigma_1^{\text{int}}[<, +, \equiv]$, by building semi-decidars with stabilizing inputs for each K , and then using the closure properties presented in Section 3.2.

Proposition 20. Any language from $\Sigma_1^{\text{int}}[<, +, \equiv]$ is $\text{PP}[<]$ -semi-decidable.

Proof. Let $L \in \Sigma_1^{\text{int}}[<, +, \equiv]$. By Lemma 19, L is a finite union of languages of the form $L' = \{a_0 w_1 a_1 \dots w_m a_m \mid w_i \in \Sigma^*, a_i \in \Sigma, \varphi(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto |a_i|_\sigma)\}$ where $m \geq 0$ and φ is a Presburger formula. It suffices to show that L' is $\text{PP}[\mathcal{N}]$ -semi-decidable with stabilizing inputs. Indeed, by Lemma 8, that class of languages is closed under union.

Let $\underline{\sigma} = \{\sigma \mid \sigma \in \Sigma\}$, $A_i = \underline{\sigma} \times \{i\}$, $W_i = \Sigma \times \{i\}$ and $\Gamma = \bigcup_{i \in [0..m]} A_i \cup \bigcup_{i \in [1..m]} W_i$. We justify that the three following languages K , K' and K'' over alphabet Γ are semi-decided with stabilizing inputs:

- $K = A_0^* W_1^* A_1^* \dots W_m^* A_m^*$,
- $K' = \{w \in \Gamma^+ \mid \varphi(x_{i,\sigma} \mapsto |w|_{(\sigma,i)}, y_{i,\sigma} \mapsto |w|_{(\underline{\sigma},i)})\}$,
- $K'' = \left\{ w \in \Gamma^+ \mid \sum_{\gamma \in A_1} |w|_\gamma = 1 \wedge \dots \wedge \sum_{\gamma \in A_m} |w|_\gamma = 1 \right\}$ (recall $A_i \cap A_j = \emptyset$ for $i \neq j$).

Let $f(2i) = A_i$ and $f(2i+1) = W_{i+1}$. We have $f(0^* 1^* 2^* \dots (2m)^*) = K$. By Proposition 7 and Lemma 9, K is $\text{PP}[<]$ -semi-decidable with stabilizing inputs. Since $K', K'' \in \text{PP}[\emptyset]$, these two languages are $\text{PP}[\emptyset]$ -decidable with stabilizing inputs [1, 28].

By Lemma 8, the language $K \cap K' \cap K''$ is $\text{PP}[<]$ -semi-decidable with stabilizing inputs. Let $g((\sigma, i)) = g((\underline{\sigma}, i)) = \{\sigma\}$. We are done by Lemma 9 since $L' = g(K \cap K' \cap K'')$. ◀

XX:12 Population Protocols over Ordered Agents

491 ▶ **Corollary 21.** $\Delta_1^{\text{int}}[\langle, +, \equiv] \subseteq \text{PP}[\langle]$.

492 The language $(ab)^+$ over alphabet $\{a, b\}$ plays a key role in the study of DA; indeed,
 493 DA is the largest class of regular languages, closed under the so-called variety operations,
 494 that does *not* contain $(ab)^+$. Strikingly, $\Delta_1^{\text{int}}[\langle, +, \equiv]$ cannot either express this language,
 495 although it belongs to $\Pi_1^{\text{int}}[\langle, +, \equiv]$, since it is described by $(\forall x)[(x \equiv_2 1) \leftrightarrow a(x)]$.

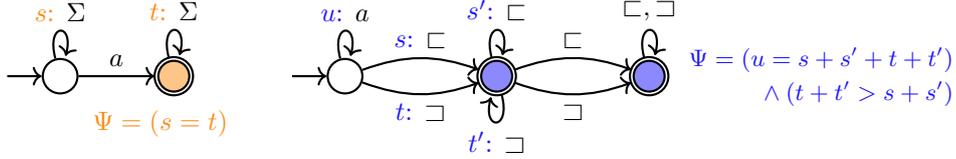
496 ▶ **Proposition 22.** *The language $(ab)^+$ does not belong to $\Sigma_1^{\text{int}}[\langle, +, \equiv]$.*

5.2 Partially-ordered Parikh automata

497 A *partially-ordered Parikh automaton* (poPA) is a tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F, \Psi)$ where:

- 499 ■ Q is a finite set of *states* equipped with a partial order \leq ;
- 500 ■ Σ is a finite *alphabet*;
- 501 ■ $\delta \subseteq Q \times \Sigma \times Q$ is the *transition relation*, that satisfies $p \leq q$ for all $(p, \sigma, q) \in \delta$; if $p = q$,
 502 the transition is dubbed a *self-loop*, and otherwise, a *progress transition*;
- 503 ■ $q_0 \in Q$ is the *initial state*, and $F \subseteq Q$ is the set of *final states*;
- 504 ■ Ψ is a Presburger formula over variables δ .

505 For all $t = (q, \sigma, q') \in \delta$, we write $q \rightarrow^t q'$ and $\pi_\Sigma(t) = \sigma$. We naturally lift these notations
 506 to sequences. A word $w \in \Sigma^*$ is *accepted* by \mathcal{A} if there exist $\rho \in \delta^*$ and $q \in F$ such that
 507 $q_0 \rightarrow^\rho q$, $\pi_\Sigma(\rho) = w$ and $\Psi(t \mapsto |\delta|_t)$ holds. The *language* of \mathcal{A} is the set $L(\mathcal{A})$ of words it
 508 accepts. Figure 2 depicts two examples of poPA.



■ **Figure 2** Example of poPA for the median language $\bigcup_{n \geq 0} \Sigma^n a \Sigma^n$ (left) and the coDyck-witness language $\{a^n v \mid n \geq 1, v \in \{\square, \sqsupset\}^{\geq n}, |v[1..n]|\square > |v[1..n]|\sqsupset\}$ (right).

509 Lemma 19 allows us to translate a formula from $\Sigma_1^{\text{int}}[\langle, +, \equiv]$ into a poPA by guessing a
 510 factorization $a_0 w_1 a_1 \cdots w_m a_m$ with a progress transition for each a_i , and a self-loop for each
 511 w_i ; and then using the Presburger acceptance formula of the automaton to verify the guess.
 512 We can also translate a poPA into a formula of $\Sigma_1^{\text{int}}[\langle, +, \equiv]$ by guessing the position of the
 513 progress transitions and verifying the validity of the path. This yields:

514 ▶ **Theorem 23.** *A language is recognized by some poPA iff it belongs to $\Sigma_1^{\text{int}}[\langle, +, \equiv]$.*

515 Since $\Sigma_2[\langle]$ is known to be equivalent to partially-ordered automata, and since the latter
 516 trivially form a subclass of poPA, it is the case that $\Sigma_2[\langle] \subseteq \Sigma_1^{\text{int}}[\langle, +, \equiv]$. As $\Sigma_2[\langle]$ is a
 517 subset of the regular languages, this implies that $\Delta_2[\langle] \subset (\Delta_1^{\text{int}}[\langle, +, \equiv] \cap \text{Reg})$. The strict
 518 inclusion follows from the fact that $\Delta_2[\langle]$ cannot express the language $E = \{a^{2n} \mid n \geq 1\}$
 519 over alphabet $\{a\}$, whereas it trivially belongs to $\Delta_1^{\text{int}}[\langle, +, \equiv]$, with the formula $\mathbf{max} \equiv_2 0$.

520 A poPA \mathcal{A} is *weakly unambiguous*³ if every $w \in L(\mathcal{A})$ is accepted by at most one path
 521 (w.r.t. F and Ψ). Note that the automata of Figure 2 are weakly unambiguous and can be

³ The nomenclature, introduced in [6], stems from prior studies [8] which called “unambiguous” the PA with an unambiguous underlying automaton. We note that the examples of Figure 2 can be shown, using the tools of [8], not to be expressible with unambiguous PA.

522 complemented. By Theorem 23, this means that both languages belong to $\Delta_1^{\text{int}}[<, +, \equiv]$.
 523 More generally, we conjecture that weakly-unambiguous poPA are closed under complement.

524 **Observation 24.** *If weakly-unambiguous poPA are closed under complement, then any*
 525 *language recognized by a weakly-unambiguous poPA belongs to $\Delta_1^{\text{int}}[<, +, \equiv]$.*

526 Recall that DA is also characterized by two-way *deterministic* partially-ordered automata.
 527 We can show that two-way poPA are equivalent to poPA. However, defining two-way
 528 deterministic PA as two-way deterministic automata with a Presburger constraint, it is
 529 known that the model is as expressive as unambiguous PA [15], which cannot express the
 530 languages of Figure 2.

531 6 Expressiveness of IO-PP[\mathcal{N}] and PP[\mathcal{N}] when successor is available

532 We now consider protocols where transitions may be fired only when two agents are adjacent.
 533 Formally, we look at IO-PP[\mathcal{N}] and PP[\mathcal{N}] with $+1 \in \mathcal{N}$, where $+1$ (sometimes written **succ**
 534 in the literature) is the predicate $\{(x, x+1) \mid x \in \mathbb{N}\}$. This study is reminiscent of the classic
 535 study of $\text{FO}^2[<, +1]$ [14, 21, 33, 36] that followed that of $\text{FO}^2[<]$, the two-variable fragment
 536 of $\text{FO}[<]$ that is as expressive as IO-PP[$<$].

537 **Example 25.** Recall that the set of all b -stable configurations of a PP[$<$] is subword-closed,
 538 and hence regular. The following PP[$+1$] has a non-regular set of \top -stable configurations:

$$539 \begin{array}{ccc} a^{\bar{x}}, b^{\bar{x}} \xrightarrow{\text{true}} a^{\bar{y}}, b^{\bar{c}} & a^{\bar{y}}, a \xrightarrow{+1} \checkmark, a^{\bar{x}} & a^{\bar{y}}, b^{\bar{c}} \xrightarrow{+1} \mathbf{x}, \mathbf{x} \\ 540 & b, b^{\bar{c}} \xrightarrow{+1} b^{\bar{x}}, \checkmark & \end{array}$$

541 Let us set the opinion of each state to \top , except for \mathbf{x} . The set of \top -stable configurations
 542 cannot be regular since, from $w \in a^{\bar{x}}a^*b^*b^{\bar{x}}$, the state \mathbf{x} will appear iff $|w|_a = |w|_b$, e.g.,

$$543 a^{\bar{x}}abb^{\bar{x}} \rightarrow a^{\bar{y}}abb^{\bar{c}} \rightarrow \checkmark a^{\bar{x}}bb^{\bar{c}} \rightarrow \checkmark a^{\bar{x}}b^{\bar{x}}\checkmark \rightarrow \checkmark a^{\bar{y}}b^{\bar{c}}\checkmark \rightarrow \checkmark \mathbf{x} \checkmark \checkmark.$$

544 The main result of this section is:

545 **Theorem 26.** *Let \mathcal{N} be a set of NSPACE(n)-decidable numerical predicates that contains*
 546 *$+1$. We have IO-PP[\mathcal{N}] = PP[\mathcal{N}] = NSPACE(n).*

547 To show this result, it is enough to prove that $\text{NSPACE}(n) \subseteq \text{IO-PP}[+1]$, since $\text{PP}[\mathcal{N}] \subseteq$
 548 $\text{NSPACE}(n)$ has been established beforehand in Theorem 4.

549 6.1 IO-PP[$+1$] = PP[$+1$]

550 We start by establishing that IO-PP[$+1$] are as expressive as PP[$+1$].

551 **Lemma 27.** *Every language semi-decided by a protocol in PP[$+1$] is also semi-decided by*
 552 *a protocol in IO-PP[$+1$].*

553 **Proof sketch.** The construction turns a PP[$+1$] protocol \mathcal{P} into an immediate-observation
 554 protocol \mathcal{P}' by replacing each transition $\delta = (a, b) \xrightarrow{+1} (c, d)$ with a short *four-step handshake*
 555 along the successor relation, using auxiliary markers a^δ and $b^{\text{ack}, \delta}$. The four rules are:

$$556 \begin{array}{ccc} a, b \xrightarrow{+1} a^\delta, b & a^\delta, b \xrightarrow{+1} a^\delta, b^{\text{ack}, \delta} \\ 557 a^\delta, b^{\text{ack}, \delta} \xrightarrow{+1} c, b^{\text{ack}, \delta} & c, b^{\text{ack}, \delta} \xrightarrow{+1} c, d \end{array}$$

558 Intuitively, one agent announces that it wants to perform δ , the neighbor acknowledges, and
 559 the two agents then update one after the other; additional clean-up rules erase incomplete
 560 handshakes so that executions cannot block when several handshakes overlap. \blacktriangleleft

561 **6.2** $\text{NSPACE}(n) \subseteq \text{PP}[+1]$ 562 Thanks to Lemma 27, we can focus on showing that $\text{NSPACE}(n) \subseteq \text{PP}[+1]$.563 **► Lemma 28.** *The language of a linear-bounded nondeterministic Turing machine is $\text{PP}[+1]$ -*
564 *semi-decidable.*565 To match the semi-decision semantics of protocols, we replace a linear-bounded Turing
566 machine M by a variant M_{rst} in which (i) from any non-accepting state the machine may
567 reset to the initial configuration on the same input, and (ii) once an accepting state is reached,
568 the machine enters a final sweeping phase that overwrites the whole tape with a symbol f .569 The protocol \mathcal{P}_M simulates a linear-bounded machine M_{rst} using the successor relation.
570 The main difficulty is that protocol inputs do *not* come with explicit endmarkers: initially,
571 every agent locally *guesses* that it is both the left and the right boundary, and stores its
572 input letter forever. The protocol then proceeds as follows. As long as neighboring boundary
573 guesses are consistent, the agents interpret their “tape component” as the current tape content
574 and faithfully simulate one step of M_{rst} (including the head marker) using local successor
575 interactions. If an inconsistency is detected, the protocol triggers a *reset* by spawning left-
576 and/or right-moving reset tokens. These tokens propagate along the successor order, restoring
577 each visited cell to its stored input symbol and thereby erasing any partial simulation. When
578 a token reaches the corresponding boundary, it restores the appropriate boundary component
579 and disappears, leaving behind a clean segment. Thanks to this reinitialization mechanism,
580 any run can be forced to reach a well-formed tape with unique boundaries, from which the
581 simulation behaves exactly like M_{rst} . Thanks to (ii), upon acceptance, the whole tape is
582 overwritten by f , yielding a \top -stable configuration.583 **► Corollary 29.** $\text{NSPACE}(n) \subseteq \text{PP}[+1]$.584 **Proof.** Let $L \in \text{NSPACE}(n)$. As $\text{NSPACE}(n)$ is closed under complement, there exist two
585 linear-bounded nondeterministic Turing machines M and \overline{M} recognizing L and $\Sigma^+ \setminus L$,
586 respectively. By the above construction, we can build two $\text{PP}[+1]$ protocols, say \mathcal{P}_M and
587 $\mathcal{P}_{\overline{M}}$, that semi-decide L and $\Sigma^+ \setminus L$, respectively. Lemma 5 allows to conclude. ◀588 **7** **Decidability of checking whether a population protocol is a decider**589 A good portion of our results assumes that a given population protocol is a decider or a
590 semi-decider in order to *construct* an object (a Turing machine in Theorem 4, a formula in
591 Theorem 10, etc.). To fully understand how constructive these proofs are, we ought to study
592 whether it is decidable, given a population protocol, to check if it is a decider. In other words,
593 *is the syntax of deciders decidable?* From the perspective of formal verification, this is also a
594 natural problem, known as the well-specification problem. It is decidable for $\text{PP}[\emptyset]$ [12], and
595 undecidable when the alphabet is infinite, except for immediate-observation protocols [34].596 We show that the problem is undecidable already for $\text{PP}[<]$ and $\text{IO-PP}[+1]$, and provide
597 a natural conjecture that would entail that the problem is decidable for $\text{IO-PP}[<]$.598 **7.1** **The syntax of $\text{PP}[<]$, $\text{IO-PP}[+1]$ and $\text{PP}[+1]$ deciders are undecidable**599 First, we consider the *emptiness problem*: given a $\text{PP}[\mathcal{N}]$ protocol \mathcal{P} , is there an execution
600 from an initial configuration u to a \top -stable one? We use a technique from [34] to show a
601 general reduction to the syntax problem.

602 ▶ **Lemma 30.** *The emptiness problem for $\text{PP}[\mathcal{N}]$ (resp. $\text{IO-PP}[\mathcal{N}]$) reduces to deciding the*
 603 *complement of the syntax of $\text{PP}[\mathcal{N}]$ (resp. $\text{IO-PP}[\mathcal{N}]$) deciders.*

604 **Proof sketch.** From a protocol $\mathcal{P} = (Q, \Sigma, O, \Delta)$, we construct a protocol \mathcal{P}' such that \mathcal{P}' is
 605 *not* a decider iff \mathcal{P} can reach a \top -stable configuration. The construction adds a fresh sink
 606 state q_{\perp} with $O(q_{\perp}) = \perp$, and for every state q with $O(q) = \perp$ we add a transition allowing
 607 an agent in q to switch to q_{\perp} .

608 As a consequence, from any configuration that contains a \perp -agent, a fair execution can
 609 drive the system to the \perp -stable consensus q_{\perp}^* . On the other hand, if a \top -stable configuration
 610 is reachable in \mathcal{P} , then the same configuration is reachable in \mathcal{P}' and remains \top -stable there:
 611 none of the added transitions to q_{\perp} are enabled from \top -states. Therefore \mathcal{P}' is not a decider.

612 Conversely, if \mathcal{P} has no reachable \top -stable configuration, then no fair run of \mathcal{P}' stabilizes
 613 to \top . Every fair run stabilizes to q_{\perp}^* and \mathcal{P}' decides the empty language. The construction
 614 preserves immediate observation, as each new transition updates at most one agent. ◀

615 This reduction, combined with the translation from linear-bounded Turing machines to
 616 $\text{IO-PP}[+1]$ from Section 6, already yields the following result.

617 ▶ **Corollary 31.** *The syntax of $\text{IO-PP}[+1]$ and $\text{PP}[+1]$ deciders are undecidable.*

618 ▶ **Theorem 32.** *The emptiness problem is undecidable for $\text{PP}[<]$. Hence the syntax of $\text{PP}[<]$*
 619 *deciders is also undecidable.*

620 **Proof sketch.** We reduce from the Post correspondence problem. Specifically, we take
 621 two homomorphisms $h_1, h_2: B^* \rightarrow A^*$ and construct a $\text{PP}[<]$ that can reach a \top -stable
 622 configuration if and only if there is a word u such that $h_1(u) = h_2(u)$. We obtain it as a
 623 product of three protocols. One makes sure that we can only reach a \top -stable configuration
 624 from an initial one of the form $\#_1 v \#_2 u \#_3$ with $v \in B^*$ and $u \in A^*$. The two others check
 625 that we can reach a \top -stable configuration only when $h_1(u) = v$ and $h_2(u) = v$, respectively.
 626 Each one does so by making agents simulate two reading heads going through u and v in
 627 lockstep and verifying that $h_i(u) = v$. Initially all agents have opinion \perp , and in order to
 628 change it they must carry the reading head at some point, hence we cannot skip any letter
 629 in order to reach a \top -stable configuration.

630 Hence, a \top -stable configuration is reachable iff there exist u, v with $h_1(u) = v = h_2(u)$. ◀

631 7.2 The syntax of $\text{IO-PP}[<]$ deciders is decidable, conditionally

632 We present a conjecture on the reachability relation of $\text{IO-PP}[<]$, then show that it entails
 633 decidability of the syntax of their deciders:

634 ▶ **Conjecture 33.** *The set of configurations reachable from a DA language in an $\text{IO-PP}[<]$*
 635 *protocol is also a DA language.*

636 ▶ **Theorem 34.** *If Conjecture 33 holds, then the syntax of $\text{IO-PP}[<]$ deciders is decidable.*

637 **Proof sketch.** We use two semi-decision procedures. The first one checks, for each length n ,
 638 if the protocol is a decider on words of length n . If the input protocol is not a decider, we
 639 will observe it for some n . The second one looks for invariants witnessing that the protocol
 640 is a decider. We look for two disjoint languages K_{\top}, K_{\perp} such that every input word is in
 641 one of the two, they are closed under transitions of the protocol, and from everywhere in K_b
 642 we can reach a b -stable configuration. Assuming the conjecture, we can assume that those
 643 invariants are in DA. We show that we can enumerate potential regular invariants and check
 644 the requirements, yielding the second semi-decision procedure. ◀

8 Open questions

The most tantalizing open question left open by this work is the characterization of $\text{PP}[\prec]$. We conjecture that the logic and automata models introduced in Section 5 are tight:

► **Conjecture 35.** $\text{PP}[\prec]$ is the class of languages recognized by $\Delta_1^{\text{int}}[\prec, +, \equiv]$ and the class of languages recognized by weakly unambiguous poPA.

The previous conjecture would entail in particular that weakly unambiguous poPA are closed under complement. We note that it is not known whether weakly unambiguous PA themselves are closed under complement, but it may be easier to show such closure under the partially ordered assumption.

Conjecture 33 is also very natural and left open.

References

- 1 Dana Angluin, James Aspnes, Melody Chan, Michael J. Fischer, Hong Jiang, and René Peralta. Stably computable properties of network graphs. In *First IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 63–74, 2005. doi:10.1007/11502593_8.
- 2 Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006. doi:10.1007/s00446-005-0138-3.
- 3 Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007. doi:10.1007/s00446-007-0040-2.
- 4 James Aspnes and Eric Ruppert. An introduction to population protocols. In *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer, 2009. doi:10.1007/978-3-540-89707-1_5.
- 5 Michael Blondin and François Ladouceur. Population protocols with unordered data. In *Proc. 50th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 115:1–115:20, 2023. doi:10.4230/LIPICS.ICALP.2023.115.
- 6 Alin Bostan, Arnaud Carayol, Florent Koechlin, and Cyril Nicaud. Weakly-unambiguous Parikh automata and their link to holonomic series. In *Proc. 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 114:1–114:16, 2020. doi:10.4230/LIPICS.ICALP.2020.114.
- 7 Olivier Bournez, Johanne Cohen, and Mikaël Rabie. Homonym population protocols. *Theory of Computing Systems*, 62(5):1318–1346, 2018. doi:10.1007/S00224-017-9833-2.
- 8 Michaël Cadilhac, Alain Finkel, and Pierre McKenzie. Unambiguous constrained automata. *International Journal of Foundations of Computer Science*, 24(7):1099–1116, 2013. doi:10.1142/S0129054113400339.
- 9 Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2014. doi:10.1007/S11047-013-9393-6.
- 10 Zoë Diamadi and Michael J Fischer. A simple game for the study of trust in distributed systems. *Wuhan University Journal of Natural Sciences*, 6(1):72–82, 2001. doi:10.1007/BF03160228.
- 11 Javier Esparza and Michael Blondin. *Automata theory: An algorithmic approach*. The MIT Press, 2023.
- 12 Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of population protocols. In *Proc. 26th International Conference on Concurrency Theory (CONCUR)*, pages 470–482, 2015. doi:10.4230/LIPICS.CONCUR.2015.470.

- 691 **13** Javier Esparza, Pierre Ganty, Rupak Majumdar, and Chana Weil-Kennedy. Verification
692 of immediate observation population protocols. In *Proc. 29th International Conference on*
693 *Concurrency Theory (CONCUR)*, volume 118, pages 31:1–31:16, 2018. doi:10.4230/LIPICS.
694 CONCUR.2018.31.
- 695 **14** Kousha Etessami, Moshe Y. Vardi, and Thomas Wilke. First-order logic with two variables
696 and unary temporal logic. *Information and Computation*, 179(2):279–295, 2002. doi:10.1006/
697 INCO.2001.2953.
- 698 **15** Emmanuel Filiot, Shibashis Guha, and Nicolas Mazzocchi. Two-way Parikh automata. In
699 *Proc. 39th IARCS Annual Conference on Foundations of Software Technology and Theoretical*
700 *Computer Science (FSTTCS)*, pages 40:1–40:14, 2019. doi:10.4230/LIPICS.FSTTCS.2019.40.
- 701 **16** Adam Ganczorz, Leszek Gasieniec, Tomasz Jurdzinski, Jakub Kowalski, and Grzegorz
702 Stachowiak. Selective population protocols. In *Proc. 26th International Symposium on*
703 *Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 225–239, 2024.
704 doi:10.1007/978-3-031-74498-3_16.
- 705 **17** Rachid Guerraoui and Eric Ruppert. Names trump malice: Tiny mobile agents can toler-
706 ate Byzantine failures. In *Proc. 36th International Colloquium Automata, Languages and*
707 *Programming (ICALP)*, pages 484–495, 2009. doi:10.1007/978-3-642-02930-1_40.
- 708 **18** Simon Halfon. *On Effective Representations of Well Quasi-Orderings. (Représentations*
709 *Effectives des Beaux Pré-Ordres)*. PhD thesis, Université Paris-Saclay, France, 2018. URL:
710 <https://tel.archives-ouvertes.fr/tel-01945232>.
- 711 **19** Graham Higman. Ordering by divisibility in abstract algebras. *Proceedings of the London*
712 *Mathematical Society*, s3-2(1):326–336, 1952. doi:10.1112/plms/s3-2.1.326.
- 713 **20** Sige-Yuki Kuroda. Classes of languages and linear-bounded automata. *Information and*
714 *Control*, 7(2):207–223, 1964. doi:10.1016/S0019-9958(64)90120-2.
- 715 **21** Kamal Lodaya, Paritosh K. Pandya, and Simoni S. Shah. Around dot depth two. In *Proc.*
716 *14th International Conference on Developments in Language Theory (DLT)*, pages 303–315,
717 2010. doi:10.1007/978-3-642-14455-4_28.
- 718 **22** M. Lothaire. *Algebraic Combinatorics on Words*. Encyclopedia of Mathematics and its
719 Applications. Cambridge University Press, 2002.
- 720 **23** Othon Michail, Ioannis Chatzigiannakis, and Paul G. Spirakis. Mediated population protocols.
721 *Theoretical Computer Science*, 412(22):2434–2450, 2011. doi:10.1016/J.TCS.2011.02.003.
- 722 **24** J. Andres Montoya. Asymptotic reasoning with two variables. In *Proc. 31st International*
723 *Workshop on Logic, Language, Information, and Computation (WoLLIC)*, pages 38–55, 2025.
724 doi:10.1007/978-3-031-99536-1_3.
- 725 **25** Jean-Éric Pin. Syntactic semigroups. In *Handbook of Formal Languages, Volume 1: Word,*
726 *Language, Grammar*, pages 679–746. Springer, 1997. doi:10.1007/978-3-642-59136-5_10.
- 727 **26** Jean-Éric Pin and Pascal Weil. Polynomial closure and unambiguous product. In *Proc. 22nd*
728 *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 348–359,
729 1995. doi:10.1007/3-540-60084-1_87.
- 730 **27** Frank P Ramsey. On a problem of formal logic. In *Classic Papers in Combinatorics*, pages
731 1–24. Springer, 1987. doi:10.1007/978-0-8176-4842-8_1.
- 732 **28** Michael Raskin. Modular population protocols. In *Proc. 20th International Symposium*
733 *on Algorithmics of Wireless Networks (ALGOWIN)*, pages 173–187, 2024. doi:10.1007/
734 978-3-031-74580-5_13.
- 735 **29** Marcel-Paul Schützenberger. Sur le produit de concaténation non ambigü. *Semigroup Forum*,
736 13(1):47–75, 1976. doi:10.1007/bf02194921.
- 737 **30** Thomas Schwentick, Denis Thérien, and Heribert Vollmer. Partially-ordered two-way automata:
738 A new characterization of DA. In *Proc. 5th International Conference on Developments in*
739 *Language Theory (DLT)*, pages 239–250, 2001. doi:10.1007/3-540-46011-X_20.
- 740 **31** Michael Sipser. *Introduction to the theory of computation*. Thomson Course Technology,
741 second edition, 2006.

XX:18 Population Protocols over Ordered Agents

- 742 **32** Pascal Tesson and Denis Thérien. Diamonds are forever: The variety DA. In *Semigroups,*
743 *algorithms, automata and languages*, pages 475–499. World Scientific, 2002. doi:10.1142/5050.
- 744 **33** Denis Thérien and Thomas Wilke. Over words, two variables are as powerful as one quantifier
745 alternation. In *Proc. 30th Annual ACM Symposium on the Theory of Computing (STOC)*,
746 pages 234–240, 1998. doi:10.1145/276698.276749.
- 747 **34** Steffen van Bergerem, Roland Guttenberg, Sandra Kiefer, Corto Mascle, Nicolas Waldburger,
748 and Chana Weil-Kennedy. Verification of population protocols with unordered data. In *Proc.*
749 *51st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages
750 156:1–156:20, 2024. doi:10.4230/LIPICS.ICALP.2024.156.
- 751 **35** Chana Weil-Kennedy. *Observation Petri Nets*. PhD thesis, Technical University of Munich,
752 Germany, 2023.
- 753 **36** Philipp Weis and Neil Immerman. Structure theorem and strict alternation hierarchy for FO²
754 on words. In *Proc. 21st International Workshop on Computer Science Logic (CSL)*, pages
755 343–357, 2007. doi:10.1007/978-3-540-74915-8_27.
- 756 **37** Hiroto Yasumi, Fukuhito Ooshita, and Michiko Inoue. Population protocols for graph class
757 identification problems. In *Proc. 25th International Conference on Principles of Distributed*
758 *Systems, (OPODIS)*, pages 13:1–13:19, 2021. doi:10.4230/LIPICS.OPODIS.2021.13.

A Appendix

A.1 Example of a PP[<] for the median language

Consider the *median language* $L = \bigcup_{n \in \mathbb{N}} \Sigma^n a \Sigma^n$. Let us describe a protocol $\mathcal{P} = (Q, \Sigma, O, \Delta)$ that decides L . The states are defined as

$$Q = \Sigma \times \{\blacktriangleleft, \blacktriangleright, \blacklozenge\} \times \{\perp, \top\}.$$

The components respectively represent the input letter; a belief on whether the center is on the right, here, or on the left; and a belief on the output.

We identify input a with state $(a, \blacktriangleright, \top)$, and each input $\sigma \neq a$ with $(\sigma, \blacktriangleright, \perp)$. We set $O((x, y, z)) = z$. The set Δ is defined by these rules, each describing a family of transitions:

<i>Population halving</i>		
(1)	$(x, \blacktriangleright, z), (x', \blacktriangleright, z')$	$\xrightarrow{\quad} (x, \blacktriangleleft, \perp), (x', \blacktriangleright, \perp)$
<i>Center finding</i>		
(2)	$(x, \blacktriangleright, z), (x', \blacktriangleleft, z')$	$\xrightarrow{\quad} (x, \blacktriangleleft, \perp), (x', \blacktriangleright, x' = a)$
(3)	$(x, \blacktriangleright, z), (x', \blacktriangleright, z')$	$\xrightarrow{\quad} (x, \blacktriangleright, x = a), (x', \blacktriangleright, \perp)$
<i>Output propagation</i>		
(4)	$(x, y, z), (x', \blacktriangleright, z')$	$\xrightarrow{\text{true}} (x, y, z'), (x', \blacktriangleright, z')$ for $y \in \{\blacktriangleleft, \blacktriangleright\}$
(5)	$(x, y, \top), (x', y', \perp)$	$\xrightarrow{\text{true}} (x, y, \perp), (x', y', \perp)$ for $y, y' \in \{\blacktriangleleft, \blacktriangleright\}$

By fairness, the first rule must be used until one or zero \blacktriangleright remains. Moreover, by fairness, the second and third rules will respectively move the \blacktriangleleft 's to the left, and the \blacktriangleright 's to the right. If some $(x, \blacktriangleright, y)$ remains, then, by fairness and the fourth rule, it will propagate its output y , which is \top iff $x = a$, by the choice of initial states and by rule (2-3). Otherwise, if the population is of even length, the fifth rule will be used to propagate \perp .

Figure 3 depicts all configurations reachable from the initial configuration $aab \in L$. Any fair run of a population protocol leads to a bottom strongly connected component of such a reachability graph. Thus, in this example, every fair runs leads to $a_{\blacktriangleleft}^{\top} a_{\blacktriangleright}^{\top} b_{\blacktriangleright}^{\top}$, which is \top -stable. Note that $a_{\blacktriangleleft}^{\perp} a_{\blacktriangleright}^{\perp} b_{\blacktriangleright}^{\perp}$ is a \perp -consensus, but is not \perp -stable.

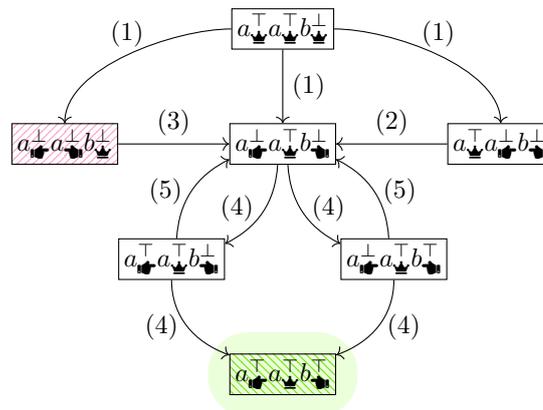


Figure 3 Configurations reachable from aab , where x_y^z stands for (x, y, z) . Self-loops arising from “no operation” transitions are omitted. The hatched nodes are consensus; the bottom one is stable.

XX:20 Population Protocols over Ordered Agents

778 Further observe that the reachability graph of Figure 3 has non-trivial cycles, for instance:

$$779 \quad a_{\perp}^{\perp} a_{\perp}^{\top} b_{\perp}^{\perp} \rightarrow a_{\perp}^{\perp} a_{\perp}^{\top} b_{\perp}^{\top} \rightarrow a_{\perp}^{\perp} a_{\perp}^{\top} b_{\perp}^{\perp} \rightarrow a_{\perp}^{\perp} a_{\perp}^{\top} b_{\perp}^{\top} \rightarrow \dots$$

780 Informally, the first and second agents are fighting to convince the third agent. However, by
781 fairness, this is not allowed to happen indefinitely. Eventually, the configuration $a_{\perp}^{\top} a_{\perp}^{\top} b_{\perp}^{\top}$, at
782 the bottom, is reached.

783 In the above specific protocol \mathcal{P} , for each initial configuration, its reachability graph has
784 a unique trivial bottom strongly connected component, made of one configuration of the
785 form $\perp^n \top^n$ or $\perp^n \perp^n$. However, in general, it needs not be unique or trivial. A fair run
786 becomes b -stable iff it visits a bottom strongly connected component whose configurations
787 are *all* b -consensuses.

788 In a decider, for a given initial configuration, *all* bottom strongly connected components
789 must consist only of stable configurations, all of the *same* output.

790 A.2 Missing proofs from Section 3.1



791 ► **Lemma 5.** *A language L is $\text{PP}[\mathcal{N}]$ -decidable iff L and its complement are $\text{PP}[\mathcal{N}]$ -semi-*
792 *decidable. The same holds for IO-PP[\mathcal{N}].*

793 **Proof.** \Rightarrow) This is immediate, since a protocol deciding L is a semi-decider for L , and the
794 complement protocol, obtained by swapping \top and \perp in the opinion function, is a semi-decider
795 for $\Sigma^+ \setminus L$.

796 \Leftarrow) This requires combining the two semi-deciding protocols $\mathcal{P}_+, \mathcal{P}_-$, for L and $\Sigma^+ \setminus L$
797 respectively, into a protocol \mathcal{P} . We let $\mathcal{P}_s = (Q_s, \Sigma, O_s, \Delta_s)$, for $s \in \{+, -\}$, and we define
798 $\mathcal{P} = (Q, \Sigma', O, \Delta)$ as follows.

799 We let $Q = Q_+ \times Q_- \times \{+, -\}$, with $\Sigma' = \{(a, a, s) \mid a \in \Sigma, s = + \text{ iff } a \in L\}$. We identify
800 $a \in \Sigma$ with the state $(a, a, s) \in Q$ where $s = +$ iff $a \in L$. By this identification, the protocol
801 \mathcal{P} recognizes a language over Σ and is correct on single agent runs by definition (since with
802 a single agent no transition can be taken). The third component is called the *belief* of the
803 agent. The function O maps each (q_+, q_-, s) to $O_+(q_+)$ if $s = +$, and $\neg O_-(q_-)$ otherwise.
804 Finally, the transitions of \mathcal{P} are of two kinds:

- 805 ■ (*Simulation.*) A transition of \mathcal{P}_+ or \mathcal{P}_- can be used, on the appropriate component Q_+
806 or Q_- . The two other components stay the same.
- 807 ■ (*Flip.*) If two agents have different beliefs, then one of them can flip its belief. The **true**
808 numerical predicate is used for these transitions. An agent in state (q_+, q_-, s) can also
809 flip its belief by itself if $O_s(q_s) = \perp$.

810 Consider $u \in \Sigma^+$ of length at least 2. Assume that $u \in L$, the case where $u \notin L$
811 being symmetric. Let $v \in \text{Post}^*(u)$. Let v_+ and v_- be the projections of v on Q_+ and Q_-
812 respectively. Observe that for both $s \in \{+, -\}$, v_s is reachable from u in \mathcal{P}_s . In consequence,
813 since $u \in L$, there is an execution of \mathcal{P}_+ from v_+ to a \top -stable w_+ . There is also an execution
814 of \mathcal{P}_- from v_- to a configuration w_- where at least one agent has opinion \perp . By following
815 both those executions one after the other, we can go from v to a configuration w in \mathcal{P} , so
816 that the projection of w on Q_+ and Q_- are w_+ and w_- .

- 817 ■ If the agent which has opinion \perp in w_- has opinion $-$, it can switch it to $+$.
- 818 ■ Then, we can flip all the other agents to $+$, by making them observe this agent. The
819 resulting configuration is \top -stable.

820 As mentioned, the case $u \notin L$ is symmetric, showing that \mathcal{P} is a decider. Furthermore,
821 observe that if both \mathcal{P}_+ and \mathcal{P}_- are *immediate observation* protocols, then so is \mathcal{P} . ◀

822 ▶ **Lemma 6.** *If $L_1, L_2 \subseteq \Sigma^+$ are PP[\mathcal{N}]-semi-decidable, then it is also the case for $L_1 \cap L_2$
823 and $L_1 \cup L_2$. This further holds for IO-PP[\mathcal{N}] and deciders.*

824 **Proof.** The proof is standard, and we will give, in Lemma 8, a slightly more intricate proof
825 with the same ideas, so we omit it here. ◀

826 A.3 Missing proofs from Section 3.2

827 ▶ **Proposition 7.** *The language $0^*1^* \dots k^*$ is IO-PP[\prec]-semi-decidable with stabilizing inputs.*

828 **Proof of correctness.** Note that \mathcal{P} is input-saving and immediate-observation. Let us show
829 that \mathcal{P} semi-decides L with stabilizing inputs. Let $u \in \Sigma^+$ and $u \rightsquigarrow^* v$. Let $v = w_0 \rightarrow w_1 \rightarrow$
830 \dots be a fair run. For the special case of $|v| = 1$, note that each state starts with output \top
831 and this remains so as no transition is ever enabled. Let us assume that $|v| \geq 2$.

832 *Case $\iota(v) \in L$.* The first rule is permanently disabled in w_0 . Hence, by fairness, the
833 second rule must swap the last component of each agent to \top . Once this happens, the
834 population has reached a \top -consensus, and the configuration cannot change anymore.

835 *Case $\iota(v) \notin L$.* Recall that the first component of each agent is fixed from w_0 onwards.
836 Let $i < j$ be positions such that $\iota(v)[i] > \iota(v)[j]$. Using the second rule, all agents can change
837 their opinion to \top , and then, using the first rule, agent i can change its opinion to \perp . By
838 fairness, this happens infinitely often, which means that the run does not stabilize. ◀

839 ▶ **Lemma 8.** *If $L_1, L_2 \subseteq \Sigma^+$ are PP[\mathcal{N}]-semi-decidable with stabilizing inputs, then it is also
840 the case for $L_1 \cap L_2$ and $L_1 \cup L_2$. This further holds for IO-PP[\mathcal{N}] and deciders.*

841 **Proof.** Let $\mathcal{P}_1 = (Q_1, \Sigma, \Delta_1, O_1)$ and $\mathcal{P}_2 = (Q_2, \Sigma, \Delta_2, O_2)$ be the protocols that respectively
842 PP[\mathcal{N}]-semi-decide L_1 and L_2 with stabilizing inputs.

843 **Intersection.** We simply take the product of both protocols; simulate them independently
844 with a common input component; and output the conjunction of their outputs.

845 Let $Q_1 = \Sigma \times R_1$ and $Q_2 = \Sigma \times R_2$. Formally, we define $\mathcal{P} = (Q, \Sigma, \Delta, O)$ by $R = R_1 \times R_2$,
846 $Q = \Sigma \times R$, $O((\sigma, q_1, q_2)) = O_1((\sigma, q_1)) \wedge O_2((\sigma, q_2))$, and the following rules:

$$847 \quad (\sigma, q_1, r), (\sigma', q_2, s) \xrightarrow{P} (\sigma, q_3, r), (\sigma', q_4, s) \quad \text{for } (\sigma, q_1), (\sigma', q_2) \xrightarrow{P} (\sigma, q_3), (\sigma', q_4) \in \Delta_1$$

$$848 \quad \text{and } r, s \in R_2,$$

$$849 \quad (\sigma, r, q_1), (\sigma', s, q_2) \xrightarrow{P} (\sigma, r, q_3), (\sigma', s, q_4) \quad \text{for } (\sigma, q_1), (\sigma', q_2) \xrightarrow{P} (\sigma, q_3), (\sigma', q_4) \in \Delta_2$$

$$850 \quad \text{and } r, s \in R_1.$$

851 We identify each $\sigma \in \Sigma$ with (σ, σ, σ) .

852 Note that \mathcal{P} is input-saving, and immediate-observation if it is the case of \mathcal{P}_1 and \mathcal{P}_2 .
853 Let us show that \mathcal{P} semi-decides $L = L_1 \cap L_2$ with stabilizing inputs. For every $w \in Q^+$,
854 let $\pi_1(w) \in Q_1^+$ be the projection of w onto its first two components, and let $\pi_2(w) \in Q_2^+$
855 be the projection of w onto its first and third components. Let $u \in \Sigma^+$ and $u \rightsquigarrow^* v$. Let
856 $v = w_0 \rightarrow w_1 \rightarrow \dots$ be a fair run from v . By definition of \mathcal{P} , the sequence $\pi_i(w_0), \pi_i(w_1), \dots$
857 is a fair run of \mathcal{P}_i . Moreover, $\iota(v) = \iota(\pi_1(v)) = \iota(\pi_2(v))$.

858 *Case $\iota(v) \in L$.* Since $\iota(v) \in L_1 \cap L_2$, the runs of \mathcal{P}_1 and \mathcal{P}_2 eventually stabilize to \top -stable
859 configurations. Since O is a conjunction, the same holds for the run of \mathcal{P} .

860 *Case $\iota(v) \notin L$.* Let $i \in \{1, 2\}$ be such that $\iota(v) \notin L_i$. The run of \mathcal{P}_i visits infinitely many
861 configurations containing a state $q \in Q_i$ with $O_i(q) = \perp$. Since O is a conjunction, the

XX:22 Population Protocols over Ordered Agents

862 same holds for the run of \mathcal{P} . Furthermore, if \mathcal{P}_i is a decider, then the run of \mathcal{P}_i eventually
863 stabilizes to \perp -stable configurations, which implies the same for the run of \mathcal{P} .

864 **Union.** We cannot simply take the previous construction and redefine the output mapping
865 to $O((\sigma, q_1, q_2)) = O_1((\sigma, q_1)) \vee O_2((\sigma, q_2))$. For example, if \mathcal{P}_1 and \mathcal{P}_2 have two agents with
866 outputs $\top\perp$ and $\perp\top$ respectively, then their disjunction yields $\top\top$, which is incorrect as
867 neither \mathcal{P}_1 nor \mathcal{P}_2 is in a \top -consensus.

868 Instead, we extend Q with an extra component $\{1, 2\}$ that indicates which output should
869 be used, and define $O((\sigma, q_1, q_2, i)) = O_i((\sigma, q_i))$. The transitions from Δ simply ignore this
870 new component. However, we need to add this (immediate-observation) rule:

$$871 \quad (\sigma, q_1, q_2, i), (\sigma', q'_1, q'_2, i') \xrightarrow{\text{true}} (\sigma, q_1, q_2, 3-i), (\sigma', q'_1, q'_2, i')$$

$$872 \quad \text{for } i \neq i' \vee O_i(q_i) = \perp \vee O_i(q'_i) = \perp.$$

874 This way, as long as agents have not agreed on a common choice, or as long as their choice is
875 not in a \top -consensus, they may change their mind.

876 Note that this is correct even for “deciders”: If \mathcal{P}_1 and \mathcal{P}_2 both reach \perp -stability, then
877 agents will indefinitely change their choice, but nonetheless steadily output \perp . ◀



878 ▶ **Lemma 9.** *Let $f: \Sigma \rightarrow 2^\Gamma$. If $L \subseteq \Sigma^+$ is PP[\mathcal{N}]-semi-decidable with stabilizing inputs,
879 then $f(L)$ is PP[\mathcal{N}]-semi-decidable with stabilizing inputs. This further holds for IO-PP[\mathcal{N}].*

880 **Proof.** Let $\mathcal{P} = (Q, \Sigma, \Delta, O)$ be the protocol that semi-decides L with stabilizing inputs.
881 Since \mathcal{P} is input-saving, its states are of the form $Q = \Sigma \times R$.

882 We provide a protocol \mathcal{P}' where, on input $u \in \Gamma^+$, the population internally starts with
883 an arbitrary word v such that $u \in f(v)$ and runs \mathcal{P} on v . With luck, it may be the case that
884 $v \in L$, but this needs not be the case. Therefore, the population may change its choice: If an
885 agent of \mathcal{P}' has input γ and encounters a \perp -state of \mathcal{P} , then it can change its internal input
886 of \mathcal{P} to any letter σ such that $\gamma \in f(\sigma)$.

887 Let us now describe the protocol formally. For every $\gamma \in \Gamma$, let σ_γ be an arbitrary letter⁴
888 of Σ such that $\gamma \in f(\sigma_\gamma)$. To handle the corner case of populations with a single agent, if
889 there is choice such that $\sigma_\gamma \in L$, then we take one. We define $\mathcal{P}' = (Q', \Gamma, \Delta', O')$ as follows:

- 890 ■ $Q' = \Gamma \times Q$;
- 891 ■ We identify each $\gamma \in \Gamma$ with (γ, σ_γ) ;
- 892 ■ $O'((\gamma, q)) = O(q)$;
- 893 ■ The transitions of Δ' are defined by

$$894 \quad (\gamma, q_1), (\gamma', q_2) \xrightarrow{P} (\gamma, q_3), (\gamma', q_4) \quad \text{for } \gamma, \gamma' \in \Gamma \text{ and } (q_1, q_2) \xrightarrow{P} (q_3, q_4) \in \Delta,$$

$$895 \quad (\gamma, q), (\gamma', q') \xrightarrow{\text{true}} (\gamma, \sigma), (\gamma', q') \quad \text{for } \gamma, \gamma' \in \Gamma, q, q' \in Q \text{ and } \sigma \in \Sigma \text{ such that}$$

$$896 \quad \gamma \in f(\sigma), \text{ and } O(q) = \perp \text{ or } O(q') = \perp.$$

897 The protocol is input-saving. The second rule is immediate-observation. Moreover, if \mathcal{P}
898 is immediate-observation, then it is also the case of the first rule. It remains to show that \mathcal{P}'
899 semi-decides $f(L)$ with stabilizing inputs. For all $w \in (Q')^+$, let $\pi(w)$ be the projection of w
900 onto its second component. Recall that $\iota(w)$ is the projection of w onto its first component.

⁴ If no such σ_γ exists, then we can simply map γ to a dummy state with opinion \perp , as no word containing γ belongs to $f(L)$. That being said, we will never invoke the lemma with such an f .

901 Let $u \in \Gamma^+$ and $u \rightsquigarrow^* v$. Let $v = w_0 \rightarrow w_1 \rightarrow \dots$ be a fair run. The case where $|v| = 1$, and
 902 hence $\iota(v) = \gamma \in \Gamma$, is trivially correct by the choice of σ_γ . Thus, suppose that $|v| \geq 2$.

903 *Case $\iota(v) \in f(L)$.* By hypothesis, there exists $v' \in L$ such that $\iota(v) \in f(v')$. For the
 904 sake of contradiction, suppose that no \top -stable configuration is visited by the fair run. By
 905 assumption, there exist indices $i_0 < i_1 < \dots$ such that $O'(w_{i_j}) = \perp$ for all $j \geq 0$. By
 906 definition of O' , this means that $O(\pi(w_{i_j})) = \perp$ for all $j \geq 0$. By fairness and the second
 907 rule, we can change the second component of the population to v' (and hence the input of \mathcal{P}
 908 to v'). Since \mathcal{P} semi-decides L with stabilizing inputs, we can reach a \top -stable configuration
 909 of \mathcal{P} in the second component. Consequently, the second rule becomes permanently disabled,
 910 which implies that \mathcal{P}' visits a \top -stable configuration, a contradiction.

911 *Case $\iota(v) \notin f(L)$.* For the sake of contradiction, suppose that there exists $j \geq 0$ such
 912 that $O'(w_j) = O'(w_{j+1}) = \dots = \top$. This means that $w_j \rightarrow w_{j+1} \rightarrow \dots$ only uses the first
 913 rule. By definition of O' , we have $O(\pi(w_j)) = O(\pi(w_{j+1})) = \dots = \top$. Since \mathcal{P} semi-decides
 914 L with stabilizing inputs, this implies that $\iota(\pi(w_j)) \in L$ and hence $\iota(w_j) \in f(L)$, which is a
 915 contradiction. \blacktriangleleft

916 A.4 Missing proofs from Section 4.1

917 \blacktriangleright **Lemma 11.** *Any protocol in $\text{PP}[\leq]$ is a well-structured transition system w.r.t. \preceq , that is,*
 918 *for all $u, u', w \in Q^*$ with $u \preceq w$ and $u \rightarrow u'$, there exists w' such that $u' \preceq w'$ and $w \rightarrow w'$.*
 919 *Further, the set of b -stable configurations is subword-closed and computable.*

920 **Proof.** In a $\text{PP}[\leq]$, we have $u \rightarrow u'$ if and only if there exist $a, b, c, d \in Q$ and $u_1, u_2, u_3 \in Q^*$
 921 such that $u = u_1 a u_2 b u_3$, $u' = u_1 c u_2 d u_3$ and $(a, b) \xrightarrow{\leq} (c, d) \in \Delta$. Let $w \in Q^*$ be such that
 922 $u \preceq w$. We can write w as $w = w_1 a w_2 b w_3$ with $u_1 \preceq w_1$, $u_2 \preceq w_2$ and $u_3 \preceq w_3$. We obtain
 923 $w \rightarrow w' = w_1 c w_2 d w_3$ and $u' \preceq w'$ as desired.

924 Let $U_0 = Q^* \setminus O^{-1}(b)^*$, the set of configurations which are *not* b -consensuses. For all i ,
 925 let $U_{i+1} = U_i \cup \{u \in Q^* \mid (\exists v \in U_i)[u \rightarrow v]\}$, the set of configurations that can reach U_0 in at
 926 most $i+1$ steps. We let $D_i = Q^* \setminus U_i$. The set of b -stable configurations is $D = Q^* \setminus \bigcup_{i \in \mathbb{N}} U_i$.

927 A key observation is that since our only predicate is $<$, all D_i are subword-closed. As a
 928 consequence, so is their intersection D . Furthermore, the sequence $(D_i)_{i \in \mathbb{N}}$ is a descending
 929 chain of subword-closed sets. Since \preceq is a well quasi-order, every such chain eventually
 930 stabilizes, hence there exists $i \in \mathbb{N}$ such that $D_i = D_{i+1}$. We can compute D_i for each index
 931 until we find one such that $D_i = D_{i+1}$. Such D_i is then the set of b -stable configurations. \blacktriangleleft

932 We first show that our protocols are robust to the insertion of repeated letters, under
 933 specific contexts:

934 \blacktriangleright **Lemma 12.** *In an IO-PP[$<$], if $u w v \rightarrow^* u' w' v'$ with $|u| = |u'|$ and $|v| = |v'|$, then for all*
 935 *z with $\alpha(z) \subseteq \alpha(w)$, there exists z' such that $u w z w v \rightarrow^* u' w' z' w' v'$ and $\alpha(z') \subseteq \alpha(w')$.*

936 **Proof.** We prove this for a single step, the result follows immediately by induction. Suppose
 937 there is a step from $u w v$ to $u' w' v'$, let $(a, b) \xrightarrow{\leq} (a', b')$ be the associated transition. Since the
 938 protocol is IO, we have $a = a'$ or $b = b'$. We assume $a = a'$, the other case being symmetric.

939 If the observing agent is in u or v , then $w = w'$ and we can go from $u w z w v$ to $u' w' z' w' v' =$
 940 $u' w z w v'$ with the same transition. If the observing agent is in w , then $u = u'$, $v = v'$ and,
 941 from $u w z w v$, we start by applying that transition for each b in z , observing the same agent
 942 labelled a as in the step from $u w v$ to $u' w' v'$. We thus turn the factor z into z' where all b are
 943 turned into b' . We can then apply that same transition twice to turn each of the two copies of
 944 w to w' . Note that $\alpha(z') = (\alpha(z) \cup \{b'\}) \setminus \{b\}$, and $\alpha(w') \supseteq (\alpha(w) \cup \{b'\}) \setminus \{b\} \supseteq \alpha(z')$ as w'

XX:24 Population Protocols over Ordered Agents

945 is obtained from w by deleting an occurrence of b and adding an occurrence of b' . We obtain
 946 $uw'z'w'v$, which is equal to $u'w'z'w'v'$ (since here $u = u'$ and $v = v'$) with $\alpha(z') \subseteq \alpha(w')$. ◀



947 ► **Lemma 13.** *Let $L \in \text{IO-PP}[\prec]$. There exists $m \geq 1$ such that, for all $w_1, \dots, w_m \in \Sigma^+$ and*
 948 *$z \in \Sigma^*$ with $\alpha(w_1) = \dots = \alpha(w_m) \supseteq \alpha(z)$, we have $w_1 \cdots w_m \equiv_L (w_1 \cdots w_m)z(w_1 \cdots w_m)$.*

949 **Proof.** Lemma 11 indicates that for both $b \in \{\top, \perp\}$, the set of b -stable configurations is
 950 subword-closed. As a consequence, writing B^ε for $B \cup \{\varepsilon\}$, it is a finite union of languages of the
 951 form $A_1^* B_1^\varepsilon \cdots A_k^* B_k^\varepsilon$ with A_i, B_i subsets of Q , the set of states of our protocol [18, Sect. 6.1.1].
 952 Let K_b be the maximal factorisation size k over all those languages, $K = \max(K_\top, K_\perp)$ and
 953 $m = 2K + 1$.

954 Let $w_1, \dots, w_m, z \in \Sigma^*$ be such that $\alpha(w_1) = \dots = \alpha(w_m) \supseteq \alpha(z)$, and let $u, v \in \Sigma^*$. We
 955 show that for both $b \in \{\top, \perp\}$, if we can reach a b -stable configuration from $uw_1 \cdots w_mv$ then
 956 it is also possible from $uw_1 \cdots w_mzw_1 \cdots w_mv$. This implies that $uw_1 \cdots w_mv$ is accepted if
 957 and only if $uw_1 \cdots w_mzw_1 \cdots w_mv$ is, proving the lemma.

958 Let $b \in \{\top, \perp\}$, suppose we can reach a b -stable configuration w' from $uw_1 \cdots w_mv$. Since
 959 steps are length-preserving, we can divide w' into $u'w'_1 \cdots w'_m v'$ where u', v' and each w'_i
 960 have the same length as u, v and each w_i , respectively.

961 Since w' is b -stable, there exists $k < K_b$ and an expression $A_1^* B_1^\varepsilon \cdots A_k^* B_k^\varepsilon$ whose language
 962 contains w' , and only contains b -stable words. Since $m > 2K_b \geq 2k$, there exist i, j such that
 963 $u'w'_1 \cdots w'_{i-1} \in A_1^* B_1^\varepsilon \cdots A_j^*$, $w'_{i+1} \cdots w'_m v' \in A_j^* B_j^\varepsilon \cdots B_k^\varepsilon$, and $w'_i \in A_j^*$.

964 Since $w = uw_1 \cdots w_mv \rightarrow^* u'w'_1 \cdots w'_m v' = w'$, by Lemma 12 there is a partial run from

$$965 \quad \tilde{w} := (uw_1 \cdots w_{i-1})w_i(w_{i+1} \cdots w_m)z(w_1 \cdots w_{i-1})w_i(w_{i+1} \cdots w_mv)$$

966 to

$$967 \quad \tilde{w}' := (u'w'_1 \cdots w'_{i-1})w'_i z' w'_i (w'_{i+1} \cdots w'_m v')$$

968 for some z' with $\alpha(z') \subseteq \alpha(w'_i)$, and $|z'| = |(w_{i+1} \cdots w_m)z(w_1 \cdots w_{i-1})|$.

969 Furthermore, since we have $u'w'_1 \cdots w'_{i-1} \in A_1^* B_1^\varepsilon \cdots A_j^*$, and $w'_{i+1} \cdots w'_m v' \in A_j^* B_j^\varepsilon \cdots B_k^\varepsilon$,
 970 and $z' \in \alpha(w'_i)^* \subseteq A_j^*$, we infer $\tilde{w}' \in A_1^* B_1^\varepsilon \cdots A_k^* B_k^\varepsilon$, hence \tilde{w}' is b -stable.

971 We have shown that if there is an execution from $uw_1 \cdots w_mv$ to a b -stable configuration,
 972 then there is also one from $uw_1 \cdots w_mzw_1 \cdots w_mv$ to a b -stable configuration, for all $u, v \in \Sigma^*$
 973 and both $b \in \{\top, \perp\}$. By definition of the language of \mathcal{P} , this means that $w_1 \cdots w_m$ and
 974 $w_1 \cdots w_mzw_1 \cdots w_m$ are equivalent with respect to \equiv_L . ◀



975 ► **Lemma 15.** *Every IO-PP[\prec]-decidable language is regular.*

976 **Proof.** We start by showing that some patterns, reminiscent of *sesquipowers* [22], are
 977 unavoidable in long strings. An m -reducible pattern is a word $w_1 \cdots w_mzw_1 \cdots w_m$ such that
 978 $\alpha(w_1) = \dots = \alpha(w_m) \supseteq \alpha(z)$. We have:

979 ▷ **Claim 36.** For all alphabet Q of size ℓ and all $m \in \mathbb{N}$, there exists $B(\ell, m) \in \mathbb{N}$ such that
 980 for all $w \in Q^*$, if $|w| \geq B(\ell, m)$ then w contains an m -reducible pattern.

981 **Proof.** We proceed by induction on the size ℓ of the alphabet Q . If Q is a singleton, then
 982 the result is clear with $B(1, m) = 2m + 1$.

983 Otherwise, define $B(\ell, m) = mB(\ell - 1, m)\ell^{mB(\ell-1, m)}$. Let $w \in Q^*$ be such that $|w| \geq$
 984 $B(\ell, m)$, let u be an infix of w of maximal length such that $|\alpha(u)| < \ell$, and let $M = |u|$.

985 ■ If $M \geq B(\ell - 1, m)$, then by induction hypothesis u contains an m -reducible pattern,
 986 thus so does w .

987 ■ Otherwise, we can split w into $r = B(\ell, m)/B(\ell - 1, m) = m^{\ell^{mB(\ell-1, m)}}$ factors of length
 988 $B(\ell - 1, m)$, plus a suffix v : $w = u_1 \cdots u_r v$. Since $M < B(\ell - 1, m)$, by definition of M ,
 989 each u_k contains all letters in Q . Since $r \geq m^{\ell^{mB(\ell-1, m)}}$, there must exist indices i, j such
 990 that $i+m < j$ and $u_i \cdots u_{i+m-1} = u_j \cdots u_{j+m-1}$. As all u_k contain the same set of letters,
 991 we obtain an m -reducible pattern $(u_i \cdots u_{i+m-1})(u_{i+m} \cdots u_{j-1})(u_j \cdots u_{j+m-1})$. ◀

992 By Claim 36 applied with m taken from Lemma 13, there is a bound B such that every
 993 word of length more than B is equivalent under \equiv_L to a shorter word. As a consequence,
 994 every word is equivalent to a word of length at most B .

995 This means that \equiv_L has finitely many equivalence classes, yielding the result by the
 996 Myhill-Nerode theorem. ◀

997 A.5 Missing proofs from Section 4.2

998 ▶ **Proposition 16.** *The language $\{w \in \Sigma^+ \mid |w|_a = 1\}$ is IO-PP[\emptyset]-semi-decidable with*
 999 *stabilizing inputs.*

1000 **Proof.** We define $\mathcal{P} = (Q, \Sigma, \Delta, O)$ by $Q = \Sigma \times \Sigma \times \{\top, \perp\}$, $O((\sigma, \sigma', o)) = (o \wedge \sigma = \sigma') \vee ((\sigma =$
 1001 $a \neq \sigma'))$, and the following rules:

$$1002 \quad (\sigma, a, o), q \xrightarrow{\text{true}} (\sigma, \sigma, \perp), q \quad \text{for } \sigma \neq a, \quad (2)$$

$$1003 \quad (a, \sigma, o), q \xrightarrow{\text{true}} (a, a, \top), q, \quad (3)$$

$$1004 \quad (\sigma, \sigma, o), (a, a, \top) \xrightarrow{\text{true}} (\sigma, \sigma, \top), (a, a, \top) \quad \text{for } \sigma \neq a, \quad (4)$$

$$1005 \quad (\sigma, \sigma, o), (\sigma', \sigma', \perp) \xrightarrow{\text{true}} (\sigma, \sigma, \perp), (\sigma', \sigma', \perp), \quad (5)$$

$$1006 \quad (a, a, o), (a, a, o') \xrightarrow{\text{true}} (a, a, \perp), (a, a, o'). \quad (6)$$

1007 We identify a with (a, a, \top) , and each $\sigma \in \Sigma \setminus \{a\}$ with (σ, σ, \perp) .

1008 Note that \mathcal{P} is input-saving and immediate-observation. Let us show that \mathcal{P} decides
 1009 the language of the statement with stabilizing inputs. Let $u \in \Sigma^+$ and $u \rightsquigarrow^* v$. Let
 1010 $v = w_0 \rightarrow w_1 \rightarrow \cdots$ be a fair run.

1011 If $|v| = 1$, then v is of the form (a, a, \top) , (a, σ', \perp) , (σ, a, \top) , or (σ, σ', \perp) where $\sigma, \sigma' \neq a$.
 1012 By definition of O , the first (resp. last) two states have output \top (resp. \perp) as desired.

1013 Suppose $|v| \geq 2$. By fairness and rules (2)–(3), we can assume without loss of generality
 1014 that each agent has its first two components equal and immutable along the fair run.

1015 *Case $|\iota(v)|_a = 0$.* Since each $\sigma \neq a$ is identified with (σ, σ, \perp) , and since an a can only
 1016 disappear via rule (2), there must exist a position i and $\sigma \neq a$ such that $v[i] = (\sigma, \sigma, \perp)$. By
 1017 fairness and rule (5), a \perp -consensus can be reached. From there, all rules are disabled.

1018 *Case $|\iota(v)|_a = 1$.* The single agent whose first component is a can set its third component
 1019 to \top with rule (3), and set the third component of other agents to \top with rule (4). From
 1020 there, all rules are disabled and hence a \top -stable configuration has been reached.

1021 *Case $|\iota(v)|_a \geq 2$.* By fairness, rules (3) and (6) cause the third component of a -agents to
 1022 alternate indefinitely between \top and \perp . ◀

1023 A.6 Missing proofs from Section 5.1

1024 ▶ **Lemma 19.** *Any language from $\Sigma_1^{\text{int}}[<, +, \equiv]$ is a finite union of languages of the form*
 1025 $K = \{a_0 w_1 a_1 \cdots w_m a_m \mid w_i \in \Sigma^*, a_i \in \Sigma, \varphi(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto |a_i|_\sigma)\}$ *where $m \geq 0$ and φ*
 1026 *is a Presburger formula over variables $\{x_{i,\sigma} \mid i \in [1..m], \sigma \in \Sigma\} \cup \{y_{i,\sigma} \mid i \in [0..m], \sigma \in \Sigma\}$.*

XX:26 Population Protocols over Ordered Agents

1027 **Proof.** Let $\varphi = (\exists x_1, \dots, x_m)[\psi] \in \Sigma_1^{\text{int}}[<, +, \equiv]$ where ψ is quantifier-free. Without loss of
 1028 generality, ψ may assume that $1 = x_0 < x_1 < \dots < x_m < x_{m+1} = \mathbf{max}$. Indeed,

- 1029 ■ For the first and last positions, we add the two extra variables with these two constraints,
 1030 and replace each other occurrence of 1 and \mathbf{max} by x_0 and x_{m+1} ;
- 1031 ■ For the ordering, we can change the formula so that it tests all orderings:

$$1032 \bigvee_{\pi \in S_m} (\exists x_0, x_1, \dots, x_m, x_{m+1})[x_1 \leq \dots \leq x_m \wedge \psi_\pi],$$

- 1033 where S_m is the set of permutations over $[1..m]$, and ψ_π is ψ with x_i replaced by $x_{\pi(i)}$;
- 1034 ■ For distinctness, we can similarly enumerate all possible equivalence classes of “=” and
 1035 associate a variable to each class, e.g., if we guess that $x_0 < x_1 = x_2 < x_3 = x_4 < x_5$,
 1036 then we change the subformula to $(\exists y_0, y_1, y_2, y_3)[\psi']$ where ψ' is the formula obtained by
 1037 renaming variables in ψ as follows: $x_0 \mapsto y_0$, $\{x_1, x_2\} \mapsto y_1$, $\{x_3, x_4\} \mapsto y_2$ and $x_5 \mapsto y_3$.

1038 Once variables are strictly ordered, we can further assume that each expression $\#_a(x_i, x_j)$
 1039 satisfies $j \in \{i, i+1\}$. Indeed, if $i < j$, then $\#_\sigma(x_i, x_j)$ can be substituted with

$$1040 \sum_{\ell=i}^{j-1} \#_\sigma(x_\ell, x_{\ell+1}) - \sum_{\ell=i+1}^{j-1} \#_\sigma(x_\ell, x_\ell).$$

1041 Let ψ' denote the formula ψ where each expression $\#_\sigma(x_i, x_{i+1})$ is replaced with $x_{i,\sigma}$,
 1042 and each $\#_\sigma(x_i, x_i)$ is replaced with $y_{i,\sigma}$. Consider $K = \{a_0 w_1 a_1 \dots w_n a_m \mid w_i \in \Sigma^*, a_i \in \Sigma,$
 1043 $\psi'(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto |a_i|_\sigma)\}$. We claim that $v \models (\exists x_0, \dots, x_m)[\psi]$ iff $v \in K$.

1044 \Rightarrow) We take $a_i = v[x_i]$ for each $i \in [0..m]$, and $w_i = v[x_{i-1} + 1..x_i - 1]$ for each $i \in [1..m]$.
 1045 By the assumptions on ψ , it is the case that $\psi'(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto |a_i|_\sigma)$ holds.

1046 \Leftarrow) We have $v = a_0 w_1 a_1 w_2 \dots w_m a_m$ where $w_i \in \Sigma^*$, $a_i \in \Sigma$ and $\psi'(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto$
 1047 $|a_i|_\sigma)$ holds. By taking $x_i = |a_0 w_1 \dots a_i|$ for $i \in [0..m]$, $\psi(x_0, \dots, x_m)$ holds. \blacktriangleleft

1048 \uparrow **Corollary 21.** $\Delta_1^{\text{int}}[<, +, \equiv] \subseteq \text{PP}[<]$.

1049 **Proof.** Let $L \in \Delta_1^{\text{int}}[<, +, \equiv]$. Since $L \in \Sigma_1^{\text{int}}[<, +, \equiv]$, by Proposition 20, L is $\text{PP}[<]$ -semi-
 1050 decidable. As the complement of L also belongs to $\Sigma_1^{\text{int}}[<, +, \equiv]$, the same applies. Thus, we
 1051 are done by Lemma 5. \blacktriangleleft

1052 \uparrow **Proposition 22.** *The language $(ab)^+$ does not belong to $\Sigma_1^{\text{int}}[<, +, \equiv]$.*

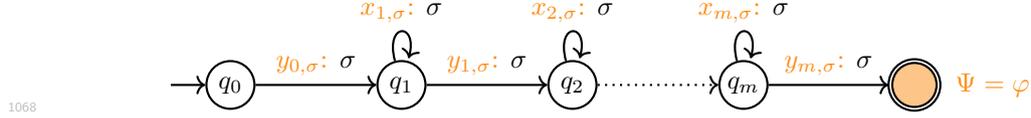
1053 **Proof.** To derive a contradiction, suppose that $(ab)^+$ belongs to $\Sigma_1^{\text{int}}[<, +, \equiv]$. By Lemma 19,
 1054 $(ab)^+ = K_1 \cup \dots \cup K_\ell$ where K_j is of the form $\{a_{j,0} w_{j,1} a_{j,1} \dots w_{j,m_j} a_{j,m_j} \mid w_{j,i} \in \Sigma^*, a_{j,i} \in$
 1055 $\Sigma, \varphi_j(x_{i,\sigma} \mapsto |w_{j,i}|_\sigma, y_{i,\sigma} \mapsto |a_{j,i}|_\sigma)\}$, $m_j \geq 0$ and φ_j is a Presburger formula.

1056 Let $v = (ab)^{n+1}$ where $n = \max(m_1, \dots, m_\ell)$. We have $v \in L$ and so $v \in K_j$ for some $j \in$
 1057 $[1..\ell]$. Let $v = a_{j,0} w_{j,1} a_{j,1} \dots w_{j,m_j} a_{j,m_j}$ be decomposed as in K_j . By $|v| = 2n + 2 > 2m_j + 1$
 1058 and the pigeonhole principle, there is $i \in [1..m_j]$ with $|w_{j,i}| \geq 2$. Let $w'_{j,i}$ be a word obtained
 1059 from $w_{j,i}$ by swapping two adjacent letters. Let v' be obtained from v by this change. As
 1060 the letter counts within $w_{j,i}$ and $w'_{j,i}$ are the same, we have $v' \in K_j$, a contradiction. \blacktriangleleft

1061 A.7 Missing proofs from Section 5.2

1062 \uparrow **Theorem 23.** *A language is recognized by some poPA iff it belongs to $\Sigma_1^{\text{int}}[<, +, \equiv]$.*

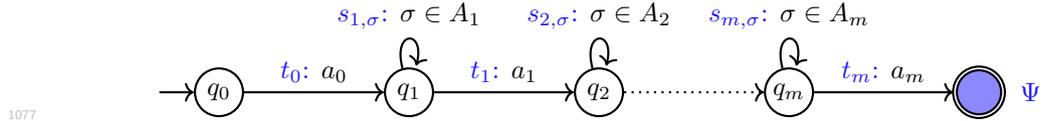
1063 **Proof.** \Leftarrow) Let $L \in \Sigma_1^{\text{int}}[\langle, +, \equiv]$. By Lemma 19, L can be written as a finite union of
 1064 languages of the form $K = \{a_0 w_1 a_1 \cdots w_m a_m \mid w_i \in \Sigma^*, a_i \in \Sigma, \varphi(x_{i,\sigma} \mapsto |w_i|_\sigma, y_{i,\sigma} \mapsto$
 1065 $|a_i|_\sigma)\}$ where $m \geq 0$ and φ is a Presburger formula. Since poPA are nondeterministic, they
 1066 are trivially closed under union. Thus, it suffices to give a poPA for K , which we do below.
 1067 Here, each transition labeled with “ σ ” stands for $|\Sigma|$ distinct transitions:



1069 \Rightarrow) Let \mathcal{A} be a poPA. Let K be the words of $L(\mathcal{A})$ of length at most one. We can
 1070 construct a formula for each word of K :

$$1071 \quad \varphi_\varepsilon = \bigwedge_{\sigma \in \Sigma} \#_\sigma(1, \mathbf{max}) = 0 \text{ and } \varphi_a = \#_a(1, \mathbf{max}) = 1 \wedge \bigwedge_{\sigma \neq a} \#_\sigma(1, \mathbf{max}) = 0.$$

1072 Let us now consider $K' = L(\mathcal{A}) \setminus K$. As \mathcal{A} is partially ordered, its language can be
 1073 described as the finite union of poPA organized as straight lines alternating between self-loops
 1074 and progress transitions. As we consider K' (and so words of length at least 2), we can
 1075 explicitly read the first and last letters, which yields automata of the following form, where
 1076 $a_i \in \Sigma$ and $A_i \subseteq \Sigma$:



1078 We convert such an automaton into this formula:

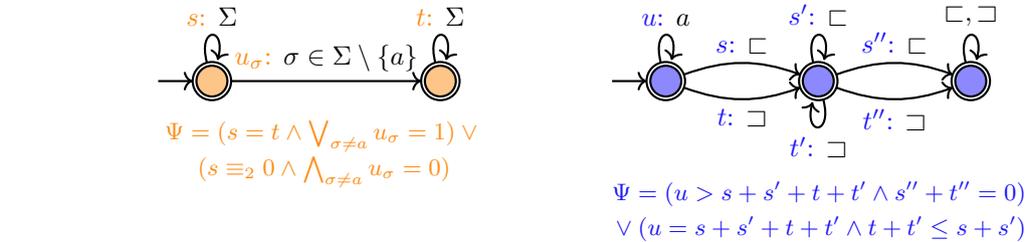
$$1080 \quad (\exists x_0, x_1, \dots, x_m)[1 = x_0 < x_1 < \dots < x_m = \mathbf{max}] \wedge \bigwedge_{i=0}^m \#_{a_i}(x_i, x_i) = 1 \wedge$$

$$1081 \quad \bigwedge_{i=1}^m \sum_{\sigma \in \Sigma \setminus A_i} \#_\sigma(x_{i-1} + 1, x_i - 1) = 0 \wedge \Psi(t_i \mapsto 1, s_{i,\sigma} \mapsto \#_\sigma(x_{i-1} + 1, x_i - 1)).$$

1082 We are done by taking the disjunction of all formulas. \blacktriangleleft

1083 \blacktriangleright **Observation 24.** If weakly-unambiguous poPA are closed under complement, then any
 1084 language recognized by a weakly-unambiguous poPA belongs to $\Delta_1^{\text{int}}[\langle, +, \equiv]$.

1085 **Proof.** First, observe that the two automata of Figure 2 for the median and coDyck-witness
 1086 languages can be complemented as follows, while remaining weakly unambiguous:



1088 In general, let \mathcal{A} and \mathcal{A}' be poPA with $L(\mathcal{A}') = \overline{L(\mathcal{A})}$. By Theorem 23, $L(\mathcal{A}), L(\mathcal{A}') \in$
 1089 $\Sigma_1[\langle, +, \equiv]$. Thus, $L(\mathcal{A}) = \overline{L(\mathcal{A}')} \in \Pi_1[\langle, +, \equiv]$ and so $L(\mathcal{A}) \in \Delta_1[\langle, +, \equiv]$. \blacktriangleleft

1090 **A.8 Missing proofs from Section 6.1**



1091 ► **Lemma 27.** *Every language semi-decided by a protocol in $\text{PP}[+1]$ is also semi-decided by*
 1092 *a protocol in $\text{IO-PP}[+1]$.*

1093 **Proof.** Let $\mathcal{P} = (Q, \Sigma, O, \Delta)$ be a protocol in $\text{PP}[+1]$. We construct a protocol $\mathcal{P}' =$
 1094 $(Q', \Sigma, O', \Delta')$ in $\text{IO-PP}[+1]$ such that $L(\mathcal{P}') = L(\mathcal{P})$.

1095 *States and outputs.* For each transition $\delta = (a, b) \rightarrow (c, d)$ in Δ , we introduce auxiliary states
 1096 a^δ and $b^{\text{ack},\delta}$. Formally, $Q' = Q \cup \{a^\delta, b^{\text{ack},\delta} \mid \delta = (a, b) \rightarrow (c, d) \in \Delta\}$. The output function
 1097 ignores the auxiliary markers: $O'(q) = O(q)$, $O'(a^\delta) = O(a)$ and $O'(b^{\text{ack},\delta}) = O(b)$.

1098 *Transitions.* For each $\delta = (a, b) \rightarrow (c, d) \in \Delta$, the set Δ' contains these four transitions:

$$\begin{aligned} 1099 & a, b \xrightarrow{+1} a^\delta, b \\ 1100 & a^\delta, b \xrightarrow{+1} a^\delta, b^{\text{ack},\delta} \\ 1101 & a^\delta, b^{\text{ack},\delta} \xrightarrow{+1} c, b^{\text{ack},\delta} \\ 1102 & c, b^{\text{ack},\delta} \xrightarrow{+1} c, d. \end{aligned}$$

1103 Each of these steps updates at most one of the two agents at a time, hence \mathcal{P}' is an immediate-
 1104 observation protocol. To avoid blocking in the presence of overlapping partial simulations,
 1105 we also include the following “clean-up” rules:

$$\begin{aligned} 1106 & a^\delta, x \xrightarrow{+1} a, x && \text{for } x \neq b^{\text{ack},\delta}, \\ 1107 & y, b^{\text{ack},\delta} \xrightarrow{+1} y, d && \text{for } y \neq a^\delta. \end{aligned}$$

1108 *Step simulation.*

1109 ► **Claim 37.** If $u, v \in Q^*$ and $u \rightarrow v$ in \mathcal{P} , then $u \rightarrow^* v$ in \mathcal{P}' .

1110 **Proof.** Let $\delta = (a, b) \xrightarrow{+1} (c, d)$ be the transition used in \mathcal{P} . Let $u = u_1 a b u_2$ and $v = u_1 c d u_2$.
 1111 In \mathcal{P}' , the four transitions associated with δ yield $u_1 a b u_2 \rightarrow u_1 a^\delta b u_2 \rightarrow u_1 a^\delta b^{\text{ack},\delta} u_2 \rightarrow$
 1112 $u_1 c b^{\text{ack},\delta} u_2 \rightarrow u_1 c d u_2$. ◀

1113 *Projection to \mathcal{P} .* We define a projection $f: Q'^* \rightarrow Q^*$ that interprets a \mathcal{P}' -configuration as a
 1114 \mathcal{P} -configuration. For $w = w_1 \cdots w_n$, define $f(w) = w'_1 \cdots w'_n$ by:

- 1115 ■ if $w_i \in Q$, then $w'_i = w_i$;
- 1116 ■ if $w_i = a^\delta$ for some $\delta = (a, b) \rightarrow (c, d)$ and $i < n$ with $w_{i+1} = b^{\text{ack},\delta}$, then $w'_i = c$;
- 1117 ■ if $w_i = a^\delta$ but the above condition fails, then $w'_i = a$;
- 1118 ■ if $w_i = b^{\text{ack},\delta}$ for some $\delta = (a, b) \rightarrow (c, d)$, then $w'_i = d$.

1119 ► **Claim 38.** Let $u \in Q^*$ and $v \in Q'^*$. If $u \rightarrow^* v$ in \mathcal{P}' , then $u \rightarrow^* f(v)$ in \mathcal{P} .

1120 **Proof.** We prove the claim by induction on the length of $u \rightarrow^* v$ in \mathcal{P}' . The only step that
 1121 can change the projected configuration is the acknowledgement step $(a^\delta, b) \xrightarrow{+1} (a^\delta, b^{\text{ack},\delta})$,
 1122 in which case f precisely applies the corresponding transition $\delta = (a, b) \xrightarrow{+1} (c, d)$ on the
 1123 projected word. All other rules leave $f(\cdot)$ unchanged. ◀

1124 ► **Claim 39.** For every $v \in Q'^*$ reachable in \mathcal{P}' , it is the case that $v \rightarrow^* f(v)$ in \mathcal{P}' .

1125 **Proof.** Starting from v , repeatedly apply the clean-up rules to remove every pending marker
 1126 a^δ not followed by the matching $b^{\text{ack},\delta}$, and then complete each matching pair $a^\delta b^{\text{ack},\delta}$ by the
 1127 two remaining simulation steps $(a^\delta, b^{\text{ack},\delta}) \xrightarrow{+1} (c, b^{\text{ack},\delta}) \xrightarrow{+1} (c, d)$. This yields $f(v)$. ◀

1128 We are now ready to prove the lemma. Let us show that $w \in L(\mathcal{P})$ iff $w \in L(\mathcal{P}')$.
 1129 \Rightarrow) Let $w \rightarrow^* u$ in \mathcal{P}' . Thanks to Claim 38, we know that $w \rightarrow^* f(u)$ in \mathcal{P} . Since
 1130 $w \in L(\mathcal{P})$, configuration $f(u)$ can reach a \top -stable configuration v in \mathcal{P} . Hence, thanks to
 1131 Claim 39 and Claim 37, we conclude that $u \rightarrow^* f(u) \rightarrow^* v$ in \mathcal{P}' . To derive a contradiction,
 1132 suppose that v is not \top -stable in \mathcal{P}' . Configuration v can reach a non- \top -consensus v' in
 1133 \mathcal{P}' . Consider v' to be such a configuration reachable in a minimal number of steps. The
 1134 last step must be a transition of the form $(a^\delta, b^{\text{ack},\delta}) \xrightarrow{+1} (c, b^{\text{ack},\delta})$ with $O'(c) = O(c) = \perp$,
 1135 or $(c, b^{\text{ack},\delta}) \xrightarrow{+1} (c, d)$, with $O'(d) = O(d) = \perp$. Hence, $f(v')$ is not a \top -consensus in \mathcal{P} .
 1136 Therefore, $f(v) = v \rightarrow^* f(v')$ in \mathcal{P} by Claim 38, contradicting the \top -stability of v in \mathcal{P} .
 1137 \Leftarrow) Let $w \rightarrow^* u$ in \mathcal{P} . By Claim 37, $w \rightarrow^* u$ in \mathcal{P}' . As $w \in L(\mathcal{P}')$, configuration u can
 1138 reach a \top -stable configuration v in \mathcal{P}' . From Claim 39, we have $v \rightarrow^* f(v)$ in \mathcal{P}' . Note that
 1139 $f(v)$ is \top -stable in \mathcal{P}' by construction. Moreover, from Claim 38, we have $u \rightarrow^* f(v)$ in \mathcal{P} .
 1140 To derive contradiction, suppose that $f(v)$ is not \top -stable in \mathcal{P} . We have $f(v) \rightarrow^* v'$ in \mathcal{P} for
 1141 some non- \top -consensus v' . Thanks to Claim 37, we obtain $f(v) \rightarrow^* v'$ in \mathcal{P}' , contradicting
 1142 the \top -stability of $f(v)$ in \mathcal{P}' . Hence, $f(v)$ is \top -stable in \mathcal{P} , and so $w \in L(\mathcal{P})$. \blacktriangleleft

1143 A.9 Missing proofs from Section 6.2

1144 \uparrow **► Lemma 28.** *The language of a linear-bounded nondeterministic Turing machine is PP[+1]-*
 1145 *semi-decidable.*

1146 **Proof.** Let M be a linear-bounded nondeterministic Turing machine over an input alphabet
 1147 Σ containing fixed endmarkers L and R , which are never overwritten. Thus, any accepted
 1148 input has the form LwR with no endmarker within w . We define M_{rst} as M equipped with a
 1149 *reset* operation. Each non-endmarker cell stores a pair (x, a) consisting of its current symbol
 1150 x and its initial input symbol a . Initially, M_{rst} replaces each $a \in \Sigma \setminus \{L, R\}$ by (a, a) . From
 1151 any non-accepting state, the machine may enter a reset state, scan the tape and restore every
 1152 cell (x, a) to (a, a) , return to L , and restart in the initial state. Moreover, upon reaching
 1153 an accepting state, M_{rst} performs a final sweep that overwrites the tape with a designated
 1154 accepting symbol f .

1155 **► Lemma 40.** *For every $w \in \Sigma^*$, w is accepted by M if and only if for every configuration*
 1156 *u reachable from w in M_{rst} , there exists a run of M_{rst} from u to an accepting configuration.*

1157 **Notations.** Let S be the set of states of M_{rst} , let Γ be its tape alphabet, and let $\Sigma \subseteq \Gamma$
 1158 be its input alphabet. Write s_0 for the initial state, $F \subseteq S$ for the set of accepting states,
 1159 and $T \subseteq S \times \Gamma \times \{\triangleright, \triangleleft\} \times S \times \Gamma$ for the transition relation, where “ \triangleright ” and “ \triangleleft ” respectively
 1160 denote left and right head movements. Let $\Gamma_w = \Gamma \setminus \{L, R\}$ denote the set of non-endmarker
 1161 symbols. An M_{rst} -configuration is a word over Γ with exactly one head marker, i.e., an
 1162 element of

$$1163 \bigcup_{s \in S} (L, s)\Gamma_w^*R \cup \bigcup_{x \in \Gamma_w, s \in S} L\Gamma_w^*(x, s)\Gamma_w^*R \cup L\Gamma_w^* \bigcup_{s \in S} (R, s).$$

1164 **Protocol \mathcal{P}_M .** Let us first define a protocol \mathcal{P}'_M in PP[+1] which, for all $w = w_1 \cdots w_n \in$
 1165 Σ^* and $w' = (\underline{L}, w_1, w_1)(w_2, w_2) \cdots (w_{n-1}, w_{n-1})(w_n, w_n, \bar{R})$, accepts w' iff M accepts w .
 1166 Intuitively, \mathcal{P}'_M takes as input a word w where the leftmost and rightmost agents know that
 1167 they are, and are in charge of simulating both themselves and the endmarker cells.

1168 We will later expand \mathcal{P}'_M into a protocol \mathcal{P}_M , with rules which allow for agents to guess
 1169 if they are first or last; in \mathcal{P}_M , a word w will keep resetting itself until agents have rightfully
 1170 guessed if they are first or last.

XX:30 Population Protocols over Ordered Agents

1171 *State space.* Let $TS = \Gamma \cup (S \times \Gamma)$ be the state of possibly head-marked tape symbols. We
 1172 use the following state forms:

- 1173 ■ the machine input alphabet Σ ;
- 1174 ■ internal cell states (b, γ) where $b \in \Sigma$ is the stored input letter and $\gamma \in TS$ is the current
 1175 tape symbol;
- 1176 ■ two-sided boundary states $(\underline{\gamma}, b, \gamma', \overline{\gamma''})$; an agent in this state guesses that it is both the
 1177 first and only agents. It stores its internal cell state as b, γ' , and the state of the left
 1178 and right endmarker as γ and γ'' . The notations $\underline{\gamma}$ and $\overline{\gamma''}$ allows us to distinguish easily
 1179 where the endmarker states are stored.
- 1180 ■ left-boundary states $(\underline{\gamma}, b, \gamma')$ and right-boundary states $(b, \gamma, \overline{\gamma'})$ should be interpreted in
 1181 the same way.
- 1182 ■ reset tokens rst-left^b and rst-right^b .

1183 *Encoding.* We encode the machine configuration $x_L x_1 \cdots (x_i, s) \cdots x_n x_R$ reachable on input
 1184 $b_1 \cdots b_n$ as the protocol configuration

$$1185 \quad (\underline{x_L}, b_1, x_1) (b_2, x_2) \cdots (b_i, (x_i, s)) \cdots (b_n, x_n, \overline{x_R}).$$

1186 Rules of \mathcal{P}'_M .

1187 *Starting the simulation.* An agent currently carrying the left boundary component may
 1188 initiate the simulation by placing the head on its tape symbol:

$$1189 \quad (\underline{L}, *, *) , * \xrightarrow{\text{true}} ((L', s_0), *, *) , *.$$

1190 We introduce symbol L' to prevent a same left endmarker to start several machine execution.

1191 *Simulation of the machine.* We simulate M_{rst} on the tape components, leaving the stored
 1192 input letters unchanged. Below, $*$ stands for an arbitrary boundary component (possibly
 1193 absent). For each transition $(q, x, \triangleright, p, y) \in T$, the set Δ contains:

$$1194 \quad ((q, x), a, z, *) , * \xrightarrow{\text{true}} (\underline{y}, a, (p, z), *) , *$$

$$1195 \quad (*, a, (q, x), \overline{z}), * \xrightarrow{\text{true}} (*, a, y, \overline{(p, z)}), *$$

$$1196 \quad (*, a, (q, x)), (b, z) \xrightarrow{+1} (*, a, y), (b, (p, z)).$$

1197 For transitions of the form $(q, x, \triangleleft, p, y) \in T$, the set Δ contains analogous transitions.

1198 *Final sweeping.* Let $Acc = \{f\} \cup (F \times \Gamma)$. We further include these rules:

$$1199 \quad (*, a, \gamma), (b, \gamma', *) \xrightarrow{+1} (*, a, f), (b, f, *) \quad \text{if } \gamma \text{ or } \gamma' \in Acc,$$

$$1200 \quad (\underline{\gamma}, a, \gamma', *) , * \xrightarrow{\text{true}} (\underline{f}, a, f, *) , * \quad \text{if } \gamma \text{ or } \gamma' \in Acc,$$

$$1201 \quad (*, a, \gamma, \overline{\gamma'}), * \xrightarrow{\text{true}} (*, a, f, \overline{f}), * \quad \text{if } \gamma \text{ or } \gamma' \in Acc.$$

1202 *Opinion function.* We define O as follows:

$$1203 \quad O(b, \gamma) = \top \iff \gamma \in Acc,$$

$$1204 \quad O(\underline{\gamma}, a, \gamma') = \top \iff \gamma, \gamma' \in Acc,$$

$$1205 \quad O(a, \gamma, \overline{\gamma'}) = \top \iff \gamma, \gamma' \in Acc,$$

$$1206 \quad O(\underline{\gamma}, a, \gamma', \overline{\gamma''}) = \top \iff \gamma, \gamma', \gamma'' \in Acc,$$

1207 where all other states have opinion \perp .

1208 **Expanding \mathcal{P}'_M into \mathcal{P}_M .**

1209 *Initialization.* On input letter $a \in \Sigma \setminus \{L, R\}$, an agent initially guesses it is both the first
1210 and the last cell, and stores its input symbol:

$$1211 \quad a, * \xrightarrow{\text{true}} (\underline{L}, a, a, \overline{R}), *$$

1212 *Detecting an inconsistency.* Below, “ \square ” stands for an arbitrary tape symbol in TS . Whenever
1213 two adjacent agents carry incompatible boundary information, we enter a reset mode:

$$\begin{array}{ll}
 1214 & (b, \square), (\square, a, \square) \xrightarrow{+1} \text{rst-left}^b, \text{rst-right}^a & \text{(mismatch on the left)} \\
 1215 & (\square, b, \square), (\square, a, \square) \xrightarrow{+1} (\underline{L}, b, b), \text{rst-right}^a & \text{(two left-boundaries)} \\
 1216 & (a, \square, \overline{\square}), (b, \square) \xrightarrow{+1} \text{rst-left}^a, \text{rst-right}^b & \text{(mismatch on the right)} \\
 1217 & (a, \square, \overline{\square}), (b, \square, \overline{\square}) \xrightarrow{+1} \text{rst-left}^a, (b, b, \overline{R}) & \text{(two right-boundaries)} \\
 1218 & (\square, a, \square, \overline{\square}), (\square, b, \square, \overline{\square}) \xrightarrow{+1} (\underline{L}, a, a), (b, b, \overline{R}) & \text{(two-sided vs two-sided)} \\
 1219 & (\square, a, \square, \overline{\square}), (\square, b, \square) \xrightarrow{+1} (\underline{L}, a, a), \text{rst-right}^b & \text{(two-sided vs left)} \\
 1220 & (a, \square, \overline{\square}), (\square, b, \square, \overline{\square}) \xrightarrow{+1} \text{rst-left}^a, (b, b, \overline{R}) & \text{(right vs two-sided)} \\
 1221 & (a, \square, \overline{\square}), (\square, b, \square) \xrightarrow{+1} \text{rst-left}^a, \text{rst-right}^b & \text{(right vs left)}.
 \end{array}$$

1222 *Reset propagation.* A right-moving token restores the stored input symbol and moves one
1223 step to the right:

$$\begin{array}{ll}
 1224 & \text{rst-right}^a, (*, b, \square) \xrightarrow{+1} (a, a), \text{rst-right}^{(b)} & \text{(propagate right)} \\
 1225 & \text{rst-right}^a, \text{rst-left}^b \xrightarrow{+1} \text{rst-left}^b, \text{rst-right}^a & \text{(tokens cross)}.
 \end{array}$$

1226 The left-moving reset token rst-left is symmetric.

1227 *End of propagation.* When rst-right (resp. rst-left) reaches a right (resp. left) boundary, it
1228 restores the right-boundary (resp. left-boundary) component:

$$\begin{array}{l}
 1229 \quad \text{rst-right}^a, (*, b, \square, \overline{\square}) \xrightarrow{+1} (a, a), (b, b, R) \\
 1230 \quad (\square, b, \square, *) \xrightarrow{+1} \text{rst-left}^a \xrightarrow{+1} (\underline{L}, b, b), (a, a).
 \end{array}$$

1231 ► **Lemma 41.** \mathcal{P}_M is a semi-decider of $L(M)$.

1232 **Proof.** Let $w \in \Sigma^*$. It suffices to show that $w \in L(M)$ if and only this condition holds:

1233 (P1) For every partial run of \mathcal{P}_M from w reaching a configuration c ,
there exists a run from c to a \top -stable configuration.

1234 On the other hand, by Lemma 40, we have that $w \in L(M)$ if and only if:

1235 (P2) For every configuration u reachable from w in M_{rst} , there
exists a run of M_{rst} from u to an accepting configuration.

1236 We prove that (P1) and (P2) are equivalent.

1237 **(P1) \Rightarrow (P2).** Assume (P1), and let u be any configuration reachable from w in M_{rst} . Write
1238 $w = a_1 \cdots a_n$. Consider the following run of \mathcal{P}_M :

$$\begin{array}{l}
 1239 \quad a_1 \cdots a_n \rightarrow^+ (\underline{L}, a_1, a_1, \overline{R}) \cdots (\underline{L}, a_n, a_n, \overline{R}) \\
 1240 \quad \rightarrow (\underline{L}, a_1, a_1) (a_2, a_2, \overline{R}) \cdots (\underline{L}, a_n, a_n, \overline{R}) \\
 1241 \quad \rightarrow (\underline{L}, a_1, a_1) \text{rst-left}^{a_2} (a_3, a_3, \overline{R}) \cdots (\underline{L}, a_n, a_n, \overline{R})
 \end{array}$$

XX:32 Population Protocols over Ordered Agents

$$\begin{aligned} &\rightarrow^+ (\underline{L}, a_1, a_1) \text{rst-left}^{a_2} \text{rst-left}^{a_3} \cdots \text{rst-left}^{a_{n-1}} (a_n, \overline{R}) \\ &\rightarrow^+ (\underline{L}, a_1, a_1) (a_2, a_2) \cdots (a_{n-1}, a_{n-1}) (a_n, a_n, \overline{R}). \end{aligned}$$

Let w' denote the last configuration above. From it, we follow the simulation rules of \mathcal{P}'_M so as to reproduce a machine run from w to u in M_{rst} , thus reaching a protocol configuration \tilde{u} that encodes u . By (P1), there exists a partial run from \tilde{u} to a \top -stable configuration.

By construction, from \tilde{u} no inconsistency can be detected (the boundary components are consistent and no reset token is present), and the simulation has already been started on the left. Hence, every enabled transition from \tilde{u} is a simulation step of M_{rst} . Therefore, the above fair run projects to a run of M_{rst} starting from u . Since the target configuration is \top -stable, the machine must eventually enter an accepting state and perform its final sweeping phase, yielding a configuration in which all tape symbols are f . In particular, u reaches an accepting configuration in M_{rst} , proving (P2).

(P2) \Rightarrow (P1). Assume (P2). Let $c = c_1 \cdots c_n$ be any protocol configuration reachable from the input $w = a_1 \cdots a_n$. We show that there exists a fair continuation from c to a \top -stable configuration. First, if c still contains some agent in a state $a \in \Sigma$, we turn them into symbols $(\underline{L}, a, a, \overline{R})$. Observe that, once initialized, the leftmost (resp. rightmost) agent can never lose its left-boundary (resp. right-boundary) information. Furthermore, using repeatedly the reset-propagation rules (namely the rules moving **rst-right** one step to the right and **rst-left** one step to the left, together with the token-crossing rule), we can reach a configuration with no reset token.

Case 1: c contains wrong boundary components. We first eliminate them by repeated resets. Let $N(c)$ be the number of *wrong* boundary components, namely

$$\begin{aligned} N(c) = & \left| \{1 < i \leq n \mid c_i \text{ carries a left-boundary component } \underline{\gamma}\} \right| + \\ & \left| \{1 \leq i < n \mid c_i \text{ carries a right-boundary component } \overline{\gamma}\} \right|. \end{aligned}$$

We claim that if $N(c) > 0$, then c has a continuation to some configuration c' such that $N(c') < N(c)$, and, moreover, the whole segment between the closest left boundary to the left, and the closest right boundary to the right is reset (i.e., its tape components become of the form (a, a)).

To prove the claim, pick a wrong left boundary component; that is, choose a position $i > 1$ such that c_i carries $\underline{\gamma}$. Then c_i can interact with its predecessor at position $i - 1$, yielding a successor configuration

$$c_1 \cdots c_{i-2} c'_{i-1} c'_i c_{i+1} \cdots c_n.$$

By inspection of the inconsistency rules of \mathcal{P}_M , the pair (c'_{i-1}, c'_i) is necessarily of one of the following forms:

- **rst-left^y, rst-right^z**: we propagate **rst-left** to the nearest left boundary and **rst-right** to the nearest right boundary, resetting the whole intermediate segment;
- **(\underline{s}, b, b), rst-right^z**: here position $i - 1$ is the nearest left boundary to the left of i , and propagating **rst-right** to the nearest right boundary resets the segment to its right;
- **(\underline{s}, a, a), ($b, b, \overline{s'}$)**: in this case both boundaries are restored immediately around the cut, and the enclosed segment is reset.

In all cases, after the (finite) propagation of reset tokens, we reach a configuration c' such that $N(c') < N(c)$ and the affected segment has been reset to (a, a) symbols. Repeating this

1284 process, we eventually reach a configuration with $N(\cdot) = 0$, hence with correct boundary
1285 placement. In particular, we can reach a configuration of the form

$$1286 \quad c_{\text{init}} = (\underline{L}, a_1, a_1) (a_2, a_2) \cdots (a_{n-1}, a_{n-1}) (a_n, a_n, \overline{R}).$$

1287 From c_{init} , the protocol can simulate, step by step, any run of M_{rst} on input w (using
1288 only simulation rules, as no inconsistency rule is enabled anymore). By (P2), from the corre-
1289 sponding machine configuration there exists a run to an accepting configuration; simulating
1290 it yields a protocol run to a configuration in which the machine is accepting and the sweeping
1291 phase has written f everywhere, i.e., a configuration of the form

$$1292 \quad (\underline{\gamma}_0, a_1, \gamma_1) (a_2, \gamma_2) \cdots (a_{n-1}, \gamma_{n-1}) (a_n, \gamma_n, \overline{\gamma_{n+1}}) \quad \text{with } \gamma_i \in \text{Acc for all } i.$$

1293 From such a configuration, no inconsistency can be detected and no simulation step is enabled,
1294 hence it is terminal; since every agent has opinion \top , it is \top -stable. Hence, $w \in L(\mathcal{P}_M)$.

1295 *Case 2: c has correct boundary placement.* Assume c has a unique left-boundary component
1296 at position 1 and a unique right-boundary component at position n . We can write

$$1297 \quad c = (\underline{\gamma}_0, a_1, \gamma_1) c' (a_n, \gamma_n, \overline{\gamma_R}),$$

1298 where the middle factor c' carries no boundary component. Using repeatedly the reset-
1299 propagation rules (namely the rules moving rst-right one step to the right and rst-left one
1300 step to the left, together with the token-crossing rule), we can reach a configuration with no
1301 reset token, of the form

$$1302 \quad v = (\underline{\gamma}_L, a_1, \gamma_1) (a_2, \gamma_2) \cdots (a_{n-1}, \gamma_{n-1}) (a_n, \gamma_n, \overline{\gamma_R}).$$

1303 We claim that v contains *at most one* head-marked tape symbol (q, x) . Assume for
1304 contradiction that v contains two head markers. Since the only rules that *create* a head
1305 marker are the machine-start rules at the left boundary (of the form $(\underline{L}, *) \rightarrow ((L', s_0), *)$
1306 and the subsequent simulation rules), this can only happen in one of the following ways:

- 1307 1. two *distinct* positions carrying a left-boundary component have started a machine execu-
1308 tion (possibly before some boundary components were later removed), or
- 1309 2. the *same* left-boundary position has started the execution twice.

1310 We call them *starters*.

1311 (1) *Two distinct starters.* Let p be the rightmost starter. The corresponding head marker
1312 must lie at or to the right of p (it is created at p and then moves by simulation rules). Since v
1313 has correct boundary placement, the left-boundary component that was present at p at start
1314 time must have been removed at some point. By the definition of the inconsistency-detection
1315 rules, removing a left-boundary component that is not at position 1 can only occur via one
1316 of the following rule patterns:

- 1317 a) a rule that creates a right-moving reset token, i.e., a transition whose right-hand side
1318 contains rst-right , or
- 1319 b) a rule that restores a right boundary locally, i.e., a transition whose right-hand side
1320 contains a state of the form $(\square, \square, \overline{R})$.

1321 In case (a), the created rst-right token is propagated until it reaches the next right boundary
1322 using the reset-propagation rules, and the end-of-propagation rule applies at that boundary.
1323 Along this propagation, the whole traversed segment is reset, and in particular any head
1324 marker within that segment is removed (the simulated head never crosses a right boundary

XX:34 Population Protocols over Ordered Agents

1325 component). In case (b), the starter position becomes (or is adjacent to) a right boundary,
 1326 hence the head marker cannot move past it by any simulation rule, and it is eliminated when
 1327 that boundary-restoration rule fires. In both subcases, the head marker created by p cannot
 1328 survive up to v , a contradiction.

1329

1330 (2) *The same starter starts twice.* For the same left-boundary position to start twice, it must
 1331 have been reset in between (otherwise the machine-start rule would not be enabled again at
 1332 that position). By construction, any reset affecting the left boundary can only occur through
 1333 one of the following rule families:

- 1334 a) creation of a **rst-right** token at or next to the left boundary (a rule whose right-hand side
 1335 contains **rst-right** ^{a}),
- 1336 b) a two-sided inconsistency rule producing both restored boundaries in one step (a rule
 1337 whose right-hand side contains both $(\underline{L}, \square, \square)$ and $(\square, \square, \overline{R})$),
- 1338 c) interaction between the left boundary and a left-moving token, i.e., a rule of the form
 1339 $(\square, \square, \square), \text{rst-left}^a \xrightarrow{+1} (\underline{L}, \square, \square), (a, a)$.

1340 We show that in each case, any existing head marker is deleted.

1341 (a) **Creation of rst-right.** Consider the step that creates **rst-right** ^{a} at the left boundary. Imme-
 1342 diately before this step, the head marker cannot lie strictly to the left of that boundary. If it
 1343 lay strictly to the right while the boundary component still existed, then to reach the left side
 1344 it would have to cross a left boundary component by a simulation rule, which is impossible
 1345 since simulation rules never move the head across boundary components. Therefore the head
 1346 marker must be at the boundary position (or in the reset segment adjacent to it) and is
 1347 removed when the reset is created and propagated.

1348

1349 (b) **Two-sided inconsistency.** If a two-sided inconsistency rule fires and produces a restored
 1350 left boundary $(\underline{L}, \square, \square)$ and a restored right boundary $(\square, \square, \overline{R})$ in one step, then the head
 1351 marker cannot lie beyond either boundary (simulation rules never cross boundary compo-
 1352 nents). Hence the head marker lies inside the segment that is reset by that rule and is deleted.

1353

1354 (c) **Interaction with rst-left.** Before reaching the left boundary, a token **rst-left** is created by
 1355 an inconsistency involving a right boundary component (i.e., by one of the inconsistency-
 1356 detection rules whose right-hand side contains **rst-left** ^{a}). During the subsequent propagation
 1357 to the left (by the reset-propagation rules), the token resets every visited position. Since
 1358 the head marker cannot cross a right boundary component, it must lie within the segment
 1359 traversed by **rst-left**, and is therefore deleted when the token passes.

1360 In all cases we reach a contradiction. Therefore, v contains at most one head marker.

1361 *Restarting a clean simulation.* From v , we can use the internal reset mechanism of M_{rst} (as
 1362 simulated by \mathcal{P}_M) to reset the tape contents to the stored input and restart from the initial
 1363 machine configuration

$$1364 \quad ((\underline{L}, s_0), a_1, a_1) (a_2, a_2) \cdots (a_{n-1}, a_{n-1}) (a_n, a_n, \overline{R}),$$

1365 provided the head is not already in an accepting state. If it is already accepting, then by
 1366 our assumption on M_{rst} , the machine performs its sweeping phase and writes f everywhere,
 1367 yielding \top -stable configuration of the form

$$1368 \quad (\underline{\gamma}_0, a_1, \gamma_1) (a_2, \gamma_2) \cdots (a_{n-1}, \gamma_{n-1}) (a_n, \gamma_n, \overline{\gamma_{n+1}}) \quad \text{with } \gamma_i \in \text{Acc for all } i$$

1369 Otherwise, starting from the clean initial configuration, we simulate (step by step) a
 1370 run of M_{rst} that reaches an accepting configuration, which exists by (P2). Simulating its
 1371 sweeping phase again yields a \top -stable configuration as above. Hence, $w \in L(\mathcal{P}_M)$. ◀

1372

1373 A.10 Missing proofs from Section 7.1

↑ 1374 ▶ **Lemma 30.** *The emptiness problem for $\text{PP}[\mathcal{N}]$ (resp. $\text{IO-PP}[\mathcal{N}]$) reduces to deciding the
 1375 complement of the syntax of $\text{PP}[\mathcal{N}]$ (resp. $\text{IO-PP}[\mathcal{N}]$) deciders.*

1376 **Proof.** Take a protocol $\mathcal{P} = (Q, \Sigma, O, \Delta)$. We build a protocol \mathcal{P}' such that \mathcal{P}' is not a
 1377 decider if and only if a \top -stable configuration is reachable in \mathcal{P} .

1378 If Σ contains a letter a with opinion \top , then the configuration a is \top -stable and reachable.
 1379 In that case, we set \mathcal{P}' as an arbitrary protocol which is not a decider.

1380 Otherwise, we build $\mathcal{P}' = (Q', \bar{\Sigma}, O', \Delta')$ from \mathcal{P} as follows: For each initial state $a \in \Sigma$
 1381 in \mathcal{P} , we add another state \bar{a} with opinion \perp . We denote $\bar{\Sigma}$ the set of those additional states.
 1382 We also add another state q_{\perp} with opinion \perp . The set of initial states of \mathcal{P}' is $\bar{\Sigma}$. For each
 1383 $\bar{a} \in \bar{\Sigma}$ and $q \in Q'$ we add a transition $(\bar{a}, q) \xrightarrow{\text{true}} (a, q)$. As a consequence, from every initial
 1384 configuration $\bar{a}_1 \cdots \bar{a}_k$ with $k \geq 2$, we can reach its counterpart in $a_1 \cdots a_k$ in Σ^* .

1385 Further, for every pair of states q_1, q_2 such that at least one of the two has opinion
 1386 \perp , we add a transition $(q_1, q_2) \xrightarrow{\text{true}} (q_1, q_{\perp})$. Note that by this last transition, from every
 1387 configuration of length ≥ 2 which is not a \top -consensus, we can reach a configuration of q_{\perp}^* .

1388 Suppose that there is a run in \mathcal{P} from an initial configuration $u = a_1 \cdots a_k \in \Sigma^+$ to a
 1389 \top -stable configuration v in \mathcal{P} . We must have $k \geq 2$ since we assumed that all $a \in \Sigma$ have
 1390 opinion \perp . In \mathcal{P}' , the configuration $\bar{u} = \bar{a}_1 \cdots \bar{a}_k$ is initial, and it can reach both q_{\perp}^k and v ,
 1391 respectively a \perp and a \top -stable configuration. As a result, the protocol is not a decider.

1392 Now suppose that a \top -stable configuration is unreachable in \mathcal{P} . Let u be an initial
 1393 configuration and let $u \rightarrow^* v$. Since v is not \top -stable, we have $v \rightarrow^* w$ for some configuration
 1394 w containing a state with opinion \perp . Thus $w \rightarrow^* q_{\perp}^*$ and the latter is a \perp -consensus. As a
 1395 result, the protocol is a decider (and recognizes the empty language).

1396 Finally, observe that if \mathcal{P} is immediate-observation then so is \mathcal{P}' . ◀

↑ 1397 ▶ **Corollary 31.** *The syntax of $\text{IO-PP}[+1]$ and $\text{PP}[+1]$ deciders are undecidable.*

1398 **Proof.** In Section 6, we presented an effective construction translating a linear-bounded
 1399 Turing machine into an $\text{IO-PP}[+1]$ protocol with the same language. Since emptiness is
 1400 undecidable for linear-bounded Turing machines [31, Thm. 5.10], it must also be undecidable
 1401 for $\text{IO-PP}[+1]$ protocols. By Lemma 30, so is the syntax of $\text{IO-PP}[+1]$ deciders. ◀

↑ 1402 ▶ **Theorem 32.** *The emptiness problem is undecidable for $\text{PP}[<]$. Hence the syntax of $\text{PP}[<]$
 1403 deciders is also undecidable.*

1404 **Proof.** We reduce from the Post correspondence problem (PCP), which is undecidable. We
 1405 are given two finite alphabets A, B , and two homomorphisms $h_1, h_2: B^* \rightarrow A^*$, and must
 1406 determine whether there exists $w \in B^*$ such that $h_1(w) = h_2(w)$.

1407 We construct three protocols, each in charge of verifying a property of the initial word:

- 1408 ■ \mathcal{P}_0 checks that the input word is in $\#_1 A^* \#_2 B^* \#_3$;
- 1409 ■ \mathcal{P}_1 checks that $h_1(u) = v$ for a word $\#_1 v \#_2 u \#_3$;
- 1410 ■ \mathcal{P}_2 checks that $h_2(u) = v$ for a word $\#_1 v \#_2 u \#_3$.

XX:36 Population Protocols over Ordered Agents

1411 For the first one, we simply use the fact that $\#_1 A^* \#_2 B^* \#_3$ is in DA, and hence recognized
 1412 by an (immediate-observation) protocol \mathcal{P}_0 . In particular an input word can reach a \top -stable
 1413 configuration if and only if it is in that language.

We now construct \mathcal{P}_i for $i \in \{1, 2\}$. The set of states of both protocol is

$$\{a, a^{head}, a^{next} \mid a \in A\} \cup B \cup \{(b, \pi) \mid b \in B, \pi \text{ prefix of } h(b)\} \cup \{\#_1, \#_2, \overline{\#_2}, \#_3, q_\top\}.$$

1414 The opinion of q_\top is \top , all other states have opinion \perp .

1415 Protocol rules are as follows, for all $a, c \in A, b, d \in B$ and prefix π of $h_i(b)$:

<i>Start</i>	<i>Head</i>
$\#_1, a \xrightarrow{\leq} q_\top, a^{head}$	$a^{head}, (b, \pi) \xrightarrow{\leq} a^{next}, (b, \pi a)$
$\#_2, b \xrightarrow{\leq} \overline{\#_2}, (b, \varepsilon)$	if πa is a prefix of $h_i(b)$
<i>Next (A)</i>	<i>Next (B)</i>
$a^{next}, c \xrightarrow{\leq} q_\top, c^{head}$	$(b, h_i(b)), d \xrightarrow{\leq} q_\top, (d, \varepsilon)$
$a^{next}, \overline{\#_2} \xrightarrow{\leq} q_\top, q_\top$	$(b, h_i(b)), \#_3 \xrightarrow{\leq} q_\top, q_\top$

1417 Our final protocol \mathcal{P} is simply the product of those three protocols, the opinion of a
 1418 triple of states being the conjunction of their opinions. This way, we can reach a \top -stable
 1419 configuration from an initial configuration if and only if we can reach a \top -stable configuration
 1420 from it in all three protocols.

1421 We now show that this is the case if and only if that initial configuration is of the form
 1422 $\#_1 u \#_2 v \#_3$ with $u = h_1(v) = h_2(v)$.

1423 If the initial configuration $\#_1 h(a_1) \cdots h(a_k) \#_2 a_1 \cdots a_k \#_3$ satisfies those conditions, then
 1424 it is straightforward to build an execution from it to q_\top^* in both \mathcal{P}_1 and \mathcal{P}_2 . Furthermore, it
 1425 can reach a \top -stable configuration in \mathcal{P}_0 since it is in $\#_1 A^* \#_2 B^* \#_3$.

1426 Now consider an initial configuration from which we can reach a \top -consensus in all
 1427 three protocols. By definition of \mathcal{P}_0 , it must be of the form $\#_1 a_1 \cdots a_k \#_2 b_1 \cdots b_\ell \#_3$ with
 1428 $a_1, \dots, a_k \in A$ and $b_1, \dots, b_\ell \in B$.

1429 Now consider protocol \mathcal{P}_i , with $i \in \{1, 2\}$. We call the states $\{a^{head}, a^{next} \mid a \in A\} \cup \{\#_1\}$
 1430 *special states*. It is easy to see that while we do not use the last **Next (A)** rule, there is
 1431 exactly one special state in the configuration, and there is no special state afterwards. Also
 1432 observe that only letters to the right of the special state can be modified. Since all agents
 1433 must reach q_\top , this imposes that all agents that have a state in A initially must hold the
 1434 special state, one by one from left to right.

1435 A symmetric argument shows that all agents with a state in B initially must transform
 1436 into q_\top using **Next (B)** rules, one by one, from left to right. Now let us have a look at the
 1437 **Head** rule: $a^{head}, (b, \pi) \xrightarrow{\leq} a^{next}, (b, \pi a)$. Let w be the word obtained by taking the letter a
 1438 associated with each of those transitions, in the order in which they appear in the run.

1439 By the arguments above, we must have $w = a_1 \cdots a_k$, and $w = h_i(b_1) \cdots h_i(b_\ell)$. As a
 1440 result, the input word is indeed of the form $\#_1 h_i(b_1) \cdots h_i(b_\ell) \#_2 b_1 \cdots b_\ell \#_3$.

1441 As we have shown this for both $i \in \{1, 2\}$, we obtain that the initial configuration must
 1442 be of the form $\#_1 u \#_2 v \#_3$ with $u = h_1(v) = h_2(v)$.

1443 In conclusion, there is a reachable \top -stable configuration if and only if the PCP instance
 1444 has a solution. ◀

1445 A.11 Missing proofs from Section 7.2

1446 ↑ ▶ **Theorem 34.** *If Conjecture 33 holds, then the syntax of IO-PP[<] deciders is decidable.*

1447 **Proof.** Let $\text{Post}(u) = \{v \mid u \rightarrow v\}$ and $\text{Post}^*(u) = \{v \mid u \rightarrow^* v\}$. Similarly, let $\text{Pre}(u) =$
 1448 $\{v \mid v \rightarrow u\}$ and $\text{Pre}^*(u) = \{v \mid v \rightarrow^* u\}$. We extend these to sets, e.g., $\text{Post}^*(L) =$
 1449 $\bigcup_{w \in L} \text{Post}^*(w)$.

1450 We reuse the machinery of Section 4.1, specifically Lemma 11, Lemma 13, and the
 1451 combinatorial Claim 36.

1452 If a IO-PP[<] protocol $\mathcal{P} = (Q, \Sigma, O, \Delta)$ is not a decider, then it is witnessed by the
 1453 configuration space of words of length n for some $n \geq 2$. This can be checked for each n
 1454 incrementally. We only have to provide a semi-decision procedure which returns yes if and
 1455 only if the protocol is a decider.

1456 Observe that assuming Conjecture 33 holds, a protocol \mathcal{P} is a decider if and only if there
 1457 exist languages K_{\top}, K_{\perp} in DA over Q^* such that:

- 1458 ■ K_{\top} and K_{\perp} are disjoint;
- 1459 ■ $\Sigma^* \subseteq K_{\top} \cup K_{\perp}$;
- 1460 ■ For both b , for all $u \in K_b$ and $v \in \text{Post}(u)$, it is the case that $v \in K_b$;
- 1461 ■ For both b , $K_b \subseteq \text{Pre}^*(C_b)$ with C_b the set of b -stable configurations.

1462 The right-to-left direction is clear: if such invariants exist, then the protocol recognizes
 1463 $K_{\top} \cap \Sigma^*$. For the other direction, by Theorem 10, if \mathcal{P} is a decider then its language L is in
 1464 DA, hence so is $\Sigma^* \setminus L$. Take $K_{\top} = \text{Post}^*(L)$ and $K_{\perp} = \text{Post}^*(\Sigma^* \setminus L)$. Further, assuming
 1465 Conjecture 33, K_{\top} and K_{\perp} are also in DA. They clearly satisfy all the conditions.

1466 Given \mathcal{P} , we enumerate automata $\mathcal{A}, \mathcal{B}_{\top}, \mathcal{B}_{\perp}$ recognizing languages in DA.

1467 We check whether:

- 1468 ■ $L(\mathcal{A}) \subseteq L(\mathcal{B}_{\top})$, $\Sigma^* \setminus L(\mathcal{A}) \subseteq L(\mathcal{B}_{\perp})$;
- 1469 ■ For both b , for all $u \in L(\mathcal{B}_b)$ and $v \in \text{Post}(u)$, it is the case that $v \in L(\mathcal{B}_b)$;
- 1470 ■ For both b , $L(\mathcal{B}_b) \subseteq \text{Pre}^*(C_b)$.

1471 The first condition is easy to check. For the second one, we can build a non-deterministic
 1472 automaton recognizing $\bigcup_{u \in L(\mathcal{B}_b)} \text{Post}(u)$, by making it guess two positions at which it applies
 1473 a transition of the protocol, while checking that the resulting word is in $L(\mathcal{B}_b)$. The last
 1474 condition is more subtle. To check it, we argue that if there is a word in $L(\mathcal{B}_b) \setminus \text{Pre}^*(C_b)$
 1475 then there is one of bounded length. It then suffices to check inclusion over words of length
 1476 below this bound.

1477 First we prove a property of language of DA by a classical use of Ramsey's theorem.

1478 ▷ **Claim 42.** For all language $L \subseteq Q^*$ in DA, there exists $m \in \mathbb{N}$, computable from an
 1479 automaton recognizing L , such that for all $n \geq m$ and u_1, \dots, u_n, z with $\emptyset \neq \alpha(u_1) = \dots =$
 1480 $\alpha(u_n) \supseteq \alpha(z)$, we have $u_1 \cdots u_n \equiv_L u_1 \cdots u_n z u_1 \cdots u_n$.

1481 **Proof.** Let $(M, \cdot, 1_M)$ be the syntactic monoid of L , $\varphi: Q^* \rightarrow M$ a morphism and $F \subseteq M$
 1482 such that $L = \varphi^{-1}(F)$.

1483 Given a sequence of words u_1, \dots, u_n , consider the complete undirected graph (without
 1484 loops) over $\{0, \dots, n\}$, where for all $i < j$ the edge $\{i, j\}$ is colored $\varphi(u_{i+1} \cdots u_j)$. By
 1485 Ramsey's theorem [27], there exists a uniform bound $m \in \mathbb{N}$ such that whenever $n \geq m$ this
 1486 graph contains a monochromatic 3-clique, i.e., there are $i < j < k$ such that $\varphi(u_{i+1} \cdots u_j) =$
 1487 $\varphi(u_{j+1} \cdots u_k) = \varphi(u_{i+1} \cdots u_k)$.

1488 Let u_1, \dots, u_n, z with $n \geq m$ and $\emptyset \neq \alpha(u_1) = \dots = \alpha(u_n) \supseteq \alpha(z)$, and let $i < j < k$ be
 1489 as described above. Since φ is a morphism, $\varphi(u_{i+1} \cdots u_k) = \varphi(u_{i+1} \cdots u_j) \varphi(u_{j+1} \cdots u_k) =$
 1490 $\varphi(u_{i+1} \cdots u_k)^2$. In other words, $u_{i+1} \cdots u_k \equiv_L (u_{i+1} \cdots u_k)^2$. By Equation (1), we have
 1491 $u_{i+1} \cdots u_k \equiv_L u_{i+1} \cdots u_k (u_{k+1} \cdots u_n z u_1 \cdots u_i) u_{i+1} \cdots u_k$.

XX:38 Population Protocols over Ordered Agents

1492 By appending a prefix $u_1 \cdots u_i$ and a suffix $u_{k+1} \cdots u_n$ to both sides of the equivalence,
1493 we obtain $u_1 \cdots u_n \equiv_L (u_1 \cdots u_n)z(u_1 \cdots u_n)$. \triangleleft

1494 We are now ready to show the bound on witnesses for (the negation of) the third condition.

1495 \triangleright **Claim 43.** We can compute a bound k such that if $L(\mathcal{B}_b) \setminus \text{Pre}^*(C_b)$ is non-empty then it
1496 contains a word of length at most k .

1497 *Proof.* Since $L(\mathcal{B}_b)$ belongs to DA, we can compute m from \mathcal{B}_b as given by Claim 42.

1498 By Lemma 11, the set of b -stable configurations is subword-closed and computable. There
1499 is therefore a computable k such that it is made of a finite union of languages described by
1500 expressions of the form $A_1^* B_1^\epsilon \cdots A_\ell^* B_\ell^\epsilon$ with $\ell \leq k$. Let $n = \max(m, k)$.

1501 Let w be a word in $L(\mathcal{B}_b) \setminus \text{Pre}^*(C_b)$ of length $> B(|\Sigma|, n)$ as given by Claim 36. The
1502 word w is of the form $xu_1 \cdots u_n z u_1 \cdots u_n y$ with $\emptyset \neq \alpha(u_1) = \cdots = \alpha(u_n) \supseteq \alpha(z)$.

1503 The word $w' = xu_1 \cdots u_n y$ belongs to $L(\mathcal{B}_b)$ since $n \geq m$. It is not in $\text{Pre}^*(C_b)$ since
1504 otherwise w would be in $\text{Pre}^*(C_b)$, analogously to the proof of Lemma 13. \triangleleft

1505 In conclusion, to check the third condition we can compute the bound k given by the
1506 claim above, and check if there is a word of length at most k in $L(\mathcal{B}_b) \setminus \text{Pre}^*(C_b)$. \blacktriangleleft

B Supplementary constructions and explorations

B.1 Alternative construction for $\Delta_2[\prec] \subseteq \text{IO-PP}[\prec]$

► **Lemma 44.** *Any language in $\Sigma_2[\prec]$ is IO-PP $[\prec]$ -semi-decidable.*

Proof. We know that any language in $\Sigma_2[\prec]$ is a finite union languages of the form $L_0 a_1 L_1 \dots a_n L_n$ where L_i is a downward-closed language for all $1 \leq i \leq n$. Furthermore, any downward closed language can be written as a finite union of languages of the form $A_0^* a_1 A_1^* \dots a_n A_n^*$ where $A_i \subseteq \Sigma$. Therefore, since semi-decidable languages are closed under finite union, it is sufficient to build a protocol recognizing L of the form $A_0^* a_1 A_1^* \dots a_{n-1} A_{n-1}^* a_n$.

The protocol is designed so that an input word u can reach a \top -stable configuration if and only if $u \in L$. It proceeds in three steps: (1) Each agent guesses an index in $\{0, \dots, n\}$ corresponding to a position in the pattern. (2) Once an agent has observed at least one agent for each odd position (in an order that is consistent with its own view), it can “lock” itself in a \top -state. (3) Any agent that is not locked can always reinitialize (i.e., restart from its input symbol and try a different guess). Even locked agents may reinitialize if they encounter an agent whose guess is incompatible with theirs.

Let us give explicitly the protocol rules.

(1) **Guessing rules.** Each agent can make a guess, without observing any neighbor.

$$a \rightarrow a(i)(-1) \text{ if } \begin{cases} a \in A_i \text{ and } i \text{ is even,} \\ a = a_i \text{ and } i \text{ is odd.} \end{cases}$$

In state $a(i)(\ell)$, i is the guessed position of the letter in the language structure, whereas $\ell \in \{-1, 1, 3, \dots\}$ is its level in the lock process. A letter is fully locked when $\ell = n - 1$, i.e., the opinion function maps each symbol in $\{\sigma(i)(n - 1) \mid \sigma \in \Sigma, i \leq n\}$ to \top , and every other symbol to \perp .

(2) **Locking rules.** When one agent with a guess i is at lock level ℓ , it can increase its lock level to $\ell + 2$ when either its guess is $\ell + 2$, (and it is the agent $a_{\ell+2}$ of the pattern), or it encounters one agent with guess ℓ to its right (resp. left) and this is coherent with its view, i.e. $i < \ell + 2$ (resp $i > \ell + 2$):

$$\begin{aligned} (a(i)(\ell), b(\ell + 2)(\ell')) &\xrightarrow{\prec} (a(i)(\ell + 2), b(\ell + 2)(\ell')) && \text{with } i < \ell + 2 \\ (b(\ell + 2)(\ell'), a(i)(\ell)) &\xrightarrow{\prec} (b(\ell + 2)(\ell'), a(i)(\ell + 2)) && \text{with } i > \ell + 2 \\ a(\ell + 2)(\ell) &\rightarrow a(\ell + 2)(\ell + 2) \end{aligned}$$

(3) **Unlocking rules.**

$$(a(i)(\ell), b(j)(\ell')) \xrightarrow{\prec} (a, b(j)(\ell')) \text{ if } \begin{cases} i > j \\ i = j \text{ and } i \text{ is odd and greater than } -1. \end{cases}$$

$$a(i)(\ell) \rightarrow a \text{ if } \ell < n - 1$$

$$(a, b(i)(\ell)) \rightarrow (a, b)$$

$$(b(i)(\ell), a) \rightarrow (b, a)$$

We claim that a word $u = u_1 \dots u_k$ can reach a word $u_\top = u_1(x_1)(n - 1) \dots u_k(x_k)(n_1)$ in Σ_\top only if for all odd i between 1 and $n - 1$, there exists $j \leq k$ such that $x_j = i$. Indeed, consider the last letter to go from state $a(x)(i - 2)$ to $a(x)(i)$ in the execution from u to

XX:40 Population Protocols over Ordered Agents

1545 u_{\top} : to do that, $a(x)(i-2)$ must observe a letter $b(i)(\ell)$ (maybe itself). For the sake of
 1546 contradiction, imagine that b wears a different label in u_{\top} : it means that after being observed
 1547 by a , b uses an unlocking rule to reinitialize itself, and then later guesses a new position
 1548 and locked itself; hence a is not the last letter to do the i -th locking step, and we reach a
 1549 contradiction.

1550 Thanks to the first unlocking rule, u_{\top} is a \top -consensus iff $x_1 \leq \dots \leq x_k$, and for each
 1551 odd i between 1 and $n-1$, there exist at most one j such that $x_j = i$. Thus, if u is not in L ,
 1552 u cannot reach a \top -consensus, because there is no labelling of u through the guessing rule
 1553 producing an unlockable sequence of x_i .

1554 On the other hand, if u is in L , there exists such a labelling, hence there is a run producing
 1555 a \top -consensus from u . Now it only remains to show that any *fair* run reaches a \top -consensus.
 1556 Observe that, from any configuration u' which is not a \top -consensus, there is an execution
 1557 from u' to u using the unlocking rules: any partially locked letter can reinitialize itself, and
 1558 a fully locked letter can reinitialize by observing a position contradicting with its own from
 1559 a letter at its right. Furthermore, any un-labelled letter a can reinitialize any other letter,
 1560 even fully locked. Thus, any infinite fair run from u would enter and leave u infinitely often,
 1561 and explore every configuration accessible from u , including a \top -consensus. Thus every fair
 1562 run from u is finite and reaches a \top -consensus.

1563 Therefore, the protocol semi-decides L . ◀

1564 B.2 Alternative construction for $\Delta_1^{\text{int}}[\langle, +, \equiv] \subseteq \text{PP}[\langle]$

1565 ▶ **Proposition 45.** $\Delta_1^{\text{int}}[\langle, +, \equiv] \subseteq \text{PP}[\langle]$.

1566 **Proof.** We show that any language in $\Sigma_1^{\text{int}}[\langle, +, \equiv]$ is semi-decided by a protocol in $\text{PP}[\langle]$.
 1567 Such a language is a finite union of intersections $K = R \cap S$, where R is a language of the
 1568 form $A_0^* a_1 A_2^* \dots a_{n-1} A_n^*$ and S is a semilinear constraint protocol. It suffices to semi-decide
 1569 one such K .

iv: it is more
 than an intersec-
 tion, otherwise
 we know how to
 intersect proto-
 cols

1570 We construct a protocol \mathcal{P} with two components: \mathcal{P}_{struct} from Lemma 44 that semi-decides
 1571 R and \mathcal{P}_{count} .

1572 A fair run of \mathcal{P}_{struct} on input w annotates every agent with a position label (k) and a
 1573 locked/unlocked label. When an agent is labelled with position k , it means this agent believe
 1574 they are part of the k -th block of R . Furthermore, for any $w \rightarrow^* w'$:

- 1575 ■ If at least one letter is unlocked in w' , there exists a run $w' \rightarrow^* w$.
- 1576 ■ If w' is fully locked, but an agent observes that their position belief is in contradiction
 1577 with an other agent belief, they can unlock themselves.
- 1578 ■ If w' is fully locked with coherent position beliefs (which is possible iff $w \in R$) then no
 1579 more transition rules can be taken; \mathcal{P}_{struct} stabilizes.

1580 This stabilizing property of \mathcal{P}_{struct} is what will allow us to combine it with \mathcal{P}_{count} .
 1581 Formally, following Angluin et al. [1], an input process is *stabilizing* if it remains invariant
 1582 after some time T_{stab} . By theorem 4 in [1], there exists a protocol \mathcal{P}_{count} that, given stabilizing
 1583 inputs, eventually stabilizes to output \top if the inputs satisfy S , and \perp otherwise.

iv: there is a
 leap from the
 way Theorem 4
 is given and that

1584 Protocol \mathcal{P}_{count} runs only on agents locked by \mathcal{P}_{struct} with a specific feedback loop: if an
 1585 agent has opinion \perp in \mathcal{P}_{count} , they can unlock themselves. We show that a stable \top -consensus
 1586 is reachable $\iff w \in K$.

- 1587 ■ Case $w \in K$: There exists a correct structural guess. Consider a fair path the agents
 1588 guess this structure and lock. Since $w \in S$, \mathcal{P}_{count} eventually stabilizes to output \top . Once

output \top is reached, no unlocking can occur, and the system remains in a permanently locked structural state with \top output, i.e., a \top -consensus.

- Case $w \notin K$: If $w \notin R$, \mathcal{P}_{struct} never locks stably by definition. If $w \in R$ but $w \notin S$, suppose \mathcal{P}_{struct} locks in position w' . Since w' represent a vector violating S . Thus \mathcal{P}_{count} can always unlock an agent. ◀

Proof. Any language in $\Sigma_1^{\text{int}}[\langle, +, \equiv]$ is a finite union of languages L where L is defined by

- a finite set of special positions x_1, \dots, x_n .
- Semi-linear conditions on the infixes delimited by these special positions.

Formally, L is characterized a number n and a finite set \mathcal{S} of constraints of the form $\sum_{\sigma \in A} \#_{\sigma}(x_i, x_j) \bowtie \sum_{\sigma \in B} \#_{\sigma}(x_k, x_l)$ with $A, B \subseteq \Sigma$ and $\bowtie \in \{=, <, \leq, >, \geq\}$ and $0 \leq i \leq j \leq n+1$ and $0 \leq k \leq l \leq n+1$. For any $u \in \Sigma^*$, $u \in L$ iff there exist positions x_1, \dots, x_n in u such that u verifies all constraints in \mathcal{S} (positions x_0 and x_{n+1} denotes the first and last positions of the word).

Let us build \mathcal{P} that semi-decides such a L . We do that by composing protocols \mathcal{P}_{struct} and $\mathcal{P}(S)$ for all $S \in \mathcal{S}$.

First, \mathcal{P}_{struct} factorizes u as $u_0 a_1 u_2 \dots a_{n-1} u_n$ by annotating every agent with a position belief: either (x_k) , if the agent believe themselves to be special position x_k , or $(k+)$ if they believe to be after x_k (but before x_{k+1} if exists), or (0) for agent before x_0 . Then every agent “locks” themselves by checking they can see every special position. Any agent can unlock themselves if they observe a contradiction between their beliefs and someone else, or someone not fully locked. A \top -stable configuration is a configuration where every agent is fully locked with a belief coherent with some factorisation $u_0 a_1 u_2 \dots a_{n-1} u_n$ of u .

For more details on how this protocol works, see Proof of Lemma 44.

Now, let us give a protocol \mathcal{P}_S which verifies a constraint S of the form $\sum_{\sigma \in A} \#_{\sigma}(x_i, x_j) \bowtie \sum_{\sigma \in B} \#_{\sigma}(x_k, x_l)$ with $A, B \subseteq \Sigma$ and $\bowtie \in \{=, <, \leq, >, \geq\}$ and $0 \leq i \leq j \leq n+1$ and $0 \leq k \leq l \leq n+1$. Assume that $\bowtie = <$ (other cases works similarly):

Here are the transition rules of \mathcal{P}_S , for any agents a, b, c, d fully locked such that agent a is fully locked in A with position belief between i and j , agent b is fully locked in B with position belief between k and l , and agent c is in one of these configurations. Agent d is fully locked with irrelevant position belief.

- $a, b \rightarrow a^2, b^2$ for $a \in A, b \in B$
- $b \rightarrow b^{1-\top}$
- $b^{1-\top}, a \rightarrow b^2, a^2$
- $b^{1-\top}, c_2 \rightarrow b^1, c^{2-\top}$
- $d \rightarrow d^{\top}$

\mathcal{P} is a composition of protocol \mathcal{P}_{struct} , and protocols $\mathcal{P}(S)$ for all $S \in \mathcal{S}$. The labels from protocol $\mathcal{P}(S)$ carry a S -identifier, and can co-exist together. A *end-configuration* in \mathcal{P} is a word u' where each agent is locked and has a position belief, and has a \top label from $\mathcal{P}(S)$ for all $S \in \mathcal{S}$ Furthermore, all position beliefs must be coherent. Σ_{\top} is the minimal alphabet such that every end-configuration is a \top -consensus. We will show later that end-configurations are exactly \top -stable configurations.

We add the following feedback rules, between the composed protocols:

- **Free-to-unlock rule** Letters that are locked but have no label can always unlock themselves. Furthermore, letters that are locked cannot be unlocked unless they bear no label from any \mathcal{P}_S .

IV: I don't think this proof can work without giving the details of Pcount. If Pcount works for some time then unlock a letter, we cannot force this letter to completely reset w, there could be some leftovers from the Pcount execution

IV: Here is my attempt at a proof. To make it work, I had to add some precision on how Pstruct and Pcount compose together, and how to carefully clean up when doing feedback

MB: How do you get rid of coefficients from linear combinations? And from modulus? This seems crucial; agents cannot take the role of a position outside of the range [1..n].

XX:42 Population Protocols over Ordered Agents

- 1636 ■ **Prepare-to-change-belief rule.** If \mathcal{P}_{struct} can unlock an agent, but this agent is
1637 labelled, it can receive an “un-label yourself” order.
- 1638 ■ **Unlabel rules.** An agent with such an “unlabel yourself” order can reverse any
1639 transition from any protocol S involving itself and another agent (even if this agent has
1640 received no order). It then loses the “unlabel yourself”.
- 1641 ■ **Unhappy rule.** If an agent has some labels but not all needed, it can give itself the
1642 “unlabel yourself” order.

1643 These feedback rules ensure the following properties:

- 1644 ■ If $u \in L$, u can reach an end-configuration. Furthermore, end-configuration are \top -stable:
1645 no feedback or unlock rules can be applied.
- 1646 ■ For any run $u \rightarrow^* u'$ such that u' is not an end-configuration, there exists a run $u' \rightarrow^* u$.
1647 ■ Either some agent is not fully labelled: they can use the **Unhappy rule**, then unlabel
1648 themselves through the Unlabel rules, then unlock themselves with the **Free-to-unlock**
1649 **rule**. From there, they can use the **Prepare-to-change-belief rule** on every labelled
1650 agent. They all unlabel themselves then unlock themselves.
- 1651 ■ Or there is a contradiction between beliefs. The agent observing this contradiction can
1652 use the **Prepare-to-change-belief rule** on themselves, then apply the same process as
1653 above.
- 1654 ■ The unlabel rules, combined to the rules of any protocol \mathcal{P}_S , ensure that, at any point of an
1655 execution, all agents labelled by \top form a word that would verify S if their position belief
1656 were valid. Therefore, if there is an execution $u \rightarrow u'$ where u' is an end-configuration,
1657 then $u \in L$. Indeed, position beliefs are coherent in u' coherent, and for any $S \in \mathcal{S}$ every
1658 agent is labelled \top by \mathcal{P}_S . ◀

1659 B.3 Alternative proof of $\Delta_2[<] \subseteq \Delta_1^{\text{int}}[<, +, \equiv]$

1660 ▶ **Proposition 46.** $\Delta_2[<] \subseteq (\Delta_1^{\text{int}}[<, +, \equiv] \cap \text{Reg})$.

1661 **Proof.** Let $L \in \Sigma_2[<]$. As mentioned in Section 4.2, L is a finite union of languages of the
1662 form $L' = A_0^* a_1 A_1^* \cdots a_k A_k^*$ where $A_i \subseteq \Sigma$. Let $\chi_{i,j}$ be 1 if $a_i \notin A_j$, and 0 otherwise. The
1663 language L' can be expressed by

$$\begin{aligned}
 & \exists x_1, \dots, x_k (x_1 < \cdots < x_k) \wedge a_1(x_1) \wedge \cdots \wedge a_k(x_k) \wedge \\
 & \sum_{\sigma \in \Sigma \setminus A_0} \#_{\sigma}(1, x_1) = \chi_{1,0} \wedge \sum_{\sigma \in \Sigma \setminus A_k} \#_{\sigma}(x_k, \mathbf{max}) = \chi_{k,k} \wedge \\
 & \bigwedge_{1 \leq i < k} \left(\sum_{\sigma \in \Sigma \setminus A_i} \#_{\sigma}(x_i, x_{i+1}) = \chi_{i,i} + \chi_{i+1,i} \right).
 \end{aligned}$$

1668 Thus, $L' \in \Sigma_1^{\text{int}}[<, +, \equiv]$ and so $L \in \Sigma_1^{\text{int}}[<, +, \equiv]$. Since $L \in \Delta_2[<]$, we have $\bar{L} \in \Sigma_2[<]$ and
1669 hence $\bar{L} \in \Sigma_1^{\text{int}}[<, +, \equiv]$ by the same reasoning. Therefore, $L \in \Delta_1^{\text{int}}[<, +, \equiv]$. ◀

1670 B.4 Partially-ordered two-way Parikh automata

1671 A *partially-ordered two-way Parikh automaton* (*p.o. 2-PA*) is a tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F, \Psi_{\text{acc}})$
1672 where:

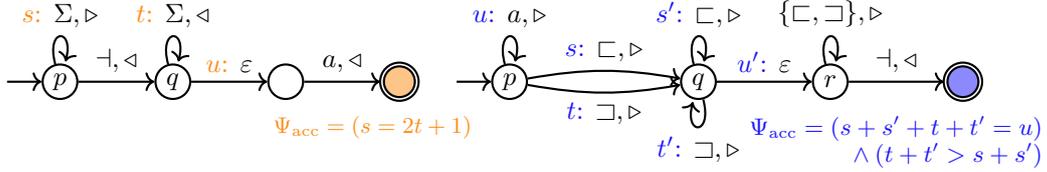
- 1673 ■ Q is a finite set of *states* equipped with a partial order \leq ;
- 1674 ■ Σ is a finite *alphabet*;

- 1675 ■ $\delta \subseteq Q \times (\Sigma_{\vdash\lrcorner} \cup \{\varepsilon\}) \times \{\triangleright, \triangleleft\} \times Q$ is the *transition relation*, that satisfies the following,
 1676 for every $(p, \sigma, d, q), (p', \sigma', d', q') \in \delta$:
- 1677 ■ $p \leq q$ (cycles are self-loops);
 - 1678 ■ if $p = q = p' = q'$, then $d = d'$ (a self-loop reads in a single direction);
 - 1679 ■ if $p = p'$ and $\sigma, \sigma' \in \Sigma_{\vdash\lrcorner}$, then $\sigma \neq \sigma'$ (nondeterminism arises from ε);
 - 1680 ■ if $p = q$, then $\sigma \neq \varepsilon$ (no ε -self-loops);
 - 1681 ■ $\sigma = \vdash$ implies $\sigma \neq \triangleleft$, and $\sigma = \lrcorner$ implies $\sigma \neq \triangleright$ (the head never goes out of bound);
- 1682 ■ $q_0 \in Q$ is the *initial state*;
- 1683 ■ $F \subseteq Q$ is the set of *final states*, from which there is no transition;
- 1684 ■ Ψ_{acc} is a Presburger formula over variables δ , called the *acceptance condition*.

1685 A *configuration* of \mathcal{A} is a triple (q, i, \mathbf{c}) where $q \in Q$ is the current state; $i \in \mathbb{N}$ is the
 1686 head position; and $\mathbf{c}: \delta \rightarrow \mathbb{N}$ maps transitions to their visit count. For every $w \in \Sigma^n$, let
 1687 $w[0] = \vdash$ and $w[n+1] = \lrcorner$. On input w , a step $(q, i, \mathbf{c}) \xrightarrow{t} (q', i', \mathbf{c}')$ can occur if transition
 1688 $t = (q, \sigma, \tau, q') \in \delta$ satisfies the following:

- 1689 ■ If $\sigma \in \Sigma_{\vdash\lrcorner}$, then $w[i] = \sigma$, and $i' = i + 1$ if $\tau = \triangleright$, or $i' = i - 1$ if $\tau = \triangleleft$;
- 1690 ■ If $\sigma = \varepsilon$, then $i' = i$;
- 1691 ■ $\mathbf{c}'(t) = \mathbf{c}(t) + 1$ and $\mathbf{c}'(s) = \mathbf{c}(s)$ for $s \neq t$.

1692 We say that \mathcal{A} *accepts* the input w if $(q_0, 1, \mathbf{0}) \xrightarrow{*} (q, i, \mathbf{c})$ where $q \in F$ and $\Psi_{\text{acc}}(\mathbf{c})$ holds.
 1693 The *language* of \mathcal{A} is the set of words it accepts. Figure 4 depicts two examples of p.o. 2-PA.



1694 ■ **Figure 4** Example of p.o. 2-PA for the median language $\bigcup_{n \geq 0} \Sigma^n a \Sigma^n$ (left) and the coDyck-
 1695 witness language $\{a^n v \mid n \geq 1, v \in \{\square, \sqsupset\}^{\geq n}, |v[1..n]|_{\square} > |v[1..n]|_{\sqsupset}\}$ (right).

1694 For all $t = (p, \sigma, d, q)$, let $\text{out}(t) = p$, $\text{letter}(t) = \sigma$, $\text{dir}(t) = d$ and $\text{in}(t) = q$. We distinguish
 1695 *looping transitions* $\delta_{\circlearrowleft} = \{t \in \delta \mid \text{out}(t) = \text{in}(t)\}$ and *progress transitions* $\delta_{\rightarrow} = \delta \setminus \delta_{\circlearrowleft}$.

1696 ► **Proposition 47.** *A language is recognized by some p.o. 2-PA iff it belongs to $\Sigma_1^{\text{int}}[\langle, +, \equiv]$.*

1697 A resolver is a mechanism that indicates what to do when a nondeterministic choice
 1698 arises. We introduce a type of resolver that takes a progress transition as soon as some linear
 1699 constraint holds, and only uses looping transitions otherwise.

1700 Let τ be a function that associates to each progress transition a Presburger formula over
 1701 Q , such that $(t \neq t' \wedge \text{out}(t) = \text{out}(t') \wedge \varepsilon \in \{\text{letter}(t), \text{letter}(t')\})$ implies $\tau(t) \wedge \tau(t') \equiv \text{false}$.

1702 We will interpret formulas over Q as formulas over δ by mapping each q to $\sum_{t \in \delta \mid \text{in}(t)=q} \mathbf{c}(t)$.
 1703 Given $t \in \delta$, we write $(q, i, \mathbf{c}) \xrightarrow{t}_{\tau} (q', i', \mathbf{c}')$ iff $(q, i, \mathbf{c}) \xrightarrow{t} (q', i', \mathbf{c}')$ and either

- 1704 ■ $t \in \delta_{\rightarrow}$ and $\tau(t)$ holds; or
- 1705 ■ $t \in \delta_{\circlearrowleft}$ and $\neg \tau(s)$ holds for each $s \in \delta_{\rightarrow}$ such that $\text{out}(s) = q$.

1706 We write $L_{\tau}(\mathcal{A})$ to denote the language recognized by \mathcal{A} using $\xrightarrow{*}_{\tau}$ instead of $\xrightarrow{*}$. We say
 1707 that τ is a *semilinear resolver* if $L_{\tau}(\mathcal{A}) = L(\mathcal{A})$.

XX:44 Population Protocols over Ordered Agents

1708 ► **Example 48.** Reconsider automaton $\mathcal{A}_{\text{left}}$ from Figure 4. By assigning $\tau(u) = (p - 2q = 1)$,
 1709 the automaton accepts the same language, even with $\Phi_{\text{acc}} = \top$. For automaton $\mathcal{A}_{\text{right}}$ from
 1710 Figure 4, setting $\tau(u') = (p - q = 0)$ yields the same language, even with $\Phi_{\text{acc}} = (t + t' > s + s')$.
 1711 The advantage here is that it makes both automata “deterministic”.

1712 ► **Proposition 49.** *Given a p.o. 2-PA \mathcal{A} and a semilinear resolver τ , it is possible to construct*
 1713 *a p.o. 2-PA \mathcal{A}' and a semilinear resolver τ' such that $L_{\tau'}(\mathcal{A}') = \overline{L_{\tau}(\mathcal{A})}$.*

1714 **Proof.** Let $\mathcal{A} = (Q, \Sigma, \delta, q_0, F, \Psi_{\text{acc}})$. Since $(t \neq t' \wedge \text{out}(t) = \text{out}(t') \wedge \varepsilon \in \{\text{letter}(t), \text{letter}(t')\})$
 1715 implies $\tau(t) \wedge \tau(t') \equiv \text{false}$, two progress transitions of \mathcal{A} are never simultaneously enabled.
 1716 As \rightarrow_{τ}^* prioritize progress transitions, this means that each input word has a single run in \mathcal{A} .
 1717 Since \mathcal{A} is partially ordered, it cannot loop forever. Therefore, a word is rejected by \mathcal{A} iff its
 1718 single run reaches F with $\neg\Psi_{\text{acc}}$, or reaches a dead-end in $Q \setminus F$. We construct \mathcal{A}' and τ' so
 1719 that they detect these two outcomes.

1720 For every $p \in Q$, let

$$1721 \quad A_p = \{a \in \Sigma_{\neq \perp} \mid a = \text{letter}(t) \text{ for some } t \in \delta \text{ with } \text{out}(t) = p\},$$

$$1722 \quad S_p = \{s \in \delta_{\rightarrow} \mid \text{out}(s) = p \wedge \text{letter}(s) = \varepsilon\}.$$

1723 We construct $\mathcal{A}' = (Q', \Sigma, \delta', q'_0, F', \Psi'_{\text{acc}})$ and τ' as follows:

- 1724 ■ $Q' = Q \cup \{q_{\text{acc}}\}$, $q'_0 = q_0$ and $F' = \{q_{\text{acc}}\}$;
- 1725 ■ δ' and τ' obtained by extending δ and τ as follows:
 - 1726 ■ For each $t = (p, \sigma, d, q) \in \delta_{\rightarrow}$ with $\sigma \neq \varepsilon$, we add $t' = (p, \sigma, d, q_{\text{acc}})$ with $\tau'(t') =$
 1727 $\neg\tau(t) \wedge \bigwedge_{s \in S_p} \neg\tau(s)$. Indeed, as t is the only transition that can read σ from p ,
 1728 automaton \mathcal{A} rejects unless an ε -transition is enabled by τ (in which case it is taken).
 - 1729 ■ For each $a \in \Sigma_{\neq \perp} \setminus A_p$, we add $t_{p,a} = (p, a, d, q_{\text{acc}})$ with $\tau'(t_{p,a}) = \bigwedge_{s \in S_p} \neg\tau(s)$, and
 1730 $d = \triangleright$ iff $a \neq \perp$. Indeed, since no transition can read a from p , automaton \mathcal{A} rejects
 1731 unless an ε -transition is enabled by τ (in which case it is taken);
- 1732 ■ $\Psi'_{\text{acc}} = \bigwedge_{t \in \delta' \setminus \delta} (t > 0) \vee \neg\Psi_{\text{acc}}$.

1734 ► **Proposition 50.** *For every p.o. 2-PA \mathcal{A} and linear resolver τ , $L_{\tau}(\mathcal{A}) \in \Delta_1^{\text{int}}[\langle, +, \equiv]$.*

1735 **Proof.** By Proposition 49, the class of languages recognized by these automata is closed
 1736 under complement. Thus, it suffices to give a formula in Σ_1^{int} .

1737 Since \mathcal{A} is partially ordered, every accepting run is of the form

$$1738 \quad (q_0, \cdot, \cdot) \xrightarrow{\rho_0} (q_0, \cdot, \cdot) \xrightarrow{t_1} (q_1, \cdot, \cdot) \xrightarrow{\rho_1} (q_1, \cdot, \cdot) \xrightarrow{t_2} \cdots \xrightarrow{\rho_{m-1}} (q_{m-1}, \cdot, \cdot) \xrightarrow{t_m} (q_m, \cdot, \cdot),$$

1739 where each ρ_i only uses looping transitions, each t_i is a progress transition, and $q_m \in F$.

1740 We will construct a formula that guesses t_1, \dots, t_m and checks that the run is valid. We
 1741 introduce existential variables $\mathbf{x} = (x_1, \dots, x_m)$ and $\mathbf{y} = (y_0, y_1, \dots, y_m)$ where x_i and y_i are
 1742 the head position respectively before and after taking t_i . We set $y_0 = 1$.

1743 The validity of ρ_i requires that the interval between x_i and x_{i+1} contains no letter outside
 1744 of $A_i = \{\text{letter}(s) \mid s \in \delta_{\cup}, \text{out}(s) = \text{out}(t_i)\}$. It is expressed as

$$1745 \quad \varphi_{\text{loop},i} = \bigwedge_{a \in \Sigma_{\neq \perp} \setminus A_i} (\#_a(u, v) = 0) \text{ where } u = \min(y_i, x_{i+1}) \text{ and } v = \max(y_i, x_{i+1}).$$

1746 The validity of the positions is expressed as follows:

$$1747 \quad \varphi_{y,i} = \begin{cases} y_i = x_i & \text{if } \text{letter}(t_i) = \varepsilon, \\ y_i = x_i + 1 & \text{if } \text{letter}(t_i) \neq \varepsilon \text{ and } \text{dir}(t_i) = \triangleright, \\ y_i = x_i - 1 & \text{if } \text{letter}(t_i) \neq \varepsilon \text{ and } \text{dir}(t_i) = \triangleleft. \end{cases}$$

1748 It remains to express that transition t_i is valid. Given a Presburger formula φ over Q ,
1749 and a state q , let us define

$$1750 \quad \varphi_{\text{disabled},i} = \forall x ((0 \leq x \leq \text{out}(t_i)) \implies \neg \varphi(\text{out}(t_i) \mapsto x)).$$

1751 Since Presburger arithmetic admits quantifier elimination, $\varphi_{\text{disabled},i}$ can be made quantifier-
1752 free. We define

$$1753 \quad \varphi_{\text{progress},i} = ((\text{letter}(t_i) \neq \varepsilon) \rightarrow \#_{\text{letter}(t)}(x_i, x_i) = 1) \wedge \tau(t_i) \wedge \bigwedge_{\substack{s \in \delta_{\rightarrow} \setminus \{t\} \\ \text{out}(s) = \text{out}(t) \\ \text{letter}(s) = \varepsilon}} \tau(s)_{\text{disabled},i}.$$

1754 Recall that τ yields Presburger formulas over Q , which technically does not typecheck in
1755 $\varphi_{\text{progress},i}$. Formally, we consider that q_i stands for $|y_i - x_{i+1}|$. ◀

1756 ▶ **Corollary 51.** *Any language decided by a p.o. 2-PA with linear resolving belongs to PP[<].*

1757 **B.5 Alternative proof for $\text{NSPACE}(n) \subseteq \text{IO-PP}[+1]$**

 1758 Let us fix a linear-bounded nondeterministic automaton $M = (P, \Sigma, \Sigma', \delta, p_0, p_{\text{acc}}, p_{\text{rej}})$ where
 1759 P is the set of states; Σ is the input alphabet; $\Sigma' \supseteq \Sigma$ is the tape alphabet; $\delta \subseteq (P \times$
 1760 $\Sigma') \times (P \times \Sigma' \times \{\triangleleft, \triangleright\})$ is the transition relation; and $p_0, p_{\text{acc}}, p_{\text{rej}}$ are respectively the initial,
 1761 accepting and rejecting states.

 1762 Traditionally, on input $w \in \Sigma$, a linear-bounded automaton has access to $\vdash w \dashv$; has
 1763 no left-move on \vdash ; has no right-move on \dashv ; and cannot overwrite the endmarkers. For
 1764 convenience, we consider instead that there are no endmarkers, but that the head stays put
 1765 when moving left (resp. right) from the first (resp. last) cell. It is simple to construct such a
 1766 machine with the same language. We further assume that M may only enter q_{acc} with its
 1767 head on the rightmost cell, and that there is no transition from q_{acc} .

 1768 Let $\Gamma = (P \cup \{-\}) \times \Sigma' \times \{\text{fst}, -, \text{lst}\}$. For every $s = (p, a, x) \in \Gamma$, let $\text{sta}(s) = p$ and
 1769 $\text{pos}(s) = x$. We say that a word $w \in \Gamma^+$, with $|w| \geq 2$, is a *valid configuration* of M if

- 1770 ■
- $\text{pos}(w[i]) = \text{fst}$
- iff
- $i = 1$
- ,
-
- 1771 ■
- $\text{pos}(w[i]) = \text{lst}$
- iff
- $i = |w|$
- ,
-
- 1772 ■ there is a unique position
- i
- with
- $\text{sta}(w[i]) \in P$
- .

 1773 For example, $w = (-, a, \text{fst})(p, a, -)(-, b, \text{lst})$ represents the configuration where M is in
 1774 state p , with its head on the second cell, and its tape containing aab .

 1775 ► **Proposition 52.** *The following language is $\text{PP}[+1]$ -semi-decidable with stabilizing inputs:*
 1776 $L = \{w \in \Gamma^{\geq 2} \mid w \text{ is a valid configuration of } M \text{ and it leads to acceptance}\}$.

 1777 **Proof.** Let us construct a protocol $\mathcal{P} = (Q, \Gamma, O, \Delta)$. We define $Q = \Gamma \times \Gamma \times \Gamma \times \{\perp, \top, \bullet\}$
 1778 and $O((x, y, z, o)) = (x = y \wedge o = \top)$. We identify each input $\gamma \in \Gamma$ with $(\gamma, \gamma, \gamma, \perp)$.

 1779 Intuitively, the state (x, y, z, o) indicates that the agent has input x ; has been acting as if
 1780 its input was y ; currently holds z ; and has output belief o . Whenever $x \neq y$, the agent has
 1781 changed its mind on its input, and must consequently reset the whole population.

 1782 The transitions of Δ are defined by these rules:

 1783

<i>Simulation of M</i>		
(1)	$(*, *, (p, a, *), *), (*, *, (-, *, *), *) \xrightarrow{+1} (*, *, (-, b, *), *), (*, *, (q, *, *), *)$	for $((p, a), (q, b, \triangleright)) \in \delta$
(2)	$(*, *, (-, *, *), *), (*, *, (p, a, *), *) \xrightarrow{+1} (*, *, (q, *, *), *), (*, *, (-, b, *), *)$	for $((p, a), (q, b, \triangleleft)) \in \delta$
(3)	$(x, *, (p, a, *), *) \xrightarrow{\text{true}} (x, *, (q, b, *), *)$	for $((p, a), (q, b, \triangleright)) \in \delta$ and $\text{pos}(x) = \text{lst}$
(4)	$(x, *, (p, a, *), *) \xrightarrow{\text{true}} (x, *, (q, b, *), *)$	for $((p, a), (q, b, \triangleleft)) \in \delta$ and $\text{pos}(x) = \text{fst}$
<i>Belief propagation</i>		
(5)	$(x, *, *, *), (x', *, z', *) \xrightarrow{\text{true}} (x, *, *, \top), (x', *, z', *)$	if $\text{pos}(x) = \text{fst} \wedge \text{pos}(x') = \text{lst} \wedge \text{sta}(z') = p_{\text{acc}}$
(6)	$(*, *, *, \top), (*, *, *, *) \xrightarrow{\text{true}} (*, *, *, \top), (*, *, *, \top)$	
<i>Reset on input changes</i>		
(7)	$(x, y, *, *), (x', y', *, *) \xrightarrow{\text{true}} (x, y, *, *), (x', y', *, \bullet)$	if $\text{pos}(x') = \text{lst} \wedge (x \neq y \vee x' \neq y')$
(8)	$(*, *, *, *), (x', *, *, \bullet) \xrightarrow{+1} (*, *, *, \bullet), (x', x', x', \perp)$	
(9)	$(x, *, *, \bullet) \xrightarrow{\text{true}} (x, x, x, \perp)$	if $\text{pos}(x) = \text{fst}$
<i>Erroneous configuration detection</i>		
(10)	$(x, *, *, *), (x', *, *, *) \xrightarrow{\text{true}} (x, *, *, \perp), (x', *, *, \perp)$	if $\text{sta}(x), \text{sta}(x') \in P$
(11)	$(x, *, *, *), (x', *, *, *) \xrightarrow{\text{true}} (x, *, *, \perp), (x', *, *, \perp)$	if $\text{pos}(x) = \text{pos}(x') \in \{\text{fst}, \text{lst}\}$
(12)	$(x, *, *, *), (x', *, *, *) \xrightarrow{+1} (x, *, *, \perp), (x', *, *, \perp)$	if $\text{pos}(x) = \text{lst} \vee \text{pos}(x') = \text{fst}$

 1785 Let us show that \mathcal{P} semi-decides L with stabilizing inputs. Let $u \in \Gamma^n$. For the special
 1786 case of $n = 1$, note that each state starts with output \perp and this remains so as no transition
 1787 is ever enabled. Thus, let us assume that $n \geq 2$.

1788 Let $u = u_0 \rightsquigarrow u_1 \rightsquigarrow \dots \rightsquigarrow u_\ell = v$, and let $v = w_0 \rightarrow w_1 \rightarrow \dots$ be a fair run. Without
 1789 loss of generality, we may assume $\ell = 0$ or $\iota(u_\ell) \neq \iota(u_{\ell-1})$. Indeed, it suffices to take ℓ as the
 1790 last moment where the input changes.

1791 We make a case distinction, where each case assumes that the previous ones do not hold.

1792 *Case 1: $\iota(v)$ is valid.* By validity, rules (10–12) are permanently disabled from v onwards.

1793 If $\ell > 0$, then there exists a position i such that $v[i] = (x, y, z, o)$ with $x \neq y$. By validity
 1794 of $\iota(v)$, we have $\text{pos}(\iota(v)[1]) = \text{fst}$ and $\text{pos}(\iota(v)[n]) = \text{lst}$. Thus, by fairness, the population
 1795 must eventually use rule (7) to assign \clubsuit to agent n ; use rule (8) to reset the population from
 1796 right to left; and use rule (9) to complete this reset with agent 1. From there, rules (7–9) are
 1797 permanently disabled. Thus, by fairness, M is simulated faithfully with rules (1–4), and \top
 1798 eventually spreads with rules (5–6).

1799 The case of $\ell = 0$ is the same except for the fact that rules (7–9) are disabled from the
 1800 very beginning, where the simulation of M already begins.

1801 *Case 2: $\iota(v)$ has several “fst”, “lst” or states from P ; or a wrongly placed “fst” or “lst”.* Since
 1802 the first component of each agent stops changing from v onwards, fairness and rules (10–12)
 1803 guarantee that at least one \perp -belief must occur infinitely often.

1804 *Case 3: $\ell = 0$ and $\iota(v)$ is invalid.* Rule (5) is never enabled and hence no \top ever appears.

1805 *Case 4: $\ell > 0$ and $\iota(v)$ is invalid.* By assumption, there is a position i such that $v[i] =$
 1806 (x, y, z, o) with $x \neq y$. By the latter, we have $O(v[i]) = \perp$. We consider three subcases.

1807 *Case 4-A: $\iota(v)$ has no “lst”.* By the lack of “lst”, rule (9) is permanently disabled and hence
 1808 agent i is stuck with output \perp .

1809 *Case 4-B: $\iota(v)$ has no “fst”.* The only way the population can possibly reach a \top -consensus
 1810 is by resetting agent i . This requires creating at least one \clubsuit . Since $\iota(v)$ contains no “fst”,
 1811 rule (9) can never be used. Therefore, if \clubsuit appears, then at least one \clubsuit must remain forever.
 1812 As $O((*, *, *, \clubsuit)) = \perp$, this means that there is always at least one agent with output \perp .

1813 *Case 4-C: $\iota(v)$ has no state from P .* Since the previous (sub)cases do not hold, agents 1 and n
 1814 respectively hold “fst” and “lst”. As agent i has a mismatch on its first two components, the
 1815 only way to obtain \top is by carrying a full reset from agent n to agent 1 using rules (7–9). If
 1816 this happens, then by the lack of state from P , rules (1–5) are permanently disabled. \blacktriangleleft

1817 **► Proposition 53.** *The language $L(M)$ is PP[+1]-semi-decidable.*

1818 **Proof.** Let K denote the set of valid initial configurations of M . Formally, let $K = \Gamma_0 \Gamma_1^* \Gamma_2$
 1819 where $\Gamma_0 = \{p_0\} \times \Sigma \times \{\text{fst}\}$, $\Gamma_1 = \{-\} \times \Sigma \times \{-\}$ and $\Gamma_2 = \{-\} \times \Sigma \times \{\text{lst}\}$.

1820 Let $f(i) = \Gamma_i$ and $U_i = \{w \in \{0, 1, 2\}^+ : |w|_i = 1\}$. We have

$$1821 K = f(0^* 1^* 2^* \cap U_0 \cap U_2).$$

1822 Thus, by Propositions 7 and 16 and by Lemmas 8 and 9, the language K is IO-PP[+1]-semi-
 1823 decidable with stabilizing inputs.

1824 Let $K' = \{w \in \Gamma^{\geq 2} \mid w \text{ is a valid configuration of } M \text{ and it leads to acceptance}\}$. By
 1825 Proposition 52, K' is PP[+1]-semi-decidable with stabilizing inputs. Let $g((x, y, z)) = y$. We
 1826 have

$$1827 L(M) \cap \Sigma^{\geq 2} = g(K \cap K').$$

1828 So, by Lemmas 8 and 9, the language $L(M) \cap \Sigma^{\geq 2}$ is PP[+1]-semi-decidable.

1829 We can trivially show that $L(M) \cap \Sigma$ is PP[\emptyset]-decidable. Thus, we are done by taking
 1830 the union, with Lemma 6. \blacktriangleleft

MB: The current
 Proposition 7
 is for IO-PP[<],
 but it is trivial
 to adapt it to
 IO-PP[+1].