

Population Protocols that Correspond to Symmetric Games[☆]

Olivier Bournez^a, Jérémie Chalopin^b, Johanne Cohen^c, Xavier Koegler^d

^a*Ecole Polytechnique & Laboratoire d'Informatique (LIX),
91128 Palaiseau Cedex, France*

^b*CNRS & Laboratoire d'Informatique Fondamentale de Marseille, CNRS &
Aix-Marseille Université,
39 rue Joliot Curie, 13453 Marseille Cedex 13, France*

^c*CNRS & PRiSM,
45 Avenue des Etats Unis, 78000 Versailles, France*

^d*École Normale Supérieure & Université Paris Diderot - Paris 7,
Case 7014, 75205 Paris Cedex 13, France*

Abstract

Population protocols have been introduced by Angluin et al. as a model of networks consisting of very limited mobile agents that interact in pairs but with no control over their own movement. A collection of anonymous agents, modeled by finite automata, interact pairwise according to some rules that update their states.

The model has been considered as a computational model in several papers. Input values are initially distributed among the agents, and the agents must eventually converge to the the correct output. Predicates on the initial configurations that can be computed by such protocols have been characterized under various hypotheses. The model has initially been motivated by sensor-networks, but it can be seen more generally as a model of networks of anonymous agents interacting pairwise. This includes sensor networks, ad-hoc networks, or models from chemistry.

In an orthogonal way, several distributed systems have been termed in literature as being realizations of games in the sense of game theory. In this paper, we investigate under which conditions population protocols, or

[☆]This work and all authors were partly supported by ANR Project SOGEA and by ANR Project SHAMAN, Xavier Koegler was partly supported by COST Action 295 DYNAMO and ANR Project ALADDIN

more generally pairwise interaction rules, can be considered as the result of a symmetric game. We prove that not all rules can be considered as symmetric games. We characterize the computational power of symmetric games. We prove that they have very limited power: they can count until 2, compute majority, but they can not even count until 3.

As a side effect of our study, we also prove that any population protocol can be simulated by a symmetric one (but not necessarily a game).

Key words: Population Protocols, Computation Theory, Distributed Computing, Algorithmic Game Theory

1. Introduction

The computational power of networks of anonymous resource-limited mobile agents has been investigated recently.

In particular, Angluin et al. proposed in [2] a model of distributed computations. In this model, called *population protocols*, finitely many finite-state agents interact in pairs chosen by an adversary. Each interaction has the effect of updating the state of the two agents according to a joint transition function.

A protocol is said to (*stably*) *compute* a predicate on the initial states of the agents if, in any fair execution, after finitely many interactions, all agents reach a common output that corresponds to the value of the predicate.

The model was originally proposed to model computations realized by sensor networks in which passive agents are carried along by other entities. The canonical example of [2] corresponds to sensors attached to a flock of birds and that must be programmed to check some global properties, like determining whether more than 5% of the population has elevated temperature. Motivating scenarios also include models of the propagation of trust [9].

Much of the work so far on population protocols has concentrated on characterizing which predicates on the initial states can be computed in different variants of the model and under various assumptions. In particular, the predicates computable by the unrestricted population protocols from [2] have been characterized as being precisely the semi-linear predicates, that is those predicates on counts of input agents definable in first-order Presburger arithmetic [20]. Semi-linearity was shown to be sufficient in [2] and necessary in [3].

Variants considered so far include restriction to one-way communications, restriction to particular interaction graphs, to random interactions, with possibly various kind of failures of agents. Solutions to classical problems of distributed algorithmics have also been considered in this model. Refer to [4] for a survey and a complete discussion.

The population protocol model shares many features with other models already considered in the literature. In particular, models of pairwise interactions have been used to study the propagation of diseases [14], or rumors [8]. In chemistry, the chemical master equation has been justified using (stochastic) pairwise interactions between the finitely many molecules [17, 13]. In that sense, the model of population protocols may be considered as fundamental in several fields of study, since it appears as soon as anonymous agents interact pairwise.

In an orthogonal way, pairwise interactions between finite-state agents are sometimes motivated by the study of the dynamics of particular two-player games from game theory. For example, paper [10] considers the dynamics of the so-called *PAVLOV* behavior in the iterated prisoner lemma. Several results about the time of convergence of this particular dynamics towards the stable state can be found in [10], and [11], for rings, and complete graphs.

The purpose of this article is to better understand whether and when pairwise interactions, and hence population protocols, can be considered as the result of a game. We want to understand if restricting to rules that come from a symmetric game is a limitation, and in particular whether restricting to rules that can be termed *PAVLOV* in the spirit of [10] is a limitation. We do so by giving solutions to several basic problems using rules of interactions associated to a symmetric game, and by characterizing their power: We prove that they can count until 2, they can compute *MAJORITY*, but they can not even count until 3.

As such protocols must also be symmetric, we are also discussing whether restricting to symmetric rules in population protocols is a limitation. We prove that any population protocol can be simulated by a symmetric one (but not necessarily a game).

In Section 2, we briefly recall population protocols. In Section 3, we recall some basics from game theory. In Section 4, we discuss how a game can be turned into a dynamics, and introduce the notion of Pavlovian population protocol. In Section 5 we prove that any symmetric deterministic 2-states population protocol is Pavlovian, and that the problem of computing the OR, AND, as well as the leader election and majority problem admit Pavlovian

solutions. We then characterize their power by proving that they can count until 2, but they can not count until 3 in Section 6. We prove that symmetric population protocols, unlike the restricted class of Pavlovian population protocols can compute all semi-linear predicate in Section 7.

Related work. Population protocols have been introduced in [2], and proved to compute all semi-linear predicates. They have been proved not to be able to compute more in [3]. Various restrictions on the initial model have been considered up to now. An (almost) up to date survey can be found in [4].

Variants include discussions about the influence of removing the assumption of two-way interaction: One-way interaction models include variants where agents communicate by anonymous message-passing, with immediate delivery or delayed delivery, or where agents can record it has sent a message, or queue incoming messages [1]. However, as far as we know, the constraint of restricting to symmetric rules has not been yet explicitly considered, nor restricting to rules that correspond to games in the population protocol literature.

More generally, population protocols arise as soon as populations of anonymous agents interact in pairs. Our original motivation was to consider rules corresponding to two-players games, and population protocols arose quite incidentally. The main advantage of the [2] settings is that it provides a clear understanding of what is called a computation by the model. Many distributed systems have been described as the result of games, but as far as we know there has not been attempts to characterize what can be computed by games in the spirit of this computational model.

In this paper, we turn two players games into dynamics over agents, by considering *PAVLOV* behavior. This is inspired by [10, 11, 16] that consider the dynamics of a particular set of rules termed the *PAVLOV* behavior in the iterated prisoner lemma. The *PAVLOV* behavior is sometimes also termed *WIN-STAY, LOSE-SHIFT* [18, 5]. Notice, that we extended it from two-strategies two-players games to n-strategies two-players games, whereas above references only talk about two-strategies two-players games, and mostly of the iterated prisoner lemma.

This is clearly not the only way to associate a dynamic to a game. There are several famous classical approaches: The first consists in repeating games: see for example [19, 6]. The second in using models from evolutionary game theory: refer to [15, 21] for a presentation of this latter approach. The approach considered here falls in method that consider dynamics obtained by

selecting at each step some players and let them play a fixed game. Alternatives to *PAVLOV* behavior could include *MYOPIC* dynamics (at each step each player chooses the best response to previously played strategy by its adversary), or the well-known and studied *FICTIOUS – PLAYER* dynamics (at each step each player chooses the best response to the statistics of the past history of strategies played by its adversary). We refer to [12, 6] for a presentation of results known about the properties of the obtained dynamics according to the properties of the underlying game. This is clearly non-exhaustive, and we refer to [5] for an incredible zoology of possible behaviors for the particular iterated prisoner lemma game, with discussions of their compared merits in experimental tournaments.

We obtain a characterization of the power of Pavlovian population protocols in terms of closure properties that show that they can count until 2, but not until 3. Notice that several variants of (one-way) population protocols have been characterized in [1] in a $COUNT_k$ hierarchy. The class obtained here seems close to the $COUNT_2$ level of this latter hierarchy [1]. However, on one hand, this is not exactly this class (for example *MAJORITY* is computable but not in the $COUNT_2$ level), and on the other hand, as no class there is formally proved to correspond to $COUNT_2$, this shows that the class considered here is different, and not reducible to the variants of [1].

Notice that a preliminary version of this article has been presented in *Complexity of Simple Programs CSP'08*. Compared to this conference version, we simplified some constructions, we added a few protocols, we extended deeply related work discussions, and mainly, we solved the statements conjectured there: we provide here a characterization of the power of Pavlovian population protocols, whereas it was open at the time of the presentation of this preliminary version.

2. Population Protocols

A protocol [2, 4] is given by $(Q, \Sigma, \iota, \omega, \delta)$ with the following components. Q is a finite set of *states*. Σ is a finite set of *input symbols*. $\iota : \Sigma \rightarrow Q$ is the initial state mapping, and $\omega : Q \rightarrow \{0, 1\}$ is the individual output function. $\delta \subseteq Q^4$ is a joint transition relation that describes how pairs of agents can interact. Relation δ is sometimes described by listing all possible interactions using the notation $(q_1, q_2) \rightarrow (q'_1, q'_2)$, or even the notation $q_1 q_2 \rightarrow q'_1 q'_2$, for $(q_1, q_2, q'_1, q'_2) \in \delta$ (with the convention that $(q_1, q_2) \rightarrow (q_1, q_2)$ when no rule is specified with (q_1, q_2) in the left-hand side). The protocol

is termed *deterministic* if for all pairs (q_1, q_2) there is only one pair (q'_1, q'_2) with $(q_1, q_2) \rightarrow (q'_1, q'_2)$. In that case, we write $\delta_1(q_1, q_2)$ for the unique q'_1 and $\delta_2(q_1, q_2)$ for the unique q'_2 .

Notice that, in general, rules can be non-symmetric: if $(q_1, q_2) \rightarrow (q'_1, q'_2)$, it does not necessarily follow that $(q_2, q_1) \rightarrow (q'_2, q'_1)$.

Computations of a protocol proceed in the following way. The computation takes place among n *agents*, where $n \geq 2$. A *configuration* of the system can be described by a vector of all the agents' states. The state of each agent is an element of Q . Because agents with the same states are indistinguishable, each configuration can be summarized as an unordered multiset of states, and hence of elements of Q .

Each agent is given initially some input value from Σ : Each agent's initial state is determined by applying ι to its input value. This determines the initial configuration of the population.

An execution of a protocol proceeds from the initial configuration by interactions between pairs of agents. Suppose that two agents in state q_1 and q_2 meet and have an interaction. They can change into state q'_1 and q'_2 if (q_1, q_2, q'_1, q'_2) is in the transition relation δ . If C and C' are two configurations, we write $C \rightarrow C'$ if C' can be obtained from C by a single interaction of two agents: this means that C contains two states q_1 and q_2 and C' is obtained by replacing q_1 and q_2 by q'_1 and q'_2 in C , where $(q_1, q_2, q'_1, q'_2) \in \delta$. An *execution* of the protocol is an infinite sequence of configurations C_0, C_1, C_2, \dots , where C_0 is an initial configuration and $C_i \rightarrow C_{i+1}$ for all $i \geq 0$. An execution is *fair* if for all configurations C that appear infinitely often in the execution, if $C \rightarrow C'$ for some configuration C' , then C' appears infinitely often in the execution.

At any point during an execution, each agent's state determines its output at that time. If the agent is in state q , its output value is $\omega(q)$. The configuration output is 0 (respectively 1) if all the individual outputs are 0 (respectively 1). If the individual outputs are mixed 0s and 1s then the output of the configuration is undefined.

Let p be a predicate over multisets of elements of Σ . Predicate p can be considered as a function whose range is $\{0, 1\}$ and whose domain is the collection of these multisets. The predicate is said to be computed by the protocol if, for every multiset I , and every fair execution that starts from the initial configuration corresponding to I , the output value of every agent eventually stabilizes to $p(I)$.

Multisets of elements of Σ are in clear bijection with elements of $\mathbb{N}^{|\Sigma|}$: a

multiset over Σ can be identified by a vector of $|\Sigma|$ components, where each component represents the multiplicity of the corresponding element of Σ in this multiset. It follows that predicates can also be considered as functions whose range is $\{0, 1\}$ and whose domain is $\mathbb{N}^{|\Sigma|}$.

The following was then proved in [2, 3].

Theorem 1 ([2, 3]). *A predicate is computable in the population protocol model if and only if it is semilinear.*

Recall that semilinear sets are known to correspond to predicates on counts of input agents definable in first-order Presburger arithmetic [20].

We will use the following notation as in [1]: the set of all functions from a set X to a set Y is denoted by Y^X . Let Σ be a finite non-empty set. For all $f, g \in \mathbb{R}^\Sigma$, we define the usual vector space operations

$$\begin{aligned} (f + g)(\sigma) &= f(\sigma) + g(\sigma) && \text{for all } \sigma \in \Sigma \\ (f - g)(\sigma) &= f(\sigma) - g(\sigma) && \text{for all } \sigma \in \Sigma \\ (cf)(\sigma) &= cf(\sigma) && \text{for all } \sigma \in \Sigma, c \in \mathbb{R} \\ (f \cdot g)(\sigma) &= \sum_{\sigma} f(\sigma)g(\sigma). \end{aligned}$$

Abusing notation as in [1], we will write σ for the function $\sigma(\sigma') = [\sigma = \sigma']$, for all $\sigma' \in \Sigma$, where $[\text{condition}]$ is 1 if condition is true, 0 otherwise.

3. Game Theory

We now recall the simplest concepts from Game Theory. We focus on non-cooperative games, with complete information, in extensive form.

The simplest game is made up of two players, called I and II , with a finite set of options, called *pure strategies*, $Strat(I)$ and $Strat(II)$. Denote by $A_{i,j}$ (respectively: $B_{i,j}$) the score for player I (resp. II) when I uses strategy $i \in Strat(I)$ and II uses strategy $j \in Strat(II)$.

The scores are given by $n \times m$ matrices A and B , where n and m are the cardinality of $Strat(I)$ and $Strat(II)$. The game is termed *symmetric* if A is the transpose of B : this implies that $n = m$, and we can assume without loss of generality that $Strat(I) = Strat(II)$.

In this paper, we will restrict to symmetric games.

Example 1 (Prisoner's dilemma). *The case where A and B are the following matrices*

$$A = \begin{pmatrix} R & S \\ T & P \end{pmatrix}, B = \begin{pmatrix} R & T \\ S & P \end{pmatrix}$$

with $T > R > P > S$ and $2R > T + S$, is called the prisoner's dilemma. We denote by C (for cooperation) the first pure strategy, and by D (for defection) the second pure strategy of each player.

As the game is symmetric, matrix A and B can also be denoted by:

		<i>Opponent</i>	
		C	D
<i>Player</i>	C	R	S
	D	T	P

A strategy $x \in \text{Strat}(I)$ is said to be a best response to strategy $y \in \text{Strat}(II)$, denoted by $x \in BR(y)$ if

$$A_{z,y} \leq A_{x,y} \tag{1}$$

for all strategies $z \in \text{Strat}(I)$.

A pair (x, y) is a (*pure*) *Nash equilibrium* if $x \in BR(y)$ and $y \in BR(x)$. A pure Nash equilibrium does not always exist.

In other words, two strategies (x, y) form a Nash equilibrium if in that state neither of the players has a unilateral interest to deviate from it.

Example 2. *On the example of the prisoner's dilemma, $BR(y) = D$ for all y , and $BR(x) = D$ for all x . So (D, D) is the unique Nash equilibrium, and it is pure. In it, each player has score P . The well-known paradox is that if they had played (C, C) (cooperation) they would have had score R , that is more. The social optimum (C, C) , is different from the equilibrium that is reached by rational players (D, D) , since in any other state, each player fears that the adversary plays C .*

We will also introduce the following definition: Given some strategy $x' \in \text{Strat}(I)$, a strategy $x \in \text{Strat}(I)$ is said to be a best response to strategy $y \in \text{Strat}(II)$ among those different from x' , denoted by $x \in BR_{\neq x'}(y)$ if

$$A_{z,y} \leq A_{x,y} \tag{2}$$

for all strategy $z \in \text{Strat}(I), z \neq x'$.

Of course, the roles of II and I can be inverted in the previous definition.

There are two main approaches to discuss dynamics of games. The first consists in repeating games [19, 6]. The second in using models from evolutionary game theory. Refer to [15, 21] for a presentation of this latter approach.

Repeating Games. Repeating k times a game, is equivalent to extending the space of choices into $Strat(I)^k$ and $Strat(II)^k$: player I (respectively II) chooses his or her action $x(t) \in Strat(I)$, (resp. $y(t) \in Strat(II)$) at time t for $t = 1, 2, \dots, k$. Hence, this is equivalent to a two-player game with respectively n^k and m^k choices for players.

To avoid confusion, we will call *actions* the choices $x(t), y(t)$ of each player at a given time, and *strategies* the sequences $X = x(1), \dots, x(k)$ and $Y = y(1), \dots, y(k)$, that is to say the strategies for the global game.

If the game is repeated an infinite number of times, a strategy becomes a function from integers to the set of actions, and the game is still equivalent to a two-player game¹.

Behaviors. In practice, player I (respectively II) has to solve the following problem at each time t : given the history of the game up to now, that is to say

$$X_{t-1} = x(1), \dots, x(t-1)$$

and

$$Y_{t-1} = y(1), \dots, y(t-1)$$

what should I (resp. II) play at time t ? In other words, how to choose $x(t) \in Strat(I)$? (resp. $y(t) \in Strat(II)$?)

It is natural to suppose that this is given by some behavior rules:

$$x(t) = f(X_{t-1}, Y_{t-1}),$$

$$y(t) = g(X_{t-1}, Y_{t-1})$$

for some particular functions f and g .

¹but whose matrices are infinite.

The Specific Case of the Prisoner's Lemma. The question of the best behavior rule to use for the prisoner lemma gave birth to an important literature. In particular, after the book [5], that describes the results of tournaments of behavior rules for the iterated prisoner lemma, and that argues that there exists a best behavior rule called *TIT – FOR – TAT*. This consists in cooperating at the first step, and then do the same thing as the adversary at subsequent times.

A lot of other behaviors, most of them with very picturesque names have been proposed and studied: see for example [5].

Among possible behaviors there is *PAVLOV* behavior: in the iterated prisoner lemma, a player cooperates if and only if both players opted for the same alternative in the previous move. This name [5, 16, 18] stems from the fact that this strategy embodies an almost reflex-like response to the payoff: it repeats its former move if it was rewarded by R or T points, but switches behavior if it was punished by receiving only P or S points. Refer to [18] for some study of this strategy in the spirit of Axelrod's tournaments.

The *PAVLOV* behavior can also be termed *WIN-STAY, LOSE-SHIFT* since if the play on the previous round results in a success, then the agent plays the same strategy on the next round. Alternatively, if the play resulted in a failure the agent switches to another action [5, 18].

Going From 2 Players to N Players. *PAVLOV* behavior is Markovian: a behavior f is *Markovian*, if $f(X_{t-1}, Y_{t-1})$ depends only on $x(t-1)$ and $y(t-1)$.

From such a behavior, it is easy to obtain a distributed dynamic. For example, let's follow [10], for the prisoner's dilemma.

Suppose that we have a connected graph $G = (V, E)$, with N vertices. The vertices correspond to players. An instantaneous configuration of the system is given by an element of $\{C, D\}^N$, that is to say by the state C or D of each vertex. Hence, there are 2^N configurations.

At each time t , one chooses randomly and uniformly one edge (i, j) of the graph. At this moment, players i and j play the prisoner dilemma with the *PAVLOV* behavior. It is easy to see that this corresponds to executing the following rules:

$$\left\{ \begin{array}{l} CC \rightarrow CC \\ CD \rightarrow DD \\ DC \rightarrow DD \\ DD \rightarrow CC. \end{array} \right. \quad (3)$$

What is the final state reached by the system? The underlying model is a very large Markov chain with 2^N states. The state $E^* = \{C\}^N$ is absorbing. If the graph G does not have any isolated vertex, this is the unique absorbing state, and there exists a sequence of transformations that transforms any state E into this state E^* . As a consequence, from well-known classical results in Markov chain theory, whatever the initial configuration is, with probability 1, the system will eventually be in state E^* [7]. The system is *self-stabilizing*.

Several results about the time of convergence towards this stable state can be found in [10], and [11], for rings, and complete graphs.

What is interesting in this example is that it shows how to go from a game, and a behavior to a distributed dynamic on a graph, and in particular to a population protocol when the graph is a complete graph.

4. From Games To Population Protocols

In the spirit of the previous discussion, to any symmetric game, we can associate a population protocol as follows.

Definition 1 (Associating a Protocol to a Game). *Assume a symmetric two-player game is given. Let Δ be some threshold.*

The protocol associated to the game is a population protocol whose set of states is Q , where $Q = \text{Strat}(I) = \text{Strat}(II)$ is the set of strategies of the game, and whose transition rules δ are given as follows:

$$(q_1, q_2, q'_1, q'_2) \in \delta$$

where

- $q'_1 = q_1$ when $M_{q_1, q_2} \geq \Delta$
- $q'_1 \in BR_{\neq q_1}(q_2)$ when $M_{q_1, q_2} < \Delta$

and

- $q'_2 = q_2$ when $M_{q_2, q_1} \geq \Delta$
- $q'_2 \in BR_{\neq q_2}(q_1)$ when $M_{q_2, q_1} < \Delta$,

where M is the matrix of the game.

Remark 1. By subtracting Δ to each entry of the matrix M , we can assume without loss of generality that $\Delta = 0$. We will do so from now on.

Definition 2 (Pavlovian Population Protocol). *A population protocol is Pavlovian if it can be obtained from a game as above.*

Remark 2. Clearly a Pavlovian population protocol must be symmetric: indeed, whenever $(q_1, q_2, q'_1, q'_2) \in \delta$, one has $(q_2, q_1, q'_2, q'_1) \in \delta$.

5. Some Specific Pavlovian Protocols

We now discuss whether assuming protocols Pavlovian is a restriction. We start by an easy consideration.

Theorem 2. *Any symmetric deterministic 2-states population protocol is Pavlovian.*

Proof. Consider a deterministic symmetric 2-states population protocol. Note $Q = \{+, -\}$ its set of states. Its transition function can be written as follows:

$$\left\{ \begin{array}{ll} ++ & \rightarrow \alpha_{++}\alpha_{++} \\ +- & \rightarrow \alpha_{+-}\alpha_{-+} \\ -+ & \rightarrow \alpha_{-+}\alpha_{+-} \\ -- & \rightarrow \alpha_{--}\alpha_{--} \end{array} \right. \quad (4)$$

for some $\alpha_{++}, \alpha_{+-}, \alpha_{-+}, \alpha_{--}$.

This corresponds to the symmetric game given by the following pay-off matrix M

		Opponent	
		+	-
Player	+	β_{++}	β_{+-}
	-	β_{-+}	β_{--}

where for all $q_1, q_2 \in \{+, -\}$,

- $\beta_{q_1 q_2} = 1$ if $\alpha_{q_1 q_2} = q_1$,
- $\beta_{q_1 q_2} = -1$ otherwise.

□

Unfortunately, not all rules correspond to a game.

Proposition 1. *Some symmetric population protocols are not Pavlovian.*

Proof. Consider for example a deterministic 3-states population protocol with set of states $Q = \{q_0, q_1, q_2\}$ and a joint transition function δ such that $\delta_1(q_0, q_0) = q_1$, $\delta_1(q_1, q_0) = q_2$, $\delta_1(q_2, q_0) = q_0$.

Assume by contradiction that there exists a 2-player game corresponding to this 3-states population protocol. Consider its payoff matrix M . Let $M(q_0, q_0) = \beta_0$, $M(q_1, q_0) = \beta_1$, $M(q_2, q_0) = \beta_2$. We must have $\beta_0 \geq \Delta = 0$, $\beta_1 \geq \Delta = 0$ since all agents that interact with an agent in state q_0 must change their state. Now, since q_0 changes to q_1 , q_1 must be a strictly better response to q_0 than q_2 : hence, we must have $\beta_1 > \beta_2$. In a similar way, since q_1 changes to q_2 , we must have $\beta_2 > \beta_0$, and since q_2 changes to q_0 , we must have $\beta_0 > \beta_1$. From $\beta_1 > \beta_2 > \beta_0$ we reach a contradiction. \square

This indeed motivates the following study, where we discuss which problems admit a Pavlovian solution.

5.1. Basic Protocols

Proposition 2. *There is a Pavlovian protocol that computes the logical OR (resp. AND) of input bits.*

Proof. Consider the following protocol to compute *OR*,

$$\left\{ \begin{array}{l} 01 \rightarrow 11 \\ 10 \rightarrow 11 \\ 00 \rightarrow 00 \\ 11 \rightarrow 11 \end{array} \right. \quad (5)$$

and the following protocol to compute *AND*,

$$\left\{ \begin{array}{l} 01 \rightarrow 00 \\ 10 \rightarrow 00 \\ 00 \rightarrow 00 \\ 11 \rightarrow 11 \end{array} \right. \quad (6)$$

Since they are both deterministic 2-states population protocols, they are Pavlovian. \square

Remark 3. Notice that OR (respectively AND) protocol corresponds to the predicate $[x.0 \geq 1]$ (resp. $[x.0 = 0]$), where x is the input. A simple change of notation yields a protocol to compute $[x.\sigma \geq 1]$ and $[x.\sigma = 0]$ for any input symbol σ .

Remark 4. All previous protocols are “naturally broadcasting” i.e., eventually all agents agree on some (the correct) value. With previous definitions (which are the classical ones for population protocols), the following protocol does not compute the XOR or input bits, or equivalently does not compute predicate $[x.1 \equiv 1 \pmod{2}]$.

$$\left\{ \begin{array}{l} 01 \rightarrow 01 \\ 10 \rightarrow 10 \\ 00 \rightarrow 00 \\ 11 \rightarrow 00 \end{array} \right. \quad (7)$$

Indeed, the answer is not eventually known by all the agents. It computes the XOR in a weaker form i.e., eventually, all agents will be in state 0, if the XOR of input bits is 0, or eventually only one agent will be in state 1, if the XOR of input bits is 1.

Proposition 3. There is a Pavlovian protocol that computes the threshold predicate $[x.\sigma \geq 2]$, which is true when there are at least 2 occurrences of input symbol σ in the input x .

Proof. The following protocol is a solution taking

- $\Sigma = \{0, \sigma\}$, $Q = \{0, \sigma, 2\}$,
- $\omega(0) = \omega(\sigma) = 0$,
- $\omega(2) = 1$.

$$\left\{ \begin{array}{l} 00 \rightarrow 00 \\ 0\sigma \rightarrow 0\sigma \\ \sigma 0 \rightarrow \sigma 0 \\ 02 \rightarrow 22 \\ 20 \rightarrow 22 \\ \sigma\sigma \rightarrow 22 \\ \sigma 2 \rightarrow 22 \\ 2\sigma \rightarrow 22 \\ 22 \rightarrow 22 \end{array} \right. \quad (8)$$

Indeed, if there is at least two σ , then by fairness and by the rule number 6, they will ultimately be changed into two 2s. Then 2s will turn all other agents into 2s. Now, this is the only way to create a 2.

This is a Pavlovian protocol since it corresponds to the following payoff matrix.

		Opponent		
		0	σ	2
Player	0	0	0	-1
	σ	0	-1	-1
	2	1	1	1

□

Hence, Pavlovian population protocols can count until 2. We will prove later on that they can not count until 3.

5.2. Leader Election

The classical solution [2] to the leader election problem (starting from a configuration with ≥ 1 leaders, eventually exactly one leader survives) is the following:

$$\left\{ \begin{array}{l} LL \rightarrow LN \\ LN \rightarrow LN \\ NL \rightarrow NL \\ NN \rightarrow NN \end{array} \right. \quad (9)$$

Notice that we use the terminology “leader election” as in [2] for this protocol, but that it may be considered more as a “mutual exclusion” protocol.

Unfortunately, this protocol is non-symmetric, and hence non-Pavlovian.

Remark 5. *Actually, the problem is with the first rule, since one wants two leaders to become only one. If the two leaders are identical, this is clearly problematic with symmetric rules.*

However, the leader election problem can actually be solved by a Pavlovian protocol, at the price of a less trivial protocol.

Proposition 4. *The following Pavlovian protocol solves the leader election problem, as soon as the population is of size ≥ 3 .*

$$\left\{ \begin{array}{l} L_1L_2 \rightarrow L_1N \\ L_1N \rightarrow NL_2 \\ L_2N \rightarrow NL_1 \\ NN \rightarrow NN \\ L_2L_1 \rightarrow NL_1 \\ NL_1 \rightarrow L_2N \\ NL_2 \rightarrow L_1N \\ L_1L_1 \rightarrow L_2L_2 \\ L_2L_2 \rightarrow L_1L_1 \end{array} \right. \quad (10)$$

Proof. Indeed, starting from a configuration containing not only N s, eventually after some time configurations will have exactly one leader, that is one agent in state L_1 or L_2 .

Indeed, the first rule and the fifth rule decrease strictly the number of leaders whenever there are more than two leaders. Now the other rules, preserve the number of leaders, and are made such that an L_1 can always be transformed into an L_2 and vice-versa, and hence are made such that a configuration where first or fifth rule applies can always be reached whenever there are more than two leaders. The fact that it solves the leader election problem then follows from the hypothesis of fairness in the definition of computations.

This is a Pavlovian protocol, since it corresponds to the following payoff matrix.

		Opponent		
		L_1	L_2	N
Player	L_1	-3	0	-3
	L_2	-1	-3	-3
	N	-2	-3	0

□

5.3. Majority

Proposition 5. *The majority problem (given some population of input symbols σ and σ' , determine whether there are more σ than σ') can be solved by a Pavlovian population protocol.*

Remark 6. *If one prefers, the predicate $[x.\sigma \geq x.\sigma']$, where σ and σ' are two input symbols, and x is the input, is computable by a Pavlovian population protocol.*

Proof. We claim that the following protocol outputs 1 if there are more σ than σ' in the initial configuration and 0 otherwise,

$$\left\{ \begin{array}{l} NY \rightarrow YY \\ YN \rightarrow YY \\ N\sigma \rightarrow Y\sigma \\ \sigma N \rightarrow \sigma Y \\ Y\sigma' \rightarrow N\sigma' \\ \sigma' Y \rightarrow \sigma' N \\ \sigma\sigma' \rightarrow NY \\ \sigma'\sigma \rightarrow YN \end{array} \right. \quad (11)$$

taking

- $\Sigma = \{\sigma, \sigma'\}, Q = \{\sigma, \sigma', Y, N\}$,
- $\omega(\sigma) = \omega(Y) = 1$,
- $\omega(\sigma') = \omega(N) = 0$.

In this protocol, the states Y and N are “neutral” elements for our predicate but they should be understood as *Yes* and *No*. They are the “answers” to the question: are there more 0s than 1s.

This protocol is made such that the numbers of σ and σ' are preserved except when a σ meets a σ' . In that latter case, the two agents are deleted and transformed into a Y and a N .

If there are initially strictly more σ than σ' , from the fairness condition, each σ' will be paired with a σ and at some point no σ' will left. By fairness and since there is still at least a σ , a configuration containing only σ and Y s will be reached. Since in such a configuration, no rule can modify the state of any agent, and since the output is defined and equals to 1 in such a configuration, the protocol is correct in this case

By symmetry, one can show that the protocol outputs 0 if there are initially strictly more σ' than σ .

Suppose now that initially, there are exactly the same number of σ and σ' . By fairness, there exists a step when no more agents in the state σ or σ'

left. Note that at the moment where the last σ is matched with the last σ' , a Y is created. Since this Y can be “broadcasted” over the N s, in the final configuration all agents are in the state Y and thus the output is correct.

This protocol is Pavlovian, since it corresponds to the following payoff matrix.

		Opponent			
		N	Y	σ	σ'
Player	N	1	-1	-1	1
	Y	0	1	1	-1
	σ	0	0	0	-1
	σ'	0	0	-1	0

□

6. Bounds on the Power of Pavlovian Population Protocols

We proved that predicates $[x.\sigma = 0]$, $[x.\sigma \geq 1]$, $[x.\sigma \geq 2]$ can be computed by some Pavlovian population protocols, as well as $[x.\sigma \geq x.\sigma']$.

It is clear that the subset of the predicates computable by Pavlovian population protocols is closed by negation: just switch the value of the individual output function of a protocol computing a predicate to get a protocol computing its negation.

Notice that, unlike what happens for general population protocols, composing Pavlovian population protocols into a Pavlovian population protocol is not easy. It is not clear whether Pavlovian computable predicates are closed by conjunctions: classical constructions for general population protocols can not be used directly.

The power of Pavlovian population protocols is actually rather limited as they can count up to 2, but not 3.

Theorem 3. *There is no Pavlovian protocol that computes the threshold predicate $[x.\sigma \geq 3]$, which is true when there are at least 3 occurrences of input symbol σ in the input x .*

Proof. We will prove this by contradiction. Assume there exists such a Pavlovian protocol. Without loss of generality we may assume that $\Sigma = \{0, \sigma\}$ is a subset of the set of states Q .

As the protocol is Pavlovian, and hence symmetric, any rule $qq \rightarrow q'q''$, is such that $q' = q''$, that is to say of the form $qq \rightarrow q'q'$ for all $q \in Q$.

Let then consider the sequence of rules such that $\sigma\sigma \rightarrow q_1q_1 \rightarrow q_2q_2 \rightarrow \dots \rightarrow q_kq_k \rightarrow \dots$ where $\sigma, q_1q_2, q_3, \dots, q_k \in Q$.

Since Q is finite, there exist two distinct integers k and ℓ such that $q_k = q_\ell$ and $k < \ell$.

The case $k + 1 = \ell$ is not possible. Indeed, we would have the rule $q_kq_k \rightarrow q_kq_k$. Consider the inputs x_3 and x_4 such that $x_3 = \{\sigma, \sigma\}$ and $x_4 = \{\sigma, \sigma, \sigma, \sigma\}$. x_4 must be accepted. From x_4 there is a derivation $x_4 \rightarrow \{q_1, q_1, \sigma, \sigma\} \rightarrow \{q_1, q_1, q_1q_1\} \rightarrow^* \{q_k, q_k, q_k, q_k\}$. This latter configuration is terminal from the above rule. Since x_4 must be accepted, we must have $\omega(q_k) = 1$. However, from x_3 there is a derivation $x_3 \rightarrow \{q_1, q_1\} \rightarrow^* \{q_k, q_k\}$, where the last configuration is also terminal. We reach a contradiction, since its output would be $\omega(q_k) = 1$, whereas x_3 must be rejected.

Hence, $k + 1 < \ell$, and $q_kq_k \rightarrow q_{k+1}q_{k+1} \rightarrow \dots \rightarrow q_\ell q_\ell \rightarrow q_kq_k$. Let T be then the set of states $T = \{q_i : k \leq i \leq \ell\}$.

Since $q_iq_i \rightarrow q_{i+1}q_{i+1}$ is among the rules, since the protocol is Pavlovian with a matrix M , and by definition of Pavlov behavior, we must have $q_{i+1} = BR_{\neq q_i}(q_i)$ (with the convention that $q_{\ell+1}$ is q_k). So, $BR(q_i)$ can be q_{i+1} or q_i .

Let then discuss the rules

$$q_iq_j \rightarrow q'_iq'_j \tag{12}$$

for $q_i, q_j \in T$.

There are three possibilities for the value of q'_i :

1. $q'_i = q_i$ if $M_{q_iq_j} \geq \Delta$
2. $q'_i = q_j$, if $M_{q_iq_j} < \Delta$ and if $q_j = BR_{\neq q_i}(q_j)$
3. $q'_i = q_{j+1}$ if $M_{q_iq_j} < \Delta$ and if $q_{j+1} = BR_{\neq q_i}(q_j)$

In any case, we see that the value of q'_i is in T .

Symmetrically, we have three possibilities for q'_j , all of them in T .

Hence, all rules of the form (12) preserve T : we have $q'_i, q'_j \in T$, as soon as $q_i, q_j \in T$.

Consider still then the inputs x_3 and x_4 such that $x_3 = \{\sigma, \sigma\}$ and $x_4 = \{\sigma, \sigma, \sigma, \sigma\}$. From x_4 there is a derivation $x_4 \rightarrow \{q_1, q_1, \sigma, \sigma\} \rightarrow \{q_1, q_1, q_1q_1\} \rightarrow^* \{q_k, q_k, q_k, q_k\}$. From this last configuration, by above remark, the state of all agents will be in T . As x_4 must be accepted, ultimately all agents will be in states that belong to T whose image by ω is 1. Consider now x_3 . From x_3 there is a derivation $x_3 \rightarrow \{q_1, q_1\} \rightarrow^* \{q_k, q_k\}$ that then will go through all configurations $\{q_iq_i\}$, for the $q_i \in T$ in turn. This can

not eventually stabilize to elements whose image by ω is 0, as some of the elements of T have image 1 by ω , and hence x_3 is not accepted. This yields a contradiction, and hence such a Pavlovian protocol can not exist. \square

7. The Power of Symmetric Population Protocols

Pavlovian Population protocols are symmetric. We just proved that they have a very limited computational power. However, assuming population protocols symmetric (not-necessarily Pavlovian) is not truly a restriction.

Proposition 6. *Any population protocol can be simulated by a symmetric population protocol, as soon as the population is of size ≥ 3 .*

Before proving this proposition, we state the (immediate) main consequence.

Corollary 1. *A predicate is computable by a symmetric population protocol if and only if it is semilinear.*

Proof. To a population protocol $(Q, \Sigma, \iota, \omega, \delta)$, with $Q = \{q_1, \dots, q_n\}$ associate population protocol $(Q \cup Q', \Sigma, \iota, \omega, \delta')$ with $Q' = \{q'_1, \dots, q'_n\}$, $\omega(q') = \omega(q)$ for all $q \in Q$, and for all rules

$$qq \rightarrow \alpha\beta$$

in δ , the following rules in δ' :

$$\left\{ \begin{array}{l} qq' \rightarrow \alpha\beta \\ q'q \rightarrow \beta\alpha \\ qq \rightarrow q'q' \\ q'q' \rightarrow qq \\ q\gamma \rightarrow q'\gamma \\ q'\gamma \rightarrow q\gamma \\ \gamma q \rightarrow \gamma q' \\ \gamma q' \rightarrow \gamma q \end{array} \right.$$

for all $\gamma \in Q \cup Q', \gamma \neq q, \gamma \neq q'$, and for all pairs of rules

$$\left\{ \begin{array}{l} qr \rightarrow \alpha\beta \\ rq \rightarrow \delta\epsilon \end{array} \right.$$

with $q, r \in Q$, the following rules in δ' :

$$\left\{ \begin{array}{l} qr' \rightarrow \alpha\beta \\ r'q \rightarrow \beta\alpha \\ rq' \rightarrow \delta\epsilon \\ q'r \rightarrow \epsilon\delta. \end{array} \right.$$

The obtained population protocol is clearly symmetric. Now the first set of rules guarantees that a state in Q can always be converted to its primed version in Q' and vice-versa. By fairness, whenever a rule $qq \rightarrow \alpha\beta$ (respectively $qr \rightarrow \alpha\beta$) can be applied, then the corresponding two first rules of the first set of rules (resp. of the second set of rules) can eventually be fired after possibly some conversions of states into their primed version or vice-versa. \square

References

- [1] D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
- [2] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. In *Twenty-Third ACM Symposium on Principles of Distributed Computing*, pages 290–299. ACM Press, July 2004.
- [3] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *PODC '06: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 292–299, New York, NY, USA, 2006. ACM Press.
- [4] James Aspnes and Eric Ruppert. An introduction to population protocols. In *Bulletin of the EATCS*, volume 93, pages 106–125, 2007.
- [5] Robert M. Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.
- [6] Ken Binmore. *Jeux et Théorie des jeux*. DeBoeck Université, Paris-Bruxelles, 1999. Translated from ‘Fun and Games: a text on game theory’ by Francis Bismans and Eulalia Damaso.

- [7] Pierre Brémaud. *Markov Chains, Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, New York, 2001.
- [8] DJ Daley and DG Kendall. Stochastic Rumours. *IMA Journal of Applied Mathematics*, 1(1):42–55, 1965.
- [9] Z. Diamadi and M.J. Fischer. A simple game for the study of trust in distributed systems. *Wuhan University Journal of Natural Sciences*, 6(1-2):72–82, 2001.
- [10] Martin E. Dyer, Leslie Ann Goldberg, Catherine S. Greenhill, Gabriel Istrate, and Mark Jerrum. Convergence of the iterated prisoner’s dilemma game. *Combinatorics, Probability & Computing*, 11(2), 2002.
- [11] Laurent Fribourg, Stéphane Messika, and Claudine Picaronny. Coupling and self-stabilization. In Rachid Guerraoui, editor, *Distributed Computing, 18th International Conference, DISC 2004, Amsterdam, The Netherlands, October 4-7, 2004, Proceedings*, volume 3274 of *Lecture Notes in Computer Science*, pages 201–215. Springer, 2004.
- [12] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. Number 624. December 1996. available at <http://ideas.repec.org/p/cla/levarc/624.html>.
- [13] D.T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188(1-3):404–425, 1992.
- [14] Herbert W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, December 2000.
- [15] J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 4:479–519, 2003.
- [16] D. Kraines and V. Kraines. Pavlov and the prisoner’s dilemma. *Theory and Decision*, 26:47–79, 1988.
- [17] James Dickson Murray. *Mathematical Biology. I: An Introduction*. Springer, third edition, 2002.
- [18] M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game. *Nature*, 364(6432):56–58, 1993.

- [19] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [20] M. Presburger. Über die Vollständigkeit eines gewissen systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. *Comptes-rendus du I Congres des Mathematicians des Pays Slaves*, pages 92–101, 1929.
- [21] Jörgen W. Weibull. *Evolutionary Game Theory*. The MIT Press, 1995.