



Dynamic Computation of Population Protocols

Marin Bertier, Yann Busnel, Anne-Marie Kermarrec

► **To cite this version:**

Marin Bertier, Yann Busnel, Anne-Marie Kermarrec. Dynamic Computation of Population Protocols. IEEE. the 17th IEEE International Conference on Telecommunications - Ad-hoc and Sensor Communications track (ICT 2010), Apr 2010, Doha, Qatar. pp.100-107, 2010.

HAL Id: hal-00480988

<https://hal.archives-ouvertes.fr/hal-00480988>

Submitted on 5 May 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamic Computation of Population Protocols

Marin Bertier
IRISA / INSA Rennes
Campus universitaire de Beaulieu
35042 Rennes Cedex – FRANCE
Email: Marin.Bertier@irisa.fr
Tel: +33 2 99 84 22 06

Yann Busnel
LINA / University of Nantes
2, rue de la Houssinière – BP 92208
44322 Nantes Cedex 03 – FRANCE
Email: Yann.Busnel@univ-nantes.fr
Tel: +33 2 51 12 58 95

Anne-Marie Kermarrec
INRIA Rennes - Bretagne Atlantique
Campus universitaire de Beaulieu
35042 Rennes Cedex – FRANCE
Email: Anne-Marie.Kermarrec@inria.fr
Tel: +33 2 99 84 25 98

Abstract—Population protocols provide theoretical foundations for mobile tiny device networks in which global behavior emerges from a set of simple interactions between anonymous agents. The works in this area mostly focus on studying the computational power of the model. Results hold as long as a fair scheduler, which governs the interactions between nodes, ensures that all reachable system states may eventually happen.

This paper studies for the first time the impact of the agents' mobility model on the convergence speed of population protocols, emphasizing the dynamic of the computation. We propose an augmented population protocol model where each edge of the interaction graph is weighted, representing the probability of two agents to interact. This model enables to define the behavior of the scheduler under various mobility models. We have empirically shown that mobility models have a significant impact on the convergence speed of the protocols [10]. In fact, we observed that the uniform distribution always provides the best convergence time. Such a model is representative of the well-known Random Way Point model used to evaluate most of mobile ad-hoc network protocols.

In this paper, we formally prove that a uniform distribution of weights provides the lowest bound of average convergence speed for a large class of population protocols. Therefore, this analysis reveals that the Random Way Point model, following this distribution, provides the best case scenario questioning its relevance as a reference model.

I. INTRODUCTION

Sensor networks are composed of tiny computation units able to communicate and collect data from their environment. This leads to a whole class of applications, where sensors are embedded on human or animals, aiming at observing their behavior or computing global properties. Such settings impose intermittent and arbitrary communications between sensors, which are specifically studied in the context of *Delay Tolerant Network (DTN)* [9] and *Population Protocols* [1]. Delay Tolerant Network may be seen as an evolution of *Mobile Ad-hoc NETWORK (MANET)* [21]. In the DTN area, the challenge consists in designing distributed applications able to cope with the uncertainty on the connectivity and the dynamicity induced by the mobility. Population Protocols propose a formalism to study the convergence of distributed algorithms based on a succession of interactions between nodes. These two approaches are complementary. In this paper, we leverage these complementarities and propose an enriched population protocol model to deal with mobility patterns.

MANET and DTN: Consider a network composed of mobile nodes equipped with wireless networking capabilities and able to communicate with each other only when they are within transmission range. A common DTN scenario considers a network suffering from frequent connectivity disruptions, making the topology only intermittently and partially connected. Most DTN algorithms propose probabilistic solutions, where the probabilities are strongly dependent on the mobility model. Many works in this area focus on establishing the correct mobility model to represent human's movement behaviors in order to determine how to optimize distributed algorithms.

Population protocols: Population protocols, introduced in [1], model the interactions between mobile agents with very limited power. Population protocols provide common theoretical foundations for distributed systems in which a global behavior emerges from a set of simple interaction between nodes. Population protocols consist of finite sets of states, inputs, outputs and a transition function. The set of possible node interactions is represented by a graph of interactions. An interaction represents the fact that two agents are sufficiently close, for a sufficiently long time, to interact by exchanging their local information.

The power of population protocols lies in the simplicity of the model. No specific assumption is made on the agents' synchrony, the system infrastructure or the order of the interactions. A *scheduler* which is only assumed to be *fair* guides the way the interactions actually take place. A fair scheduler simply ensures that every possible evolution of the system eventually happens.

On the impact of mobility: Most of the works in population protocols focus on the computational power of such protocols [3], [2], [4], [5], [15]. No specific assumption is usually made on the model of interactions between agents. When a specific interaction model is considered, it is usually uniform [3]. Yet, in reality, it is unlikely that the mobility pattern of a set of mobile entities is uniform. Although this has no impact on the actual power of the model and the asymptotic convergence, it is very likely that the convergence speed of such a system is impacted by the mobility patterns of agents. In mobile networks, it has been shown that mobility models strongly impact the outcome of a protocol [16]. As a consequence, characterizing realistic mobility models [12], [14] is an active area of research.

Contributions: In this paper, we study the impact of the agents’ mobility model on the convergence speed of population protocols. To this end, we introduce *MAPP (Mobility Applied to Population Protocol)*, a population protocol model augmented to take into account the probability distribution of agents’ interactions. To this end, to each edge of the interaction graph is assigned a weight reflecting the interaction frequency between the two agents linked by that edge.

In the former empirical study [10], we consider a uniform interaction distribution as well as a set of non-uniform mobility patterns identified by the mobile network community [12], [14]. We observed a significant impact of the mobility patterns on the speed of convergence. More specifically, a uniform distribution turns out to consistently achieve the best average convergence time.

Then, in this paper, we formally prove that this uniform distribution of weights provides the lowest bound of the average number of steps to reach convergence in a large class of population protocols.

Finally, we also demonstrate that the Random Way Point mobility model, which is extensively used to evaluate mobile networks, implies a uniform distribution in MAPP. Interestingly enough, the Random Way Point model is recognized as non-realistic but it is usually used as a neutral setting. In this paper, we show that it actually provides the *best-case* scenario. This may raise questions about the relevance of using such a model as a representative mobility model.

Roadmap: The rest of the paper is organized as follows. Section II provides the computation models of population protocols as well as MAPP, our extension of this model. Section III introduces the mathematical background. Section IV presents the proof that the uniform distribution leads to the lowest bound in term of average number of steps required to converge in population protocols. Finally, in Section V, we prove that the Random Way Point model can be modeled by a uniform distribution. We conclude and list some open issues in Section VI.

II. COMPUTATION MODELS

A. Population protocols

The original population protocol model [1], [6] considers a collection of agents with an associated input value. Pairwise interaction of agents is governed by a fair scheduler (*cf.* Section II-B). An agent in this model is represented as a finite state machine, for which state is updated upon interaction only. Updates are defined by a transition function δ . Agents compute an output value related to their current state. This value eventually converges to the expected correct output value.

More formally, a population protocol is composed of:

- an interaction graph $\Lambda(\Upsilon, \Theta)$ where Υ represents a set of $n \geq 2$ anonymous agents and Θ the set of all possible pairwise interactions between these agents. In the basic model, $\Theta = \{(v, v') \in \Upsilon^2 | v \neq v'\}$ (Λ is complete);
- a finite input alphabet Σ ;
- a finite output alphabet Y ;

- a finite set of possible agent’s states Q ;
- an input function $\iota : \Sigma \rightarrow Q$ mapping inputs to states;
- an output function $\omega : Q \rightarrow Y$ mapping states to outputs;
- a transition relation $\delta : Q \times Q \rightarrow Q \times Q$ on pair of states.

In the following, we call $(p, q) \mapsto (p', q')$ a transition if $[(p, q) \mapsto (p', q')] \in \delta$. A transition can occur between two agents’ states only if these two agents have an interaction. The protocol is deterministic if δ is a function (*i.e.* at most one possible transition for each pair in Q^2).

A configuration of the system corresponds to a mapping vector of all agents’ states. As agents are anonymous, two agents with the same state are indistinguishable. Then, each configuration can be viewed as an unordered multiset of states. We denote $C \rightarrow C'$ the fact that a configuration C' can be obtained from C in one step (*i.e.* with only one transition for one existing interaction $\theta \in \Theta$). An execution of the protocol is a finite or infinite sequence of population configurations C_0, C_1, C_2, \dots such that $\forall i, C_i \rightarrow C_{i+1}$.

In brief, a population protocol stably computes a function $f : \Sigma^+ \rightarrow Y$ if $\forall n \in \mathbb{N}, \forall \sigma \in \Sigma^n$, every fair execution, with n agents initialized with the elements of σ , eventually stabilizes to output $f(\sigma)$ (*i.e.* output value of every agent eventually stabilizes to $f(\sigma)$).

The MAPP extension affects only the interaction scheduler without jeopardizing its fairness, nor the computable predicates (which corresponds to the Presburger arithmetic¹). In fact, putting a probability distribution on schedules may allow for computation of more powerful predicates with high probability (*e.g.* [3]). This problem is avoided here as we are still requiring convergence in *all* executions, even those that occur with low probability.

B. Mobility Applied to Population Protocols

In order to formally study the impact of the mobility model on the convergence speed of protocols, we propose an extension of this model.

Modeling the scheduler: In the population protocol model, a fair scheduler determines the order of interactions. This fairness assumption ensures that an attainable state can be effectively reached. More formally, considering a given configuration C , for every configuration C' obtained from C with a single interaction of two agents (*i.e.* $C \rightarrow C'$), if C appears infinitely often during the execution, then the configuration C' must also appear infinitely often.

The main objective of our extension is to model the scheduler, to define its heuristics in term of pair selection during a schedule, while ensuring this fairness condition, and then capture mobility patterns. In *Mobility Applied to Population Protocol (MAPP)*, a weight is assigned to each edge of the interaction graph, reflecting the probability for an interaction to happen at the next step of the execution between the agents connected by this edge. As we consider the interaction as atomic, two concurrent interactions can be viewed as sequential. Moreover, the order of these interactions

¹For more details, see [2].

does not matter as all the agents involved in the latter are distinct.

In the following, we extend the interaction graph definition as $\Lambda(\Upsilon, \Theta)$ such that

$$\forall \theta \in \Theta, \exists v, v' \in \Upsilon, v \neq v' \wedge \theta = v \xrightarrow{p_{v,v'}} v'.$$

We denote this probability $p_{v,v'}$ or p_θ . Without loss of generality, we can assume that $\sum_{\theta \in \Theta} p_\theta = 1$.

From a practical point of view, MAPP allows to model the choices of the scheduler during the execution of a given protocol. In fact, given a population in a configuration C , the scheduler chooses the protagonists of the following interaction according to the probability of those agents to meet.

In the following, we denote as *out-weight distribution of v* (OWD) the set of all the out-degree weights of a given agent v . More formally, given an agent $v \in \Upsilon$, its OWD is defined by the following multiset: $\mathcal{P}_v = \{p_{v,\psi} | \psi \in \Upsilon - \{v\}\}$, that contains all the interaction probabilities in which v is involved. For a population of n agents, it is obvious that $|\mathcal{P}_v| = n - 1$.

On the relevance of this extension: We must first check that our model preserves the fairness condition imposed on the scheduler.

Lemma 1: For all schedulers following the probability distribution of a given MAPP, this scheduler also respects the fairness assumption.

Proof: $\forall \theta \in \Theta$, we consider that $p_\theta > 0$. Then, the probability that this interaction occurs instead of any other is not null. Then, given two specific configurations C and C' such that $C \rightarrow C'$ with the interaction θ , the probability that this transition will be chosen by the scheduler is not null either. So, if C appears infinitely in the execution then C' will also appear infinitely in the execution. In fact, any execution permitted by the MAPP model and not fair, happens with probability 0. ■

C. Related works and integration

Recently, an important number of studies has been conducted on the former model, or its various extensions [6]. For instance, adding some constrains on the possible interaction scheduler (while obviously ensuring the fairness assumption) [3] permit to analyze the real-time evolution of the system states. We argues that MAPP is more generic than [3] as the latter paper only studies the uniform OWD.

Considering side works according to our approach, self-stabilization of population protocols has also be studied, without initial state knowledge, based strictly on the input set [5]. Also, Delporte-Gallet *et al.* introduced how agent failure could affect the computation power of population protocols [15]. Notably, they shown that, starting from any protocol computing a function in the former model, it exist a generic transformation which provide a $O(1)$ -fault tolerant protocol. However, this method requires some weaker characteristic of the problem. In [4], Angluin *et al.* claimed that the unidirectional communication model (which is equivalent to a directed interaction graph) implies a restricted computation power, in comparison with the classical bidirectional one.

One of the main advantages of our extension consists in its simple integration among all variant extension starting from the population protocols model. For instance, we show that the computational power of MAPP is unchanged from the former model. Thus, any result according to the self-organization of these protocols [5] and to take into account failures [15] remain valid in MAPP.

Moreover, using MAPP, we can model most of the extensions. Consequently, considering a *restricted interaction graph* [1] can be viewed as a MAPP with part of edges of interaction graph labeled by 0. As well, all result proposed in the context of *random interactions* [1], [3] are valid for all MAPP that are set using a uniform OWD. Finally, different models introduced in [4] concerning the modeling of one-way communications remain valid in MAPP, using an directed graph.

III. THEORETICAL ANALYSIS OF CONVERGENCE

Many papers [1], [3], [2], [4], [5], [7], [15], [17] have investigated the power of population protocols and some extensions. Yet, very few address the convergence speed (in the context of population protocols, it corresponds to the average number of steps or the time needed to reach the stabilized configuration). MAPP provides a simple model to analyze theoretically the system evolution the over time.

In this paper, we present some results issued from these analyses. We introduce below the necessary mathematic tools. Therefore, due to space constraints, we only focus on fundamental theoretical results but we do not develop on estimating the behavior of population according to a specific protocol. However, some of these empirical analyses are available in [10].

Definitions

Using MAPP, it is possible to model the behavior of the scheduler, and then, formally study the state evolution of a population. At each step, the scheduler chooses a specific pair of agents according to the probability given in the interaction graph Λ . Obviously, the interaction chosen at a specific step is independent of the past choices. Therefore, a given MAPP can be seen as a *Markov chain with a finite state space*, as the number of agents and the size of Q (the set of possible state of an agent) are finite. Then, the transition probability distribution of the Markov chain can be represented by a stochastic matrix, in which the evolution of the system can be extracted *a priori*. Effectively, the Markov chain describing the system evolution is composed of the set of all configurations of the given MAPP, and the set of transitions is directly extracted from the probability distribution available in the interaction graph.

In the following, we consider T an ordering index set, such as the natural numbers \mathbb{N} , the non-negative real numbers $[0, +\infty)$, or a subset of these. Elements $t \in T$ can be thought of as “times”. From this time notion, given a stochastic process, it is possible to infer the time corresponding to the “occurrence” of the system in a specific state. Consequently,

given that the Markov chain is ergodic, we can define the expectation of this time, corresponding to a mean hitting time:

Definition 1 (First hit time): Given a probability space (Ω, Σ, Pr) and a measurable state space S , let $X : \Omega \times T \rightarrow S$ be a stochastic process, and let \mathcal{A} be a measurable subset of the state space S . The *first hit time* $\tau_{\mathcal{A}} : \Omega \rightarrow [0, +\infty]$ is the random variable defined by

$$\tau_{\mathcal{A}}(\omega) = \inf\{t \in T \mid X_t(\omega) \in \mathcal{A}\}$$

Definition 2 (Mean & variance of hitting time): Given a state \mathcal{I} and a set of first time hit times $\{\tau_{\mathcal{A}}^i\}_{i \in \{n \in \mathbb{N} : n < N\}}$, the *mean hitting time* corresponds to the expected value of $\tau_{\mathcal{A}}$ starting on state \mathcal{I} :

$$\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{A}}) = \lim_{N \rightarrow +\infty} \frac{1}{N} \cdot \sum_{i=1}^N \tau_{\mathcal{A}}^i$$

and the *variance hitting time* is defined by:

$$\sigma_{\mathcal{I}}^2(\tau_{\mathcal{A}}) = \lim_{N \rightarrow +\infty} \frac{1}{N-1} \cdot \sum_{i=1}^N (\tau_{\mathcal{A}}^i - \mathbb{E}(\tau_{\mathcal{A}}))^2$$

In the following, *convergence* refers to the point at which the stationary state in the Markov chain associated to a MAPP is reached (i.e. the stable configuration of a given population). Thus, the term *convergence speed* represents the mean hitting time of the stationary distribution.

IV. A LOWER BOUND: UNIFORMITY IS ALWAYS BEST

We observe from the simulations in [10] that the mobility patterns of agents might have a significant impact on the convergence speed of population protocols. More specifically, we observe that the uniform distribution consistently leads to higher convergence speed than the other distributions. Although this seems natural for some operations (as for the *sum* operation in which all agents have to interact with each other several times to converge – cf. [1]), it is actually counter-intuitive for some others e.g. flooding. Indeed, for the *or* operation, one might think that a global OWD, in which the source agent is involved in almost all the interactions², converges faster than the uniform OWD.

In this section, we propose a theorem that rebuts this intuition. Consider a class of population protocols, which contains resolution for all possible semi-linear predicates. One way to compute predicates definable in Presburger arithmetic is to reduce them to a particular class of base protocols that involve leader agents wandering around picking up values (as shown in [1]). This class, so called *reducePP*, has the same size than the class of computable population protocol. Below, we prove, for the latter class, that a uniform distribution of weights always achieves the best convergence speed, regardless of the considered operation in population protocols. That means that the uniform OWD corresponds to the lower bound of the mean

hitting time for *reducePP* (i.e. the average number of steps needed to reach the stabilized state of a population).

Theorem 2: For any function computable by a population protocol in *reducePP*, the lower bound of the convergence speed is reached using a uniform OWD in MAPP.

Proof: Roughly speaking, in this proof, we first (i) characterize any predicate computable by a population protocol. Then, (ii) by characterizing how to compute them using a combination of population protocols, we will show that (iii) any population protocol in *reducePP* has a polynomial mean hitting time. Finally, we prove that (iv) any polynomial mean hitting time have a lower bound for a uniform OWD.

Characterization of computable functions: In [1], Theorem 5 states that every predicate belonging to the Presburger arithmetic is stationary computable by a population protocol. This arithmetic fully characterizes the wholeness of population protocols' computable function [2]. Thus, these two domains of function share the same equivalence class.

First, consider a factorization of all Presburger arithmetic's predicate. Let $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ be an arbitrary input alphabet, and $A_i, c, m \geq 2$ integer constants. Then, it has been shown that the entire Presburger arithmetic can be stationary computable using a combination of the following predicates on non-negative integers x_1, \dots, x_k :

- $\sum_i a_i x_i < c$;
- $\sum_i a_i x_i \equiv_m c$ (i.e. $\sum_i a_i x_i \equiv c$ modulo m);
- Any 2-place Boolean function ξ .

We present below the population protocols that compute each element of these last generator set of predicates. We prove below that for any combination of this predicates, a uniform OWD in MAPP corresponds to the optimal convergence speed of this combination. Thus, any predicate of the Presburger arithmetic, computed by a population protocol from *reducePP*, has an optimal convergence speed with the uniform OWD.

Generator population protocols' definition: We now present in details the population protocols, introduced in [1], which compute the two first aforementioned predicates.

Let $s = \max(|c| + 1, m, \max_i |a_i|)$. In both protocols, the set of space Q is the set $\{0, 1\} \times \{0, 1\} \times \{u \in \mathbb{Z} \mid -s \leq u \leq s\}$, and the associated function ι corresponds to $\sigma_i \mapsto (1, 0, a_i)$. The first *bit* of the state is denoted the *leader bit* and is used to elect an unique leader, which aggregates the value of the linear combination. The second bit is denoted the *output bit* which stores, for each agent, the output value computed by the last encounter leader. The third entry of a state is a *counter* used for collecting the linear combination of x_i (left-hand side of the previous predicates). The output function ω simply maps (\cdot, b, \cdot) to b .

We now describe the transition rules for each of the two protocols (the correction of these protocol is proved in [1]).

- Consider, for all integers u, u' such that $-s \leq u, u' \leq s$, the two following functions:

$$\begin{cases} q(u, u') = \max(-s, \min(s, u + u')) \\ r(u, u') = u + u' - q(u, u') \end{cases}$$

²For instance, consider the following distribution: one agent has a huge probability to interact with any other agent, and all other possible interactions have a tiny probability to happen.

It is obvious that $q(u, u'), r(u, u') \in [-s, s]$ and that $q(u, u') + r(u, u') = u + u'$. We define $b(u, u') = 1$ if $q(u, u') < c$ and 0 otherwise. The δ function is define as follow if at least ℓ or ℓ' are equal to 1:

$$\begin{aligned} & (\ell, \cdot, u), (\ell', \cdot, u') \\ & \rightarrow (1, b(u, u'), q(u, u')), (0, b(u, u'), r(u, u')). \end{aligned}$$

In the case where both ℓ and ℓ' are null, the interaction has no effect.

- Consider now that $b(u, u') = 1$ if $u + u' \equiv_m c$ and 0 otherwise. The second protocol works using the following class of transition:

$$\begin{aligned} & (\ell, \cdot, u), (\ell', \cdot, u') \\ & \rightarrow (1, b(u, u'), (u + u') \bmod m), (0, b(u, u'), 0). \end{aligned}$$

if at least ℓ or ℓ' are equal to 1. Otherwise, when both ℓ and ℓ' are null, the interaction has still no effect.

Let consider the last item of aforementioned predicates. Let ξ a 2-place Boolean function. In the proof of Lemma 3 in [1], computation of a ξ 's combination of two stationary computable predicates F and G is done as follow. Let \mathcal{A} (respectively \mathcal{B}) a protocol that stationary computes F (respectively G); we assume that \mathcal{A} and \mathcal{B} share the same input set Σ . Consider the protocol \mathcal{C} which stationary computes $\xi(F, G)$ by parallel composition of \mathcal{A} and \mathcal{B} (the population runs protocols \mathcal{A} and \mathcal{B} in parallel and outputs the value of ξ applies to the outputs of the two computed predicates F et G).

In more details, let $Q_{\mathcal{A}}$ and $Q_{\mathcal{B}}$ the state set respectively of \mathcal{A} and \mathcal{B} . The set of states of \mathcal{C} is defined as $Q_{\mathcal{C}} = Q_{\mathcal{A}} \times Q_{\mathcal{B}}$. The associated input map $\iota_{\mathcal{C}}$ corresponds to $\sigma \in \Sigma \mapsto (\iota_{\mathcal{A}}(\sigma), \iota_{\mathcal{B}}(\sigma))$ and the transition function is defined as $\delta_{\mathcal{C}}((p_1, p_2), (q_1, q_2)) = ((p'_1, p'_2), (q'_1, q'_2))$ with $\delta_{\mathcal{A}}(p_1, q_1) = (p'_1, q'_1)$ and $\delta_{\mathcal{B}}(p_2, q_2) = (p'_2, q'_2)$. Finally, the output map applies ξ to the both protocols' outputs:

$$\omega_{\mathcal{C}}((q_1, q_2)) = \xi(\omega_{\mathcal{A}}(q_1), \omega_{\mathcal{B}}(q_2)).$$

The convergence speed of \mathcal{C} directly depends on the speed of \mathcal{A} and \mathcal{B} . In more details, the convergence speed of \mathcal{C} is exactly the same than the lowest one between \mathcal{A} and \mathcal{B} . Without loss of generality, we assume that \mathcal{A} has a mean hitting time lower than \mathcal{B} . Then, the mean hitting time of \mathcal{C} is the same than \mathcal{B} . Thus, the optimal OWD distribution for \mathcal{B} will be also optimal for \mathcal{C} .

Intermediate summary: Let us get a general vision of the remaining part of this proof. We show below that the optimal distribution for each of the two aforementioned protocols is a uniform OWD. Thereby, using a parallel combination, in *reducePP*, any predicate coming from Presburger arithmetic is stationary computable in an optimal mean hitting time with a uniform OWD. This extension result is given by the fact that any mean hitting time function is characterized by polynomial function of p_{θ} .

Mean hitting time is polynomial: Let \mathcal{P} be a protocol and $\mathcal{M}_{\mathcal{P}}$ its associated Markov chain. Let \mathcal{S} be the state corresponding to the stationary distribution. Considering an initial state \mathcal{I} , we have to solve the following simultaneous equations, containing $m = \frac{n(n-1)}{2}$ variables:

$$\begin{cases} \mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}}) = f(p_{\theta_1}, \dots, p_{\theta_m}) \\ \sum_{\theta \in \Theta} p_{\theta} = 1 = g(p_{\theta_1}, \dots, p_{\theta_m}) \end{cases} \quad (1)$$

We define a path in $\mathcal{M}_{\mathcal{P}}$ by a sequence of states of $\mathcal{M}_{\mathcal{P}}$: $\langle k_1, k_2, \dots, k_s \rangle$ with $s \in \mathbb{N} \setminus \{0\}$. Let us define $\mathcal{C}_s(\mathcal{I}, \mathcal{S})$ the set of all paths from \mathcal{I} to \mathcal{S} with a length equals to s and in which \mathcal{S} not appears but the last:

$$\mathcal{C}_s(\mathcal{I}, \mathcal{S}) =$$

$$\{k_1, k_2, \dots, k_s \mid k_1 = \mathcal{I}, k_s = \mathcal{S}, \forall i \in \{1, \dots, s-1\}, k_i \neq \mathcal{S}\}$$

So, it is possible to infer a formal expression of the mean hitting time:

$$\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}}) = \sum_{s \in \mathbb{N} \setminus \{0\}} \sum_{c \in \mathcal{C}_s(\mathcal{I}, \mathcal{S})} \mathbb{E}_{\mathcal{I}}[\tau_{\mathcal{S}} \mid c] \cdot \mathbb{P}[c]$$

As the path in $\mathcal{M}_{\mathcal{P}}$ is determined for a specific $c \in \mathcal{C}_s(\mathcal{I}, \mathcal{S})$, we have $\mathbb{E}_{\mathcal{I}}[\tau_{\mathcal{S}} \mid c] = s$. Hence:

$$\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}}) = \sum_{s \in \mathbb{N} \setminus \{0\}} s \cdot \mathbb{P}[\mathcal{C}_s(\mathcal{I}, \mathcal{S})]$$

Let $q_{A,B}$ the probability to go from A to B in the Markov chain $\mathcal{M}_{\mathcal{P}}$. Thus, we have:

$$\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}}) = \sum_{s \in \mathbb{N} \setminus \{0\}} s \cdot \sum_{c \in \mathcal{C}_s(\mathcal{I}, \mathcal{S})} q_{\mathcal{I}, k_2} \cdot q_{k_2, k_3} \cdot \dots \cdot q_{k_{s-1}, \mathcal{S}}$$

In the Markov chain $\mathcal{M}_{\mathcal{P}}$, every transition probability label only depends on the potential interaction which makes the system evolve from a state k to another state k' , and, consequently, only depends on the sum of the p_{θ} corresponding to these interactions. Moreover, to analyze the advancement of the system state, the chain $\mathcal{M}_{\mathcal{P}}$ associated with each of the two aforementioned protocols can be simplified to one that only takes into account the first bit of the triplet state (the leader bit). Formally, let E be the states' set of $\mathcal{M}_{\mathcal{P}}$ where $E = \{0, 1\}^n$ for n agents in the population. As agents are anonymous, all states with the same distribution of the leader bits are clustered into a common *meta-state* in the resulting Markov chain. For each possible transition of δ , the number of leader bits, which are set to 1, cannot be increased. It can only be strictly decreased in case of both agent which act in the interaction own a 1 leader bit, or remain constant in case of one interacting agent owns a 1 leader bit, and the other owns a 0 leader bit.

Thus, there are two kinds of transitions in $\mathcal{M}_{\mathcal{P}}$:

$$\forall e, e' \in E, q_{e,e'} = \begin{cases} \mathbb{P}[i_1 \rightarrow i_2] = p_{i_1, i_2} & \text{if } \begin{cases} \forall j \notin \{i_1, i_2\}, e_j = e'_j \\ e'_{i_1} = e_{i_2} = 1 \\ e_{i_1} = e'_{i_2} = 0 \end{cases} \\ \sum_{j \neq i_0 \wedge e_j = 1} \mathbb{P}[j \rightarrow i_0] (= p_{j, i_0}) & \text{if } \begin{cases} \forall j \neq i_0, e_j = e'_j \\ e_{i_0} = 1 \\ e'_{i_0} = 0 \end{cases} \\ 0 & \text{otherwise} \end{cases}$$

Then, for any transition q in $\mathcal{M}_{\mathcal{P}}$, q is a linear application of p_{θ} . This infers that $\mathbb{E}_{\mathcal{I}}(\tau_S)$ and, by definition in Equation (1), the function f is a polynomial on p_{θ} .

On the lower-bound characterization: To find the optimal distribution for $p_{\theta} \in]0, 1[$, we are looking for a minimization of $f(p_{\theta_1}, \dots, p_{\theta_m})$, according to variables $(p_{\theta_1}, \dots, p_{\theta_m})$, under the constraints of Equation (1): $\sum_{\theta \in \Theta} p_{\theta} = 1$ (Moreover, by this constraint, it is possible to deduct p_{θ_m} from the other variables $(p_{\theta_1}, \dots, p_{\theta_{m-1}})$).

Consider an OWD in the interaction graph such that $\forall \theta, p_{\theta} > 0$ (complete interaction graph). Let D , an open subset of \mathbb{R}^n , defined as follow :

$$D = \left\{ (p_{\theta})_{\theta \in \Theta} \in]0, 1[^m \mid \sum_{\theta \in \Theta} p_{\theta} = 1 \right\}.$$

In topology, using the Taylor formula, we can argue that if a function in C^1 admits a minima in an open set, the derivative is null at this point. Then, as f is a polynomial, f and its derivative of first order are continuous. So, $f \in C^1$. Consequently, the minimal value of f on the closure of D is reached either on one of its limit point, or on a vector p^* inside D such that $\nabla f(p^*) = 0$ (∇ is the gradient of f and is defined as $\nabla f(p) = \left(\frac{\partial f}{\partial p_{\theta_1}}(p), \dots, \frac{\partial f}{\partial p_{\theta_m}}(p) \right)$). In the population protocol model, the interaction graph is complete and thus, $\forall \theta \in \Theta, p_{\theta} \neq 0$. So, the minimal value of f is reached on $p^* \in D$ if $\nabla(f)(p^*) = 0$.

As the minimal value is reached on such a p^* vector, using the weak Lagrangian principle, we can infer that $\nabla(g)(p^*) = 0$ (where g correspond to the function defined in Equation 1). Beyond, by definition, g is a constant function. Thus, all partial derivatives of g are identical: $\forall \theta, \theta' \in \Theta, \frac{\partial g}{\partial p_{\theta}}(p^*) = \frac{\partial g}{\partial p_{\theta'}}(p^*)$. Then, p^* is the equidistributed point on D (i.e. $\forall \theta, \theta' \in \Theta, p_{\theta} = p_{\theta'}$).

Finally, for any given population protocol, the minimum of $\mathbb{E}_{\mathcal{I}}(\tau_S)$ is reached for a uniform distribution of $p_{\theta}, \theta \in \Theta$, namely, the uniform OWD. ■

We can conclude that, for any population protocol in *reducePP*, it is impossible to have a mean convergence speed better than the one obtained using a totally uniform interaction graph.

V. ON THE RELEVANCE OF THE RANDOM WAY POINT MODEL

The Random Way Point mobility model is extensively used to evaluate mobile ad-hoc networks, though often criticized

for its lack of realism [19]. In this section, we show that the uniform distribution of MAPP is in fact equivalent to the random waypoint model. Beyond the theoretical interest of this proof, our objective is to emphasize the fact that, while the use of the random way point model is often justified in an attempt to provide a neutral setting, this actually provides the *best setting* with respect to convergence time.

Theorem 3: The Random Way Point mobility model is equivalent to a uniform OWD in MAPP.

Proof: The Random Way Point model can be formally defined as follows. Consider n mobile agents, with initial position $p_1^{(0)}, \dots, p_n^{(0)}$. At the beginning of the experiment, each agent picks a random destination point and a space velocity, according to law P_{pos} and P_{vit} respectively, common to all agents in the network. As a guideline, usually, the destination point distribution is uniform according to the given moving area and the space velocity is picked uniformly among a given interval. We show our result in a more general way and only assume that each sample is an independent ergodic system.

For every agent x , we define $\gamma_x(t)$ as the trajectory of this agent inside the moving area. This trajectory only depends of the sequence of pair $\langle (p_x^{(i)}, v_x^{(i)}) \rangle_{i \in \mathbb{N}^*}$. As the spatial distribution of nodes is not uniform [8] in case of using a reflecting random waypoint (bounded moving area – not a tore – and bouncing on side), we prove below that the contact probability for two agents is uniform (agents' moves are i.i.d.³).

Let x, y be two agents. Consider the following equation during a fixed period T that gives a formal expression of the temporal mean contact period:

$$\frac{1}{T} \int_0^T \mathbf{1}_{|\gamma_x(t) - \gamma_y(t)| \leq \varepsilon}(t) \cdot dt. \quad (2)$$

In this equation, $\mathbf{1}_{|\gamma_x(t) - \gamma_y(t)| \leq \varepsilon}(t)$ represents the *indicator function* which is defined as 1 if x and y are separated by a distance lower or equal to ε , and 0 otherwise. If Equation 2 tends toward a constant value completely independent of the (x, y) pair choice, then this constant is the same for any pair of agents in the system.

Consider $\gamma_x(t)$ the generic trajectory probability law for a given agent x :

$$P_x = \left(\bigotimes_{i=1}^{\infty} p_x^{(i)} \right) \otimes \left(\bigotimes_{i=1}^{\infty} v_x^{(i)} \right).$$

Consider now the pair of trajectories (γ_x, γ_y) . This last pair is a random variable according to the law $P_x \otimes P_y$. We can define the following system:

$$\left((p_x^{(i)}, v_x^{(i)}, p_y^{(i)}, v_y^{(i)})_{i \in \mathbb{N}^*}, P_x \otimes P_y, \lambda_{\otimes} \right)$$

where λ_{\otimes} is the invariant product Lebesgue measure. By definition, this system is a product of independent ergodic systems, and consequently, it is itself an ergodic system. Then, it is possible to apply the Ergodic Theorem on it (also

³i.i.d. means *independent and identically-distributed*.

named *large numbers law* which argues that the temporal mean converges toward the spatial mean). Let us introduced the spatial mean as follow:

$$\mathbb{E}[\mathbf{1}_{|\gamma_x - \gamma_y| \leq \varepsilon}] = \iint \mathbf{1}_{|\gamma_x - \gamma_y| \leq \varepsilon} \cdot dP_x(\gamma_x) \cdot dP_y(\gamma_y) \quad (3)$$

For *a.e.*⁴ trajectories γ_x and γ_y , when T tends toward ∞ , Equation 2 tends toward Equation 3. This last equation corresponds to an integral among all possible trajectories of x and y agents. So, this expression does not depend on x and y , but only on ε and on the trajectory probability laws P . These last laws are identical for every agents in the system.

In short, for any pair of agents, the average number of contacts is the same. Then, every pair of agents has the same probability to come in contact at time t . This means that we can simulate the random waypoint model by using a uniform OWD in MAPP. ■

VI. CONCLUSION AND OPEN ISSUES

a) Contributions: In this paper, we studied the impact of the agents' mobility model on the convergence speed of population protocols. We introduced *MAPP (Mobility Applied to Population Protocol)*, a population protocol model augmented to take into account the probability distribution of agent interactions. In this model, each edge of the interaction graph is weighted by a probability value reflecting the interaction frequency between the two agents linked by that edge.

We observe from previous empirical studies [10] a significant impact of various mobility models on the convergence speed population protocols.

From this empirical study, we also observed that the uniform interaction distribution turns out to consistently achieve the best average convergence time. Thus, in this paper, we formally proved that the uniform OWD provides the lowest bound of the average number of steps to reach convergence for population protocols in a larger class than *reducePP*.

Finally, we also demonstrated that the Random Way Point mobility model, which is the most extensively used to evaluate mobile networks, infers a uniform distribution in MAPP. Interestingly enough, not only the Random Way Point model is recognized as non-realistic, but we show that it actually provides the *best-case* scenario. This may question the relevance of using such a model as a representative mobility model.

b) Open Issues: This work is only a base step about stochastic analyses of population protocols. Several questions remain open. For instance, we plan to analyze MAPP on more realistic mobility models by taking into account the correlation between proximity of agents and the probability of interaction. To achieve these analyses, we will compare our former results with other realistic mobility model as proposed in [11], [18], [20]. It may be also interesting to compare and merge outcomes issues from another probabilistic study of population protocols [13].

⁴*a.e.* means *almost every*: one says that a property holds almost everywhere if the set of elements for which the property does not hold is a null set, i.e. is a set with measure zero.

Finally, the range of population protocols that accept their mean hitting time as polynomial, is obviously larger than *reducePP*. Any of these specific population protocol falls into the scope of Theorem 2. Searching for the exact set of *polynomial PP* is a challenging open question.

REFERENCES

- [1] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing (Special Issue: PODC'04)*, 18(4):235–253, March 2006.
- [2] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semilinear. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing (PODC '06)*, pages 292–299, New York, NY, USA, 2006. ACM Press.
- [3] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing (Special Issue: DISC'07)*, 21:183–199, September 2008.
- [4] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing (Special Issue: PODC'06)*, 20:279–304, November 2007.
- [5] Dana Angluin, James Aspnes, Michael J. Fischer, and Hong Jiang. Self-stabilizing population protocols. In *9th International Conference Principles of Distributed Systems (OPDIS'05)*, volume 3974 of *Lecture Notes in Computer Science*, pages 103–117, Pisa, Italy, December 2005.
- [6] James Aspnes and Eric Ruppert. An introduction to population protocols. *Bulletin of the European Association for Theoretical Computer Science, Distributed Computing Column*, 93:98–117, October 2007.
- [7] Marin Bertier, Yann Busnel, and Anne-Marie Kermarrec. On gossip and populations. In *the 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2009)*, Piran, Slovenia, May 2009. Springer-Verlag.
- [8] Christian Bettstetter, Giovanni Resta, and Paolo Santi. The node distribution of the random waypoint mobility model for wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 2(3):257–269, 2003.
- [9] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss. Delay-tolerant networking: an approach to interplanetary internet. *IEEE Communications Magazine*, 41(6):128–136, June 2003.
- [10] Yann Busnel, Marin Bertier, and Anne-Marie Kermarrec. On the impact of the mobility on convergence speed of population protocols. Rapport de Recherche RR-6580, Institut National de Recherche en Informatique et Application (INRIA), Rennes, France, juillet 2008.
- [11] Tracy Camp, Jeff Boleng, and Vanessa Davies. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing*, 2(5):483–502, 2002.
- [12] Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing*, 6(6):606–620, 2007.
- [13] Ioannis Chatzigiannakis and Paul G. Spirakis. BA: The dynamics of probabilistic population protocols. In *22nd international symposium on Distributed Computing (DISC '08)*, pages 498–499, Arcachon, France, 2008. Springer-Verlag.
- [14] Vania Conan, Jérémie Leguay, and Timur Friedman. Characterizing pairwise inter-contact patterns in delay tolerant networks. In *Conference on Autonomic Computing and Communication Systems (ACM Autonomics 2007)*, Rome, Italy, October 2007.
- [15] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Eric Ruppert. When birds die: Making population protocols fault-tolerant. In *Second IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'06)*, pages 51–66, San Francisco, CA, USA, June 2006.
- [16] Laurent Fribourg, Stéphane Messika, and Claudine Picaronny. Coupling and self-stabilization. *Distributed Computing (Special Issue: DISC'04)*, 18(3):221–232, February 2006.
- [17] Rachid Guerraoui and Eric Ruppert. Even small birds are unique: Population protocols with identifiers. Technical Report Technical Report CSE-2007-04, Dept of Computer Science and Engineering, York University, September 2007.

- [18] Seongik Hong, Injong Rhee, Seong Joon Kim, Kyunghan Lee, and Song Chong. Routing performance analysis of human-driven delay tolerant networks using the truncated levy walk model. In *1st ACM SIGMOBILE workshop on Mobility models (MobilityModels '08)*, pages 25–32, Hong Kong, China, 2008.
- [19] Amit Jardosh, Elizabeth M. Belding-Royer, Kevin C. Almeroth, and Subhash Suri. Towards realistic mobility models for mobile ad hoc networks. In *9th annual international conference on Mobile computing and networking (MobiCom '03)*, pages 217–229, San Diego, CA, USA, 2003.
- [20] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, Li Shiuan Peh, and Daniel Rubenstein. Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. In *10th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, volume 5, pages 96–107, 2002.
- [21] Samba Sesay, Zongkai Yang, and Jianhua He. A survey on mobile ad hoc wireless network. *Information Technology Journal*, 3(2):168–175, 2004.

Marin Bertier is actually an Assistant Professor at INSA Rennes (France). He received his PhD from University of Paris 6 (France) in 2004. During his PhD, he was member of the INRIA project REGAL on Resources management in large scale distributed systems. Then he was member of the LRI and the INRIA project Grand Large during one year. Finally, he is an assistant professor at the INSA of Rennes since september 2005.



Mail : Marin.Bertier@irisa.fr

Yann Busnel is currently an Assistant Professor in University of Nantes (France), working in the ATLAS-GDD team. He obtained his PhD in Computer Science at the University of Rennes (France) in November 2008, defended the thesis “Self-organized collaborative information system for large-scale wireless sensor networks: Moving theory into practice”. Then, in 2008-09, he spent one year as a Researcher at MIDLAB, the distributed systems research group of the University of Rome “La Sapienza” (Italy). He works essentially on very large scale dynamic distributed system and data management. He also an external reviewer of several international journals and conferences.



Mail : Yann.Busnel@univ-nantes.fr

Anne-Marie Kermarrec is actually a Senior Researcher at INRIA Rennes – Bretagne Atlantique (France). Before joining INRIA in February 2004, she was with Microsoft Research in Cambridge as a Researcher since March 2000. Before that, she obtained her Ph.D. from the University of Rennes (France) in October 1996. She also spent one year (1996-1997) in the Computer Systems group of Vrije Universiteit in Amsterdam (The Netherlands) working in the GLOBE project in collaboration with Maarten van Steen and Andrew. S. Tanenbaum. She acts usually as a member of a large number of Program Committee, Program Chair and Invited Editor.



Mail : Anne-Marie.Kermarrec@inria.fr