



HAL
open science

On the Impact of the Mobility on Convergence Speed of Population Protocols

Yann Busnel, Marin Bertier, Anne-Marie Kermarrec

► **To cite this version:**

Yann Busnel, Marin Bertier, Anne-Marie Kermarrec. On the Impact of the Mobility on Convergence Speed of Population Protocols. [Research Report] RR-6580, INRIA. 2008, pp.30. inria-00296584v2

HAL Id: inria-00296584

<https://inria.hal.science/inria-00296584v2>

Submitted on 27 Nov 2008

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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*On the Impact of the Mobility on Convergence Speed
of Population Protocols*

Yann Busnel — Marin Bertier — Anne-Marie Kermarrec

N° 6580

Juillet 2008

Thème NUM

A large, light gray stylized letter 'R' is positioned to the left of the text.

*R*apport
de recherche



On the Impact of the Mobility on Convergence Speed of Population Protocols

Yann Busnel*, Marin Bertier[†], Anne-Marie Kermarrec[‡]

Thème NUM — Systèmes numériques
Projets Asap

Rapport de recherche n° 6580 — Juillet 2008 — 27 pages

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, ensuring 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. 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 models the behaviour of the scheduler with respect to various mobility models. We empirically show that mobility models do have a significant impact on the convergence speed of the protocols. In fact, we observe 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.

Finally, we formally prove that a uniform distribution of weights provides the lowest bound of average convergence speed for any population protocol. Therefore, this analysis reveals that the Random Way Point model, following this distribution, provides the best case scenario. This may question its relevance as the most representative model.

Key-words: Population protocols, Mobility model, Markov chain, Theoretical analysis, Stochastic process, Lower bound.

* IRISA / Université de Rennes 1

[†] IRISA / INSA Rennes

[‡] IRISA / INRIA Rennes – Bretagne Atlantique

Sur l'impact de la mobilité sur la vitesse de convergence des protocoles de population

Résumé : Dans ce rapport de recherche, notre étude se focalise sur les protocoles de population. Nous présentons ici l'impact de la mobilité des agents sur la vitesse de convergence de ces protocoles. Après un court développement de nos motivations, nous étendons les modèles de protocoles de population présenté au paragraphe 2.1 afin de formaliser le comportement de l'ordonnanceur équitable. De ce modèle, nous pouvons ainsi analyser formellement la vitesse de convergence de n'importe quel protocole. Cependant, comme l'obtention d'équation s'avère être une tâche laborieuse, voire inaccessible dans la plupart des cas.

Nous présentons alors un ensemble de résultats empirique sur les trois protocoles introduits en section 2.2. Fort de ces observations, nous démontrons l'existence d'une borne inférieure de la vitesse de convergence pour tout protocole des deux classes sus-cités, ainsi que la configuration de l'ordonnanceur pour l'atteindre.

Issu de ce résultat, nous achèverons ce rapport de recherche par une réflexion sur la pertinence du modèle de mobilité classique dit de "points tournants aléatoires" (ou *Random Way-Point* pour les anglo-saxons).

Mots-clés : Protocoles de population, Modèle de mobilité, Chaîne de Markov, Analyse théorique, Processus stochastique, Borne inférieure.

1 Introduction

Sensor networks are composed of tiny computation units able to communicate and collect data from their environment. This led to a whole class of applications, where sensors are embedded on human or animals, aiming at observing their behaviour 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) [10] and *Population Protocols* [1]. Delay Tolerant Network may be seen as an evolution of *Mobile Ad-hoc Network* (MANET) [17]. 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 dependant on the mobility model. Many works in this area focus on establishing the correct mobility model to represent human's movement behaviours 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 global behaviour 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 and 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 only assumption is to be *fair* guides the way the interaction actually takes place. A fair scheduler simply ensures that every possible evolution of the system may eventually happen.

On the impact of mobility Most of the works in population protocol focus on the computational power of such protocols [2, 3, 4, 5, 13]. No specific assumption is usually made on the model of interactions between agents. When a specific interaction model is considered, it is usually uniform [2]. Yet, in reality, the mobility pattern of a set of mobile entities is not 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 [14]. As a consequence, characterizing realistic mobility models [11, 12] is an active area.

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 agent 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.

We first empirically study the impact of various mobility models on the convergence speed of the most classical population protocols namely the *or* and *sum modulo x* operations. We integrate the mobility aspect of such networks along two lines: (i) the distribution of interactions (this is modelled by the weight on the interaction graph) and (ii) the distribution of pairwise inter-contact time¹ for a given distribution of interactions.

In the empirical study, we consider a uniform interaction distribution as well as a set of non-uniform mobility patterns identified by the mobile network community [11, 12]. We observe 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.

We then formally prove that this uniform distribution of weights provides the lowest bound of the average number of steps to reach convergence in population protocols.

Finally, we also demonstrate 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 also we show that it actually provides the *best-case* scenario. This may lead to question the relevance of using such a model as a representative mobility model.

Roadmap The rest of the paper is organized as follow. Section 2 provides the computation models of population protocols as well as MAPP, our extension of this model. In Section 3 and 4, we present respectively the empirical context and study of the impact of mobility on the convergence speed of a few traditional protocols (*or* and *sum modulo 4*). Section 5 presents the proof that the uniform distribution leads to the lowest bound in term of average number of step required to convergence in population protocols. Finally, in Section 6, we prove that the Random Way Point model can be modelled by a uniform distribution. We conclude and list some open issues in Section 7.

2 Computation models

2.1 Population protocols

The seminal 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 2.3). An agent in this model is represented as a finite state machine, 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:

¹Note that Manet and DTN work usually focus only on inter-contact time property.

- 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, $\Theta = \{(v, v') \in \Upsilon^2 \mid v \neq v'\}$ (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 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 stabilize to output $f(\sigma)$ (*i.e.* output value of every agent eventually stabilizes to $f(\sigma)$).

Population protocols are able to compute any predicate of the Presburger arithmetic, namely the semilinear predicates [3, 4], and only these. In brief, a set $L \subseteq \mathbb{N}^k$ is *linear* if $\exists v_0, v_1, \dots, v_m \in \mathbb{N}^k$ such that $L = \{v_0 + \kappa_1 v_1 + \dots + \kappa_m v_m \mid \kappa_1, \dots, \kappa_m \in \mathbb{N}\}$. As a *semilinear set* is a union of linear sets, a *semilinear predicate* is defined as the one true precisely on a semilinear set. The MAPP extension affects only the interaction scheduler without jeopardizing its fairness, nor the computable predicates.

2.2 Illustrating protocols

In the following, we illustrate our results using some classical population protocols. Then, we introduce in this paragraph three of these.

Or operation This protocol is equivalent to an epidemic dissemination or a *flooding* operation. This simple protocol consists, for any agent with input 0, to output 1 as soon as it encounters an agent with input 1. More formally, $\Sigma = Y = Q = \{0, 1\}$. ι and ω correspond to the identity function and δ is the singleton $\{(0, 1) \rightarrow (1, 1)\}$. All other transitions are defined to leave the pair of states unchanged.

Majority This operation consists in computing the majority of the set of agents. Consider two kinds of agents in the system (0 and 1, male and female for example) depicted by \perp_0 and \perp_1 . The majority output 1 (*i.e.* each agent outputs 1 eventually) if the initial configuration contains a strict majority of \perp_1 and 0 otherwise. As counting

anonymous objects in a distributed way is impossible, the idea of the protocol is to consider agent pairwise and observe if any agent remains “*isolated*”. More formally, $\Sigma = \{\perp_0, \perp_1\}$ in which the first type of entity is represented by \perp_0 and the second one by \perp_1 . $Y = \{0, 1\}$ and $Q = \{\perp_0, \perp_1, 0, 1\}$. ι corresponds to the identity function and ω maps \perp_0 and 0 to 0 and \perp_1 and 1 to 1. δ contains the four following transitions: $(\perp_0, \perp_1) \mapsto (0, 0)$, $(\perp_0, 1) \mapsto (\perp_0, 0)$, $(\perp_1, 0) \mapsto (\perp_1, 1)$, $(0, 1) \mapsto (0, 0)$. All other transitions are defined to leave the pair of states unchanged.

Sum modulo 4 This protocol computes in a distributed manner the sum modulo 4 of all inputs picked in $\Sigma = \{0, 1, 2, 3\}$ (as well Y is the same set than Σ). All the values can be collected by a single agent, which eventually stabilizes to the sum, modulo 4. In order to remove the yet computed input values, each agent’s value becomes \perp labelled with the value of the eventually unique agent with a non- \perp value. Then, we have $Q = \{0, 1, 2, 3, \perp_0, \perp_1, \perp_2, \perp_3\}$. Let ι corresponds to the identity function and $\omega(v) = \omega(\perp_v) = v$. The only transition rules of δ , which do not leave the pair of states unchanged, are the following: $\forall v, w \in \{0, 1, 2, 3\}, (v, w) \mapsto (v + w, \perp_{v+w})$ and $(v, \perp_w) \mapsto (v, \perp_v)$ in which the addition is made modulo 4.

2.3 Mobility Applied to Population Protocols

In order to formally study the impact of the mobility model on the convergence speed of protocols, we proposed 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 all 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, while ensuring this fairness condition, and then capture mobility patterns. In *Mobility Applied to Population Protocol (MAPP)*, a weight is assigned to all edges 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. 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 . Obviously, we have $\sum_{\theta \in \Theta} p_\theta \leq 1$. 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 meeting probability of those agents.

On the relevance of this extension We must first check that our model verifies 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. $\forall \theta \in \Theta$, if $p_\theta > 0$ then the probability that this interaction occurs instead 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 appears infinitely in the execution.

Conversely, if $p_\theta = 0$, without loss of generality, we can consider a restricted interaction graph such that $\theta \notin \Theta$. In this case, it cannot exist $C \rightarrow C'$ with the interaction θ . Then, the scheduler does not break the fairness assumption. \square

2.3.1 Integration with related works

One of the main advantages of our extension consist in its easy integration among all variant models starting from the population protocols. In fact, our proposition are not decreasing the powerfulness of the initial models.

For instance, the computational power of former model are unchanged as for population protocol, it remains as semilinear predicates [3, 4]). Thus, any result according to the self-organization of these protocols [5] and to take into account failures [13] remain valid.

Then, using MAPP, we can model most of the model which extend the population protocol one. 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, 2] are valid for all MAPP set using an uniform distribution of probabilities of the graph's edges. Finally, different models introduced in [4] concerning the model of one-way communications remain valid in a MAPP using an oriented graph.

3 Theoretical analysis of convergence

Many papers [1, 2, 3, 4, 5, 13, 15] have investigated the power of population protocols and its extensions. Yet, very few are concerned with the convergence speed (in population protocol context, 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 behaviour of the evolution of the system over time.

In this section, we present the method used to proceed these analyses. First, we introduce the necessary mathematic tools. Then, we show that it may be unfeasible to extract formal expressions, even for simple protocols and small populations. Finally, we introduce the method used to estimate the behavior of population according to a specific protocol.

3.1 Background 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 only 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.

From now, it is necessary to introduce the mathematical background used in our study. In the sequel, 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 happening of the system in a specific state. Consequently, we can define the expecting 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*

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

Definition 2 (Mean and 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.

3.2 On the difficulty to analyze a simple example

Unfortunately, the number of configurations of a MAPP is growing exponentially according to the size of Q and Υ . More formally, Appendix A.1 presents the evolution analysis of a simple example and shows that the number of states of the Markov chain is $|Q|^{|\Upsilon|}$ ($= |Q|^n$).

While it might be useful and simple to understand precisely the evolution of a system with a small number of agents, it becomes impracticable to study it formally when n grows significantly. In that case, it is still interesting to study empirically these systems.

3.3 Estimation of stabilization time using Markov chain

For large systems, the mean hitting time formal calculation of the Markov chain associated with a given configuration of MAPP is prohibitively computational intensive. However, using probabilistic methods, it is possible to obtain some empirical results

on the behavior of these protocols. Despite the growing size of the Markov chain, it is always feasible to compute the possible transitions from a given configuration, and their probability to happen. Thus, it is conceivable to estimate the followed path across a Markov chain using its probability distribution. Nevertheless, it remains the problem to determine how many steps are needed to converge to the stationary distribution, within an acceptable error margin. To achieve this task, by using the *Ergodic Theorem*² and among a large number of samples, which simulate the evolution of the system before reaching the stationary distribution, the average number of step may represent a correct estimation of the mean hitting time of convergence.

In order to estimate precisely this mean convergence time and errors, we use the Markov Chain Monte Carlo (MCMC) method [8]. A Monte Carlo method is a computational algorithm which relies on repeated random sampling to compute its results [16]. Similar methods have been used to extract some empirical result in order to conjecture formal outcomes [7]. In fact, from the observation of the behaviour of a system in a large number of simulations, it is possible to extrapolate some common characteristics of all these executions. They permit to guide the following formal study.

In this paper, we denote the same way the theoretical mean hitting time and the estimated one with the Monte Carlo method (and respectively for the variance). Thus, in the MCMC method, we consider the same definitions of Section 3.1, but with a huge value of N instead of $N \rightarrow \infty$. The potential error of estimation is given by a confidence interval. This represents, with a known error percentage, an interval supposed to host the theoretical value estimated. This definition leads that greater the N , smaller the interval. More formally, we have:

Definition 3 (Confidence interval) *Given a state \mathcal{I} and a set of first time hit times $\{\tau_{\mathcal{A}}\}_{i \in \{n \in \mathbb{N}: n < N\}}$, the confidence interval of the estimation is defined as follow:*

$$\left\{ \begin{array}{l} \text{with 5\% of error:} \\ \left[\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{A}}) - 1.96 \cdot \frac{\sigma_{\mathcal{I}}(\tau_{\mathcal{A}})}{\sqrt{N}}, \mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{A}}) + 1.96 \cdot \frac{\sigma_{\mathcal{I}}(\tau_{\mathcal{A}})}{\sqrt{N}} \right] \\ \text{with 1\% of error:} \\ \left[\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{A}}) - 2.5758 \cdot \frac{\sigma_{\mathcal{I}}(\tau_{\mathcal{A}})}{\sqrt{N}}, \mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{A}}) + 2.5758 \cdot \frac{\sigma_{\mathcal{I}}(\tau_{\mathcal{A}})}{\sqrt{N}} \right] \end{array} \right.$$

4 Impact of mobility models on convergence

In this section, we study the influence of mobility model on convergence speed of population protocols, using several probability distributions in MAPP.

4.1 Sampling of probability distributions on Γ

In order to simplify the system context, and thus preserve the spirit of population protocols, in this section, we assume the unicity of *out-weight distribution of v* (OWD). More formally, given an agent $v \in \Upsilon$, the 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 $|\mathcal{P}_v| = n - 1$. Thus,

²Also known as *Large numbers law*: Given a sample of independent and identically distributed random variables with a finite population mean, the average of these observations will eventually approach and stay close to the population mean.




Figure 1: Comparison of different OWDs for one anonymous agent in MAPP

unicity means that each agent in a population owns the same OWD for a given MAPP. Formally, we have:

$$\forall v, \psi \in \Upsilon, \mathcal{P}_v = \mathcal{P}_\psi. \quad (1)$$

For instance, consider a 100 agent population. Figure 1 presents, for any agent $v \in \Upsilon$, five different OWDs used in the stochastic samples of MCMC. Each of these OWD is equivalent to a density function, represented in a decreasing way. Then, given that $\sum_{\theta \in \Theta} p_\theta = 1$, the area located below each OWD curve is equal to $\frac{2}{n}$ (the interaction graph is assumed undirected). Moreover, we consider a complete interaction graph, in which potentially, each agent may interact with any other. More formally, we have $\forall v \in \Upsilon, \forall p \in \mathcal{P}_v, p \neq 0$. The name of each OWD is given according to the striking resemblance with well-known probability distributions.

Uniform Each agent in the population has the same probability to interact with any other agent;

Pareto - Flat This distribution is related to each agent having a bias towards a few other agents;

Normal - Flat In this distribution, each agent has a bias towards a group of agents and another one privileged relation with few agents (and therefore their group);

Normal - 2 friends In this OWD, each agent has a higher probability to interact with two specific other agents;

Normal - 1 friend In this OWD, each agent has a high probability to interact with one other agent.

As the interaction graph is considered as undirected, the two last OWD are symmetric, *i.e.* if an agent a has a high probability to interact with another agent b , then b has the same high probability to interact with a .

Starting from these OWD, our empirical observations can be classified in two sets. First studies using discrete time, the topology impact of interaction graph. Secondly, we observed the impact of inter-contact time distribution model, in order to model continuous time.

4.2 Estimation in discrete time: Contact models

In the following, we study the impact of such mobility models on the convergence speed of two protocols (*or* and *sum modulo 4*).

Let observe the impact of these OWDs and of graph characteristics on the mean convergence time, and so, on the convergence speed of protocols.

4.2.1 Complete Interaction Graph

Considering in first hand, a complete intergation graph, as in the former model. Each agent has then a not-null probability to interact with any other agent of the population. Formally, we have: $\forall v \in \Upsilon, \forall p \in \mathcal{P}_v, p > 0$.

Figures 2, 3 and 4 depict, for each aforementioned OWD (*cf.* Figure 1), the average convergence speed in term of number of interactions required to reach stabilization. They represent respectively the impact of graph size, which varying from 0 to 100, for the *or*, *majority* and *sum* operations. To all of these average values is associated the confident interval with 1%-error.

We observe that the OWD has a significant impact on the convergence speed. We deep into detail for both of them.

First, the more biased the distribution, the greater the impact. It is especially verified for large graph on which the “flat” OWDs (uniform, Pareto and Normal) are usually greater than the ones with a small number of “friends”. We also observe that the uniform distribution consistently leads to faster convergence, for any protocol and any population size. We come back on this observation with more details in Section ??.

Moreover, we observe that for small-size graphs, Normal-2 and *-Flat are equivalent. Moreover, as the number of agents increases, the distribution Normal-2 gets close to Normal-1. This evolution, observed consistently over all graphs, is exacerbated in Figure 4 due to the characteristics of the sum operation. Indeed, the order of interaction for the sum operation has a stronger impact: reaching the stationary distribution required several interactions for each pair of agents.

Finally, on the context of complete interaction graph, it is interesting to study the *flushing coefficient* of highest probabilities, defined below. In fact, usely, the generated graph are regular, *i.e.* higher probabilities are uniformly spread among the whole population. However, still respecting the OWD unicity assumption, it is possible to flush the graph in order to remove highly connected area of the graph, in term of high probability (*i.e.* kind of interaction group simulating social behavior). The *flushing coefficient* represents the probability to exchange two edges in the initial regular graph, and thus, make the graph more or less randomized. Figure 5 presents for non-uniform OWD the mean number of interaction required to reach the stationary system state for the *or* function, according to the flushing coefficient. It is obvious that the clustering coefficient (*cf.* next paragraph) of highest probability edges has no impact on the convergence speed.

4.2.2 Restricted Interaction Graph

We now consider the impact of characteristics of a restricted interaction graph (see Section 2.3.1). More formally, it boil down to relax the not-null hypothesis of of probabilities: $\forall v \in \Upsilon, \forall p \in \mathcal{P}_v, p \geq 0$. Moreover, for the same clarity reason, we still consider only the OWD unicity assumption. Thus, we obtain that $\forall v \in \Upsilon, \exists p \in \mathcal{P}, p = 0$.

We only consider the *or* function in this sub-section. Indeed, the *majority* and *sum modulo* required specific topologies to converge. For instance, the *sum modulo* imply that the collecting agent has to interact with all the other agent in order to put the final value on. It is then trivial that this protocol cannot converge in the context of restricted interaction graph.

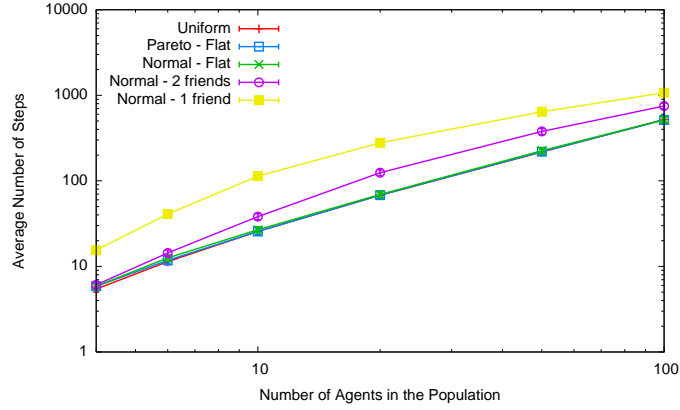


Figure 2: Average number of step and confidence interval according to the number of agents for the *or* protocol, among a complete interaction graph.

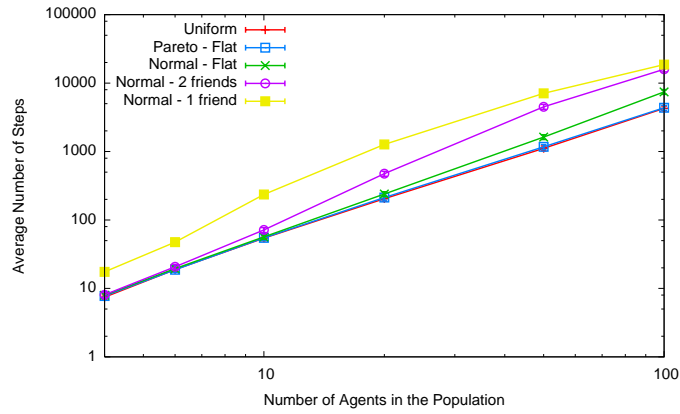


Figure 3: Average number of step and confidence interval according to the number of agents for the *majority* protocol, among a complete interaction graph.




Figure 4: Average number of step and confidence interval according to the number of agents for the *sum modulo 4* protocol, among a complete interaction graph.

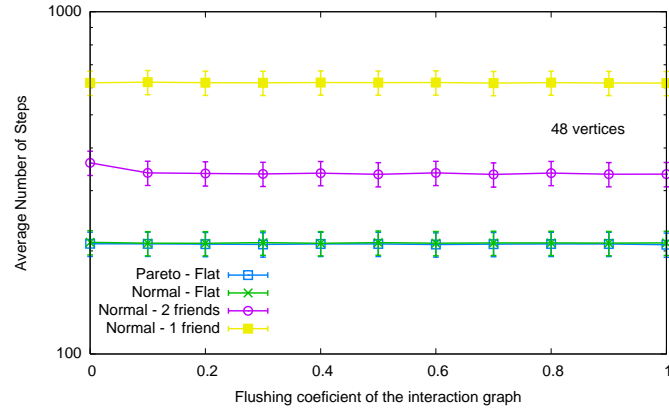


Figure 5: Average number of step and confidence interval according to the flushing coefficient for the *or* protocol, among a complete interaction graph.

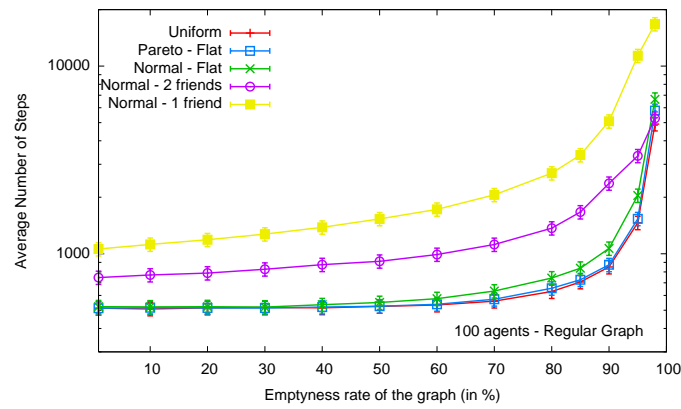


Figure 6: Average number of step and confidence interval according to the emptiness rate for the *or* protocol, among a restricted graph of 100 agents.

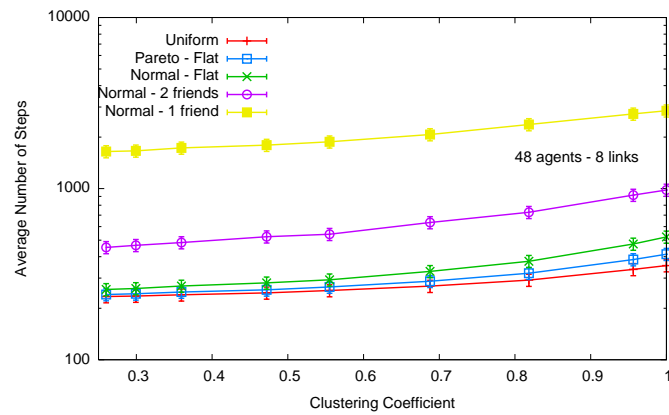


Figure 7: Average number of step and confidence interval according to the clustering coefficient for the *or* protocol, among a restricted graph of 48 agents.

Observations revealed above remains obviously valid in the context of restricted interaction graph. Two characteristics have a non-negligeable impact on the convergence speed of considered protocols: (1) the number of probability in the OWD equals to 0 and (2) the coefficient of clustering of the graph for a given OWD. Then, we focus in this section only on these two characteristics: *graph emptiness rate* and *clustering coefficient*.

The *graph emptiness rate* corresponding to the inverse of the density in term of edge quantity in any graph (*i.e.* it represents the ratio of *void* in a graph). In our context, the emptiness rate corresponds to the ratio of null probability edges.

Figure 6 depict for the various OWD, the average number of interactions in the *or* operation depending on the emptiness rate. These result are extracted from a 100 agent population on a regular graph. Despite a light impact for a low emptiness rate (complete graph to 40%), the consequences on convergence speed for high rate are significant. Even though this observation seems intuitive, ones can think that for a regular graph with an emptiness rate of 98% (*i.e.* a graph with a ring topology), the convergence speed have to be increased. On the contrary, for any given OWD (except “Normal - 1 friend” in which an agent interact with almost only one specific other), the number of interaction required grows significantly and the differences are merging in a common behavior. This speaks certainly on the fact that all OWD have only two non-null probabilities, these last are then all equivalents to the uniform one.

In other hand, we consider the clustering coefficient (CC). Initially introduced in [18], it represents the rate of neighbors vertices which are also directly link by another edge. It is often correlated to the characteristic path length in a general graph $G(V, E)$. In the following, we denote adj_i the set containing all the successors of a vertex i in a graph. More formally, we have:

Definition 4 (Clustering coefficient) *Let a vertex $v \in V$ of a graph $G(V, E)$. Given $K = \frac{k_v \cdot (k_v - 1)}{2}$ where $k_v = |adj_G(v)|$, the clustering coefficient of vertex v is:*

$$C_v = \frac{|\{(x, y) \in E | x, y \in adj_G(v)\}|}{K}.$$

Then, in global view, the clustering coefficient of graph G is $C_G = \frac{\sum_{v \in V} C_v}{|V|}$.

Definition 5 (Characteristic path length) *Let a graph $G(V, E)$. the shortest path length $L_{v, v'}$ between v and v' is the number of edges in the shortest path for $v, v' \in V^2$. Then, in global view, the characteristic path length of a graph G is define as follow:*

$$L_G = \frac{\sum_{(v, v') \in V^2} L_{v, v'}}{|V^2|}.$$

As the characteristic path coefficient is strongly correlated to the clustering coefficient [18], they have a similar impact on the convergence. Therefore, we do not consider the impact of the path coefficient impact in this paper.

Figure 7 presents, for each considered OWD, the convergence speed depending on the clustering coefficient in a 48 agent system and 8 edges, with a 83% emptiness rate. These values has been chosen regarding on sparseness values between the OWDs related by the previous observations. It is observable here that the clustering coefficient has an impact on the number of interactions required to converge, but conversely as the common knowledge, this impact remains slight.

All these measure on graph characteristics have been conducted in discrete time. The convergence time has been then formulated in number of interactions required to reach the stationary distribution of the associated Markov chain. In this context, two interactions made in parallel is considered as sequential. This hypothesis is not more valid in the continuous time context. Observing the impact in this context is the aim of the next section.

4.3 Estimation in continuous time: Inter-contact models

While the contact model provides the probability for each agent to interact with other, it is not possible to directly infer the convergence time necessary to converge. Thus, for a continuous time environment, it is not sufficient to specify only the OWD, as this last does not reflect the frequency of such interactions. This frequency can be seen as the *inter-contact frequency*.

So, given a specific OWD, the number of steps to reach the stationary distribution remains the same, regardless of the contact and inter-contact duration. In this subsection, we estimate using stochastic simulation the convergence time in the continuous time domain, for several well-known distributions of inter-contact times.

In these experiments, we consider a complete interaction graph and the OWD unicity among the population. We use three of the extensively used inter-contact time distribution, extracted notably in [11, 12]: *Exponential*, *Pareto* (known also as *Power-Law*) and *Log-Normal* distributions. These distributions are formally defined respectively by the following function

$$\lambda \cdot e^{-\lambda x} \quad , \quad \frac{k \cdot x_m^k}{x^{k+1}} \quad \text{et} \quad \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(\ln(x)-\mu)^2}{2\sigma^2}} \quad ,$$

with parameters $\lambda > 0, x_m > 0, k > 0, \sigma > 0$ and $-\infty < \mu < \infty$.

In order to measure separately the impact of contact and inter-contact duration, each simulation has been conducted first with atomic contact (*i.e.* one unique interaction for each contact of a given pair) and second, with uniform contact length (*i.e.* one interaction per second during the whole length of the contact, which duration is picked uniformly between 1 and 100 seconds). For each simulation, the settings of these distributions are comparable: the average inter-contact length is set to 15 minutes, and, in case of non-atomic contact, the maximum contact duration is set to 100 seconds.




Figure 8: Average convergence time and confidence interval according to the number of agents for the *or* protocol with atomic contact.




Figure 9: Average convergence time and confidence interval according to the number of agents for the *majority* protocol with atomic contact.

Two simulation sets has been conducted for different sizes of the population: one for the *or* and another for the *majority* primitive. Figures 8 and 9 present the results obtained using these both sets in case of atomic contact, as Figures 10 and 11 present the ones in case of uniform contact length. Each simulation set contains an estimated convergence date using MCMC in the continuous time domain, for the three aforementioned inter-contact distribution among OWD “Uniform”, “Normal - Flat” et “Normal - 2 friends”.

The observation of these figures permit to obviously conclude on the strong impact of inter-contact distribution on convergence speed of a protocol, for a given population size. Moreover, it seems clear that on these examples, the duration of contact has a small influence on the convergence speed (*cf.* the aspect of Figures 8 and 10, as well as the Figures 9 and 11).

However, we can raise two singularities of these 4 figures. While the population size has a small influence on convergence time for inter-contact distribution in *Pareto* and *LogNormal*, this one has a strong impact for the *exponential* distribution. The almost constant execution time using the two first distributions can be explained by the multiplication of parallel interaction as the population growing. At contrary, the third distribution have an influence, in gain or loss, on convergence for very large scale population. This observation infers the second characteristic of these diagrams. This last concerns also only the *exponential* distribution. As the convergence time decrease according to the population growing for OWD “Uniform” and “Normal - Flat”, it is increasing significantly in case of the OWD “Normal - 2 friends” among a large graph. This arise from the unbalancing of contact probabilities, infer by the given OWD. On the other hand, understanding why the fact that this disproportion do not have an influence in the context of *Pareto* inter-contact distribution remains open.

To put in a nutshell, the impact of the inter-contact model is meaningful but keeps the same order of magnitude and surprisingly, is independent to the duration of contact. Following these set of observation, we continuously observed that using an uniform OWD permit to reach the best average convergence time, for any graph configuration and using discrete or continuous time. Inspired by this conjecture, we propose to formalize and to prove this outcome, for any given configuration and any population protocol.




Figure 10: Average convergence time and confidence interval according to the number of agents for the *or* protocol with uniform contact length.

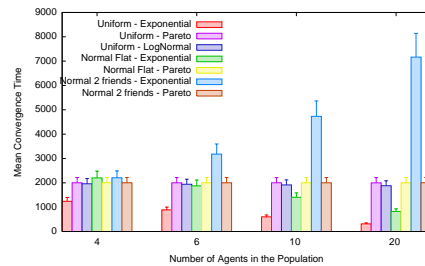


Figure 11: Average convergence time and confidence interval according to the number of agents for the *majority* protocol with uniform contact length.

5 A lower bound: Uniformity is always best

We observed from the simulations that the mobility patterns of agents might have a significant impact on the convergence speed of population protocols. Yet, the uniform distribution consistently leads to obtain a better 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), it is actually counter-intuitive for some others like flooding. Indeed, for the *or* operation, by removing the unicity of OWDs hiterto used, one might think that a global OWD, in which the source agent is involved in almost all the interaction³, converge faster than the uniform OWD.

In this section, we propose a theorem that rebuts this intuition. We prove 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 (*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, the lower bound of the convergence speed is reach using an 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 combination of population protocol, we will infer that (*iii*) any population protocol owns a polynomial mean hitting time. Finally, we prove that (*iv*) any polynomial mean hitting time accept a lower bound for an 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 characterize the wholeness of population protocols' computable function [3]. Thus, these two domains of function share the same equivalence class.

First, consider a factorization of all Presburger arithmetic's predicate. Let $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ an arbitrary input alphabet, and A_i, c, m integer constants such that $m \geq 2$. 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 that for any combination of this predicates, an uniform OWD in MAPP corresponds to the optimal convergence speed of this combination. Thus, any predicate of the Presburger arithmetic, computed by combination of these population protocols, has an optimal convergence speed with the uniform OWD.

³For 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.

As this class of combination has the same equivalence class than population protocols [1], it is obvious to conclude that if any combination of generator population protocols achieves an optimal convergence speed for a given OWD, this OWD will also be optimal for all the population protocols.

Generator population protocols' definition We first present both population protocols introduced in [1] that 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 a 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:

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

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: $(\ell, \cdot, u), (\ell', \cdot, u')$

$$\rightarrow (1, b(u, u'), q(u, u')), (0, b(u, u'), r(u, u')).$$

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: $(\ell, \cdot, u), (\ell', \cdot, u')$

$$\rightarrow (1, b(u, u'), (u + u') \bmod m), (0, b(u, u'), 0).$$

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

Inspired by the proof of Lemma 3 in [1], we still need to present the computation of a boolean function ξ on two stationary computable predicates F and G . 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 combination of \mathcal{A} and \mathcal{B} (the population runs protocols \mathcal{A} and \mathcal{B} in parallel and outputs the value of ξ applied to the outputs of the two computed predicates F and G).

More formally, 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 of 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 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 parallel combination, any predicate coming from Presburger arithmetic is stationary computable in an optimal mean hitting time with an uniform OWD. This extension result is given by the fact that any mean hitting time function is characterized by a polynomial function of p_i .

Mean hitting time is polynomial Let \mathcal{P} a protocol and $\mathcal{M}_{\mathcal{P}}$ its associated Markov chain. Let \mathcal{S} 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:

$$\begin{cases} \mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}}) = f(p_1, \dots, p_m) \\ \sum_{p_i} p_i = 1 = g(p_1, \dots, p_m) \end{cases} \quad (2)$$

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

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

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

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

As the path in $\mathcal{M}_{\mathcal{P}}$ is determined for a specific $c \in \mathcal{C}_s(\mathcal{I}, \mathcal{S})$, we have $\mathbb{E}_{\mathcal{I}}[\tau_{\mathcal{S}} | c] = s$. 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_i corresponding to these interactions. Moreover, to analyze the advancement of the system state, the chain $\mathcal{M}_{\mathcal{P}}$ associated respectively to each of the both aforementioned protocols can be simplified to only take 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, it exists two kinds of transitions in $\mathcal{M}_{\mathcal{P}}$: $\forall e, e' \in E, q_{e,e'} =$

$$\left\{ \begin{array}{ll} \mathbb{P}[i_1 \rightarrow i_2] = p_{i_1, i_2} & \text{if } \left\{ \begin{array}{l} \forall j \notin \{i_1, i_2\}, e_j = e'_j \\ e'_{i_1} = e_{i_2} = 1 \\ e_{i_1} = e'_{i_2} = 0 \end{array} \right. \\ \sum_{j \neq i_0 \wedge e_j = 1} \mathbb{P}[j \rightarrow i_0] (= p_{j, i_0}) & \text{if } \left\{ \begin{array}{l} \forall j \neq i_0, e_j = e'_j \\ e_{i_0} = 1 \\ e'_{i_0} = 0 \end{array} \right. \\ 0 & \text{otherwise} \end{array} \right.$$

Then, for any transition q in $\mathcal{M}_{\mathcal{P}}$, q is a linear application of p_i . This infers that $\mathbb{E}_{\mathcal{I}}(\tau_S)$ and, by definition in Equation (2), the function f is a polynomial on p_i .

On the lower-bound characterization To find the optimal distribution for $p_i \in]0, 1[$, we are looking for a minimization of $f(p_1, \dots, p_m)$, according to variables (p_1, \dots, p_m) , under the constraints of Equation (2): $\sum_{i=1}^m p_i = 1$ (Moreover, by this constraint, it is possible to deduct p_m from the other variables (p_1, \dots, p_{m-1})).

Consider an OWD in the interaction graph such that $\forall p_i, p_i > 0$. Let $D = \{(p_i)_{1 \leq i \leq n} \in]0, 1[^n, \sum_{i=1}^n p_i = 1\}$. It is an open subset of \mathbb{R}^n .

In topology, using the Taylor formula, we can argue that if a function in C^1 admits 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 $\nabla f(p) = \left(\frac{\partial f}{\partial p_1}(p), \dots, \frac{\partial f}{\partial p_m}(p) \right)$). In the population protocol model, the interaction graph is complete and thus, $\forall i \in \{1, \dots, n\}, p_i \neq 0$. 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 $\nabla(g)(p^*) = 0$. Beyond, by definition, g is a constant function. Thus, all partial derivatives of g are identical: $\forall i, j \in \llbracket 1, m \rrbracket, \frac{\partial f}{\partial p_i}(p^*) = \frac{\partial f}{\partial p_j}(p^*)$. Then, p^* is the equidistributed point on D (i.e. $\forall i, j \in \llbracket 1, m \rrbracket, p_i = p_j$).

Finally, for any given population protocol, the minimum of $\mathbb{E}_{\mathcal{I}}(\tau_S)$ is reached for a uniform distribution of $p_i, i \in \llbracket 1, m \rrbracket$, namely, the uniform OWD. \square

From this theorem, we can conclude that, for any population protocol, it is impossible to have a mean convergence speed better than the one obtained using a totally uniform interaction graph.

6 On the relevance of the Random Way Point model

The Random Way Point mobility model is extensively used to evaluate mobile ad-hoc networks whereas often criticized for its lack of realism. In this section, we show that the uniform distribution of MAPP is actually equivalent to the random waypoint. 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 an 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 at random a 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 [9] 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.).

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

$$\frac{1}{T} \int_0^T \mathbf{1}_{|\gamma_x(t) - \gamma_y(t)| \leq \varepsilon}(t) \cdot dt. \quad (3)$$

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

$$\mathbb{E}[\mathbf{1}_{|\gamma_x - \gamma_y| \leq \varepsilon}] = \iint \mathbf{1}_{|\gamma_x - \gamma_y| \leq \varepsilon} \cdot dP_x(\gamma_x) \cdot dP_y(\gamma_y) \quad (4)$$

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

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.

In short, for any pair of agents, the average number of contact 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.

□

7 Conclusion

In this paper, we study 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 first empirically studied the significant impact of various mobility models on the convergence speed of two classical population protocols (*or* and *sum modulo 4*) and extracted some interesting bias according to the sociability and meeting models.

From the empirical study, we observed that the uniform interaction distribution turns out to consistently achieve the best average convergence time. We then formally proved that the uniform OWD provides the lowest bound of the average number of steps to reach convergence in population protocols.

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 lead to question the relevance of using such a model as a representative mobility model.

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. Fast computation by population protocols with a leader. In *20th International Symposium on Distributed Computing (DISC'06)*, pages 61–75, Stockholm, Sweden, September 2006.
- [3] 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.
- [4] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *21th International Symposium on Distributed Computing (DISC'07)*, 20(4):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 (OPODIS'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] Joffroy Beauquier, Colette Johnen, and Stéphane Messika. Brief announcement: Computing automatically the stabilization time against the worst and the best schedulers. In *Proceedings of Distributed Computing (DISC'06)*, volume LNCS 4167, pages 443–447, Stockholm, Sweden, September 2006.
- [8] Bernd A. Berg. *Markov Chain Monte Carlo Simulations and Their Statistical Analysis*. World Scientific, 2004.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] Laurent Fribourg, Stéphane Messika, and Claudine Picaronny. Coupling and self-stabilization. *Distributed Computing (Special Issue: DISC'04)*, 18(3):221–232, February 2006.
- [15] 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.
- [16] N. Metropolis and S. Ulam. The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341, 1949.
- [17] Samba Sesay, Zongkai Yang, and Jianhua He. A survey on mobile ad hoc wireless network. *Information Technology Journal*, 3(2):168–175, 2004.
- [18] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

A Appendix

A.1 Impossibility of simple formula

In this appendix, we show that, even for a very simple context, it is prohibitively computational intensive to obtain a formal expression of the mean hitting time.

Let consider a very simple example: the *Or operation* for a generic MAPP. In this protocol, we have a minimum size of Q (here, $|Q| = 2$). Moreover, the Markov Chain is aperiodic as an "infected" agents cannot be heal. below, we show that, as simple this protocol is, the formal analysis of the convergence speed is not computable for arbitrary number of agents.

Let consider a very small population with only for 4 agents. Top of Figure 12 presents the interaction graph of this population. In order to simplify the context, we consider here that all agents share the same out-weight distribution (*i.e* each agent have exactly one edge weighted by p , another by p' and a last by p''). Thus, agents are undistinguishable by their state but also by their out-weight distribution.

Then, the bottom part of Figure 12 represents the associated Markov chain \mathcal{M} associated to this population for the *or* protocol. Agents in state 0 are represent by white circle at contrary to black circle representing agents in state 1. To identify the different configurations, all state of the Markov chain has been labeled with an integer of $[1, 8]$. Finally, in order to present a finer graph, we do not represent the loop edges that keeps the population in the same state. All of this edges are labeled by 1 minus the sum of all out-edges probability labels:

$$p_{i,i} = 1 - \left(\sum_{j \in [1, i-1] \cup [i+1, 8]} p_{i,j} \right)$$

In the following, we will denote Adj_i the set containing all the successors of a vertex i in a graph. Using stochastic analysis, by considering the starting point in state 1, we can directly extract the formal expression of the mean hitting time in that case:

$$\begin{aligned} \mathbb{E}_1(\tau_8) &= \sum_{k \in Adj_1} p_{1,k} \cdot \mathbb{E}_1[t_{1 \rightarrow k} + t_{k \rightarrow 8} | 1 \rightarrow k] \\ &= \left(\sum_{k \in Adj_1} p_{1,k} \cdot \mathbb{E}_1[t_{1 \rightarrow k} | 1 \rightarrow k] \right) \\ &\quad + \left(\sum_{k \in Adj_1} p_{1,k} \cdot \mathbb{E}_1[t_{k \rightarrow 8} | 1 \rightarrow k] \right) \\ &= \left(\sum_{k \in Adj_1} p_{1,k} \right) + \left(\sum_{k \in Adj_1} p_{1,k} \cdot \mathbb{E}_k(\tau_8) \right) \\ &= 1 + \sum_{k \in Adj_1} p_{1,k} \cdot \mathbb{E}_k(\tau_8) \end{aligned}$$

We can extract the following equation:

$$\mathbb{E}_1(\tau_8) = 2 + \sum_{q \in \{p, p', p''\}} 2 \cdot q \cdot \frac{3 - 4 \cdot q}{1 - 2 \cdot q} \quad (5)$$

Let now consider the extension of this Markov chain to n agents, slightly represented in Figure 13. In this figure, we represent the state of the population by vectors




Figure 12: Interaction graph with 4 agent shared the same OWD per agent and the associated Markov chain for the *or* operation.

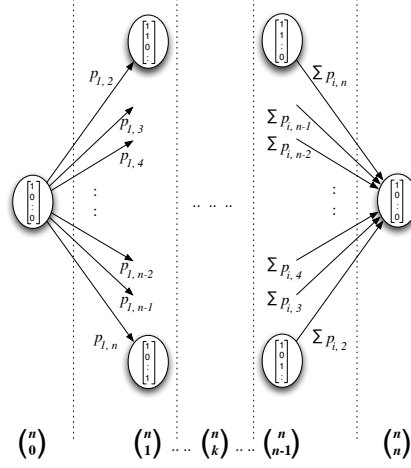


Figure 13: Markov chain for the *or* operation with a population of n agents sharing the same OWD.

containing each current agent's state. The combination presented below each column of states in this chain represents the number of states contained in this column.

In case of uniform interaction (*i.e.* $\forall \theta, \theta' \in \Theta, p_\theta = p_{\theta'}$), the Markov chain is a strait line. Then, the mean hitting time equation is quite simple. But in general case, it seems to difficult to obtain a formal equation of the mean hitting time due to the number of state in the Markov chain:

Lemma 4 *Given $|\mathcal{X}|$ agents and a generic out-weight distribution, the number of states of the associated Markov chain is $|\mathcal{Q}|^{|\mathcal{X}|}$.*

Proof. Each agent pick their state into \mathcal{Q} . Then, if we consider the system state as a vector of each agent's state, it exist at most $|\mathcal{Q}| \times \dots \times |\mathcal{Q}|$ possible states. In case of generic OWD, we cannot consider this vector as a multiset because the position of the agnet in the interaction graph is not independent. Then, we have $|\mathcal{Q}|^{|\mathcal{X}|}$ differents states, which any of them can corresponds to the initial system state. Then, the number of states of the Markov chain is $|\mathcal{Q}|^{|\mathcal{X}|}$. \square

Contents

1	Introduction	3
2	Computation models	4
2.1	Population protocols	4
2.2	Illustrating protocols	5
2.3	Mobility Applied to Population Protocols	6
2.3.1	Integration with related works	7
3	Theoretical analysis of convergence	7
3.1	Background definitions	7
3.2	On the difficulty to analyze a simple example	8
3.3	Estimation of stabilization time using Markov chain	8
4	Impact of mobility models on convergence	9
4.1	Sampling of probability distributions on Γ	9
4.2	Estimation in discrete time: Contact models	10
4.2.1	Complete Interaction Graph	11
4.2.2	Restricted Interaction Graph	11
4.3	Estimation in continuous time: Inter-contact models	15
5	A lower bound: Uniformity is always best	17
6	On the relevance of the Random Way Point model	20
7	Conclusion	22
A	Appendix	24
A.1	Impossibility of simple formula	24



Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399