



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, 2008, pp.26. inria-00296584v1

HAL Id: inria-00296584

<https://inria.hal.science/inria-00296584v1>

Submitted on 12 Jul 2008 (v1), last revised 27 Nov 2008 (v2)

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*

Preliminary version

Yann Busnel — Marin Bertier — Anne-Marie Kermarrec

N° 6580

Juillet 2008

Thème NUM

 *Rapport
de recherche*



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

Preliminary version

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

Thème NUM — Systèmes numériques
Projets Asap

Rapport de recherche n° 6580 — Juillet 2008 — vi 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, ensuring that all reachable system states are endlessly reached, governs the interactions between nodes. This assumption is crucial to ensure that the protocols eventually converge.

This paper studies for the first time the impact of the agents' mobility model on the convergence speed of population protocols. We perform our study by considering several mobility models traditionally used in the ad-hoc network community. 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.

We then formally prove that a uniform distribution of weights provides the lowest bound of average convergence speed for any population protocol.

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

Version préliminaire

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 3.1. 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 chapitre 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 to 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 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 mobile 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 Aware 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 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 extensively used in the mobile network community. 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 prove that the 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 these models. In Section 3 and 4, we present an empirical study of the impact of mobility on the convergence speed of a few traditional protocols (*or* and *sum modulo 4*). Section 5 presents the proofs that the uniform distribution leads to the lowest bound of the average number of steps to reach 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 2.2). 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 f . 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.

- a complete 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, $\Theta = \{(v, v') \in \Upsilon^2 \mid v \neq v'\}$;
- 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 $(p, q) \rightarrow (p', q')$ a transition if $[(p, q) \rightarrow (p', q')] \in \delta$. 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 $f(\sigma)$. The 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 Mobility-Aware Population Protocols

In the population protocol model, a fair scheduler determines the interactions order. This assumption ensures that a reachable state will be endlessly reached. 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 appears infinitely often.

Mobility Aware Population Protocol (*MAPP*) extends the population protocol model to capture mobility patterns. In *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 $\sum_{\theta \in \Theta} p_\theta = 1$.

Lemma 1 *For all schedulers respecting 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 $p_\theta = 0$, 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

3 How to measure the impact of mobility on the convergence speed

The goal of this section is to provide the empirical study setting of the impact of mobility on the convergence speed of population protocols. We conducted the study on two basic functions of population protocols, namely the *or* and the *sum modulo x*. In this section, we first give a brief description of the two operations chosen for the simulation study. We then describe the simulation setting, namely the Monte Carlo Markov Chain model. Two mobility characteristics may have an impact in the convergence speed: (i) the distribution of probability of two agents to interact, we call this model the *sociability model* in analogy with society²; (ii) the frequency of such interactions for a given distribution, characterized by the inter-contact time, this is called the *meeting model*. We successively present the results of these simulations using a complete interaction graphs in the following Section 4.

3.1 Illustrating algorithms

The **Or operation** 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.

The **Sum modulo 4** 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: $\forall v, w \in \{0, 1, 2, 3\}, (v, w) \rightarrow (v + w, \perp_{v+w})$ and $(v, \perp_w) \rightarrow (v, \perp_v)$ in which the addition is made modulo 4.

Eventual correctness The fairness rule is crucial to ensure that all agents eventually stabilize to the correct output value. For instance, the idea of correctness for the flooding protocol is that the number of 1 cannot decrease. If no agent has input 1, then the system is stabilized from start. Conversely, if at least one agent has an input 1 at the

²Two agents having a high probability to get in contact with each other are considered friends.

beginning, then the fairness condition ensures that the number of agents with input 1 will eventually increase to n , and then stabilize.

3.2 Estimation of stabilization time using Markov chain

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 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 MAPP, more constraints are imposed to the scheduler than in the basic population protocol model. At each step t , 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. A system is considered as converged when the stabilized configuration, or a stationary distribution in the associated Markov chain, is reached.

Unfortunately, the number of configuration 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 the following, *convergence speed* refers to the time required to reach the stationary distribution. Then, in order to estimate precisely the convergence time, we use the Markov chain Monte Carlo (MCMC) method [8]. We present in Appendix A.2 the way to compute the mean convergence time and the eventual error using MCMC.

4 Estimation of the mobility model impact

In the following, we introduce first the sampling parameters of our empirical study, then, we present the observed impact of sociability and meeting models.

Sampling environment and out-weight distribution As, in the basic population protocol model, agents are anonymous and undistinguishable, the MAPP model must verify the following assumptions. Given an agent $v \in \Upsilon$, consider the multiset $\mathcal{P}_v = \{p_{v,\psi} | \psi \in \Upsilon - \{v\}\}$, called *out-weight distribution of v* (OWD). All agents of a population must be defined with the same OWD. Otherwise, a specific agent can be choose by the scheduler for an interaction more than any other agent. That leads to a bias of the system, as agents can be distinguished from each other. The following property ensures that agents in MAPP are undistinguishable:

$$\forall v, \psi \in \Upsilon, \mathcal{P}_v = \mathcal{P}_\psi \tag{1}$$

Following this latter assumption, the OWD of an agent is not necessarily uniform, but is the same for all agents in the population.

In this section, we consider a complete interaction graph, in which potentially, each agent may interact with any other. More formally, we have $\forall v \in \Upsilon, \forall p \in \mathcal{P}_v, p \neq 0$.

Mobility Consider that the population is composed of n agents, then $|\mathcal{P}_v| = n - 1$. Each considered OWD, presented in Figure 1 (page 9), is equivalent to a probability density function. As we have assumed that $\sum_{\theta \in \Theta} p_\theta = 1$, the area located below each OWD curve is equal to $\frac{2}{n}$. For each OWD presented in this figure, we introduce below the potential realistic sociability models. Each OWD is named after 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.

Given an agent v in a 100-agent population, Figure 1 (page 9) illustrates the OWD used in the stochastic samples for the MCMC method. Each \mathcal{P}_v is represented in a decreasing way.

4.1 Impact of sociability model

In the following, we study the impact of such mobility models on the convergence speed of the two considered protocols (*or* and *sum modulo 4*) considering a complete interaction graph. As we introduced above, the OWD represents the sociability model considered.

Figures 2 and 3 (page 9) depict, for each aforementioned OWD (Figure 1 – page 9), the average number of interactions required for the *or* and *sum* operations respectively, to reach a stationary state for graph sizes varying from 0 to 100. We observe that the OWD has a significant impact on the convergence speed, the more biased the distributed, the greater the impact. We also observe that the uniform distribution consistently outperforms all other ones.

In Figure 3, we observe that for small-size graphs, Normal-2 and flat are equivalent and as the number of agents increases, the distribution Normal-2 gets close to Normal-1. This evolution, observed consistently over all graphs, is exacerbated here 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 meeting for each pair of agents.




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




Figure 2: Average number of step and confidence interval according to the number of agent n for the flooding protocol




Figure 3: Average number of step and confidence interval according to the number of agent n for the sum modulo 4 protocol




Figure 4: Average convergence time and Confidence interval according to the number of agents for the Flooding protocol with a complete interaction graph and atomic contact




Figure 5: Average convergence time and confidence interval according to the number of agents for the Flooding protocol with a complete interaction graph and uniform contact length

4.2 Simulation of meeting models

While the sociability model provides the probability for each agent to interact with other, it is not possible to directly infer the duration necessary to converge. Thus, if we considered continuous time environment, given a specific OWD does not reflect the frequency of such interactions. This last can be seen as the inter-contact frequency.

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

In these experiments, we use three of the extensively used inter-contact time distribution, extracted notably in [11, 12]: *Exponential*, *Pareto* (known also as *Power-Law*, *Log-Normal* distributions. For each simulation, the setting 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 1 minute and 40 seconds.

Figure 4 and 5 (page 9) illustrate the impact of the meeting model according to size of the interaction graph and the sociability model given. First, in Figure 4, the interactions are atomics *i.e.* only one interaction during a specific contact. On this plot, we can observe that for the same OWD (*Uniform*, *Normal-flat* or *Normal-2 friends*), given a meeting model aforementioned, the mean convergence time is impacted. As well, Figure 5 speaks to this significant impact in the non-atomic interaction. For each meeting of two agents, these last can interaction endlessly during a length picked uniformly between 1 and 100 seconds. Finally, the impacts of the meeting model are meaningful, and have the same order of magnitude, independent to the duration of meeting.

5 A lower bound: Uniformity is always the 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 outperforms other distributions. Although this seems natural for some operations, it is actually counter-intuitive for some others like flooding. For 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. One might think that for a dissemination protocol like flooding, this last distribution may imply a faster stabilization than a uniform one.

In this section we prove that a uniform distribution of weights always achieves the best convergence speed, regardless of the considered operation in population protocols.

Theorem 2 *For any predicate of the Presburger arithmetic (i.e. semilinear predicates, which can be stably computed with population protocol), the uniform OWD in MAPP corresponds to the lower bound of the mean hitting time (i.e. the average number of steps needed to reach the stabilized state).*

Proof. Roughly speaking, we first characterize any predicate computable by a population protocol. Then, by characterizing how to compute them using combination of population protocol, we will infer that any population protocol owns a polynomial mean hitting time. Finally, we prove that 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. Let $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ an arbitrary input alphabet, and A_i, c, m integer constants such that $m \geq 2$. Then, 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 .

Furthermore, for any predicate in the Presburger arithmetic, it exists an equivalent in term of computation behaviour directly derived from those three aforementioned predicates. Therefore, for any population protocol that enables the computation of those predicates, if the best mean hitting time is obtained using an uniform distribution of p_i , then the uniform OWD is the best case of any computation in Presburger arithmetic, and accordingly for any population protocols (cf. [3]).

Mapping with population protocols First of all, due to space constraint, we do not present here in details the population protocols, introduced in [1], corresponding to the aforementioned predicates, but in Appendix A.3. Anyway, based on the definition of this appendix, let \mathcal{A} (respectively \mathcal{B}) a protocol that stationary computes F (respectively G) and consider the protocol \mathcal{C} which stationary computes $\xi(F, G)$ by parallel composition of \mathcal{A} and \mathcal{B} .

The convergence speed of \mathcal{C} directly depends of the speed of \mathcal{A} and \mathcal{B} . Without loses of generality, we assume that \mathcal{A} has a mean hitting time lower than the ones of \mathcal{B} . Then, the mean hitting time of \mathcal{C} is the same than the one of \mathcal{B} . The optimal OWD distribution in term of mean hitting time for the protocol \mathcal{B} will also be the optimal for \mathcal{C} .

Intermediate summary Let us introduce in this paragraph a general sketch of the remaining part of the proof. We have to show first that the optimal distribution for each of the two aforementioned protocols is an uniform OWD. Thereby, using parallel compositions, any semilinear predicate (i.e. belonging to 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 on 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:

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

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

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

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

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

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

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

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 – cf. Appendix A.3). 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 transition in $\mathcal{M}_{\mathcal{P}}$:

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

Then, for any transition q in $\mathcal{M}_{\mathcal{P}}$, q is a linear application of p_i . This infers that $\mathbb{E}_{\mathcal{I}}(\tau_{\mathcal{S}})$ 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 $\text{grad}(f)(p^*) = 0$.

As the minimal value is reached on such a p^* vector, using the weak Lagrangian principle, we can infer that $\text{grad}(g)(p^*) = 0$. 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 random interaction graph.

6 On the relevance of 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 WayPoint mobility model is equivalent to an uniform distribution using 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.*³ 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).

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. \square

7 Conclusion

In this paper, we study the impact of the agents' mobility model on the convergence speed of population protocols. After recall the definition of this model, we introduced *MAPP (Mobility Aware 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 empirically studied the significant impact of various mobility models on the convergence speed of two classical population protocols (*or* and *sum modulo x*) and extracted some interesting bias according to the sociability and meeting models.

From the empirical study, estimations speak 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.

³*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, 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.

A Appendix

A.1 Impossibility of simple formula

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

Let consider a very small population with only for 4 agents. Top of Figure 6 presents the interaction graph of this population. We consider here the most general case with the indistinguishable assumption introduced by Equation 1 in Section 4. The bottom of this figure represents the associated Markov chain for the considered protocol. To simplified the description, all state has been labeled with an integer of $[1, 8]$. 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:

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

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 7. In this figure, we represent the state of the population by vectors 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 straight 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 *Let consider n agents and a generic OWD. The number of states of the Markov chain is $|Q|^{[n]}$.*




Figure 6: Interaction graph with 4 agent shared the same OWD per agent and the associated Markov chain for the *or* operation.

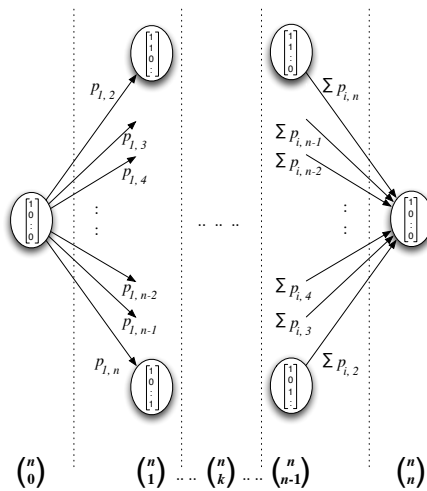


Figure 7: Markov chain for the *or* operation with a population of n agents sharing the same OWD.

Proof. Each agent pick their state into Q . Then, if we consider the system state as a vector of each agent's state, it exist at most $|Q| \times \dots \times |Q|$ 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 $|Q|^{|T|}$ differents states, which any of them can corresponds to the initial system state. Then, the number of states of the Markov chain is $|Q|^{|T|}$. \square

A.2 Monte Carlo method

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, it is possible to compute the possible transitions of a specific configuration (and associated probability of interaction) and provide empirical results on the behavior of the protocol using probabilistic methods. The more difficult problem is to determine how many steps are needed to converge to the stationary distribution within an acceptable error margin. Using the *Ergodic Theorem*⁴, among a large number of samples which stop as soon as the stationary distribution is reached, the average number of step mays 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 more formal outcome [7].

To do so, we have to introduce some definitions:

Definition 1 Let T be an ordered 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". 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 . Then, 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 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 this paper, we denote the same way the theoretical mean (and respectively variance) hitting time and the estimated one with the Monte Carlo method, corresponding to the convergence speed of a population protocol. In the MCMC method, N , the number of runs, is high but not infinite. So, some estimation errors can occur. Thus, we must finally introduce an *error estimator* to leverage the estimation:

⁴Also 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.

Definition 3 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 expectation is defined as follow:

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

A.3 Two entry-level population protocols and the Boolean combination

We present here in details the population protocols, introduced in [1], corresponding to the following predicates, 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 .

in which $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ as an arbitrary input alphabet, and A_i, c, m are integer constants such that $m \geq 2$.

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:

$$q(u, u') = \max(-s, \min(s, u + u')) \text{ and } r(u, u') = u + u' - q(u, u')$$

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.

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:

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

Contents

1	Introduction	3
2	Computation models	4
2.1	Population protocols	4
2.2	Mobility-Aware Population Protocols	5
3	How to measure the impact of mobility on the convergence speed	6
3.1	Illustrating algorithms	6
3.2	Estimation of stabilization time using Markov chain	7
4	Estimation of the mobility model impact	7
4.1	Impact of sociability model	8
4.2	Simulation of meeting models	10
5	A lower bound: Uniformity is always the best	10
6	On the relevance of Random Way Point model	13
7	Conclusion	14
A	Appendix	i
A.1	Impossibility of simple formula	i
A.2	Monte Carlo method	iii
A.3	Two entry-level population protocols and the Boolean combination	iv



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