Population Protocols with Convergence Detection

Yves Mocquard
Université de Rennes 1/IRISA,
yves.mocquard@irisa.fr

Bruno Sericola
INRIA Rennes - Bretagne Atlantique,
bruno.sericola@inria.fr

Emmanuelle Anceaume
CNRS/IRISA,
emmanuelle.anceaume@irisa.fr

***Abstract*—This paper focuses on pairwise interaction-based protocols, and proposes an universal mechanism that allows each agent to locally detect that the system has converged to the sought configuration with high probability. To illustrate our mechanism, we use it to detect the instant at which the proportion problem is solved. Specifically, let n_A (resp. n_B) be the number of agents that initially started in state A (resp. B) and $\gamma_A = n_A/n$, where n is the total number of agents. Our protocol guarantees, with a given precision $\varepsilon > 0$ and any high probability $1 - \delta$, that after $O(n \ln(n/\delta))$ interactions, any queried agent that has set the detection flag will output the correct value of the proportion γ_A of agents which started in state A, by maintaining no more than $O(\ln(n)/\varepsilon)$ integers. We are not aware of any such results. Simulation results illustrate our theoretical analysis.**

***Index Terms*—Population protocols; Detection of convergence; Large scale systems; Anonymous agents; Probabilistic analysis.**

I. Introduction

The main line of research in the population protocol model has so far been the design of pairwise interaction-based protocols that converge to a given configuration of the system as fast as possible while minimizing the number of states needed to converge to that sought configuration. Actually, since the seminal work of Aspnes [4], a considerable amount of work has been done so far to determine which properties can emerge from pairwise interactions between finite-state nodes, together with the derivation of lower bounds on the time and space needed to reach such properties (e.g., [1], [2], [3], [5], [6], [7], [8], [9], [11]).

In this paper we go a step further by proposing a mechanism that allows each agent to locally detect that the system has converged to the sought configuration with high probability. As an application, we propose to use this mechanism to detect the instant at which the proportion problem is solved. Specifically, let n_A (resp. n_B) be the number of agents that initially started in state A (resp. B) and $\gamma_A = n_A/n$ be the majority, where n is the total number of agents. Our protocol guarantees, with a given precision $\varepsilon > 0$ and any high probability $1-\delta$, that after $O(n \ln(n/\delta))$ interactions, any queried agent that has set the detection flag will output the correct value of the proportion γ_A of agents which started in state A, by maintaining no more than $O(\ln(n)/\varepsilon)$ states.

To allow each node to locally detect that its computation of the proportion has converged, we combine three algorithms, each one being run at each node of the system. The first one is dedicated to the computation of the sought property, that is the computation of the proportion γ_A. The second algorithm, run in parallel with the proportion one, aims at constructing the global clock of the system to detect the instant at which convergence is reached at all the nodes. Briefly, when the local clock of at least two nodes have reached a given threshold $T_{\max}$, this means that the number of global interactions in the system is large enough so that all the nodes of the system are able to compute the proportion value with high precision. Thus both nodes can start the propagation of a signal to inform each other node i of the system that i can derive from its local state a good approximation of the proportion γ_A. This dissemination is implemented by a randomized propagation algorithm.

We provide in Section V, a new theoretical analysis of the performance of all these three pairwise interaction-based protocols that improve upon existing ones. As will also be proven in Section V, this detection mechanism is universal in the sense that any pairwise interaction-based population protocol can be augmented with this mechanism to safely detect convergence with high probability. The only requirement that must satisfied is that an upper bound on the convergence time of that protocol must be explicitly known.

The remaining of the paper is organized as follows. Section II formalizes the addressed problem. The model of the system together with the different notations adopted in the paper are presented in Section III. The orchestration of the different ingredients of our detection mechanism are presented in Section IV. A deep theoretical analysis of the performance of our detection mechanism is presented in Section V, and a summary of the simulation results is given in Section VI. Finally, Section VII concludes the paper.

II. The addressed problem

We consider a set of n agents, interconnected by a complete graph, that asynchronously start their execution in one of two distinct states A and B. Let n_A (resp. n_B) be the number of agents whose initial state is A (resp. B), and let $\gamma_A = n_A/n$ be the proportion of the system, with n the total number of agents. We formalize the addressed problem as follows.

Definition 1 (Proportion with Convergence Detection): A population protocol ran by all the nodes of the system solves the proportion with convergence detection problem if with probability at least $1 - \delta$, for any $\delta \in (0, 1)$, any node of the system is capable of computing the proportion γ_A and detecting the instant at which the computed proportion is an ε-approximation of γ_A. The number of interactions and

This work was partially funded by the French ANR project SocioPlug (ANR-13-INFR-0003), and by the DeSceNt project granted by the Labex CominLabs excellence laboratory (ANR-10-LABX-07-01).

the number of states needed to guarantee the convergence detection must respectively be $O(\ln(n))$ and $O(\ln(n)/\varepsilon)$, where ε is the precision of the computed proportion, and n is the number of nodes.

III. Model and Notations

In this work we assume that the collection of nodes communicate through pairwise and asynchronous interactions in a failure-free environment. Initially, each node starts with an initial symbol A or B represented by ι. The input function of each node initializes its local state according to its initial symbol, and then at each interaction its state is updated using a transition function denoted by f. Interactions between nodes are random: at each discrete time, any two agents are randomly selected to interact. The notion of time in population protocols refers to as the successive steps at which interactions occur, while the parallel time refers to as the total number of interactions averaged by n, see Aspnes et al. [4]. Note that nodes do not maintain nor use identifiers, however for ease of presentation, they are numbered $1, 2, \ldots, n$.

We will denote by $(C_t^{(i)}, T_t^{(i)}, S_t^{(i)})$ the state of node i at time t, where $C_t^{(i)}$ is used by node i to evaluate the current value of the proportion, $T_t^{(i)}$ is used by node i to evaluate the global clock of the system, and $S_t^{(i)}$ is a Boolean variable used by node i to indicate whether proportion convergence has been globally reached or not. Let m and $T_{\max}$ be two systems parameters, respectively used to define the initial configuration of the nodes and to determine the global number of interactions after which convergence is reached for all the nodes. Values of both parameters are analyzed in Section V. At any time t, $C_t^{(i)}$ belongs to Q_C with $Q_C = [\![-m, m]\!]$, and we have $|Q_C| = 2m+1$; $T_t^{(i)}$ belongs to Q_T with $Q_T = [\![0, T_{\max}-1]\!]$, and we have $|Q_T| = T_{\max}$; finally $S_t^{(i)}$ belongs to Q_S with $Q_S = \{0, 1\}$, and we have $|Q_S| = 2$. Thus the size of the state set of any node is equal to $|Q_C| \times |Q_T| \times |Q_S|$.

The configuration of the system at time t is the state of each node at time t and is denoted by (C_t, T_t, S_t) where $C_t = (C_t^{(1)}, \ldots, C_t^{(n)})$, $T_t = (T_t^{(1)}, \ldots, T_t^{(n)})$, and $S_t = (S_t^{(1)}, \ldots, S_t^{(n)})$.

Interactions between nodes are orchestrated by a random scheduler: at each discrete time $t \geq 0$, two indices i and j are randomly chosen to interact with probability $p_{i,j}(t)$. The successive choices of the interacting pair of nodes are supposed to be independent and uniformly distributed, which means that we have

$$P_{i,j}(t) = \frac{1_{\{i \neq j\}}}{n(n-1)}.$$

IV. Algorithm run at each node

Each node i maintains, as its current state, a vector made of three components $(C^{(i)}, T^{(i)}, S^{(i)})$, initialized according to Algorithm 1.

Pairs of nodes interact randomly (see Algorithm 2) and during their interaction update their state by computing the average of their C values and by incrementing their clock values T

```
Function Init(i):
    if ι^(i) = A then C^(i) := m;
    if ι^(i) = B then C^(i) := −m;
    T^(i) := 0;
    S^(i) := 0;
```

Algorithm 1: Initialization of node i's state (input function)

as respectively described in Algorithms 3 and 4. The transition function f is given by $f(x, y) = (\lfloor (x+y)/2 \rfloor, \lceil (x+y)/2 \rceil)$. When two interacting nodes have both their clock equal to $T_{\max} - 1$ (i.e. the number of global interactions in the system is large enough to allow all the nodes of the system to locally compute the proportion value with high precision), they both set their signal value to 1, which indicates the starting point of the spreading (see Algorithm 5). If during an interaction, a node updates its signal value, i.e. $S = 1$, then it stops updating both its C and T values, and at any subsequent interactions, this node will only "propagate" signal S.

```
Function UpdateState(i, j):
    if Spreading (i, j) = 0 then
        if Clock (i, j) = 0 then
            Average (i, j);
        end
    end
```

Algorithm 2: Update of the state of two nodes during their interaction

```
Function Average(i,j):
    (C^(i), C^(j)) := (⌊(C^(i)+C^(j))/2⌋, ⌈(C^(i)+C^(j))/2⌉);
```

Algorithm 3: Update of C values of two interacting nodes

```
Function Clock(i, j):
    if T^(i) = T^(j) = (T_max − 1) then
        S^(i) := S^(j) := 1;
        return 1;
    end
    if T^(i) ≤ T^(j) then T^(i) := T^(i) + 1 ;
    else T^(j) := T^(j) + 1 ;
    return 0;
```

Algorithm 4: Update of S values of two interacting nodes

Upon query, a node i returns its estimation ω_A of the proportion of initial A according to its current value of $C^{(i)}$. We have

$$\omega_A(C^{(i)}) = (m + C^{(i)})/(2m)$$

Note that in addition to the proportion, node i also returns its signal $S^{(i)}$ (see Algorithm 6). As demonstrated in Section V, the proportion computed by i is an ε-approximation of γ_A with any high probability $1 - \delta$ if $S^{(i)} = 1$. Note that if $S^{(i)} = 0$ then i does not know how far its computation is from γ_A.

Function `Spreading(i, j)`:
  if $S_t^{(i)} = S_t^{(j)} = 0$ **then** **return** 0 ;
  $S^{(i)} := S^{(j)} := 1$;
  return 1;

Algorithm 5: Update of T values of two interacting nodes

Function $\omega_A(i)$:
  return $\left(\frac{m+C^{(i)}}{2m}, S^{(i)}\right)$;

Algorithm 6: Output function of the Algorithm

V. Analysis

This section is devoted to the analysis of our solution. We split the analysis into five parts, the first one devoted to the analysis of the rumor spreading function (see Section V-A), the second one to the analysis of the average function (see Section V-B) and the third one to the global clock function (see Section V-C). The analysis presented in the fourth part consists in evaluating the global behavior of our protocol by combining the previous evaluations (see Section V-D). We end this section by showing under which hypothesis our convergence detection mechanism can be applied to any pairwise interaction-based protocol (see Section V-E).

A. Analysis of the rumor spreading function

The spreading starts at the first instant t at which two interacting nodes, say i and j, have both their spreading values, $S_t^{(i)}$ and $S_t^{(j)}$, equal to 1. This instant occurs when both nodes have their clock equal to $T_{\max} - 1$. In order to analyze this spreading time, we first prove a lemma derived from Theorem 4 of [9]. let Y_t the number of informed nodes at time t and Θ_n the first instant at which all nodes know the rumor. We have

$$\Theta_n = \inf\{t \geq 0 \mid Y_t = n\}.$$

Note that in our case we have $Y_0 = 2$. This lemma gives a maximal value of spreading time with a probability less than or equal to any fixed probability $\delta \in (0,1)$ when the system initially starts with 2 nodes knowing the rumor.

Lemma *2:* For all δ, we have

$$\mathbb{P}\{\Theta_n \geq \lceil n\,(\ln(n) - \ln(\delta)/2)\rceil \mid Y_0 = 2\} \leq \delta.$$

Proof. Applying Theorem 4 of [9] with $i = 2$ leads, for every $k \geq 0$, to

$$\mathbb{P}\{\Theta_n \geq k \mid Y_0 = 2\} \leq (n-2)^2 \left(1 - \frac{2}{n}\right)^k.$$

Setting $k = \lceil n\,(\ln(n) - \ln(\delta)/2)\rceil$, we obtain

$$\left(1 - \frac{2}{n}\right)^{\lceil n(\ln(n)-\ln(\delta)/2)\rceil} \leq \left(1 - \frac{2}{n}\right)^{n(\ln(n)-\ln(\delta)/2)} = e^{n(\ln(n)-\ln(\delta)/2)\ln(1-2/n)}.$$

Using the fact that $\ln(1-x) \leq -x$, for all $x \in [0,1)$, we get

$$\left(1 - \frac{2}{n}\right)^{\lceil n(\ln(n)-\ln(\delta)/2)\rceil} \leq e^{-2\ln(n)+\ln(\delta)} = \frac{\delta}{n^2}$$

and thus

$$\mathbb{P}\{\Theta_n \geq \lceil n\,(\ln(n) - \ln(\delta)/2)\rceil \mid Y_0 = 2\} \leq \frac{(n-2)^2\delta}{n^2} \leq \delta,$$

which completes the proof. ■

Note that the proof of this lemma does not make use of the Markov inequality. The approximations have all been made with equivalents, which means that the result of this lemma is quite close to the reality. This will be illustrated in section VI-A.

B. Analysis of the average function

The average function is modelled by vector C_t. This transition function is given, for interacting nodes i and j, by

$$\left(C_{t+1}^{(i)}, C_{t+1}^{(j)}\right) = \left(\left\lfloor \frac{C_t^{(i)} + C_t^{(j)}}{2} \right\rfloor, \left\lceil \frac{C_t^{(i)} + C_t^{(j)}}{2} \right\rceil\right) \quad (1)$$
$$\text{and} \quad C_{t+1}^{(r)} = C_t^{(r)} \text{ for } r \neq i, j.$$

The following lemma, which states that the sum of the entries of vector C_t is constant, is straightforward.

Lemma *3:* For every $t \geq 0$, we have

$$\sum_{i=1}^{n} C_t^{(i)} = \sum_{i=1}^{n} C_0^{(i)}.$$

Proof. The proof is immediate since the transformation from C_t to C_{t+1} described in Relation (1) does not change the sum of the entries of C_{t+1}. Indeed, from Relation (1), we have $C_{t+1}^{(i)} + C_{t+1}^{(j)} = C_t^{(i)} + C_t^{(j)}$ and the other entries do not change their values. ■

We denote by ℓ the mean value of the sum of the entries of C_t and by L the row vector of $\mathbb{R}^n$ with all its entries equal to ℓ, that is

$$\ell = \frac{1}{n}\sum_{i=1}^{n} C_t^{(i)} \text{ and } L = (\ell, \ldots, \ell).$$

Clearly, from Lemma 3, the value of ℓ is independent of the time t. The following theorem shows that, after a given amount of time, the distance between all the $C_t^{(i)}$ and ℓ is less than $3/2$ with any high probability. Recall that the infinite norm is defined for any n-dimensional vector v by $\|v\|_\infty = \max_{i=1,\ldots,n} v_i$.

Theorem *4:* For all $\delta \in (0,1)$, if there exists a constant K such that $\|C_0 - L\| \leq K$ then, for every $t \geq n\,(2\ln(K) + 1.78\ln(n) - 7.60\ln(\delta) + 2.70)$, we have

$$\mathbb{P}\left\{\|C_t - L\|_\infty < 3/2\right\} \geq 1 - \delta.$$

Proof. See [10] ■

We now apply these results to compute the proportion γ_A of agents whose initial input was A, which is given by $\gamma_A =$

$n_A/(n_A+n_B) = n_A/n$. Recall that the output function ω_A is given, for all $x \in Q_C$, by

$$\omega_A(x) = (m+x)/(2m).$$

The following theorem shows that, after a given amount of time, the value of all the $\omega_A(C_t^{(i)})$ is an ε-approximation of the proportion γ_A with any probability.

Theorem *5:* For all $\varepsilon, \delta \in (0,1)$, setting $m = \lceil 3/(4\varepsilon) \rceil$, we have, for all $t \geq n\,(2.78\ln(n) - 2\ln(\varepsilon) - 7.60\ln(\delta) + 2.70)$,

$$P\{|\omega_A(C_t^{(i)}) - \gamma_A| < \varepsilon \text{ for all } i = 1,\ldots,n\} \geq 1-\delta.$$

Proof. We have $\|C_0 - L\| \leq m\sqrt{n}$. Applying Theorem 4, with $K = \sqrt{n}/\varepsilon \geq \lceil 3/(4\varepsilon) \rceil \sqrt{n} = m\sqrt{n}$, we obtain for all $\delta \in (0,1)$ and $t \geq n\,(2.78\ln(n) - 2\ln(\varepsilon) - 7.60\ln(\delta) + 2.70)$,

$$P\{\|C_t - L\|_\infty \geq 3/2\} \leq \delta$$

or equivalently

$$P\{|C_t^{(i)} - (\gamma_A - \gamma_B)m| < 3/2, \text{ for all } i = 1,\ldots,n\} \geq 1-\delta.$$

Since $\gamma_A + \gamma_B = 1$ we have

$$\begin{aligned} |C_t^{(i)} - (\gamma_A - \gamma_B)m| &= |C_t^{(i)} - (2\gamma_A - 1)m| \\ &= |m + C_t^{(i)} - 2m\gamma_A| \\ &= 2m|\omega_A(C_t^{(i)}) - \gamma_A|. \end{aligned}$$

Then

$$P\{|\omega_A(C_t^{(i)}) - \gamma_A| < 3/(4m), \text{ for all } i = 1,\ldots,n\} \geq 1-\delta.$$

So

$$P\{|\omega_A(C_t^{(i)}) - \gamma_A| < \varepsilon, \text{ for all } i = 1,\ldots,n\} \geq 1-\delta,$$

which completes the proof. ■

C. Analysis of the clock function

The clock function is modelled by vector T_t, and for any two interacting nodes i and j, we have

$$\left(T_{t+1}^{(i)}, T_{t+1}^{(j)}\right) = \begin{cases} \left(T_t^{(i)}+1, T_t^{(j)}\right) & \text{if } T_t^{(i)} \leq T_t^{(j)} \\ \left(T_t^{(i)}, T_t^{(j)}+1\right) & \text{if } T_t^{(i)} > T_t^{(j)}. \end{cases}$$

$$\text{and } \; T_{t+1}^{(r)} = T_t^{(r)} \text{ for } r \neq i,j.$$

The maximal difference between any two clocks at time t is also called the gap at time t. It is denoted by $\mathrm{Gap}(t)$ and is defined by

$$\mathrm{Gap}(t) = \max_{1\leq i\leq n}\left(T_t^{(i)}\right) - \min_{1\leq i\leq n}\left(T_t^{(i)}\right).$$

The following theorem gives a maximal value of the gap with any fixed probability. Note that this value is independent of the global time t.

Theorem *6:* For all $\delta \in (0,1)$, we have

$$\mathbb{P}\left\{\mathrm{Gap}(t) \geq 10\ln(n) - 10\ln(\delta) + 74\right\} \leq \delta.$$

Proof. From Relation (1) and Figure 1 of [8] in which we take $a = 10$ and $b = 74$, we obtain, for all $\sigma > 0$,

$$\mathbb{P}\left\{\mathrm{Gap}(t) \geq 10(1+\sigma)\ln(n) + 74\right\} \leq 1/n^\sigma.$$

Let $\delta \in (0,1)$. By taking $\sigma = -\ln(\delta)/\ln(n)$, we get $\sigma\ln(n) = -\ln(\delta)$ and $1/n^\sigma = \delta$, that is

$$\mathbb{P}\left\{\mathrm{Gap}(t) \geq 10\ln(n) - 10\ln(\delta) + 74\right\} \leq \delta,$$

which completes the proof. ■

The following properties will also be used in the next section. Since at each time only one node has its clock incremented by one we have

$$\sum_{i=1}^{n} T_t(i) = t.$$

It follows easily that at each instant $t \geq 0$, we have

$$\min_{1\leq i\leq n}\left(T_t^{(i)}\right) \leq \frac{t}{n} \leq \max_{1\leq i\leq n}\left(T_t^{(i)}\right) \tag{2}$$

D. Analysis of the proportion protocol with convergence detection

We now combine all the previous analyses to evaluate the behavior of our proportion protocol with convergence detection. For every $n \geq 2$ and for all $\delta \in (0,1)$, we introduce the following constants:

- $\tau_1 = \ln(n) - 0.5\ln(\delta) + 0.55$.
- $\tau_2 = 2.78\ln(n) - 2\ln(\varepsilon) - 7.60\ln(\delta) + 11.05$.
- $\tau_3 = 10\ln(n) - 10\ln(\delta) + 84.99$.

Constant τ_1 is the constant used in Lemma 2 with $\delta/3$ instead of δ. It is the maximal parallel time for the spreading protocol to converge with probability greater than $1-\delta/3$.

Constant τ_2 is the constant used in Theorem 5 with $\delta/3$ instead of δ. It is the maximal parallel time for the proportion protocol to converge with probability greater than $1-\delta/3$.

Constant τ_3 is the constant used in Theorem 6 with $\delta/3$ instead of δ. It is the maximal gap obtained with probability greater than $1-\delta/3$.

With these constants, we set $T_{\max} = \tau_2 + \tau_3$. The following theorem is the main result of the paper. It states that, after time $n(T_{\max} + \tau_1)$, all the nodes have an ε-approximation of γ_A and that the spreading is terminated, with probability greater then $1-\delta$. More practically, it also states that, if at any instant t a node has its spreading value equal to 1, then all the nodes have an ε-approximation of γ_A with probability greater than $1-\delta$.

Theorem *7:* For every $\delta \in (0,1)$ and $t \geq n(T_{\max} + \tau_1)$, we have

$$\mathbb{P}\left\{\left|\omega_A(C_t^{(i)}) - \gamma_A\right| \leq \varepsilon, S_t^{(i)} = 1 \;\forall i \in [\![1,n]\!]\right\} \geq 1-\delta.$$

Moreover, for every $\delta \in (0,1)$ and $t \geq 0$, we have

$$\mathbb{P}\left\{\left|\omega_A(C_t^{(i)}) - \gamma_A\right| \leq \varepsilon \;\forall i \in [\![1,n]\!] \,\middle|\, \exists j \text{ such that } S_t^{(j)} = 1\right\} \geq 1 - 2\delta/3.$$

Proof. The average protocol and the clock protocol both start at time 0 and run independently of each other. We consider first the clock protocol. Let Γ be the first time where two interacting nodes both have their clock value equal to $T_{\max}-1$. Applying Theorem 6 at instant Γ with $\delta/3$ instead of δ, we get

$$\mathbb{P}\left\{\text{Gap}(\Gamma) < 10\ln(n) - 10\ln(\delta/3) + 74\right\} \geq 1-\delta/3,$$

that is, by definition of τ_3, $\mathbb{P}\left\{\text{Gap}(\Gamma) < \tau_3\right\} \geq 1-\delta/3$. By definition of the gap and since that at instant Γ, we have $\max_{1\leq i\leq n}\left(T_\Gamma^{(i)}\right) = T_{\max}-1$, we obtain

$$\mathbb{P}\left\{T_{\max} - 1 - \min_{1\leq i\leq n}\left(T_\Gamma^{(i)}\right) < \tau_3\right\} \geq 1-\delta/3,$$

that is, by definition of $T_{\max}$,

$$\mathbb{P}\left\{\min_{1\leq i\leq n}\left(T_\Gamma^{(i)}\right) > \tau_2\right\} \geq 1-\delta/3.$$

From Relation (2) we have $\min_{1\leq i\leq n}\left(T_\Gamma^{(i)}\right) \leq \Gamma/n$, which leads to

$$\mathbb{P}\left\{\Gamma > n\tau_2\right\} \geq 1-\delta/3. \tag{3}$$

For what concerns the average protocol, to simplify the writing we introduce the events E_t defined by

$$E_t = \left\{\left|\omega_A(C_t^{(i)}) - \gamma_A\right| < \varepsilon \text{ for all } i \in [\![1,n]\!]\right\}.$$

Applying Theorem 5 with $\delta/3$ instead of δ we get, by definition of τ_2, for all $t \geq n\tau_2$,

$$\mathbb{P}\left\{E_t\right\} \geq 1-\delta/3.$$

The random variables $C_t^{(i)}$ and Γ being independent, we have for every $t \geq 0$,

$$\begin{aligned}
\mathbb{P}\left\{E_{\Gamma+t}, \Gamma > n\tau_2\right\} &= \sum_{s=n\tau_2+1}^{\infty} \mathbb{P}\left\{E_{s+t}, \Gamma = s\right\} \\
&= \sum_{s=n\tau_2+1}^{\infty} \mathbb{P}\left\{E_{s+t}\right\}\mathbb{P}\left\{\Gamma = s\right\} \\
&\geq (1-\delta/3)\sum_{s=n\tau_2+1}^{\infty} \mathbb{P}\left\{\Gamma = s\right\} \\
&= (1-\delta/3)\mathbb{P}\left\{\Gamma > n\tau_2\right\} \\
&\geq (1-\delta/3)^2.
\end{aligned}$$

It follows that, for every $t \geq 0$,

$$\mathbb{P}\left\{E_{\Gamma+t}\right\} \geq \mathbb{P}\left\{E_{\Gamma+t}, \Gamma > n\tau_2\right\} \geq (1-\delta/3)^2. \tag{4}$$

The starting point of the spreading period is instant $\Gamma+1$. By definition of Γ, instant $\Gamma+1$ is the first time at which exactly two agents have their spreading values equal to 1. More precisely, for every $t \geq 0$, we introduce the random variable Y_t defined by

$$Y_t = \sum_{i=1}^{n} S_t^{(i)}.$$

We have $Y_t = 0$ for every $t \leq \Gamma$ and $Y_{\Gamma+1} = 2$. The spreading time Θ_n is thus the first instant at which the spreading values of all the agents are equal to 1. It is then defined by

$$\begin{aligned}
\Theta_n &= \inf\{t \geq 0 \mid Y_t = n\} - (\Gamma+1) \\
&= \inf\{t \geq \Gamma+1 \mid Y_t = n\} - (\Gamma+1).
\end{aligned}$$

From Lemma 2 in which we use $\delta/3$ instead of δ, we have, since $Y_{\Gamma+1} = 2$,

$$\mathbb{P}\{\Theta_n < \lceil n\tau_1 \rceil\} \geq 1-\delta/3.$$

Again, to simplify the writing we introduce the events F_t defined by

$$F_t = \left\{S_t^{(i)} = 1 \text{ for all } i \in [\![1,n]\!]\right\}.$$

By definition of the Boolean $S_t^{(i)}$ and of the spreading time Θ_n, we have, for all $t \geq 0$,

$$\begin{aligned}
&\mathbb{P}\left\{F_{\Gamma+1+n\tau_1+t} \mid E_{\Gamma+1+n\tau_1+t}\right\} \\
&\quad = \mathbb{P}\left\{\Theta_n \leq n\tau_1 + t \mid E_{\Gamma+1+n\tau_1+t}\right\} \\
&\quad \geq 1-\delta/3,
\end{aligned}$$

where the last inequality follows from Lemma 2. Unconditioning and using Relation (4), we obtain

$$\begin{aligned}
&\mathbb{P}\left\{F_{\Gamma+1+n\tau_1+t}, E_{\Gamma+1+n\tau_1+t}\right\} \\
&\quad = \mathbb{P}\left\{F_{\Gamma+1+n\tau_1+t} \mid E_{\Gamma+1+n\tau_1+t}\right\}\mathbb{P}\left\{E_{\Gamma+1+n\tau_1+t}\right\} \\
&\quad \geq (1-\delta/3)(1-\delta/3)^2 = (1-\delta/3)^3.
\end{aligned}$$

Recalling that $\Gamma = \sum_{i=1}^{n} T_\Gamma^{(i)}$ and that $T_t^{(i)} \leq T_{\max}-1$, for all $t \geq 0$, we get $\Gamma \leq nT_{\max}$. Considering instant $t = s + nT_{\max} - \Gamma$ which is positive, finally leads, for all $s \geq 0$, to

$$\mathbb{P}\left\{F_{nT_{\max}+n\tau_1+s+1}, E_{nT_{\max}+n\tau_1+s+1}\right\} \geq (1-\delta/3)^3,$$

which is equivalent to say that, for all $t \geq n(T_{\max}+\tau_1)$, we have

$$\mathbb{P}\left\{F_t, E_t\right\} \geq (1-\delta/3)^3 \geq 1-\delta,$$

which completes the first part of the proof.
For the second part of the proof, note that

$$\exists j \text{ such that } S_t^{(j)} = 1 \Longleftrightarrow \Gamma \leq t.$$

We thus have applying Relation (4)

$$\begin{aligned}
\mathbb{P}\left\{E_t \mid \exists j \text{ such that } S_t^{(j)} = 1\right\} &= \mathbb{P}\left\{E_t \mid \Gamma \leq t\right\} \\
&= \mathbb{P}\left\{E_{\Gamma+t}\right\} \\
&\geq (1-\delta/3)^2 \geq 1-2\delta/3.
\end{aligned}$$

This completes the second part of the proof. ■

This theorem shows that the convergence is $O(\ln(n))$ and that then number of states needed is equal to $|Q_T \times Q_C \times Q_S| = 2\left(2\lceil 3/(4\varepsilon)\rceil + 1\right)T_{\max} = O\left(\ln(n)/\varepsilon\right)$.

E. Generalizing the convergence detection mechanism

We now show that our detection mechanism can be applied to any pairwise interaction-based protocol $\mathcal{P}$ so that any node of the system can safely detect the instant at which convergence is reached by all the nodes of the system. The only requirement for this mechanism to be applied is that the convergence time of $\mathcal{P}$ must be precisely known.

Specifically, let us consider the transition function f of the protocol $\mathcal{P}$ such that Relation (1) is replaced by

$$\left(C_{t+1}^{(i)}, C_{t+1}^{(j)}\right) = f\left(C_t^{(i)}, C_t^{(j)}\right)$$
$$\text{and} \quad C_{t+1}^{(r)} = C_t^{(r)} \text{ for } r \neq i, j.$$

Line 2 of Algorithm 3 is thus replaced by

$$\left(C^{(i)}, C^{(j)}\right) := f\left(C^{(i)}, C^{(j)}\right).$$

The initial value C_0 of vector C_t is given and the set of states Q_C of $C_t^{(i)}$ is supposed to be finite. As convergence indicator, we consider the general function ν from $(Q_C)^n$ to $\{0,1\}$. We also suppose that we have a general version of Theorem 5 stating that for all $\delta \in (0,1)$ and for all $t \geq \tau_C(n,\delta)$, we have

$$\mathbb{P}\left\{\nu(C_t) = 1\right\} \geq 1 - \delta. \tag{5}$$

Note that by taking $\tau_2 = \tau_C(n,\delta)$ and

$$\nu(C_t) = 1_{\{|\omega_A(C_t^{(i)}) - \gamma_A| < \varepsilon \text{ for all } i=1,\ldots,n\}}$$

we arrive to the previous result of Theorem 5.

Under the previous assumptions, the generalization of Theorem 7 is then the following. We set $T_{\max} = \tau_C(n,\delta/3) + \tau_3$.

Theorem *8:* For all $\delta \in (0,1)$ and for all $t \geq n(T_{\max} + \tau_1)$ we have

$$\mathbb{P}\left\{\nu(C_t) = 1\right\} \geq 1 - \delta.$$

Moreover, all $\delta \in (0,1)$ and $t \geq 0$, we have

$$\mathbb{P}\left\{\nu(C_t) = 1 \,\middle|\, \exists j \text{ such that } S_t^{(j)} = 1\right\} \geq 1 - 2\delta/3.$$

Proof. By defining $E_t = \{\nu(C_t) = 1\}$, the proof follows exactly the same lines of the proof of Theorem 7, in which τ_2 is replaced by $\tau_C(n,\delta/3)$. ■

The number of states is $|Q_T \times Q_C \times Q_S| = 2T_{\max}|Q_C|$.

VI. Simulations

In this section we first provide simulation results for the spreading, the average, and the clock functions, and then present simulation results for the full protocol.

A. Spreading rumor

This section shows how tight our bound given in Lemma 2 is to our simulation results.

For our purpose, a simulation consists in the following steps: first, all the n nodes are initialized to 0 except for two nodes which are initialized to 1. Then, at each step of the simulation, two nodes are randomly chosen to interact and update their state, by keeping the maximal value of both ones. The simulation stops when all the nodes have their values equal to 1. We have run N independent simulations and have stored and ordered the N values of the spreading times denoted by $\theta_1 \leq \ldots \leq \theta_N$. Recall that the spreading time θ_i is the total number of interactions to propagate an information to all the nodes of the system. The estimation of the instant τ such that $\mathbb{P}(\Theta_n < n\tau\} \geq 1 - \delta$ is thus given by the value $\theta_{\lceil N(1-\delta)\rceil}$.

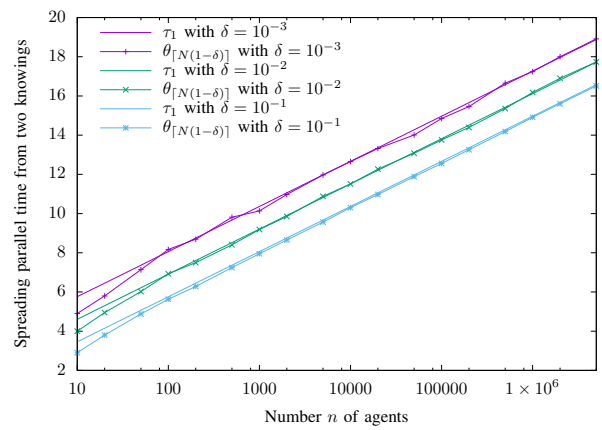

Figure 1. Parallel convergence time of Rumor Spreading as a function of n, with $N = 10^4$.

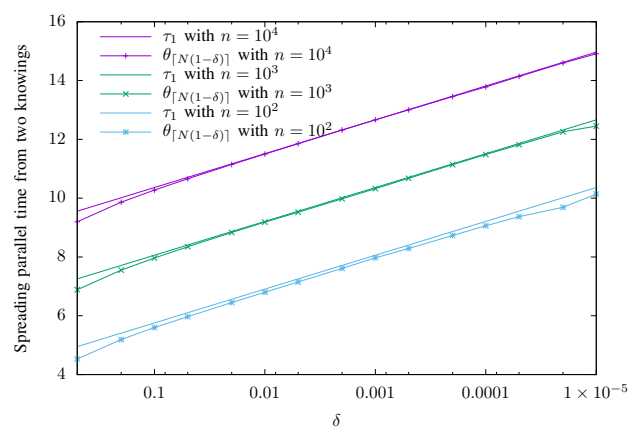

Figure 2. Parallel convergence time of Rumor Spreading as a function of δ, with $N = 10^6$.

Recall that the convergence parallel time is equal to the convergence time divided by n. Figures 1 and 2 depict the convergence parallel times $\theta_{\lceil N(1-\delta)\rceil}/n$ and τ_1 for different values of δ for the first one, and for different values n for the second one. Both figures shows that the theoretical results are quite close to the simulation ones.

B. Average

For each value of ε, we take $m = \lceil 3/(4\varepsilon)\rceil$. Next we choose $n_A = \lceil n/4m + n/2\rceil$ and $n_B = n - n_A$. A simulation consists in the steps described in Algorithm 3 and in Section V-B. The simulation stops when the difference between the minimal and the maximal values of the entries of vector C_t is less than or equal to 2. We ran N independent simulations and stored the

N values of the number of interactions performed which we ordered as $\theta_1 \leq \ldots \leq \theta_N$. The estimation of the instant τ such that, for $t \geq \tau$,

$$\mathbb{P}\left\{|\omega_A(C_t^{(i)}) - \gamma_A| < \varepsilon \text{ for all } i = 1, \ldots, n\right\} \geq 1 - \delta$$

is thus given by the value $\theta_{\lceil N(1-\delta)\rceil}$.

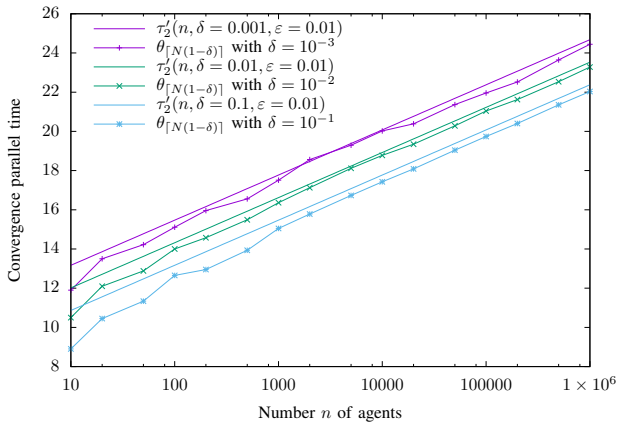

Figure 3. Parallel convergence time of Proportion Computation as a function of n, with $N = 10^4$ and $\varepsilon = 0.01$.

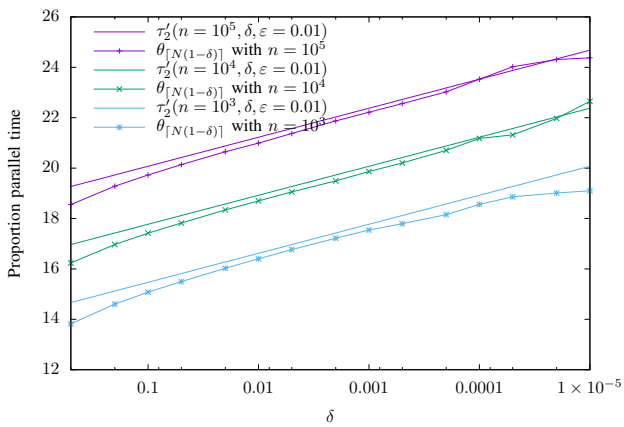

Figure 4. Parallel convergence time of Proportion Computation as a function of δ, with $N = 10^5$ and $\varepsilon = 0.01$.

Figures 3, 4 and 5 depict the convergence parallel time $\theta_{\lceil N(1-\delta)\rceil}/n$ for different values of δ in the first one, for different values of n for the second one, and for different values of ε for the third one. Note that in both the first and the second figure, we have $\varepsilon = 0.01$, that is $m = 75$. In each figure the values of $\theta_{\lceil N(1-\delta)\rceil}/n$ are compared to an intuited value $\tau_2'(n, \delta, \varepsilon)$ close to the expression of τ_2 whose coefficients have been derived from the simulation results, and given by

$$\tau_2'(n, \delta, \varepsilon) = \ln(n) - 0.5\ln(\delta) - 2\ln(\varepsilon) - 1.80. \quad (6)$$

C. The clock

For the clock protocol a simulation consists in the steps described in Algorithm 4 and in Section V-C. We start the

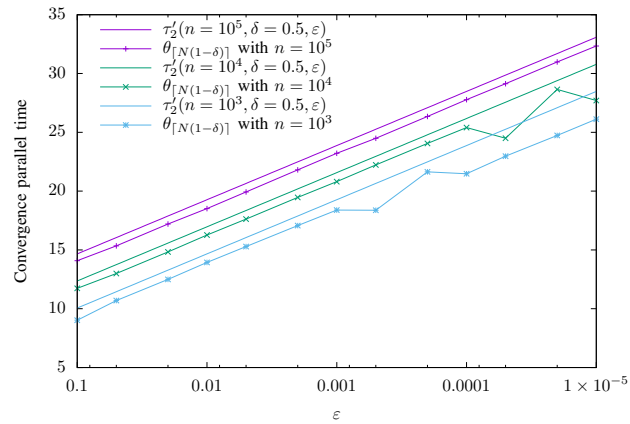

Figure 5. Parallel convergence time of Proportion Computation as a function of ε, with $N = 10^3$ and $\delta = 0.5$.

evaluation of the gap after the first $50n$ interactions. We then store the gap every 100 interactions. Wa ran x simulations and for each simulation we stored the gap y times. This means that the duration of a simulation is equal to $100y + 50n$. The number N of values of the gap obtained is thus $N = xy$. These N values are stored and reordered as $\text{Gap}_1 \leq \ldots \leq \text{Gap}_N$. The estimation of the instant τ such that

$$\mathbb{P}\left\{\text{Gap}(t) \geq \tau\right\} \leq \delta$$

is thus given by the value $\text{Gap}_{\lceil N(1-\delta)\rceil}$.

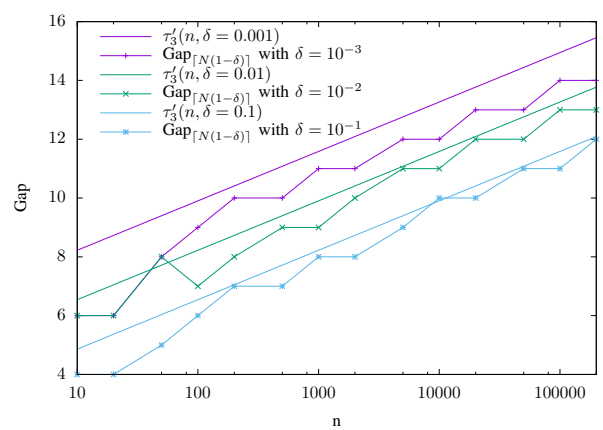

Figure 6. Gap of the clock as a function of n, with $N = 10^6$.

Figures 6 and 7 depict the gap $\text{Gap}_{\lceil N(1-\delta)\rceil}$ for different values of δ for the first one and for different values of n for the second one. In Figure 6 we chose $x = 10$ and $y = 10000$ and in Figure Figures 6 we chose $x = 100$ and $y = 10000$. In each figure the values of $\text{Gap}_{\lceil N(1-\delta)\rceil}/n$ are compared to an intuited value $\tau_3'(n, \delta)$ close to the expression of τ_3 whose coefficients have been derived from the simulation results, and given by

$$\tau_3'(n, \delta) = 0.73\ln(n) - 0.73\ln(\delta) + 1.5. \quad (7)$$

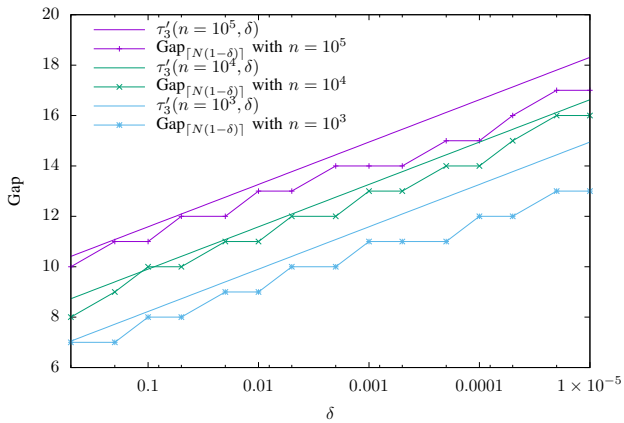

Figure 7. Gap of the clock as a function of δ, with $N = 10^7$.

D. Optimized protocol derived from simulation

From Relations (6) and (7) we derive an intuited value $T'_{\max}$ close to the expression of $T_{\max}$ for the proportion protocol with convergence detection. It is given by

$$\begin{aligned} T'_{\max} &= \tau'_2(n, \delta/3, \varepsilon) + \tau'_3(n, \delta/3) \\ &= 1.73 \ln(n) - 1.23 \ln(\delta) - 2 \ln(\varepsilon) + 1.05. \end{aligned}$$

For different values of n and ε, we ran $N = 1000$ independent simulations taking $\delta = 10^{-6}$, using the value of $T'_{\max}$ instead of $T_{\max}$. We stored the convergence times $\theta_1, \ldots, \theta_N$ defined, for $i = 1, \ldots, N$, by

$$\theta_i = \inf \left\{ t \geq 0 \mid S_t^{(j)} = 1, \text{ for all } j \in [\![1, n]\!] \right\}.$$

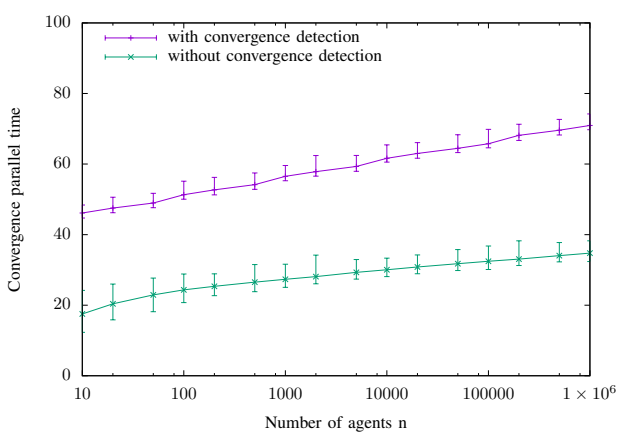

Figure 8. Comparison of the parallel convergence time of proportion computation with convergence detection and without convergence detection as a function of n, with $N = 10^3$ and $\varepsilon = 10^{-5}$.

Figure 8 compares the simulation results of the convergence parallel time (i.e., the expected value $(\theta_1 + \cdots + \theta_N)/N$, the minimal value $\min_{i=1,\ldots,N} \theta_i$ and the maximal value $\max_{i=1,\ldots,N} \theta_i$) when the convergence detection mechanism is used and when it is not used. Clearly, the cost induced by the detection mechanism is not predominant as the convergence time is of the same order of magnitude in both cases.

VII. Conclusion

In this paper we have presented how we can augment, in the population model, a proportion protocol with a convergence detection mechanism to allow each node of the system to locally detect the instant at which convergence to the sought property is reached. A deep theoretical analysis of the performance of each ingredient of our solution has been presented, and simulation results show the impressively weak impact of our detection mechanism on the convergence time of the proportion protocol. We have also shown the applicability of our convergence detection mechanism to many other pairwise interaction-based protocols. For instance, our construction can be applied to a leader election protocol provided that its convergence time is known with high probability.

References

[1] Dan Alistarh, Rati Gelashvili, and Milan Vojnovíc. Fast and exact majority in population protocols. In *Proceedings of the 34th annual ACM symposium on Principles of Distributed Computing (PODC)*, pages 47–56, 2015.

[2] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(2):183–199, 2008.

[3] Dana Angluin, James Aspnes, and David Eisenstat. A simple population protocol for fast robust approximate majority. *Distributed Computing*, 20(4):279–304, 2008.

[4] 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, 2007.

[5] Moez Draief and Milan Vojnovic. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimization*, 50(3):1087–11097, 2012.

[6] George B. Mertzios, Sotiris E. Nikoletseas, Christoforos Raptopoulos, and Paul G. Spirakis. Determining majority in networks with local interactions and very small local memory. In *Proceedings of the 41st International Colloquium (ICALP)*, pages 871–882, 2014.

[7] Yves Mocquard, Emmanuelle Anceaume, and Bruno Sericola. Optimal Proportion Computation with Population Protocols. In *Symposium on Network Computing and Applications*, Boston, United States, October 2016. IEEE.

[8] Yves Mocquard, Emmanuelle Anceaume, and Bruno Sericola. Balanced allocations and global clock in population protocols: An accurate analysis. In *SIROCCO 2018*, June 2018.

[9] Yves Mocquard, Samantha Robert, Bruno Sericola, and Emmanuelle Anceaume. Analysis of the propagation time of a rumour in large-scale distributed systems. In *Proceedings of the 15th IEEE International Symposium on Network Computing and Applications (NCA)*, 2016.

[10] Yves Mocquard, Bruno Sericola, and Emmanuelle Anceaume. Population protocols with convergence detection. Technical report, IRISA - Université Rennes 1 - CNRS - INRIA - hal-01849441 - https://hal.archives-ouvertes.fr/hal-01849441, 2018.

[11] Etienne Perron, Dinkar Vasudevan, and Milan Vojnovic. Using three states for binary consensus on complete graphs. In *Proceedings of the INFOCOM Conference*, pages 2527–2435, 2009.