Balanced allocations and global clock in population protocols: An accurate analysis

Yves Mocquard, Bruno Sericola, Emmanuelle Anceaume

► To cite this version:

Yves Mocquard, Bruno Sericola, Emmanuelle Anceaume. Balanced allocations and global clock in population protocols: An accurate analysis. SIROCCO 2018 - 25th International Colloquium on Structural Information and Communication Complexity, Jun 2018, Ma'ale HaHamisha, Israel. p.296-311, 10.1007/978-3-030-01325-7_26 . hal-01888301

HAL Id: hal-01888301

<https://hal.science/hal-01888301>

Submitted on 4 Oct 2018

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.

Balanced allocations and global clock in population protocols: An accurate analysis

Yves Mocquard^{1*}, Bruno Sericola², and Emmanuelle Anceaume³

¹ Université de Rennes 1 - IRISA (France),
yves.mocquard@irisa.fr

² INRIA Rennes - Bretagne Atlantique (France),
bruno.sericola@inria.fr

³ CNRS - IRISA (France),
emmanuelle.anceaume@irisa.fr

Abstract. The context of this paper is the two-choice paradigm which is deeply used in balanced online resource allocation, priority scheduling, load balancing and more recently in population protocols. The model governing the evolution of these systems consists in throwing balls one by one and independently of each others into n bins, which represent the number of agents in the system. At each discrete instant, a ball is placed in the least filled bin among two bins randomly chosen among the n ones. A natural question is the evaluation of the difference between the number of balls in the most loaded and the one in the least loaded bin. At time t , this difference is denoted by $\text{Gap}(t)$. A lot of work has been devoted to the derivation of asymptotic approximations of this gap for large values of n . In this paper we go a step further by showing that for all $t \geq 0$, $n \geq 2$ and $\sigma > 0$, the variable $\text{Gap}(t)$ is less than $a(1 + \sigma) \ln(n) + b$ with probability greater than $1 - 1/n^\sigma$, where the constants a and b , which are independent of t , σ and n , are optimized and given explicitly, which to the best of our knowledge has never been done before.

1 Introduction

In this paper we address the important issue of the two-choice paradigm analysis [10]. To illustrate the multi-choice paradigm, suppose that we have a set of m balls which are sequentially throws into n bins, where each ball is placed in the least filled bin among $d \geq 1$ ones randomly chosen among the n bins. Azar et al. [5] have characterized this problem by those three values (m, n, d) . A natural question is the analysis of the maximum load in any of the bins, or the maximal gap that may exist between the least loaded bin and the most loaded one. It has been proven that in the simplest case where $d = 1$ (see for example [14]), the maximum load is equal to $m/n + \Theta\left(\sqrt{(m/n) \ln n}\right)$, leading to a gap that increases with the square root of m . Now, instead of choosing a single bin at random, $d \geq 2$ bins are independently and randomly chosen, and the least loaded

* 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)

bin one among those d ones receives a ball. Then Azar et al. [5] have shown that when $m = n$ the maximum load is $\ln(\ln(n))/\ln(2) + O(1)$, and the largest gap is also equal to $\ln(\ln(n))/\ln(2) + O(1)$. These results show that by simply introducing a small choice we get a drastically improved balanced load among all the urns. Citing Mitzenmacher et al [10], "having just two random choices (i.e., $d = 2$) yields a large reduction in the maximum load over having one choice, while each additional choice beyond two decreases the maximum load by just a constant factor". Hence the name of the two-choice paradigm. Later Berenbrink et al. [7] have studied the case (m, n, d) for $d \geq 2$ and $m \gg n$, and proved that the maximum load is equal to $m/n + O(\ln(\ln(n)))$. Note that a simpler proof of this result has been recently found by Talwar and Wieder [15]. Very recently, Peres et al. [12, 13], using a measurement based on the hyperbolic cosine, have generalized the problem in the $(1 + \beta)$ -choice problem. The $(1 + \beta)$ -choice consists, with probability $1 - \beta$, in choosing one bin uniformly at random and to throw a ball in it, and with probability β , in choosing two bins uniformly at random and to throw a ball in the least loaded one. The name comes from the fact that $\mathbb{E}\{d\} = 1 + \beta$. We can note that in their model, each ball is assigned with a random weight. They found a logarithmic bound for both the gap between the maximum loaded bin and the average one [12], and for the gap between the maximum loaded bin and the minimum one [13]. In both cases the gap is $O(\log(n)/\beta)$.

The two-choice paradigm can be used in a multitude of applications, including balanced online resource allocation (where jobs need to be dynamically allocated to the least loaded processor) [1, 6, 8], priority scheduling [4], load balancing [2, 7, 9], and very recently, population protocols [3]. In the later case, the model governing the evolution of these systems consists in throwing balls one by one and independently of each others into n bins, which represents the number of agents in the system. At each discrete instant, a ball is placed in the least filled bin among two bins randomly chosen among the n ones. A natural question is the evaluation of the difference between the number of balls in the most loaded and the one in the least loaded bin. At time t , this difference is denoted by $\text{Gap}(t)$. A lot of work has been devoted to the derivation of asymptotic approximations of this gap for large values of n . In this paper we go a step further by showing that for all $t \geq 0$, $n \geq 2$ and $\sigma > 0$,

$$\mathbb{P}\{\text{Gap}(t) \geq a(1 + \sigma) \ln(n) + b\} \leq \frac{1}{n^\sigma}, \quad (1)$$

where the constants a and b , which are independent of t , σ and n , are optimized and given explicitly, which to the best of our knowledge has never been done before.

The remaining of the paper is structured as follows. In Section 2 we present the addressed problem and a simple algorithm to solve it. Section 3 is the main contribution of our work which consists in providing an accurate bound of the distribution of the gap between any two nodes. Section 4 evaluates constants a and b obtained by our analysis and compares it to constants that we derived from

the work of [4]. The gain in accuracy we obtained by our analysis is significant. Finally Section 5 provides a summary of simulations results.

2 Problem description

We consider a very large set of n nodes (also called agents), interconnected by a complete graph, that asynchronously start their execution in a given state. Agents do not maintain nor use identifiers (agents are anonymous and cannot determine whether any two interactions have occurred with the same agents or not). However, for ease of presentation the agents are numbered $1, 2, \dots, n$. Each agent keeps a local counter, initialized at 0. Agents communicate through random pairwise interactions. On each interaction, the two interacting agents compare their counters, and the one with the lower counter value increments its local counter. The objective of this simple algorithm is the construction of a global clock by guaranteeing that the values of all agent counters are concentrated according to Relation (1). As interactions are uniformly random, this can be related to the classic two choices load balancing process [13]. The goal of the paper is to evaluate the gap between any two agents, that is the maximal difference that may exist at any time t between any two local counters, by accurately evaluating constants a and b . By accurately estimating the maximal gap between any two counter nodes, other population protocols can use it as a *global clock* to perform actions in a probabilistic synchronized way.

We denote by $C_t^{(i)}$ the state of agent i at time t . The stochastic process $C = \{C_t, t \geq 0\}$, where $C_t = (C_t^{(1)}, \dots, C_t^{(n)})$, represents the vector state of the system at time t .

The choice of the two agents which interact is made using a uniform distribution. Given the pair (i, j) of agents which interact at time t , we consider the following evolution of the agents states

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

Note that in the case where agents i and j interact at time t with $C_t^{(i)} = C_t^{(j)}$ then either of two agents can be chosen to have its value increased by 1 at time $t + 1$. A particular choice is made below.

The state space of process C is thus $\mathbb{N}^n$ and a state of this process is also called a protocol configuration. At time 0, we set $C_t^{(i)} = 0$, for every $i = 1, \dots, n$. At each instant the value of only one agent is increased by 1 which means that we have, for every $t \geq 0$,

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

For every $i = 1, \dots, n$, we introduce the quantities $x_i(t) = C_t^{(i)} - t/n$, which leads, for every $t \geq 0$, to

$$\sum_{i=1}^n x_i(t) = 0.$$

The value $C_t^{(i)}$ maintained by agent i is its own view of the global clock t of the system divided by n . More precisely, the approximation of time t , provided by agent i , is $nC_t^{(i)}$.

At each discrete time $t \geq 0$, any two indices i and j are uniformly chosen to interact, independently of the vector state with probability $1/(n(n-1))$.

In order to simplify the presentation, we suppose without any loss of generality that at each instant t , the values of $x_i(t)$ are reordered in a decreasing way, assigning an arbitrary order to agents with the same value. More precisely, at time t the reordering gives

$$x_1(t) = \max_{i=1, \dots, n} (C_t^{(i)} - t/n) \geq \dots \geq x_n(t) = \min_{i=1, \dots, n} (C_t^{(i)} - t/n).$$

We denote by X the rank of the agent whose value is incremented when interaction occurs between 2 agents. In the case where two agents interacting, say i and j , are such that $C_t^{(i)} = C_t^{(j)}$, we choose to increase by 1 the one with the highest rank. If X_1 and X_2 are the ranks of the successive agents which interact, then the probability p_ℓ that agent of rank ℓ is incremented is, for $\ell = 1, \dots, n$,

$$p_\ell = \mathbb{P}\{X = \ell\} = \mathbb{P}\{X_1 = \ell, X_2 < \ell\} + \mathbb{P}\{X_1 < \ell, X_2 = \ell\} = \frac{2(\ell-1)}{n(n-1)} \quad (2)$$

As mentioned in the introduction, the goal of the paper is the evaluation of the distribution of difference between the maximum and the minimum of the entries of vector C_t . This difference is denoted by $\text{Gap}(t)$ and is given, for $t \in \mathbb{N}$, by

$$\text{Gap}(t) = \max_{1 \leq i \leq n} C_t^{(i)} - \min_{1 \leq i \leq n} C_t^{(i)} = x_1(t) - x_n(t).$$

In order to bound the complementary distribution of $\text{Gap}(t)$, we introduce the following potential functions defined, for $\alpha \in \mathbb{R}$, by

$$\Phi(t) = \sum_{i=1}^n e^{\alpha x_i(t)}, \quad \Psi(t) = \sum_{i=1}^n e^{-\alpha x_i(t)} \quad \text{and} \quad \Gamma(t) = \Phi(t) + \Psi(t).$$

The use of these two functions has been proposed in a very clever way by Y. Peres et al. in [13]. The potential function $\Gamma(t)$ is then related to function $\text{Gap}(t)$ by the following lemma.

Lemma 1 *For every $t \geq 0$, we have*

$$\Gamma(t) \geq 2e^{\alpha \text{Gap}(t)/2}. \quad (3)$$

Proof. The exponential function being convex, we have, for every $a, b \in \mathbb{R}$, $2e^{(a+b)/2} \leq e^a + e^b$. Recalling that $\text{Gap}(t) = x_1(t) - x_n(t)$, we obtain

$$\Gamma(t) = \sum_{i=1}^n e^{\alpha x_i(t)} + \sum_{i=1}^n e^{-\alpha x_i(t)} \geq e^{\alpha x_1(t)} + e^{-\alpha x_n(t)} \geq 2e^{\alpha \text{Gap}(t)/2},$$

which completes the proof. ■

This result will be used at the end of the paper for the evaluation of the distribution of $\text{Gap}(t)$ which is based on the evaluation of the one of $\Gamma(t)$, which forms the main part of the paper.

3 Analysis

We first need the two following technical lemmas which are proved in [11].

Lemma 2 *For all $x \in \mathbb{R}$, we have $1+x \leq e^x$. For all $x \in (-\infty, c]$, we have $e^x \leq 1+x+x^2$, where c is the unique positive solution to equation $e^c - 1 - c - c^2 = 0$. The value of c satisfies $1.79 < c < 1.8$.*

Lemma 3 *Let $u = (u_k)_{k \geq 1}$ and $v = (v_k)_{k \geq 1}$ be two monotonic sequences of real numbers and let m_n be the sequence of mean values of sequence v defined, for $n \geq 1$, by*

$$m_n = \frac{1}{n} \sum_{k=1}^n v_k.$$

If the sequences u and v are both non-decreasing or both non-increasing then we have

$$\sum_{k=1}^n u_k v_k \geq m_n \sum_{k=1}^n u_k.$$

If one of these two sequences is non-increasing and the other is non-decreasing then we have

$$\sum_{k=1}^n u_k v_k \leq m_n \sum_{k=1}^n u_k.$$

For every $t \geq 0$, we introduce the notation $x(t) = (x_1(t), \dots, x_n(t))$.

Lemma 4 *For all $\alpha \in (-1, 1)$, we have*

$$\mathbb{E}\{\Phi(t+1) - \Phi(t) \mid x(t)\} \leq \left(\alpha + \alpha^2 \left(1 - \frac{2}{n} \right) \right) \sum_{i=1}^n p_i e^{\alpha x_i} - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right) \Phi(t). \quad (4)$$

Proof. Since the $x_i(t)$'s are ordered, they may change value at each time. We can thus define a permutation on $\{1, 2, \dots, n\}$ named σ_t such that, for every $u = 1, \dots, n$, if $x_i(t) = C_t^{(u)} - t/n$ then $x_{\sigma_t(i)}(t+1) = C_{t+1}^{(u)} - (t+1)/n$. Suppose

that the rank of the agent (say agent u), whose value is incremented at time t , is equal to i . In this case, we have

$$\begin{aligned} x_{\sigma_t(i)}(t+1) &= C_{t+1}^{(u)} - \frac{t+1}{n} = C_t^{(u)} + 1 - \frac{t+1}{n} \\ &= C_t^{(u)} - \frac{t}{n} + 1 + \frac{t}{n} - \frac{t+1}{n} = x_i(t) + 1 - \frac{1}{n}. \end{aligned}$$

This leads, for every $i = 1, \dots, n$, to $x_{\sigma_t(i)}(t+1) = x_i(t) + 1_{\{X=i\}} - \frac{1}{n}$, where 1_A is the indicator function of event A . We then get

$$\begin{aligned} \Phi(t+1) - \Phi(t) &= \sum_{i=1}^n \left(e^{\alpha x_i(t+1)} - e^{\alpha x_i(t)} \right) = \sum_{i=1}^n \left(e^{\alpha x_{\sigma_t(i)}(t+1)} - e^{\alpha x_i(t)} \right) \\ &= \sum_{i=1}^n \left(e^{\alpha(1_{\{X=i\}} - 1/n)} - 1 \right) e^{\alpha x_i(t)}. \end{aligned}$$

Using the fact that $e^x \leq 1 + x + x^2$ for $x \leq 1$, see Lemma 2, we obtain, since $\alpha(1_{\{X=i\}} - 1/n) \leq 1$,

$$\begin{aligned} e^{\alpha(1_{\{X=i\}} - 1/n)} - 1 &\leq \alpha(1_{\{X=i\}} - 1/n) + \alpha^2(1_{\{X=i\}} - 1/n)^2 \\ &= \alpha(1_{\{X=i\}} - 1/n) + \alpha^2 \left(1_{\{X=i\}} \left(1 - \frac{2}{n} \right) + \frac{1}{n^2} \right) \\ &= \left(\alpha + \alpha^2 \left(1 - \frac{2}{n} \right) \right) 1_{\{X=i\}} - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right). \end{aligned}$$

Taking the expectation of $\Phi(t+1) - \Phi(t)$, given $x(t)$, we obtain since $\mathbb{E}\{1_{\{X=i\}}\} = p_i$,

$$\begin{aligned} \mathbb{E}\{\Phi(t+1) - \Phi(t) \mid x(t)\} &\leq \sum_{i=1}^n \left[p_i \left(\alpha + \alpha^2 \left(1 - \frac{2}{n} \right) \right) - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right) \right] e^{\alpha x_i} \\ &= \left(\alpha + \alpha^2 \left(1 - \frac{2}{n} \right) \right) \sum_{i=1}^n p_i e^{\alpha x_i} - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right) \Phi(t), \end{aligned}$$

which completes the proof. $\blacksquare$

The following relations will be frequently used in the sequel. Since, for $i = 1, \dots, n$, $p_i = 2(i-1)/(n(n-1))$, we have for all $\lambda \in (0, 1)$ with $\lambda n \in \mathbb{N}$,

$$\frac{1}{n} \sum_{i=1}^n p_i = \frac{1}{n} \tag{5}$$

$$\frac{1}{\lambda n} \sum_{i=1}^{\lambda n} p_i = \frac{\lambda n - 1}{n(n-1)} \leq \frac{\lambda}{n} \tag{6}$$

$$\frac{1}{(1-\lambda)n} \sum_{i=\lambda n+1}^n p_i = \frac{(1+\lambda)n-1}{n(n-1)} \geq \frac{1+\lambda}{n} \tag{7}$$

Corollary 5 For all $\alpha \in (0, 1)$, we have

$$\mathbb{E}\{\Phi(t+1) - \Phi(t) \mid x(t)\} \leq \frac{\alpha^2}{n} \left(1 - \frac{1}{n}\right) \Phi(t).$$

Proof. To prove this result, observe that sequence $(e^{\alpha x_i})_i$ is a non-increasing sequence and $(p_i)_i$ is a non-decreasing sequence, so using Relation (5) and applying Lemma 3 we obtain

$$\sum_{i=1}^n p_i e^{\alpha x_i(t)} \leq \frac{1}{n} \left(\sum_{i=1}^n p_i \right) \left(\sum_{i=1}^n e^{\alpha x_i(t)} \right) = \frac{\Phi(t)}{n}.$$

Putting this result in inequality (4), we get

$$\begin{aligned} \mathbb{E}\{\Phi(t+1) - \Phi(t) \mid x(t)\} &\leq \left(\alpha + \alpha^2 \left(1 - \frac{2}{n}\right) \right) \sum_{i=1}^n p_i e^{\alpha x_i} - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right) \Phi(t) \\ &\leq \left[\frac{\alpha}{n} + \frac{\alpha^2}{n} \left(1 - \frac{2}{n}\right) - \left(\frac{\alpha}{n} - \frac{\alpha^2}{n^2} \right) \right] \Phi(t) \\ &= \frac{\alpha^2}{n} \left(1 - \frac{1}{n}\right) \Phi(t), \end{aligned}$$

which completes the proof. ■

Lemma 6 For all $\alpha \in (-1, 1)$, we have

$$\mathbb{E}\{\Psi(t+1) - \Psi(t) \mid x(t)\} \leq \left(-\alpha + \alpha^2 \left(1 - \frac{2}{n}\right) \right) \sum_{i=1}^n p_i e^{-\alpha x_i} + \left(\frac{\alpha}{n} + \frac{\alpha^2}{n^2} \right) \Psi(t). \quad (8)$$

Proof. It suffices to replace α by $-\alpha$ in the proof of Lemma 4. ■

Corollary 7 For all $\alpha \in (0, 1)$, we have

$$\mathbb{E}\{\Psi(t+1) - \Psi(t) \mid x(t)\} \leq \frac{\alpha^2}{n} \left(1 - \frac{1}{n}\right) \Psi(t)$$

Proof. See [11]. ■

The two previous lemmas, which give a bound of the increase of functions $\Phi(t)$ and $\Psi(t)$, will be used to prove Theorem 12. The proof of the results follow the clever ideas of the seminal paper [13] in which the authors prove that $\text{Gap}(t)$ is less than $O(\ln(n))$ with high probability. In [4], Alistarh et al. provide a more rigorous proof from which we have extracted constants associated with this asymptotic behavior. Those constants are given at the end of Section 4. The main original idea of our paper is to parametrize as much as possible the proofs in order to obtain the smallest values of constants a and b used in Relation (1)

which is proved in Theorem 14. The numerical evaluation of these constants, obtained in Section 4, shows that they are remarkably small with respect to the ones of [4].

In the following, we introduce two variable parameters $\mu, \rho \in (0, 1/2)$ (which are fixed to $1/4$ in [13] and [4]). Since x_i 's are non-increasing we have $x_{\rho n} \geq x_{(1-\mu)n}$. Lemmas 8 and 9 deal with the balanced conditions case that is $x_{\rho n} \geq 0 \geq x_{(1-\mu)n}$. The unbalanced conditions that are the complementary cases $x_{\rho n} \geq x_{(1-\mu)n} > 0$ and $0 > x_{\rho n} \geq x_{(1-\mu)n}$ are considered respectively in Lemmas 10 and 11. Theorem 12 examines systematically each case which lead to recurrence relation for $\mathbb{E}\{\Gamma(t)\}$. Theorem 13 uses this recurrence relation to bound $\mathbb{E}\{\Gamma(t)\}$. Finally, Theorem 14 gives a precise lower bound of $\Gamma(t)$ with high probability.

Lemma 8 *Let $\alpha, \mu \in (0, 1)$ with $\mu n \in \mathbb{N}$ and $\mu > \alpha/(1 + \alpha)$. If $x_{(1-\mu)n}(t) \leq 0$ then we have*

$$\begin{aligned} \mathbb{E}\{\Phi(t+1) \mid x(t)\} &\leq \left(1 - \frac{\alpha}{n} \left[\mu - \alpha(1 - \mu) + \frac{\alpha(1 - 2\mu)}{n} \right]\right) \Phi(t) + \alpha + \alpha^2 \left(1 - \frac{2}{n}\right) \\ &\leq \left(1 - \frac{\alpha}{n} [\mu - \alpha(1 - \mu)]\right) \Phi(t) + \alpha + \alpha^2. \end{aligned} \quad (9)$$

Proof. See [11]. ■

An analogous result is obtained for $\Psi(t)$ in the following lemma.

Lemma 9 *Let $\alpha, \rho \in (0, 1)$ with $\rho n \in \mathbb{N}$ and $\rho > \alpha/(1 - \alpha)$. If $x_{\rho n}(t) \geq 0$ then we have*

$$\begin{aligned} \mathbb{E}\{\Psi(t+1) \mid x(t)\} &\leq \left(1 - \frac{\alpha}{n} \left[\rho - \alpha(1 + \rho) + \frac{\alpha(1 + 2\rho)}{n} \right]\right) \Psi(t) + \alpha\rho(1 + \rho) \\ &\leq \left(1 - \frac{\alpha}{n} [\rho - \alpha(1 + \rho)]\right) \Psi(t) + \alpha\rho(1 + \rho). \end{aligned} \quad (10)$$

Proof. See [11]. ■

Lemma 10 *Let $\alpha, \mu \in (0, 1/2)$ with $\mu n \in \mathbb{N}$ and $\mu \in (\alpha/(1 + \alpha), (1 - 2\alpha)/(1 - \alpha))$, let $\mu' \in (0, 1)$ with $\mu' n \in \mathbb{N}$ and $\mu' \in (\mu/(1 - \mu), 1/(1 + \alpha))$ and let $\gamma_1 \in (0, 1)$.*

If $x_{(1-\mu)n} > 0$ and $\mathbb{E}\{\Phi(t+1) - \Phi(t) \mid x(t)\} \geq -(1 - \mu'(\alpha + 1)) \frac{\alpha\gamma_1}{n} \Phi(t)$ and $\Phi(t) \geq \lambda_1 \Psi(t)$ then we have $\Gamma(t) \leq c_1 n$, where

$$c_1 = \left(1 + \frac{1}{\lambda_1}\right) C_1 \left(\frac{C_1}{\mu\lambda_1}\right)^{\mu/((1-\mu)\mu' - \mu)}, \quad C_1 = \frac{(1 - \mu')(2 + \alpha)}{(1 - \gamma_1)(1 - \mu'(1 + \alpha))}$$

and

$$\lambda_1 = \frac{1 - \mu - \alpha(2 - \mu)}{2\alpha}.$$

The condition $\mu < (1 - 2\alpha)/(1 - \alpha)$ is needed to assure that constant $\lambda_1 > 0$. The value of λ_1 will be used in Theorem 12. The condition $\mu' > \mu/(1 - \mu)$ is needed to assure that the power involved in constant c_1 is positive.

Proof. See [11]. ■

Lemma 11 *Let $\alpha, \rho \in (0, 1/2)$ with $\rho n \in \mathbb{N}$ and $\rho \in (\alpha/(1-\alpha), 1/(1+\alpha))$, let $\rho' \in (\rho/(1-\rho), (1-2\alpha)/(1-\alpha))$ with $\rho' n \in \mathbb{N}$ and let $\gamma_2 \in (0, 1)$.*

If $x_{\rho n} < 0$ and $\mathbb{E}\{\Psi(t+1) - \Psi(t) \mid x(t)\} \geq -[1-2\alpha-\rho'(1-\alpha)] \frac{\alpha\gamma_2}{n} \Psi(t)$ and $\Psi(t) \geq \lambda_2 \Phi(t)$ then we have $\Gamma(t) \leq c_2 n$, where

$$c_2 = \left(1 + \frac{1}{\lambda_2}\right) C_2 \left(\frac{C_2}{\rho\lambda_2}\right)^{\rho/((1-\rho)\rho'-\rho)}, \quad C_2 = \frac{(1-\rho')(2-2\alpha-\rho'(1-\alpha))}{(1-\gamma_2)(1-2\alpha-\rho'(1-\alpha))}$$

and

$$\lambda_2 = \frac{1-\rho(1+\alpha)}{2\alpha}.$$

The condition $\rho < 1/(1+\alpha)$ is needed to assure that constant $\lambda_2 > 0$. The value of λ_2 will be used in Theorem 12. The condition $\rho' > \rho/(1-\rho)$ is needed to assure that the power involved in constant c_2 is positive.

Proof. See [11]. ■

Theorem 12. *Let $\alpha, \mu, \rho \in (0, 1/2)$ with $\mu n, \rho n \in \mathbb{N}$, $\mu \in (\alpha/(1+\alpha), (1-2\alpha)/(1-\alpha))$ and $\rho \in (\alpha/(1-\alpha), 1/(1+\alpha))$. Let $\mu' \in (\mu/(1-\mu), 1/(1+\alpha))$ with $\mu' n \in \mathbb{N}$ and let $\rho' \in (\rho/(1-\rho), (1-2\alpha)/(1-\alpha))$ with $\rho' n \in \mathbb{N}$. Let $\gamma_1, \gamma_2 \in (0, 1)$. We then have*

$$\mathbb{E}\{\Gamma(t+1) \mid x(t)\} \leq \left(1 - c_4 \frac{\alpha}{n}\right) \Gamma(t) + c_3,$$

where

$$c_4 = \min \left\{ \mu - \alpha(1-\mu), \rho - \alpha(1+\rho), \gamma_1(1-\mu'(\alpha+1)), \frac{\alpha(1-\mu-\alpha(2-\mu))}{1-\mu(1-\alpha)}, \right. \\ \left. \gamma_2(1-2\alpha-\rho'(1-\alpha)), \frac{\alpha(1-\rho(1+\alpha))}{1-\rho(1-\alpha)+2\alpha} \right\}$$

and

$$c_3 = \max \left\{ \alpha(1+\alpha+\rho(1+\rho)), \alpha(1-\mu)(2-\mu), (\alpha+c_4)\alpha c_1, \alpha+\alpha^2, \right. \\ \left. (\alpha+c_4)\alpha c_2 \right\},$$

in which

$$c_1 = \left(1 + \frac{1}{\lambda_1}\right) C_1 \left(\frac{C_1}{\mu\lambda_1}\right)^{\mu/((1-\mu)\mu'-\mu)}, \quad C_1 = \frac{(1-\mu')(2+\alpha)}{(1-\gamma_1)(1-\mu'(1+\alpha))}, \\ c_2 = \left(1 + \frac{1}{\lambda_2}\right) C_2 \left(\frac{C_2}{\rho\lambda_2}\right)^{\rho/((1-\rho)\rho'-\rho)}, \quad C_2 = \frac{(1-\rho')(2-2\alpha-\rho'(1-\alpha))}{(1-\gamma_2)(1-2\alpha-\rho'(1-\alpha))}, \\ \lambda_1 = \frac{1-\mu-\alpha(2-\mu)}{2\alpha}, \quad \lambda_2 = \frac{1-\rho(1+\alpha)}{2\alpha}.$$

Proof. See [11]. ■

We are now able to give an upper bound of the expected value of $\Gamma(t)$.

Theorem 13. *For all $t \geq 0$, under the hypothesis of Theorem 12, we have $\mathbb{E}\{\Gamma(t)\} \leq c_3 n / (\alpha c_4)$.*

Proof. We prove this result by induction. For $t = 0$, we have $\Gamma(0) = 2n$. Moreover, we have

$$c_3 \geq \alpha(1 + \alpha + \rho(1 + \rho)) \geq \alpha \text{ and } c_4 \leq \mu - \alpha(1 - \mu) \leq \mu \leq 1/2,$$

which implies that $c_3 / (\alpha c_4) \geq 2$. We thus have $\mathbb{E}\{\Gamma(0)\} = 2n \leq c_3 n / (\alpha c_4)$. Suppose that the result is true for a fixed $t \geq 0$. From Theorem 12, we have

$$\begin{aligned} \mathbb{E}\{\Gamma(t+1)\} &= \mathbb{E}\{\mathbb{E}\{\Gamma(t+1) \mid x(t)\}\} \leq \mathbb{E}\left\{\left(1 - c_4 \frac{\alpha}{n}\right) \Gamma(t) + c_3\right\} \\ &\leq \left(1 - c_4 \frac{\alpha}{n}\right) \frac{c_3}{\alpha c_4} n + c_3 = \frac{c_3}{\alpha c_4} n. \end{aligned}$$

which completes the proof. ■

Theorem 14. *For all $t \geq 0$ and $\sigma > 0$, under the hypothesis of Theorem 12, we have*

$$\mathbb{P}\left\{\text{Gap}(t) \geq \frac{2(1+\sigma)}{\alpha} \ln(n) + \frac{2}{\alpha} \ln\left(\frac{c_3}{2\alpha c_4}\right)\right\} \leq \frac{1}{n^\sigma}$$

Proof. From Lemma 1 and Theorem 13, we have

$$\Gamma(t) \geq 2e^{\alpha \text{Gap}(t)/2} \text{ and } \frac{c_3 n}{\alpha c_4} \geq \mathbb{E}\{\Gamma(t)\}.$$

It follows that

$$2e^{\alpha \text{Gap}(t)/2} \geq n^\sigma \frac{c_3 n}{\alpha c_4} \implies \Gamma(t) \geq n^\sigma \frac{c_3 n}{\alpha c_4} \implies \Gamma(t) \geq n^\sigma \mathbb{E}\{\Gamma(t)\}.$$

Using Markov inequality, we obtain

$$\begin{aligned} \mathbb{P}\left\{\text{Gap}(t) \geq \frac{2(\sigma+1)}{\alpha} \ln(n) + \frac{2}{\alpha} \ln\left(\frac{c_3}{2\alpha c_4}\right)\right\} &= \mathbb{P}\left\{2e^{\alpha \text{Gap}(t)/2} \geq n^\sigma \frac{c_3 n}{\alpha c_4}\right\} \\ &\leq \mathbb{P}\{\Gamma(t) \geq n^\sigma \mathbb{E}\{\Gamma(t)\}\} \leq \frac{1}{n^\sigma}, \end{aligned}$$

which completes the proof. ■

The following corollary shows that at any time, and for any agent, its local counter approximates the global clock with high probability.

Corollary 15 *For all $t \geq 0$ and $\sigma > 0$, under the hypothesis of Theorem 12, we have*

$$\mathbb{P}\left\{\left|C_t^{(i)} - \frac{t}{n}\right| < \frac{2(1+\sigma)}{\alpha} \ln(n) + \frac{2}{\alpha} \ln\left(\frac{c_3}{2\alpha c_4}\right), \forall i = 1, \dots, n\right\} \geq 1 - \frac{1}{n^\sigma}$$

Proof. By definition, we have $x_i = C_t^{(i)} - t/n$, and since $x_n \leq 0 \leq x_1$, we have $|x_i| \leq x_1 - x_n = \text{Gap}(t)$. It follows, from Theorem 14, that

$$\begin{aligned} \mathbb{P} \left\{ \left| C_t^{(i)} - \frac{t}{n} \right| \geq \frac{2(1+\sigma)}{\alpha} \ln(n) + \frac{2}{\alpha} \ln \left(\frac{c_3}{2\alpha c_4} \right), \forall i = 1, \dots, n \right\} \\ \leq \mathbb{P} \left\{ \text{Gap}(t) \geq \frac{2(1+\sigma)}{\alpha} \ln(n) + \frac{2}{\alpha} \ln \left(\frac{c_3}{2\alpha c_4} \right) \right\} \leq \frac{1}{n^\sigma} \end{aligned}$$

which completes the proof. ■

4 Evaluation of the constants

This section is devoted to the evaluation of constants a and b of Relation (1) and, to compare them with the ones that we can derive from the analysis of Alistarh et al. [4].

From Theorem 14, we have

$$a = \frac{2}{\alpha} \text{ and } b = \frac{2}{\alpha} \ln \left(\frac{c_3}{2\alpha c_4} \right),$$

where c_3 and c_4 are given by Theorem 12. First of all, note that constraints given in Theorem 12 imply the following inequality: $\rho/(1-\rho) < (1-2\alpha)/(1-\alpha)$, that is, $\rho \leq (1-2\alpha)/(2-3\alpha)$, which combined with $\rho \geq \alpha/(1-\alpha)$, leads to $\alpha \leq (5-\sqrt{5})/10 \approx 0.276$.

For a fixed value of α , we have to determine the values of parameters $\mu, \rho, \mu', \rho', \gamma_1, \gamma_2$ that minimize constant b . This is achieved by applying a simple Monte-Carlo algorithm. Figure 1 shows several optimal values of the constants a and b , used in Theorem 14, and computed for several values of α .

α	0.17	0.18	0.19	0.20	0.21	0.22	0.23	0.24	0.25	0.26	0.27
$a = 2/\alpha$	11.77	11.12	10.53	10	9.53	9.10	8.70	8.34	8	7.70	7.41
$b = (2/\alpha) \log(c_3/(2\alpha c_4))$	59	63	68	74	82	93	109	134	179	281	739

Fig. 1. Optimal values of a and b in function of α

Let us now evaluate constants a and b obtained in the paper of Alistarh et al. [4]. Note that the goal of their work was not necessarily focused on the optimization of a and b constants. Nevertheless, as we will see, the evaluation of a and b constants is an important motivation of our work. From Relations (1) and (2) of [4] and as $\beta = 1$, we get $0 < \delta \leq \varepsilon = 1/16$ and thus we obtain, for $\gamma > 0$ and $c \geq 2$,

$$\frac{1 + \gamma + c\alpha(1 + \gamma)^2}{1 - \gamma - c\alpha(1 + \gamma)^2} \leq \frac{17}{16},$$

which gives,

$$\alpha \leq \frac{1}{33c(1+\gamma)^2} - \frac{1}{c(1+\gamma)^2} \leq \frac{1}{33c(1+\gamma)^2} \leq \frac{1}{66}.$$

Considering the difference between the lower bound and the upper bound of the inequality following (11), we obtain

$$\exp\left(\frac{\alpha B}{n} \left(3 - \frac{1}{1-\lambda}\right)\right) \leq \frac{16\lambda C(\varepsilon)}{\varepsilon},$$

which can also be written as

$$\exp\left(\frac{\alpha B}{(1-\lambda)n}\right) \leq \left(\frac{16\lambda C(\varepsilon)}{\varepsilon}\right)^{1/(2-3\lambda)}.$$

Using the last inequality obtained in the proof of Lemma 4.8, we get

$$\Gamma(t) \leq \frac{4+\varepsilon}{\varepsilon} \lambda n C(\varepsilon) \exp\left(\frac{\alpha B}{(1-\lambda)n}\right) \leq \frac{4+\varepsilon}{\varepsilon} \lambda n C(\varepsilon) \left(\frac{16\lambda C(\varepsilon)}{\varepsilon}\right)^{1/(2-3\lambda)}.$$

Using this result, we obtain from Lemma 4.11, $\mathbb{E}\{\Gamma(t)\} \leq 4Cn/(\hat{\alpha}\varepsilon)$, where

$$C = \frac{4+\varepsilon}{\varepsilon} \lambda C(\varepsilon) \left(\frac{16\lambda C(\varepsilon)}{\varepsilon}\right)^{1/(2-3\lambda)}, \quad C(\varepsilon) = \frac{(1+\delta)/\lambda - 1 + 3\varepsilon}{3\varepsilon - \varepsilon/3}$$

$$\text{and } \hat{\alpha} = \alpha(1 - \gamma - c\alpha(1 + \gamma)^2).$$

Following the same ideas we used to prove Theorem 14, we get

$$a = \frac{2}{\alpha} \text{ and } b = \frac{2}{\alpha} \ln\left(\frac{2C}{\hat{\alpha}\varepsilon}\right).$$

Since $\alpha \leq 1/66$, we have $a \geq 132$. Moreover, since $0 \leq \delta \leq \varepsilon = 1/16$, $\lambda = 2/3 - 1/54 = 35/54$, $\gamma > 0$ and $c \geq 2$, we obtain

$$C(\varepsilon) = \frac{(1+\delta)/\lambda - 1 + 3\varepsilon}{3\varepsilon - \varepsilon/3} \geq \frac{1/\lambda - 1 + 3\varepsilon}{3\varepsilon - \varepsilon/3} = \frac{1227}{280}$$

which leads to

$$C = \frac{4+\varepsilon}{\varepsilon} \lambda C(\varepsilon) \left(\frac{16\lambda C(\varepsilon)}{\varepsilon}\right)^{1/(2-3\lambda)} \geq \frac{26585}{144} \left(\frac{6544}{9}\right)^{18}.$$

Regarding $\hat{\alpha}$, we have $\hat{\alpha} = \alpha(1 - \gamma - c\alpha(1 + \gamma)^2) \leq \alpha \leq 1/66$. Therefore, we have

$$b = \frac{2}{\alpha} \ln\left(\frac{2C}{\hat{\alpha}\varepsilon}\right) \geq 132 \ln\left(\frac{1169740}{3} \left(\frac{6544}{9}\right)^{18}\right) \geq 17354.$$

It follows that constants a and b obtained from [4] satisfy $a \geq 132$ and $b \geq 17354$, which are at least two orders of magnitude larger than the ones we derived (see Figure 1).

Fig. 2. Expected proportion and gap

5 Simulations

We complete this paper by giving a summary of the experiments we have carried out to illustrate the performances of our protocol. Recall that n is the number of nodes in the system, and $T = t/n$ is the total number of interactions divided by n , which is often called the parallel time. We have conducted two types of experiments, the first one illustrates the expected proportion of nodes $Y_T(n, k)$ whose counter is equal to $T + k$ at time nT , for different values of n and k . More precisely, $Y_T(n, k)$ is defined by

$$Y_T(n, k) = \frac{1}{n} \sum_{i=1}^n 1_{\{C_{nT}^{(i)} = T+k\}}.$$

We show in Figure 2(a) the expected value of $Y_T(n, k)$, for $n = 1000$ and $k = -2, -1, 0, 1$, as a function of the parallel time T . These results have been obtained after running 10,000 independent experiments. Figure 2(a) shows that the expected value of $Y_T(n, k)$ seems to converge when T goes to infinity, and this convergence is reached very quickly. Note that for other values of k , proportions of nodes are too close to 0 to be depicted, as shown in Table 1. Table 1 shows the expected proportion of nodes $Y_T(n, k)$ whose counter is equal to $T + k$ at time $T = 50$, for different values of $n = 10^3, 10^4, 10^5, 10^6, 10^7$ and $k = -13, \dots, 4$. These results have been obtained after running 5,000 independent experiments, for each value of n . The expected value of $Y_{50}(n, k)$ seems to be almost independent of n for large values of n .

The second experiment illustrates the gaps (i.e., the maximal, average, and minimal) for different values of the size n of the system. Let $B = 2 \times 10^9$ be the total number of interactions considered. The maximal gap is computed as $\max_{100n \leq t \leq B} \text{Gap}(t)$, the minimal one is given by $\min_{100n \leq t \leq B} \text{Gap}(t)$, and the

$k \backslash n$	10^3	10^4	10^5	10^6	10^7
-13	0.0	0.0	0.0	1.4E-9	1.42E-9
-12	0.0	2.0E-8	8.0E-9	9.0E-9	6.14E-9
-11	2.0E-7	4.0E-8	2.2E-8	2.8E-8	3.048E-8
-10	2.0E-7	8.0E-8	1.88E-7	1.436E-7	1.4814E-7
-9	4.0E-7	8.0E-7	7.7E-7	7.438E-7	7.2784E-7
-8	3.0E-6	3.6E-6	3.586E-6	3.48E-6	3.6029E-6
-7	1.42E-5	1.8E-5	1.8222E-5	1.7767E-5	1.7758E-5
-6	8.98E-5	8.602E-5	8.7176E-5	8.7372E-5	8.72753E-5
-5	4.372E-4	4.2706E-4	4.2957E-4	4.2901E-4	4.29349E-4
-4	0.0021144	0.0021023	0.0021071	0.0021092	0.0021086
-3	0.0102474	0.0102890	0.0102777	0.0102800	0.0102810
-2	0.0481626	0.0483366	0.0483382	0.0483465	0.0483437
-1	0.1930704	0.1932864	0.1933165	0.1933143	0.1933182
0	0.4389352	0.4380932	0.4380715	0.4380374	0.4380346
1	0.2824746	0.2827344	0.2826797	0.2827057	0.2827070
2	0.0243744	0.0245499	0.0245973	0.0245953	0.0245949
3	7.6E-5	7.224E-5	7.2248E-5	7.27752E-5	7.27974E-5
4	0.0	0.0	0.0	4.0E-10	3.6E-10

Table 1. Expectation of $Y_{50}(n, k)$ from number of nodes n and shift k

average gap is given by

$$\frac{1}{B - 100n} \sum_{t=100n}^{B-1} \text{Gap}(t).$$

Figure 2(b) shows respectively the minimal, average and maximal gap in a system of size n over the interval $[100n, B]$ of interactions. As one may expect, the logarithmic progression of the Gap is clearly shown.

6 Conclusion

In this article we have gone a step further in the study of the two-choice paradigm by providing an accurate analysis of the gap problem. An important application of this study would be the improvement of leaderless population protocols. Indeed, we have shown in this paper that agents can construct a global clock by guaranteeing that the values of all agent counters are concentrated according to Relation (1), and thus can locally use this global clock to determine the instants at which some specific actions need to be triggered, or the instants from which all the agents of the system have converged to a given state. In the former case, this would allow agents to solve more complex problems by triggering a series of population protocols, whereas in the latter case this would allow agents to determine the instant from which all the agents have successfully computed a given feature of the population. The construction of efficient leaderless population protocols inspired from this orchestration is left for future work.

References

1. M. Adler, P. Berenbrink, and K. Schröder. Analyzing an infinite parallel job allocation process. In *Proceedings of the European Symposium on Algorithms (ESA)*, 1998.
2. M. Adler, S. Chakrabarti, M. Mitzenmacher, and L. Rasmussen. Parallel randomized load balancing. *Random Structures & Algorithms*, 13(2):159–188, 1998.
3. D. Alistarh, J. Aspnes, and R. Gelashvili. Space-optimal majority in population protocols. In A. Czumaj, editor, *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2221–2239, 2018.
4. D. Alistarh, J. Kopinsky, J. Li, and G. Nadiradze. The power of choice in priority scheduling. In *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2017.
5. Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations (extended abstract). In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 1994.
6. P. Berenbrink, A. Czumaj, T. Friedetzky, and N. D. Vvedenskaya. Infinite parallel job allocation (extended abstract). In *Proceedings of the ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 99–108, 2000.
7. P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. *SIAM Journal on Computing*, 35(6):1350–1385, 2006.
8. P. Berenbrink, F. Meyer auf der Heide, and K. Schröder. Allocating weighted jobs in parallel. *Theory of Computing Systems*, 32(3):281–300, 1999.
9. M. Mitzenmacher. Load balancing and density dependent jump Markov processes. In *Proceedings of International Conference on Foundations of Computer Science*, 1996.
10. M. Mitzenmacher, A. W. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In *Handbook of Randomized Computing*, pages 255–312. Kluwer, 2000.
11. Y. Mocquard, B. Sericola, and E. Anceaume. Balanced allocations and global clock in population protocols: An accurate analysis (Full version), Technical Report. <https://hal.archives-ouvertes.fr/hal-01790973>, 2018.
12. Y. Peres, K. Talwar, and U. Wieder. The $(1+\beta)$ -choice process and weighted balls into bins. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.
13. Y. Peres, K. Talwar, and U. Wieder. Graphical balanced allocations and the $(1 + \beta)$ -choice process. *Random Structure of Algorithms*, 47(4):760–775, 2015.
14. M. Raab and A. Steger. “Balls into bins” — a simple and tight analysis. In *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pages 159–170, 1998.
15. K. Talwar and U. Wieder. Balanced allocations: A simple proof for the heavily loaded case. In *Automata, Languages, and Programming*, pages 979–990. Springer Berlin Heidelberg, 2014.