

Resource allocation in open multi-agent systems: an online optimization analysis

Renato Vizuete, Charles Monnoyer de Galland, Julien M. Hendrickx, Paolo Frasca and Elena Panteley

Abstract—The resource allocation problem consists of the optimal distribution of a budget between agents in a group. We consider such a problem in the context of open systems, where agents can be replaced at some time instances. These replacements lead to variations in both the budget and the total cost function that hinder the overall network’s performance. For a simple setting, we analyze the performance of the Random Coordinate Descent algorithm (RCD) using tools similar to those commonly used in online optimization. In particular, we study the accumulated errors that compare solutions issued from the RCD algorithm and the optimal solution or the non-collaborating selfish strategy and we derive some bounds in expectation for these accumulated errors.

I. INTRODUCTION

We consider the optimal resource allocation problem, where a fixed amount of resource must be distributed among n agents while minimizing some separable cost function f [1]. Problems of this type can be found in many different fields of research including distributed computer systems [2], games [3], smart grids [4], etc. In some specific formulations like actuator networks [5] or power systems [6], each agent i holds a quantity d_i (which we call here the “demand” of agent i), so that the total amount of resource to be distributed is $\sum_{i=1}^n d_i$; the problem can then be written as

$$\min_{x \in \mathbb{R}^{n_p}} f(x) = \sum_{i=1}^n f_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^n x_i = \sum_{i=1}^n d_i, \quad (1)$$

where each function $f_i : \mathbb{R}^p \rightarrow \mathbb{R}$ is α -strongly convex and β -smooth, and represents the local cost held by agent i .

Problems of this type have received a lot of attention in the last years and most of them are related to a possible change in the budget due to a variation in the demand of some of the agents [7, 8]. However, in these works, only quadratic functions are considered which significantly restrict the set of potential cost functions of the agents and do not correspond

to the standard assumptions in the field of convex and smooth optimization [9, 10].

In addition to the possible variations of the budget over time in (1), the composition of the system may also change during the whole process due to the arrival, departure or replacement of agents at a time-scale comparable to that of the process, giving rise to *open multi-agent systems*. Those are motivated by the growing size of the systems that tends to slow down the process as compared to the time-scale of potential changes in the set of agents. More generally, systems naturally allowing agents to join and leave are becoming common, such as *e.g.* multi-vehicle systems or with the Plug and Play implementation [11, 12]. In the case of (1), it results in the system size n_t , the local cost functions f_i^t , and the local demands d_i^t becoming time-varying. As a consequence, the instantaneous optimum of (1), denoted $x^{*,t}$, changes with the time as well, preventing usual convergence.

Due to these possible changes in the dimension of the system, most of the related works in the field of open multi-agent systems are focused on the analysis of scalar performance indexes associated with the process, which allow overcoming the problem of time-varying dimensions. For instance, in [13, 14] the variance is proposed as a metric for the analysis of a pairwise gossip algorithm, while in [15] the mean squared error is the object of study in randomized interactions. Regarding optimization, problems such as (1) typically imply a minimization process on a long period, and hence the cost is expected to be paid on a regular basis. In such setting, a natural way of measuring the performance of an algorithm is to compute its accumulated error with respect to a given strategy over a finite number of iterations. Similar metrics occur in the context of online optimization [16], where the objective is to minimize the so-called *regret*, commonly defined as the accumulated error of the estimate x^t with respect to $x^* := \arg \min_x \sum_{t=1}^T f^t(x)$, or sometimes with respect to the time-varying solution of (1) $x^{*,t} := \arg \min_x f^t(x)$ such as *e.g.*, in [17]. Other extensions of the regret include the case of time-varying constraints, where a similar metric is used to measure the violation [18].

In this work, we analyze the performance of the Random Coordinate Descent algorithm (RCD) [19, 20] to solve (1) in open systems. We study the loss accumulated by the RCD algorithm with respect to the time-varying optimal solution $x^{*,t}$ over a finite number of iterations, and its gain with respect to the *selfish strategy* $x^{s,t}$, which consists in the absence of collaboration between the agents (*i.e.*, $x_i^{s,t} = d_i^t$), by obtaining upper bounds. Finally, we consider the case of quadratic cost functions, for which tighter results are derived.

Research supported in part by the Agence Nationale de la Recherche (ANR) via grant “Hybrid And Networked Dynamical sYstems” (HANDY), number ANR-18-CE40-0010 and by the “RevealFlight” ARC at UCLouvain, by the *Incentive Grant for Scientific Research (MIS)* “Learning from Pairwise Data” of the F.R.S.-FNRS.

R. Vizuete and C. Monnoyer de Galland equally contributed to this work. R. Vizuete and E. Panteley are with Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, 91190, Gif-sur-Yvette, France. R. Vizuete and P. Frasca are with Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, GIPSA-lab, F-38000 Grenoble, France. C. Monnoyer de Galland and J. M. Hendrickx are with the ICTEAM institute, UCLouvain, Louvain-la-Neuve, Belgium. C. Monnoyer de Galland is a FRIA fellow (F.R.S.-FNRS). (E-mail addresses: renato.vizuete@12s.centralesupelec.fr; charles.monnoyer@uclouvain.be; julien.hendrickx@uclouvain.be; paolo.frasca@gipsa-lab.fr; elena.panteley@12s.centralesupelec.fr).

II. PROBLEM FORMULATION

The set of real numbers is denoted by $\mathbb{R}$ and the set of nonnegative integers by $\mathbb{Z}_{\geq 0}$. For two vectors $x, y \in \mathbb{R}^n$, $\langle x, y \rangle = x^\top y = \sum_{i=1}^n x_i y_i$ denotes the usual Euclidean inner product and $\|x\| = \sqrt{x^\top x}$ the Euclidean norm. The set of n -dimensional vectors with nonnegative entries is denoted by $\mathbb{R}_{\geq 0}^n$. We denote the vector of size n constituted of only zeros by $\mathbf{0}_n$ and of only ones by $\mathbf{1}_n$.

A. Open resource allocation problem

We consider the problem (1), where we restrict to non-negative states $x_i \in \mathbb{R}_{\geq 0}^p$ for the agents, and where the local cost functions satisfy the following assumption.

Assumption 1 (Local cost function): The local cost function $f_i : \mathbb{R}_{\geq 0}^p \rightarrow \mathbb{R}_{\geq 0}$ of any agent i is

- continuously differentiable;
- α -strongly convex: $f_i(x) - \frac{\alpha}{2}\|x\|^2$ is convex $\forall x$;
- β -smooth: $\|\nabla f_i(x) - \nabla f_i(y)\| \leq \beta\|x - y\|, \forall x, y$;
- satisfies $\arg \min_{x \in \mathbb{R}_{\geq 0}^p} f_i(x) = \mathbf{0}_p$ and $f_i(\mathbf{0}_p) = 0$.

More generally, we use $\mathcal{F}_{\alpha, \beta}^p$ to denote the set of functions $f : \mathbb{R}_{\geq 0}^p \rightarrow \mathbb{R}_{\geq 0}$ satisfying these conditions.

Assumption 1 means that the cost paid by an agent is always nonnegative, and is zero only when the agent does not contribute at all to any activity, i.e. $x_i = \mathbf{0}_p$. It also provides an upper and lower bound on their curvature, which is rather standard in optimization [9]. It follows from Assumption 1 that the global cost satisfies $f(x) = \sum_{i=1}^n f_i(x_i) \in \mathcal{F}_{\alpha, \beta}^{np}$.

To problem (1) we associate an undirected graph $G = (\mathcal{V}, \mathcal{E})$, so that at random times a pair of agents $(i, j) \in \mathcal{E}$ is uniformly randomly chosen to interact and exchange information. Moreover, we assume the system is subject to random instantaneous arrivals and departures of agents in the system, respectively resulting in a new agent joining the system with its own local cost function and demand, or in an agent leaving the system and never coming back, with the possibility of sending a last message to their neighbours. We also consider replacements, which consist in the simultaneous occurrence of both an arrival and a departure. Hence, the system size n_t , the local cost functions f_i^t and the demands d_i^t evolve with time, and consequently the instantaneous solution of (1) is time-varying as well, and is denoted $x^{*,t}$.

B. Simplifying assumptions and reformulation

For this preliminary work, we restrict to the specific case defined by the following assumptions.

Assumption 2 (1-D functions): The local cost function of any agent at any time is one-dimensional: $f_i^t : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$.

Assumption 3 (Homogeneous demand): The demand associated with any agent i at any time t is $d_i^t = 1$.

Assumption 4: The graph $G = (\mathcal{V}, \mathcal{E})$ is complete so that any pair of agents can interact in the network, i.e., $\mathcal{E} = \mathcal{V} \times \mathcal{V}$.

Moreover, we restrict to the case where the openness of the system is solely characterized by replacements of agents (i.e., the simultaneous occurrence of an arrival and a departure), so that the system size is fixed, and $n_t = n$ for all time t . Hence,




Fig. 1. Evolution of the function value f^t evaluated with the RCD algorithm x^t defined in (8), the optimal solution $x^{*,t}$, and the selfish strategy $x^{s,t}$, in a system subject to replacements of agents (i.e., simultaneous departures and arrivals) on average once every 4 RCD steps.

the system only evolves at the instantaneous occurrences of either pairwise interactions of agents, resulting in a possible update of their states (denoted U) or replacements (denoted R). We call these occurrences “events”, and we can define the set of all the events that can possibly take place in the system, which we call “event set”, as follows:

$$\Xi := R \cup U = \left(\bigcup_{i \in \mathcal{V}} R_i \right) \cup \left(\bigcup_{(i,j) \in \mathcal{E}} U_{ij} \right), \quad (2)$$

where R_i denotes the replacement of agent i and U_{ij} a pairwise interaction between agents i and j . We assume that two distinct events never occur simultaneously, so that the system evolves in a discrete manner, where each time-step $k \in \mathbb{Z}_{\geq 0}$ corresponds to the occurrence of an event $\xi_k \in \Xi$.

Assumption 5: An event ξ_k is independent of all other events ξ_j , $j \neq k$ and of the state of the system x^t until time k , so that at each time-step either an update (i.e., an event U) happens with fixed probability p or a replacement (i.e., an event R) with fixed probability $1 - p$.

Let $S_n := \{x \in \mathbb{R}_{\geq 0}^n : \mathbf{1}^\top x = n\}$ be the feasible set, we can now express (1) in our setting under the assumptions of this section:

$$\min_{x \in S_n} f^t(x) = \sum_{i=1}^n f_i^t(x_i). \quad (3)$$

C. Performance metrics

Natural indexes for measuring the performance of an algorithm in our setting consist in evaluating its accumulated error over a finite number of iterations with respect to a given strategy. We define the two following strategies of interest in the context of the resource allocation problem:

- **Perfect collaboration:** at each time instant t the agents know the optimal solution of (3) denoted $x^{*,t}$;
- **Selfish players:** the agents do not collaborate to minimize f^t , and they operate at their individual desired point so that $x^{s,t} = \mathbf{1}_n$ at all t .

Hence, for any T , the estimate x^t obtained with a well-designed algorithm is expected to satisfy

$$\sum_{t=1}^T f^t(x^{*,t}) \leq \sum_{t=1}^T f^t(x^t) \leq \sum_{t=1}^T f^t(x^{s,t}). \quad (4)$$

The evolution of these strategies, compared with that of a given algorithm, is illustrated in Fig. 1.

We define the following performance metrics to analyze the value provided by the RCD algorithm x^t with respect to the strategies above:

$$\textbf{Dynamical Regret: } \text{Reg}_T := \sum_{t=1}^T (f^t(x^t) - f^t(x^{*,t})); \quad (5)$$

$$\textbf{Benefit: } \text{Ben}_T := \sum_{t=1}^T (f^t(x^{s,t}) - f^t(x^t)); \quad (6)$$

$$\textbf{Potential Benefit: } \text{Pot}_T := \sum_{t=1}^T (f^t(x^{s,t}) - f^t(x^{*,t})). \quad (7)$$

The “dynamical regret” and “benefit” respectively measure the accumulated error from using a given algorithm with respect to the optimal solution $x^{*,t}$ and the accumulated gain from using it instead of the selfish strategy $x^{s,t}$. The “potential benefit” is independent of the algorithm; it represents the accumulated advantage of the optimal strategy with respect to the selfish one, and satisfies $\text{Pot}_T = \text{Ben}_T + \text{Reg}_T$.

Observe that the *regret* commonly used in online optimization typically compares x^t with the overall optimal solution taken over all the iterations, *i.e.*, $x^* = \arg \min_{x \in S_n} \sum_{t=1}^T f^t(x)$. In that sense, it differs from the *dynamical regret* in (5), which compares x^t with the time-varying instantaneous optimal solution $x^{*,t} = \arg \min_{x \in S_n} f^t(x)$ at each iteration, such as *e.g.*, in [17].

D. Random Coordinate Descent algorithm and objective

We consider the Random Coordinate Descent algorithm (RCD) introduced in [19], such as whenever a pair of agents $(i, j) \in \mathcal{E}$ interact, they update their respective estimates as

$$\begin{aligned} x_i^+ &= x_i - \frac{1}{\beta} (f'_i(x_i) - f'_j(x_j)) \\ x_j^+ &= x_j - \frac{1}{\beta} (f'_j(x_j) - f'_i(x_i)). \end{aligned} \quad (8)$$

Observe that agents do not need to know α nor the properties of the network to implement (8). By contrast, β is needed for the derivation of (8), see *e.g.* [20] for details.

We moreover assume that whenever an agent *in* joins the system, it initializes its estimate as

$$x_{in} = d_{in} = 1, \quad (9)$$

and whenever an agent *out* leaves the system, it sends a last message to all its neighbours (*i.e.*, all the other agents in our setting) with its current estimate x_{out} and its demand d_{out} so the agents $i \neq out$ update their estimates as

$$x_i^+ = x_i + \frac{x_{out} - x_i}{n} = \left(1 - \frac{1}{n}\right) x_i + \frac{1}{n} x_{out}. \quad (10)$$

We show in the following proposition that RCD iterations, arrivals and departures as they are defined in (8) to (10) guarantee that as long as the initial estimate x^0 is feasible, then all the estimates remain feasible.

Proposition 1 (Well-posedness): The event set (2) guarantees that if $x^0 \in S_n$, then $x^t \in S_n$ for all t .

Proof: We first consider arrivals: the nonnegativity of x_i and preservation of the constraint is a direct consequence of (9). In the case of departures, the nonnegativity of x_i is a direct consequence of (10). Moreover, if the constraint is

satisfied at iteration k with $n_k = n$, then under the departure of the agent labelled *out* we have

$$\sum_{i \neq out} x_i^{k+1} = n - x_{out} + \frac{x_{out}}{n} (n - 1) - \frac{n - x_{out}}{n} = n - 1.$$

We finally consider iterations of the RCD algorithm. From Assumptions 1 and 2, it follows that for any $x_i \geq 0$

$$x_i f'_i(x_i) \geq \alpha |x_i|^2 \geq 0,$$

so that $f'_i(x_i) \geq 0$. Moreover, since f_i is β -smooth, one has $f'_i(x_i) \leq \beta x_i$, and therefore at each update of the RCD algorithm between agents i and j there holds

$$x_i^+ = x_i - \frac{1}{2\beta} (f'_i(x_i) - f'_j(x_j)) \geq x_i - \frac{1}{2\beta} (\beta x_i) = \frac{x_i}{2},$$

establishing the nonnegativity of x_i . A similar analysis can be used for x_j . Due to the symmetry of the update rule, the constraint is always preserved and we conclude the proof. ■

Our goal is to analyze the performance of the RCD algorithm (8) with the arrival and departure rules (9) and (10) in the setting described in Sections II-A and II-B using the metrics defined in Section II-C in expectation.

III. UPPER BOUNDS ON THE PERFORMANCE METRICS

We now derive upper bounds on the evolution of the Potential Benefit and the Dynamical Regret respectively defined in (7) and (5) in expectation. Whereas the former is only related to the problem itself, the latter actually depends on the algorithm we consider.

We first provide the following lemmas, where Lemma 1 directly follows from the equivalence of the norms.

Lemma 1: Let $x \in S_n$, then $n \leq \|x\|^2 \leq n^2$.

Lemma 2: Let $f(x) = \sum_{i=1}^n f_i(x_i)$, where all f_i satisfy Assumptions 1 and 2, then for any $x \in S_n$ there holds

$$\frac{\alpha}{2} n \leq f(x) \leq \frac{\beta}{2} n^2. \quad (11)$$

Proof: From Assumptions 1 and 2, $f(0) = f'(0) = 0$. Hence, since f is β -smooth and using Lemma 1, there holds $f(x) \leq \frac{\beta}{2} \|x\|^2 \leq \frac{\beta}{2} n^2$ which establishes the upper bound. Similarly, since f is α -strongly convex and by using Lemma 1, it follows that $f(x) \geq \frac{\alpha}{2} \|x\|^2 \geq \frac{\alpha}{2} n$, which establishes the lower bound, and concludes the proof. ■ Lemma 2 provides a global upper bound on the difference between any two solutions $x, y \in S_n$:

$$|f^t(x) - f^t(y)| \leq \frac{n}{2} (n\beta - \alpha). \quad (12)$$

This can be used to derive upper bounds on any of the metrics defined in Section II-C, *e.g.*, $\text{Ben}_T \leq \frac{n}{2} (n\beta - \alpha) T$.

A. Potential Benefit

We first obtain in the following theorem an upper bound on the expected value of the potential benefit, which we remind quantifies the accumulated advantage of using the optimal strategy rather than not collaborating at all.

Theorem 1: In the setting of Section II, there holds

$$\text{Pot}_T \leq \frac{n}{2} \alpha (\kappa - 1) T, \quad (13)$$

and in particular

$$\lim_{T \rightarrow \infty} \frac{\text{Pot}_T}{T} \leq \frac{n}{2} \alpha (\kappa - 1). \quad (14)$$

Proof: Remember that $x^{s,t} = \mathbb{1}_n$ by definition, and that $f^t(\mathbf{0}_n) = 0$ and $\nabla f^t(\mathbf{0}_n) = \mathbf{0}_n$ from Assumption 1. Hence, there holds from the β -smoothness of f^t

$$f^t(x^{s,t}) \leq \nabla f^t(\mathbf{0}_n)^\top (x^{s,t}) + \frac{\beta}{2} \|x^{s,t}\|^2 = \frac{\beta}{2} \|\mathbb{1}_n\|^2 = \frac{\beta}{2} n.$$

Similarly, since f^t is α -strongly convex, we get

$$f^t(x^{*,t}) \geq \nabla f^t(\mathbf{0}_n)^\top (x^{*,t}) + \frac{\alpha}{2} \|x^{*,t}\|^2 = \frac{\alpha}{2} \|x^{*,t}\|^2 \geq \frac{\alpha}{2} n,$$

where the last inequality follows from Lemma 1. Hence

$$f^t(x^{s,t}) - f^t(x^{*,t}) \leq \frac{\beta}{2} n - \frac{\alpha}{2} n = \frac{n}{2} (\beta - \alpha)$$

holds, and injecting it into (7) yields (13). The last result then follows from dividing (13) by T . ■

Notice that since the dynamical regret is nonnegative by definition, the bounds (13) and (14) also hold for the benefit since $\text{Ben}_T = \text{Pot}_T - \text{Reg}_T \leq \text{Pot}_T$.

B. Dynamical Regret

We now obtain an upper bound on the expected dynamical regret defined in (5), where we remind x^t is obtained with the RCD algorithm defined in Section II-D. For that purpose, we first introduce the following intermediate quantities:

$$C_t := f^t(x^t) - f^t(x^{*,t}); \quad (15)$$

$$\Delta f_t := f^{t+1}(x^{t+1}) - f^t(x^t); \quad (16)$$

$$\Delta f_t^* := f^{t+1}(x^{*,t+1}) - f^t(x^{*,t}). \quad (17)$$

Thus, C_t corresponds to the instantaneous loss of the RCD algorithm with respect to the optimal solution at iteration t , and Δf_t and Δf_t^* respectively stand for the instantaneous variation at one iteration of the total estimated cost and optimal cost, such that $C_{t+1} = C_t + \Delta f_t - \Delta f_t^*$.

In the following proposition, we study the effect or replacements on Δf_t in order to later characterize C_t in expectation, and consequently the expected dynamical regret.

Proposition 2: In the setting of Section II the replacement of an agent, denoted R , results in

$$\mathbb{E}[\Delta f_t \mid R] \leq \frac{5}{2} \beta - \frac{3}{2} \alpha. \quad (18)$$

Proof: We analyze the effects of arrivals and departures separately. Let g denote the local cost function of the joining agent at an arrival, then $f^{t+1}(x^{t+1}) = f^t(x^t) + g(1)$ and

$$\Delta f_t = f^t(x^t) + g(1) - f^t(x^t) = g(1) \leq \frac{\beta}{2}, \quad (19)$$

where the last inequality follows from Assumption 1, and in particular the β -smoothness of g .

Consider now a departure, and let ℓ denote the label of the leaving agent, such that $f^{t+1}(x^{t+1}) = \sum_{i \neq \ell} f_i^t(x_i^{t+1})$, with $x_i^{t+1} = x_i^t + \frac{x_i^t + x_\ell^t}{n}$ from (10). From the definition of departures, ℓ is uniformly selected among the n agents in the system and by taking the expected value Δf_t over the

leaving agent, one gets the following, where we omit the reference to time to lighten the notation:

$$\begin{aligned} \mathbb{E}[\Delta f_t] &= \sum_{\ell=1}^n \frac{1}{n} \left(\sum_{i \neq \ell} f_i \left(x_i + \frac{x_\ell - x_i}{n} \right) - f(x) \right) \\ &= \frac{1}{n} \sum_{\ell=1}^n \left(f \left(\frac{n-1}{n} x + \frac{x_\ell}{n} \mathbb{1}_n \right) - f_\ell(x_\ell) \right) - f(x) \\ &= \frac{1}{n} \sum_{\ell=1}^n f \left(\frac{n-1}{n} x + \frac{x_\ell}{n} \mathbb{1}_n \right) - \frac{n+1}{n} f(x). \end{aligned}$$

Since f is β -smooth from Assumption 1, one has

$$\begin{aligned} f \left(\frac{n-1}{n} x + \frac{x_\ell}{n} \mathbb{1}_n \right) &\leq \\ f(x) + \frac{1}{n} \langle \nabla f(x), x_\ell \mathbb{1}_n - x \rangle &+ \frac{\beta}{2n^2} \|x_\ell \mathbb{1}_n - x\|^2, \end{aligned}$$

and it follows that

$$\begin{aligned} \mathbb{E}[\Delta f_t] &\leq \frac{1}{n^2} \sum_{\ell=1}^n \langle \nabla f(x), x_\ell \mathbb{1}_n - x \rangle \\ &+ \frac{\beta}{2n^3} \sum_{\ell=1}^n \|x_\ell \mathbb{1}_n - x\|^2 - \frac{1}{n} f(x). \quad (20) \end{aligned}$$

From Assumption 1, in particular since f is β -smooth and α -strongly convex, it satisfies $\alpha \|x\|^2 \leq \langle \nabla f(x), x \rangle \leq \beta \|x\|^2$ for any x . Hence, reminding that $\sum_{\ell=1}^n x_\ell = n$, the first sum of (20) can be upper bounded by

$$\begin{aligned} \sum_{\ell=1}^n \langle \nabla f(x), x_\ell \mathbb{1}_n - x \rangle &= n \langle \nabla f(x), \mathbb{1}_n - x \rangle \\ &\leq n(\beta n - \alpha \|x\|^2). \end{aligned}$$

The second sum of (20) can be expressed as:

$$\begin{aligned} \sum_{\ell=1}^n \|x_\ell \mathbb{1}_n - x\|^2 &= \sum_{\ell=1}^n \sum_{i=1}^n (x_\ell^2 - 2x_\ell x_i + x_i^2) \\ &= \sum_{\ell=1}^n (n x_\ell^2 - 2n x_\ell + \|x\|^2) \\ &= 2n (\|x\|^2 - n). \end{aligned}$$

Then (20) is upper bounded by

$$\mathbb{E}[\Delta f_t] \leq \frac{1}{n} (\beta n - \alpha \|x\|^2) + \frac{\beta}{n^2} (\|x\|^2 - n) - \frac{1}{n} f(x).$$

Lemmas 1 and 2 yield $\|x\|^2 \leq n^2$ and $f(x) \geq \frac{\alpha}{2} n$, so that

$$\mathbb{E}[\Delta f_t] \leq 2\beta - \frac{3}{2} \alpha. \quad (21)$$

The conclusion follows from adding (19) and (21). ■

We can now use Proposition 2 to study the evolution of the expected dynamical regret in the following theorem.

Theorem 2: In the setting of Section II, there holds

$$\mathbb{E} \text{Reg}_T \leq C_0 \sum_{t=1}^T \eta^t + \sum_{t=0}^{T-1} \eta^t (M_f + (1-p)(T-t)\theta) \quad (22)$$

where $\eta = 1 - \frac{p}{\kappa(n-1)}$ (with p the probability that a given event is an update from Assumption 5), $M_f = \frac{n}{2}(\beta n - \alpha)$ and $\theta = \frac{5}{2}\beta - \frac{3}{2}\alpha$.

Proof: Let $\gamma = 1 - \frac{1}{\kappa(n-1)}$ denote the contraction rate of the RCD algorithm as defined in (8) [19]. Remember that there holds $C_{t+1} = C_t + \Delta f_t + \Delta f_t^*$ for all times t . Hence, at any time-step t one has

$$\mathbb{E}[C_{t+1}] = \mathbb{E}[C_t + \Delta f_t - \Delta f_t^*]. \quad (23)$$

From Assumption 5 the event at iteration t is an update, denoted U_t , with probability p , or a replacement, denoted R_t , with probability $1 - p$.

In the case of an update, we have $x^{*,t+1} = x^{*,t}$, so that $\Delta f_t^* = 0$. Hence, we have $\Delta f_t = C_{t+1} - C_t$, and since $\mathbb{E}[C_{t+1}|C_t, U_t] \leq \gamma C_t$ with the RCD algorithm from [19]:

$$\mathbb{E}[\Delta f_t|U_t] = \mathbb{E}[C_{t+1} - C_t|U_t] \leq (\gamma - 1)\mathbb{E}[C_t]. \quad (24)$$

In the replacement case, there holds

$$\mathbb{E}[\Delta f_t|R_t] \leq \theta = \frac{5}{2}\beta - \frac{3}{2}\alpha, \quad (25)$$

where θ comes from Proposition 2. Injecting (24) and (25) into (23) then yields

$$\begin{aligned} \mathbb{E}[C_{t+1}] &\leq \mathbb{E}[C_t] + p(\gamma - 1)\mathbb{E}[C_t] + (1 - p)\theta - \mathbb{E}[\Delta f_t^*] \\ &= \eta \mathbb{E}[C_t] + (1 - p)\theta - \mathbb{E}[\Delta f_t^*], \end{aligned} \quad (26)$$

where $\eta = 1 + p(\gamma - 1)$. Expression (26) actually describes the evolution of a discrete-time dynamical system of the type $v[k+1] \leq Av[k] + Bu[k]$. Standard results on that framework yield $v[k] \leq A^k v[0] + \sum_{j=0}^{k-1} A^{k-j-1} Bu[j]$, and we obtain

$$\mathbb{E}[C_t] \leq \eta^t C_0 + \sum_{j=0}^{t-1} \eta^{t-j-1} ((1 - p)\theta - \mathbb{E}[\Delta f_j^*]). \quad (27)$$

Injecting this last result into (5) then yields

$$\begin{aligned} \mathbb{E}\text{Reg}_T &= \sum_{t=1}^T \mathbb{E}[C_t] \\ &\leq C_0 \sum_{t=1}^T \eta^t + \sum_{t=1}^T \left(\sum_{j=0}^{t-1} \eta^{t-j-1} ((1 - p)\theta - \mathbb{E}[\Delta f_j^*]) \right). \end{aligned}$$

After some term re-organization, it becomes

$$\begin{aligned} \mathbb{E}\text{Reg}_T &\leq C_0 \sum_{t=1}^T \eta^t + (1 - p) \sum_{t=0}^{T-1} (T - t) \eta^t \theta \\ &\quad - \sum_{t=0}^{T-1} \eta^t \left(\sum_{j=0}^{T-t} \mathbb{E}[\Delta f_j^*] \right). \end{aligned}$$

Finally, using Lemma 2, one concludes that

$$\begin{aligned} - \sum_{j=0}^{T-t} \mathbb{E}[\Delta f_j^*] &= -\mathbb{E} \left[\sum_{j=0}^{T-t} \Delta f_j^* \right] \\ &= -\mathbb{E} [f^{T-t+1}(x^{*,T-t+1}) - f^1(x^{*,1})] \leq \frac{n}{2}(\beta n - \alpha), \end{aligned}$$

and $M_f = \frac{n}{2}(\beta n - \alpha)$ yields the conclusion. ■

We now analyze the asymptotic behavior of the averaged regret in the following corollary.

Corollary 1: Let $\rho_R := \frac{1-p}{p}$. In the same setting as that of Theorem 2, there holds

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}\text{Reg}_T}{T} \leq \rho_R(n-1)\beta \left(\frac{5\kappa-3}{2} \right). \quad (28)$$

Proof: Starting from (22), we have

$$\frac{\mathbb{E}\text{Reg}_T}{T} \leq C_0 \sum_{t=1}^T \frac{\eta^t}{T} + \sum_{t=0}^{T-1} \left(M_f \frac{\eta^t}{T} + (1-p) \left(1 - \frac{t}{T} \right) \eta^t \theta \right).$$

Remember that $\eta = 1 - \frac{p}{\kappa(n-1)} \leq 1$, so that $\sum_{t=1}^T \eta^t < T$, and $\lim_{T \rightarrow \infty} \sum_{t=1}^T \frac{\eta^t}{T} = 0$. Hence

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}\text{Reg}_T}{T} \leq \lim_{T \rightarrow \infty} (1-p) \sum_{t=0}^{T-1} \left(1 - \frac{t}{T} \right) \eta^t \theta.$$

Moreover, for $\eta < 1$, one shows $\sum_{j=0}^{\infty} \eta^j = \frac{1}{1-\eta}$ and $\sum_{j=0}^{\infty} j \eta^j = \frac{\eta}{(1-\eta)^2}$. The latter implies that $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} t \eta^t = 0$, and therefore there holds

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}\text{Reg}_T}{T} \leq (1-p) \frac{1}{1-\eta} \theta = \frac{1-p}{p} \kappa(n-1) \theta,$$

and the conclusion follows from the definitions of θ in Theorem 2, and from $\rho_R := \frac{1-p}{p}$. ■

The upper bounds for the potential benefit (13) and the dynamical regret (22) linearly scale with T . This behavior is rather natural for the former, which does not depend on the algorithm. For the latter, it is most likely unavoidable due to the introduction at each replacement of perturbations of non-decaying magnitude and without any regularity [21]. Interestingly, this behavior contrasts with standard results in online optimization, where a sublinear growth in T is desired to cancel the asymptotic averaged regret [16, Ch. 1.1]. However, these results usually apply on another definition of the regret, where x^t is compared with an overall time-independent strategy x^* computed over all T iterations, in opposition with $x^{*,t}$ which is optimal for each iteration.

Moreover, the corresponding asymptotic upper bounds linearly grow with n and α for (14), and with $n-1$ and β for (28), consistently with their expected behavior. In particular, the scaling of (28) with $n-1$ follows from the convergence rate of the RCD algorithm $\gamma = 1 - \frac{1}{\kappa(n-1)}$. Interestingly, (28) is proportional to $\rho_R(n-1) = (1-p) \frac{n-1}{p}$, and the bound can thus be seen as the ratio between the probability for a given agent to be involved in a RCD update $\frac{p}{n-1}$ (involved in γ), and the impact of replacements at the system level $1-p$, independently of n . This is consistent with the bound on the impact of replacements in (18) which is independent of n (by contrast, alternative situations such as *e.g.*, if all agents were to be reset at each replacement are expected to generate an impact growing with n). Hence, for small values of ρ_R (*i.e.*, rare replacements), the bound guarantees that the asymptotic dynamical regret remains reasonably bounded, and decays to zero when $\rho_R \rightarrow 0$, *i.e.*, for closed systems.

Finally, observe that (28) is proportional to $\frac{5\kappa-3}{2}$, consistently with the fact that a larger interval for the possible

curvature of the cost functions should generate a larger potential error at replacements. This factor is a potential source of conservatism, e.g., with respect to (14) where the scaling is in $\frac{1}{2}(\kappa - 1)$. More generally, it is not clear yet whether other algorithms than the RCD might provide tighter bounds.

Remark 1: The proofs of Theorem 2 and Corollary 1 can directly be adapted to any contraction rate $\gamma < 1$, and are thus easily generalized to any other algorithm that guarantees linear convergence; in particular $\lim_{T \rightarrow \infty} \frac{\mathbb{E}\text{Reg}_T}{T} \leq \rho_R \frac{\theta}{1-\gamma}$.

C. The case of quadratic functions

The bound on the expected dynamical regret can be refined for the particular case where all local functions are quadratic, i.e. satisfy the following additional assumption.

Assumption 6 (Quadratic functions): The local cost function of any agent i at time t is of the form

$$f_i(x_i) = \phi_i x_i^2, \quad \phi_i \in \left[\frac{\alpha}{2}, \frac{\beta}{2}\right]. \quad (29)$$

The parameter ϕ_i is randomly chosen according to a distribution with a finite support determined by the interval $\left[\frac{\alpha}{2}, \frac{\beta}{2}\right]$. Observe that functions satisfying Assumption 6 necessarily satisfy Assumption 1 as well.

Under Assumption 6, we can obtain a tighter bound than that of Proposition 2, presented in the following proposition.

Proposition 3: In the setting of Section II, and under Assumption 6, the replacement of an agent R results in

$$\mathbb{E}[\Delta f_t \mid R] \leq \frac{3n^2 - 3n + 1}{2n^2} (\beta - \alpha). \quad (30)$$

Proof: The arrival case is treated the same as in the proof of Proposition 2, resulting in $\Delta f_t \leq \frac{\beta}{2}$. The departure case follows the same first steps with $f_i(x) = \phi_i x^2$, and

$$\begin{aligned} \mathbb{E}[\Delta f_t] &= \sum_{\ell=1}^n \frac{1}{n} \left(\sum_{i \neq \ell} \phi_i \left(x_i + \frac{x_\ell - x_i}{n} \right)^2 - f(x) \right) \\ &= \frac{1}{n} \sum_{\ell=1}^n \left(\sum_{i \neq \ell} \phi_i \left(\left(x_i + \frac{x_\ell - x_i}{n} \right)^2 - x_i^2 \right) - \phi_\ell x_\ell^2 \right) \\ &= \frac{1}{n^2} \sum_{\ell=1}^n \sum_{i \neq \ell} \phi_i \left(\frac{1-2n}{n} x_i^2 + 2x_i x_\ell \frac{n-1}{n} + \frac{x_\ell^2}{n} \right) - \frac{f(x)}{n}. \end{aligned}$$

Using the fact that $\sum_{\ell=1}^n \sum_{i \neq \ell} \phi_i x_i^2 = (n-1)f(x)$ and that $\phi_i \leq \frac{\beta}{2}$ for all i , we obtain

$$\begin{aligned} \mathbb{E}[\Delta f_t] &\leq \left(\frac{1}{n^2} \frac{1-2n}{n} (n-1) - \frac{1}{n} \right) f(x) \\ &\quad + \frac{\beta}{2n^2} \sum_{\ell=1}^n \sum_{i \neq \ell} \left(2x_i x_\ell \frac{n-1}{n} + \frac{x_\ell^2}{n} \right). \end{aligned}$$

Observe that $\sum_{\ell=1}^n \sum_{i \neq \ell} x_i x_\ell = n^2 - \|x\|^2$ and $\sum_{\ell=1}^n \sum_{i \neq \ell} x_\ell^2 = (n-1)\|x\|^2$, so that a few algebraic manipulations yield

$$\mathbb{E}[\Delta f_t] \leq -\frac{3n^2 - 3n + 1}{n^3} f(x) + \beta \frac{n-1}{2n^3} (2n^2 - \|x\|^2).$$

Using $\|x\|^2 \geq n$ (from Lemma 1) and $f(x) \geq \frac{\alpha}{2}n$ (from Lemma 2) then yields

$$\mathbb{E}[\Delta f_t] \leq -\frac{3n^2 - 3n + 1}{2n^3} \alpha + \beta \frac{2n^2 - 3n + 1}{2n^3},$$

and combining with the arrival case concludes the proof. ■

The result above allows us stating the following theorem, which improves Theorem 2 and Corollary 1 respectively for the case of quadratic functions.

Theorem 3: In the setting of Section II, and under Assumption 6, there holds

$$\mathbb{E}\text{Reg}_T \leq C_0 \sum_{t=1}^T \eta^t + \sum_{t=0}^{T-1} \eta^t (M_f + (1-p)(T-t)\theta) \quad (31)$$

where $\eta = 1 - \frac{p}{\kappa(n-1)}$ (with p the probability that a given event is an update from Assumption 5), $M_f = \frac{n}{2}(\beta n - \alpha)$ and $\theta = (\beta - \alpha) \frac{3n^2 - 3n + 1}{2n^2}$. In particular,

$$\lim_{T \rightarrow \infty} \frac{\mathbb{E}\text{Reg}_T}{T} \leq \rho_R (n-1) \frac{3n^2 - 3n + 1}{2n^2} \beta (\kappa - 1). \quad (32)$$

Proof: The proof follows the exact same steps as those of Theorem 2 and Corollary 1 where Proposition 3 is used instead of Proposition 2. ■

The upper bound for the quadratic case is qualitatively better than that of the general case; it was derived based on an additional information of the cost function, thus resulting in tighter bounds. In this case, the dependence of (32) is in $\frac{3}{2}(\kappa - 1)$, which is consistent with the result derived for the potential benefit (14). In particular, for n becoming large, (32) and (28) become equivalent up to a constant β . Moreover, (32) becomes 0 when $\kappa = 1$, consistently with the expected behavior of the RCD algorithm for quadratic functions since all the cost functions would then be the same.

D. Numerical Results

To illustrate the results of Theorems 1 to 2, we consider a system of 5 agents with $\kappa = 10$ and $\rho_R = 0.0125$, the latter implies that on average there is one replacement every 80 events. We consider two possibilities: *random replacements* (RR) where the local function is randomly uniformly chosen among the set of piecewise quadratic functions satisfying Assumption 1, and *adversarial replacements* (AR), where these functions are quadratic functions $\phi_i x^2$, with $\phi_i \in \left\{\frac{\alpha}{2}, \frac{\beta}{2}\right\}$. The AR setting is expected to be less favorable than the RR setting, since replacements might result in the largest change of local functions. Notice that the bounds (13) and (22) are independent of the distribution from which the local cost functions are assigned to the agents when they join the system, so that they hold for any such assignment rule.

Fig. 2 compares the results of Theorem 1 and Corollary 1 with simulations for both random and adversarial replacements in the setting described above. Even though the theoretical bounds are conservative, they capture well the qualitative behavior of these metrics. In particular, consistently with Pot_T and $\mathbb{E}\text{Reg}_T$ that grow linearly with T , the bounds in the figure do not converge to zero, and a remaining asymptotic error is observed. Our bounds are tighter for the




Fig. 2. Evolution of the averaged asymptotic expected Potential Benefit (on the left) and dynamical regret (on the right) in a system of 5 agents with $\rho_R = 0.0125$ and $\kappa = 10$. Each plot compares the upper bounds, respectively from (14) and (28), with simulated results, either with random replacements (RR) or adversarial replacements (AR).




Fig. 3. Expected averaged regret for a system of 5 agents holding quadratic functions with $\rho_R = 0.0125$ and $\kappa = 10$. The plain blue and dashed red lines respectively correspond to the upper bounds of Theorems 2 and 3 respectively, and the dotted lines of the same colors to the corresponding asymptotic values. The dash-dotted yellow line corresponds to simulation where we consider random replacements (RR).

adversarial replacement case than the random replacement case, this suggests that our bounds might be tight for some particular choice of the joining functions at replacements, especially that on the potential benefit.

Fig. 3 compares the results of Theorems 2 and 3 with simulations in the same setting as described previously, with random replacements (RR). The figure shows that bound (32) is tighter for quadratic functions, following the fact that we have access to more information regarding the local cost functions, thus improving the estimation of the effect of replacements on the expected regret.

IV. CONCLUSION

We analyzed the performance and behavior of the Random Coordinate Descent algorithm (RCD) for solving the optimal resource allocation problem in an open system subject to replacements of agents, resulting in variations of the total cost function and of the total amount of resource to be allocated. We considered a simple preliminary setting where the budget is homogeneous and the graph is complete, and used tools inspired from online optimization to show that it is not possible to achieve convergence to the optimal solution with the RCD algorithm in expectation in open system, but that the error is expected to remain reasonable.

We have derived upper bounds on the evolution of the regret and the potential benefit in expectation and showed that due to the random choice of the new local cost function during replacements, an error is expected to be accumulated with time and cannot be compensated. A natural continuation

of this work is thus the derivation of the corresponding upper bound for the benefit, and of lower bounds for this quantities in order to validate the observed behavior. More generally, our bounds could be extended to more general settings, and their tightness can be improved to match more accurately the actual performance of the algorithm. Moreover, since our approach is based on an independent analysis of the effect of arrivals and departures, the next step of this study is to generalize it to the case where the system size changes with the time, *i.e.*, where arrivals and departures are decoupled.

REFERENCES

- [1] T. Ibaraki and N. Katoh, *Resource Allocation Problems: Algorithmic Approaches*. MIT press, 1988.
- [2] J. F. Kurose and R. Simha, "A microeconomic approach to optimal resource allocation in distributed computer systems," *IEEE Transactions on Computers*, vol. 38, no. 5, pp. 705–717, 1989.
- [3] S. Liang, P. Yi, and Y. Hong, "Distributed Nash equilibrium seeking for aggregative games with coupled constraints," *Automatica*, 2017.
- [4] P. Dai, W. Yu, and D. Chen, "Distributed Q-learning algorithm for dynamic resource allocation with unknown objective functions and application to microgrid," *IEEE Transactions on Cybernetics*, 2021.
- [5] A. Teixeira, J. Araújo, H. Sandberg, and K. H. Johansson, "Distributed actuator reconfiguration in networked control systems," *IFAC Proceedings Volumes*, vol. 46, no. 27, pp. 61–68, 2013.
- [6] P. Yi, Y. Hong, and F. Liu, "Initialization-free distributed algorithms for optimal resource allocation with feasibility constraints and application to economic dispatch of power systems," *Automatica*, 2016.
- [7] L. Bai, C. Sun, Z. Feng, and G. Hu, "Distributed continuous-time resource allocation with time-varying resources under quadratic cost functions," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 823–828.
- [8] B. Wang, S. Sun, and W. Ren, "Distributed time-varying quadratic optimal resource allocation subject to nonidentical time-varying Hessians with application to multiquadrotor hose transportation," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2022.
- [9] Y. Nesterov, *Lectures on Convex Optimization*. Springer, 2018.
- [10] S. Bubeck, "Convex optimization: Algorithms and complexity," *Foundations and Trends in Machine Learning*, vol. 8, no. 3-4, 2015.
- [11] S. Riverso, M. Farina, and G. Ferrari-Trecate, "Plug-and-play decentralized model predictive control," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 4193–4198.
- [12] M. Farina and R. Carli, "Partition-based distributed Kalman filter with plug and play features," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 1, pp. 560–570, 2018.
- [13] J. M. Hendrickx and S. Martin, "Open multi-agent systems: Gossiping with random arrivals and departures," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 763–768.
- [14] C. Monnayer de Galland, S. Martin, and J. M. Hendrickx, "Modelling gossip interactions in open multi-agent systems," *arXiv preprint arXiv:2009.02970*, 2020.
- [15] R. Vizuete, P. Frasca, and E. Panteley, "On the influence of noise in randomized consensus algorithms," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1025–1030, 2021.
- [16] E. Hazan, "Introduction to online convex optimization," *Foundations and Trends in Optimization*, vol. 2, pp. 157–325, 01 2016.
- [17] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Transactions on Automatic Control*, vol. 63, no. 3, pp. 714–725, 2018.
- [18] X. Yi, X. Li, T. Yang, L. Xie, T. Chai, and K. Johansson, "Regret and cumulative constraint violation analysis for online convex optimization with long term constraints," in *International Conference on Machine Learning*. PMLR, 2021, pp. 11 998–12 008.
- [19] I. Necoara, "Random coordinate descent algorithms for multi-agent convex optimization over networks," *IEEE Transactions on Automatic Control*, vol. 58, no. 8, pp. 2001–2012, 2013.
- [20] C. Monnayer de Galland, R. Vizuete, J. M. Hendrickx, E. Panteley, and P. Frasca, "Random coordinate descent for resource allocation in open multi-agent systems," *arXiv preprint arXiv:2205.10259*, 2022.
- [21] X. Li, L. Xie, and N. Li, "A survey of decentralized online learning," *arXiv preprint arXiv:2205.00473*, 2022.