Distributed Resource Allocation on Dynamic Networks in Quadratic Time

Thinh Thanh Doan^{a,*}, Alex Olshevsky^{b,**}

^a*Department of Electrical and Computer Engineering, University of Illinois, Urbana, IL, USA*

^b*Department of Industrial and Enterprise Systems Engineering, University of Illinois, Urbana, IL, USA*

Abstract

We consider the problem of allocating a fixed amount of resource among nodes in a network when each node suffers a cost which is a convex function of the amount of resource allocated to it. We propose a new deterministic and distributed protocol for this problem. Our main result is that the associated convergence time for the global objective scales quadratically in the number of nodes on any sequence of time-varying undirected graphs satisfying a long-term connectivity condition.

1. Introduction

We consider the problem of optimally allocating a fixed amount of resource among n agents. Each agent suffers a convex cost as a function of the amount of resource allocated to it and the goal is to distribute the resource among the agents to minimize the total cost incurred. Sometimes the problem is described in terms of utilities, with each agent having a concave utility function and the goal being to maximize the total utility.

Our goal is to develop distributed protocols for this problem, meaning that nodes are only allowed to interact with neighbors in some graph or some time-varying sequence of graphs. Our motivation comes from potential applications in sensor networks, which regularly face the problems of optimally allocating communication bandwidth and computing power [9]. Furthermore, resource allocation is a simplification of the important “economic dispatch” problem wherein geographically distributed producers of electricity must coordinate to meet a fixed demand [5, 6, 7].

The problem has an old history dating back to the classic work of Arrow and Hurwicz [1]. The first algorithm which could be implemented in a

distributed way was the “center-free” protocol of [8]. In the protocol of [8], each node increases the amount of resource allocated to it proportionally to the difference in gradients between its neighbors and itself. It was shown in [8] that, under appropriate technical conditions, this protocol will drive the amount of resource allocated to each node to the optimal value. The term “center-free” was originally meant to refer to the absence of any central coordinating authority, though in this paper we will use it to mean any update wherein nodes update the amount of resource by looking at gradient differences with neighbors. The work of [8] has spawned a number of modern follow-ups, including [10, 16, 8, 11, 14, 12] as well as the current paper.

The paper [10] considered the resource allocation problem in the context of optimal distribution of a database among the nodes of the network; some modifications of the algorithm of [8] which used not only gradient differences but also the second derivatives of the cost functions were proposed. More recently, [16] studied the case when the cost functions are strongly convex and noted that the problem of optimal weight selection for center-free methods can be cast as a semidefinite program. The work of [11] analyzed a natural class of center-free methods on time-varying networks and provided a convergence analysis. The recent paper [14, 12] studied the convergence rates of distributed protocols which repeatedly choose a random pair of neighboring nodes and perform a center-free update on that

*Corresponding author

**This research was supported by NSF grant CMMI-1463262.

Email addresses: ttdoan2@illinois.edu (Thinh Thanh Doan), aolshev2@illinois.edu (Alex Olshevsky)

pair. Finally, the work of [3] used accelerated gradient methods to design distributed protocols for a more general problem.

Our focus in this paper is on designing protocols with good convergence speed. Specifically, we are interested at how the gap to the optimal objective value scales in the worst-case with iteration k and the number of nodes n in the system.

The best previously known results were provided in the antecedent papers [11, 12]. Both papers considered the class of costs which have Lipschitz-continuous derivatives. The paper [12] considers schemes which randomly pick pairs of neighbors to perform a center-free update; if the pairs are chosen uniformly at random the convergence time implied by the results of [12] is $O(Ln^4/k)$ in expectation¹ on fixed graphs; here L is the largest of the Lipschitz constants of the derivatives of the cost functions. However, we note here that it is possible to shave off a factor of n off this bound by adjusting the probabilities in a graph-dependent way. The paper [11] does not give an explicit convergence rate for the objective, but gives a worst-case $O(LBn^3/k)$ rate for the decay of the average of *squared* gradient differences in the graph; here B is a constant which measures how long it takes for the time-varying graph sequence to reach connectivity. Improved rates were obtained in [13] and in [2] for a more general problem, but under the assumption that the graph is a fixed complete graph.

In this paper, we show a convergence rate of $O(LBn^2/k)$ for the objective under the same assumptions of Lipschitz-continuous derivatives in the more general setting of time-varying graphs. Additionally, when the costs are strongly convex, we demonstrate a geometric rate of $O\left((1 - \mu/(4Ln^2))^{k/B}\right)$ where μ is the parameter of strong convexity. For both of these rates, the number of iterations until the objective is within ϵ of its optimal value scales quadratically with the number of nodes n . This is an improvement over the results described above, though we note that our protocol involves every node contacting its neighbors and performing an update at every step (which involves $O(|E(t)|)$ messages exchanged, where $E(t)$ is the set of edges at time t , and $O(n)$ updates);

¹The convergence rate in [12] is given in terms of the eigenvalues of a certain matrix; the quartic bound above follows by putting [12] together with the well-known fact that the smallest eigenvalue of the Laplacian of a connected, undirected graph on n is $\Omega(1/n^2)$.

whereas [12] relied on only a pair of randomly chosen nodes updating at each step.

The remainder of this paper is organized as follows. We give a formal statement of the problem in Section 2. Our protocol is described in Section 3. The convergence analysis of the protocol is in Section 4. Finally, Section 5 describes the results of some simulations and we conclude in Section 6.

2. Problem Formulation

In this paper, we study distributed protocols for the following minimization problem,

$$\begin{aligned} \min \quad & \sum_{i=1}^n f_i(x_i) \\ \text{s.t.} \quad & \sum_{i=1}^n x_i = K. \end{aligned} \quad (1)$$

We assume that there are n agents or nodes which we will index as $1, \dots, n$, that $f_i : \mathbb{R} \rightarrow \mathbb{R}$ is a convex function known only to node i and $x_i \in \mathbb{R}$ is a variable stored by node i , and finally that K is some nonnegative number.

As remarked, this models a resource allocation problem among n agents: given a finite amount K of a certain resource, we must allocate it among agents $1, \dots, n$ in an optimal way.

For simplicity, we introduce notation for the total objective function $F(\mathbf{x}) = \sum_{i=1}^n f_i(x_i)$, and the feasible set $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^n x_i = K\}$.

We assume that a sequence of time-varying undirected graphs models the communication between the nodes. Specifically, we assume we are given a sequence of undirected graphs $\mathcal{G}(k) = (\mathcal{V}, \mathcal{E}(k))$ with $\mathcal{V} = \{1, \dots, n\}$; nodes i and j can send exchange messages at time k if and only if $(i, j) \in \mathcal{E}(k)$. We denote by $\mathcal{N}_i(k)$ the set of neighbors of node i at time k .

We make the following fairly standard assumption which ensures that the graph sequence $\mathcal{G}(k)$ satisfies a long-term connectivity property.

Assumption 1. There exists an integer $B \geq 1$ such that the undirected graph

$$(\mathcal{V}, \mathcal{E}(\ell B) \cup \mathcal{E}(\ell B + 1) \cup \dots \cup \mathcal{E}((\ell + 1)B - 1)) \quad (2)$$

is connected for all nonnegative integers ℓ .

We will also be assuming that each local objective function $f_i(\cdot)$ is differentiable with Lipschitz continuous derivative.

Assumption 2. For each $i = 1, \dots, n$, the function $f_i(\cdot)$ is differentiable everywhere and there exists a constant L_i such that

$$|f'_i(y_i) - f'_i(x_i)| \leq L_i |y_i - x_i|, \quad \forall x_i, y_i \in \mathbb{R}.$$

Moreover, we will be assuming that there exists at least one optimal solution.

Assumption 3. There exists a vector $\mathbf{x}^* = (x_1^*, x_2^*, \dots, x_n^*)$ with $\mathbf{x}^* \in \mathcal{S}$ which achieves the minimum in problem (1).

We will use $\mathcal{X}^*$ to denote the set of optimal solutions to problem (1); the previous assumption ensures that $\mathcal{X}^*$ is not empty.

Finally, we will be assuming that our algorithm starts from a feasible point.

Assumption 4. $\mathbf{x}(0) \in \mathcal{S}$.

For the remainder of this paper, we will be assuming that Assumptions 1, 2, 3, 4 hold without mention.

We conclude this section with a characterization of the points in the optimal set $\mathcal{X}^*$; the proof is immediate.

Proposition 1. *We have that $\mathbf{x} \in \mathcal{X}^*$ if and only if $\mathbf{x} \in \mathcal{S}$ and $f'_i(x_i) = f'_j(x_j)$ for all $i, j \in \{1, \dots, n\}$.*

3. Main Algorithm

In this section, we will introduce a distributed protocol, which we call the *gradient balancing protocol*, to solve problem (1). Before giving a statement of the algorithm, we provide some brief motivation for its form.

Previous protocols for problem (1) tended to be “center-free” updates [8, 16, 11, 12] where node i updated as

$$x_i(k+1) = x_i(k) - \sum_{j \in N_i(k)} w_{ij} (f'_i(x_i(k)) - f'_j(x_j(k))), \quad (3)$$

where w_{ij} is a collection of nonnegative weights. The protocol of [12] had a different form but proceeded in the same spirit; in that protocol, edges were repeatedly chosen according to some probability distribution and a form of the above update was performed by the incident nodes.

The protocol we propose in this paper speeds up this update by employing some local “pruning” wherein each node tries to perform a version of Eq. (3), but only with the two nodes whose derivative is largest and smallest in its neighborhood. Thus

nodes essentially ignore neighbors whose derivatives are close to their own. Intuitively, by focusing on nodes whose derivatives are far apart we increase the speed at which information propagates through the network. The idea has been previously used in [15] and is inspired by an algorithm from Chapter 7.4 of [4].

We now describe the steps node i executes at step k to update its value from $x_i(k)$ to $x_i(k+1)$. We assume that all nodes execute these steps synchronously, and furthermore that all four steps of the protocol given below can be executed before the graph changes from $\mathcal{G}(k)$ to $\mathcal{G}(k+1)$. Speaking informally, the protocol consists of each node repeatedly trying to “match” itself to the node in its neighborhood whose derivative is smallest and smaller than its own in order to perform a center-free update.

The Gradient Balancing Protocol

1. Node i broadcasts its derivative $f'_i(x_i(k))$ and Lipschitz constant L_i to its neighbors.
2. Going through the messages it has received from neighbors as a result of step 1, node i finds the neighbor with the smallest derivative that is strictly less than its own. Let p be a neighbor with this derivative; ties can be broken arbitrarily. Formally,

$$p \in \arg \min_j \{f'_j(x_j(k)) \mid j \in N_i(k), f'_j(x_j(k)) < f'_i(x_i(k))\}.$$

Node i then sends a message containing the number

$$\Delta_{ip}(k) = \frac{1}{2} \frac{f'_i(x_i(k)) - f'_p(x_p(k))}{L_i + L_p}$$

to node p .

If no neighbor of i has a derivative strictly less than $f'_i(x_i(k))$, node i does nothing during this step.

3. Node i goes through any $\Delta_{ji}(k)$ it has just received from its neighbors $j \in N_i(k)$ as a result of step 2, and finds the largest among them; ties can be broken arbitrarily. Let us suppose this is $\Delta_{qi}(k)$.

Node i then sets

$$y_i(k) = x_i(k) + \Delta_{qi}(k).$$

Furthermore, node i sends an “accept” message to node q and a “reject” message to any other

neighbor j that sent it a $\Delta_{ji}(k)$ in step 2.

If node i did not receive any $\Delta_{ji}(k)$ as a result of step 2, it sets $y_i(k) = x_i(k)$.

4. If node i did not send out $\Delta_{ip}(k)$ during step 2, or if it received a “reject” from the node p to whom it sent $\Delta_{ip}(k)$, it sets $x_i(k+1) = y_i(k)$.
If node i has received an “accept” from node p , it sets

$$x_i(k+1) = y_i(k) - \Delta_{ip}(k).$$

Informally, we will refer to the numbers $\Delta_{ij}(k)$ as “offers.” We may summarize the gradient balancing protocol as follows. Each node i makes an offer to the node with the smallest derivative (below its own) in its neighborhood, and the size of the offer is proportional to the difference of the derivatives normalized by the sum of the respective Lipschitz constants. Each node then accepts the largest offer it has received and rejects the rest. Note that each node “accepts” at most one offer and “makes” at most one offer. The final result is something like Eq. (3), except that the graph has been pruned to be of degree at most two and contain only edges between nodes whose derivatives are “far apart.”

We remark that an immediate consequence of Assumption 4 is that $\mathbf{x}(k) \in \mathcal{S}$ for all $k \geq 0$, since every accepted offer involves an increase at the receiving node and a decrease at the offering node of the same magnitude.

For concreteness, we provide an example of our protocol; see Fig. 1. The top part of the figure shows $x_i(k)$ and $f'_i(x_i(k))$ for each node in parenthesis, respectively. We assume that $L_i = 1/2$ for all $i = 1, \dots, n$. The bottom part of the figure shows the new values $x_i(k+1)$. As we can see that node B and node C send offers to node D but node D only accepts node B 's offer. Node D also sends an offer to node E and node E accepts since it is the only offer it receives. Nodes A and C do not end up participating in any accepted offers and consequently for those nodes $x_i(k+1) = x_i(k)$.

4. Convergence Analysis

We now turn to the convergence analysis of the gradient balancing protocol. We will prove upper bounds on $F(\mathbf{x}(k)) - F(\mathbf{x}^*)$ which imply that the time until this quantity shrinks below ϵ is quadratic in the number of nodes n .

Figure 1: A step of the gradient balancing protocol.

We start with the observation that the gradient balancing protocol may be rewritten in a particularly convenient way. Let us define $\bar{\mathcal{E}}(k)$ to be the set of pairs (i, j) such that either i accepts an offer from j at time k or vice versa. We can then write

$$x_i(k+1) = x_i(k) - \sum_{j \mid (i,j) \in \bar{\mathcal{E}}(k)} \frac{f'_i(x_i(k)) - f'_j(x_j(k))}{2(L_i + L_j)}, \quad (4)$$

We now begin with a series of lemmas which lead the way to our main convergence result. Our first lemma shows the monotonicity of the largest and smallest derivatives in the network.

Lemma 1. *The function $\min_{i=1,\dots,n} f'_i(x_i(k))$ is nondecreasing in k and the function $\max_{i=1,\dots,n} f'_i(x_i(k))$ is nonincreasing in k .*

Proof. Consider node j . We will show that there is always some node q such that $f'_j(x_j(k+1)) \geq f'_q(x_q(k))$. This will prove the monotonicity of the smallest derivative; the monotonicity of the largest derivative is proved analogously.

If j does not make any offers during step 2 of the gradient balancing protocols, or if it makes an offer which is rejected, then we must have $x_j(k+1) \geq x_j(k)$; since $f_j(\cdot)$ is convex this implies that $f'_j(x_j(k+1)) \geq f'_j(x_j(k))$. Thus we may take $q = j$ in this case.

On the other hand, suppose j makes an offer during step 2 which is accepted, say by node m . From Eq. (4),

$$x_j(k+1) \geq x_j(k) - \frac{f'_j(x_j(k)) - f'_m(x_m(k))}{2(L_m + L_j)},$$

and since $f_j(\cdot)$ is convex and differentiable with L_j -

Lipschitz continuous derivative,

$$\begin{aligned} f'_j(x_j(k+1)) &\geq f'_j\left(x_j(k) - \frac{f'_j(x_j(k)) - f'_m(x_m(k))}{2(L_m + L_j)}\right) \\ &\geq f'_j(x_j(k)) - \frac{L_j (f'_j(x_j(k)) - f'_m(x_m(k)))}{2(L_m + L_j)} \\ &> f'_m(x_m(k)), \end{aligned}$$

so that we may take $q = m$ in this case. $\square$

We will often need to be making statements about the d largest derivatives at time k . To avoid overburdening the reader with notation, we will begin many of our lemmas with a variation on the words “let us relabel the vertices so that the sequence $f'_1(x_1(k)), f'_2(x_2(k)), \dots, f'_n(x_n(k))$ is nonincreasing.” Under this assumption, the d largest derivatives may be taken to be $f'_1(x_1(k)), \dots, f'_d(x_d(k))$.

Furthermore, assuming the nodes have been relabeled as above, we will say that *edge* (i, j) *crosses the cut* d if one of i, j belongs to $\{1, \dots, d\}$ while the other belongs to $\{d+1, \dots, n\}$.

An example of the use of these definitions is in the following corollary, which is an immediate consequence of Lemma 1.

Corollary 1. *Let us relabel the nodes so that the sequence $f'_1(x_1(k)), f'_2(x_2(k)), \dots, f'_n(x_n(k))$ is nonincreasing. Suppose that during times $t = k, k+1, \dots, k+T$ we have that $\bar{\mathcal{E}}(t)$ did not include any edges crossing the cut d . Then for $i = 1, \dots, d$, we have that*

$$f'_i(x_i(t+T+1)) \geq f'_d(x_d(t)),$$

while for $i = d+1, \dots, n$,

$$f'_i(x_i(t+T+1)) \leq f'_{d+1}(x_{d+1}(t)).$$

Our next lemma essentially says that cuts in the graph which separate larger derivatives from smaller derivatives must have edges in $\bar{\mathcal{E}}(k)$ which cross them eventually. The proof follows from Assumption 1 on B -connectivity and is the same as the proof of Lemma 3 in [15], so we omit it.

Lemma 2. *Let $\ell \geq 0$ and let us relabel the nodes so that the sequence $f'_1(x_1(\ell B)), f'_2(x_2(\ell B)), \dots, f'_n(x_n(\ell B))$ is nonincreasing. Then for every $d \in \{1, \dots, n-1\}$, either $f'_d(x_d(\ell B)) = f'_{d+1}(x_{d+1}(\ell B))$, or there exist some time $k \in \{\ell B, \dots, (\ell+1)B-1\}$ when an edge in $\bar{\mathcal{E}}(k)$ crosses the cut d .*

We now proceed to our first substantial lemma, which shows that the gradient balancing protocol is a descent protocol (meaning that $F(\mathbf{x}(k))$ is nonincreasing).

Lemma 3.

$$F(\mathbf{x}(k+1)) \leq F(\mathbf{x}(k)) - \sum_{(i,j) \in \bar{\mathcal{E}}(k)} \frac{(f'_i(x_i(k)) - f'_j(x_j(k)))^2}{4(L_i + L_j)}. \quad (5)$$

Proof. Assumption 2 immediately implies that for all $x_i, y_i \in \mathbb{R}$

$$f_i(y_i) \leq f_i(x_i) + f'_i(x_i)(y_i - x_i) + \frac{L_i}{2}(y_i - x_i)^2.$$

Summing up both sides over $i = 1, \dots, n$, we obtain

$$F(\mathbf{y}) \leq F(\mathbf{x}) + \sum_{i=1}^n f'_i(x_i)(y_i - x_i) + \sum_{i=1}^n \frac{L_i}{2}(y_i - x_i)^2.$$

Replacing $\mathbf{x}$ by $\mathbf{x}(k)$, $\mathbf{y}$ by $\mathbf{x}(k+1)$, we obtain

$$\begin{aligned} F(\mathbf{x}(k+1)) &\leq F(\mathbf{x}(k)) \\ &\quad + \sum_{i=1}^n f'_i(x_i(k))(x_i(k+1) - x_i(k)) \\ &\quad + \sum_{i=1}^n \frac{L_i}{2}(x_i(k+1) - x_i(k))^2. \end{aligned} \quad (6)$$

On the other hand, one consequence of Eq. (4) is that

$$\begin{aligned} &\sum_{i=1}^n f'_i(x_i(k))(x_i(k+1) - x_i(k)) \\ &= - \sum_{i=1}^n \sum_{j \mid (i,j) \in \bar{\mathcal{E}}(k)} \frac{f'_i(x_i(k))}{2(L_i + L_j)} (f'_i(x_i(k)) - f'_j(x_j(k))) \\ &= - \sum_{(i,j) \in \bar{\mathcal{E}}(k)} \frac{(f'_i(x_i(k)) - f'_j(x_j(k)))^2}{2(L_i + L_j)}. \end{aligned} \quad (7)$$

Furthermore, another consequence of Eq. (4) is that

$$\begin{aligned} &\sum_{i=1}^n \frac{L_i}{2}(x_i(k+1) - x_i(k))^2 \\ &= \sum_{i=1}^n \frac{L_i}{2} \left(\sum_{j \mid (i,j) \in \bar{\mathcal{E}}(k)} \frac{f'_i(x_i(k)) - f'_j(x_j(k))}{2(L_i + L_j)} \right)^2 \\ &\leq \sum_{i=1}^n \sum_{j \mid (i,j) \in \bar{\mathcal{E}}(k)} \frac{2L_i(f'_i(x_i(k)) - f'_j(x_j(k)))^2}{8(L_i + L_j)^2} \\ &= \sum_{(i,j) \in \bar{\mathcal{E}}(k)} \frac{(f'_i(x_i(k)) - f'_j(x_j(k)))^2}{4(L_i + L_j)}, \end{aligned} \quad (8)$$

where the inequality follows because node i is incident to at most two edges in $\bar{\mathcal{E}}(k)$, which allows us to use the inequality $(a+b)^2 \leq 2(a^2+b^2)$. Finally, we substitute Eq. (7) and Eq. (8) into Eq. (6) to obtain the statement of the lemma. $\square$

Glancing at Eq. (5) which we have just derived, we might guess that the second term on the right is ultimately going to determine how fast the gradient balancing protocol will converge. Our next two lemmas provide useful lower bounds for this quantity over the time interval $k = \ell B, \dots, (\ell+1)B-1$.

Lemma 4. *Let us relabel the nodes so that the sequence $f'_1(x_1(\ell B)), f'_2(x_2(\ell B)), \dots, f'_n(x_n(\ell B))$ is in nonincreasing order. Then,*

$$\begin{aligned} & \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2 \\ & \geq \sum_{d=1}^{n-1} (f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)))^2. \end{aligned}$$

Proof. We begin the proof by introducing some notation. For all $k \in \{\ell B, \ell B+1, \dots, (\ell+1)B-1\}$, we use $D(k)$ to denote the set of $d \in \{1, \dots, n-1\}$ such that time k is the first time in $\{\ell B, \ell B+1, \dots, (\ell+1)B-1\}$ with an edge $(i, j) \in \bar{\mathcal{E}}(k)$ crossing the cut d . Note that $D(k)$ may be empty. Furthermore, given the edge $(i, j) \in \bar{\mathcal{E}}(k)$ we will use $F_{ij}(k)$ to denote all the cuts $d \in D(k)$ crossed by (i, j) at time k . Likewise, it may be the case that $F_{ij}(k)$ is empty.

We begin with the following observation. Suppose $F_{ij}(k) = \{d_1, \dots, d_q\}$ where $d_1 < d_2 < \dots < d_q$. Then

$$\begin{aligned} & (f'_i(x_i(k)) - f'_j(x_j(k)))^2 \\ & \geq \sum_{d \in F_{ij}(k)} (f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)))^2. \end{aligned} \quad (9)$$

We now justify Eq. (9). Indeed, since $d_1 \in F_{ij}(k)$, we have $d_1 \in D(k)$. By definition of $D(k)$ there were no edges (i, j) during times $\ell B, \dots, k-1$ which crossed the cut d_1 . Applying Corollary 1, we have that $f'_i(x_i(k)) \geq f'_{d_1}(x_{d_1}(\ell B))$ and that $f'_j(x_j(k)) \leq f'_{d_q+1}(x_{d_q+1}(\ell B))$. Therefore,

$$\begin{aligned} & \geq f'_{d_1}(x_{d_1}(\ell B)) - f'_{d_q+1}(x_{d_q+1}(\ell B)) \\ & \geq \sum_{d \in F_{ij}(k)} f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)). \end{aligned}$$

This implies that

$$\begin{aligned} & [f'_i(x_i(k)) - f'_j(x_j(k))]^2 \\ & \geq \sum_{d \in F_{ij}(k)} [f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B))]^2 \end{aligned}$$

A consequence of this last inequality is that

$$\begin{aligned} & \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2 \\ & \geq \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} \sum_{d \in F_{ij}(k)} (f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)))^2 \\ & = \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{d \in D(k)} (f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)))^2 \\ & = \sum_{d=1}^{n-1} (f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)))^2, \end{aligned}$$

where the final equality used the fact that every $d \in \{1, \dots, n-1\}$ such that $f'_d(x_d(\ell B)) - f'_{d+1}(x_{d+1}(\ell B)) \neq 0$ appears in some $D(k)$, which is a restatement of Lemma 2. $\square$

Our next lemma is a refinement of Lemma 4 which yields more convenient bounds.

Lemma 5.

$$\begin{aligned} & \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2 \\ & \geq \frac{1}{n^2} \sum_{i=1}^n (f'_i(x_i(\ell B)) - f'_i(x_i^*))^2. \end{aligned} \quad (10)$$

Proof. By Lemma 4, if we relabel the nodes so that $f'_1(x_1(\ell B)), f'_2(x_2(\ell B)), \dots, f'_n(x_n(\ell B))$ is non-increasing,

$$\begin{aligned} & \frac{\sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2}{\sum_{i=1}^n (f'_i(x_i(\ell B)) - f'_i(x_i^*))^2} \\ & \geq \frac{\sum_{i=1}^{n-1} (f'_i(x_i(\ell B)) - f'_{i+1}(x_{i+1}(\ell B)))^2}{\sum_{i=1}^n (f'_i(x_i(\ell B)) - f'_i(x_i^*))^2}. \end{aligned}$$

Let $q = f'_1(x_1^*)$; by Proposition 1, we have that $q = f'_i(x_i^*)$ for all $i = 1, \dots, n$. Define $g_i(z) = f'_i(z) - qz$. We can then rewrite the above inequality as

$$\begin{aligned} & \frac{\sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2}{\sum_{i=1}^n (f'_i(x_i(\ell B)) - f'_i(x_i^*))^2} \\ & \geq \frac{\sum_{i=1}^{n-1} (g'_i(x_i(\ell B)) - g'_{i+1}(x_{i+1}(\ell B)))^2}{\sum_{i=1}^n g_i'^2(x_i(\ell B))}. \end{aligned}$$

Clearly, the sequence $g'_1(x_1(\ell B)), \dots, g'_n(x_n(\ell B))$ is in nonincreasing order. It therefore follows that

$$\begin{aligned} & \frac{\sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \bar{\mathcal{E}}(k)} (f'_i(x_i(k)) - f'_j(x_j(k)))^2}{\sum_{i=1}^n (f'_i(x_i(\ell B)) - f'_i(x_i^*))^2} \\ & \geq \min_{s_1 \geq s_2 \geq \dots \geq s_n} \frac{\sum_{i=1}^{n-1} (s_i - s_{i+1})^2}{\sum_{i=1}^n s_i^2}. \end{aligned} \quad (11)$$

Lemma 5 of [15] shows that the right hand side is at least $1/n^2$. This immediately implies the lemma. $\square$

We now turn to the statement and proof of our main result. We will use R_0 to denote a measure of initial distance to an optimal solution defined as,

$$R_0 = \sup_{\mathbf{x} \in \mathcal{S}: F(\mathbf{x}) \leq F(\mathbf{x}^*)} \sup_{\mathbf{x}^* \in \mathcal{X}^*} \|\mathbf{x} - \mathbf{x}^*\|.$$

In words, R_0 is the largest distance to the set of optimal solutions from any point whose objective not larger than the objective at $\mathbf{x}(0)$. Note that R_0 may not be finite, in which case part of our result below will be vacuously true.

Our main result is then the following theorem.

Theorem 1.

$$F(\mathbf{x}(k)) - F(\mathbf{x}^*) \leq \frac{8LR_0^2 n^2}{\lfloor k/B \rfloor}, \quad (12)$$

where $L = \max_{i=1, \dots, n} L_i$ and $\lfloor z \rfloor$ denotes the largest integer which is at most z .

Furthermore, if all $f_i(\cdot)$ are μ -strongly convex² for some $\mu \geq 0$, then we also have

$$F(\mathbf{x}(k)) - F(\mathbf{x}^*) \leq \left(1 - \frac{\mu}{4Ln^2}\right)^{\lfloor k/B \rfloor} (F(\mathbf{x}(0)) - F(\mathbf{x}^*)). \quad (13)$$

Proof. By Lemma 3 we have,

$$\begin{aligned} & F(\mathbf{x}((\ell+1)B)) \\ & \leq F(\mathbf{x}(\ell B)) - \sum_{k=\ell B}^{(\ell+1)B-1} \sum_{(i,j) \in \mathcal{E}(k)} \frac{(f'_j(x_j(k)) - f'_i(x_i(k)))^2}{4(L_i + L_j)} \\ & \leq F(\mathbf{x}(\ell B)) - \sum_{i=1}^n \frac{(f'_i(x_i(\ell B)) - f'_i(x_i^*))^2}{8Ln^2}, \end{aligned} \quad (14)$$

where the last step is due to Lemma 5 and the inequality $L_i + L_j \leq 2L$.

Next, since F is convex we have

$$\begin{aligned} F(\mathbf{x}^*) - F(\mathbf{x}(\ell B)) & \geq \langle \nabla F(\mathbf{x}(\ell B)), \mathbf{x}^* - \mathbf{x}(\ell B) \rangle \\ & = \langle \nabla F(\mathbf{x}(\ell B)) - \nabla F(\mathbf{x}^*), \mathbf{x}^* - \mathbf{x}(\ell B) \rangle, \end{aligned}$$

where the last equality follows since, by Proposition 1, the components of $\nabla F(\mathbf{x}^*)$ are identical and since $\mathbf{x}(\ell B), \mathbf{x}^* \in \mathcal{S}$, we have that the entries of $\mathbf{x}^* - \mathbf{x}(\ell B)$ sum to zero. Next, negating both sides of the above equation and using Cauchy-Schwarz,

$$F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*) \leq R_0 \|\nabla F(\mathbf{x}(\ell B)) - \nabla F(\mathbf{x}^*)\|_2, \quad (15)$$

²A function $w : \mathbb{R}^m \rightarrow \mathbb{R}$ is called μ -strongly convex if $w(u) \geq w(v) + w'(v)^T(u-v) + (\mu/2)\|u-v\|_2^2$ for all $u, v \in \mathbb{R}^m$.

where we used that $F(\mathbf{x}(k))$ is nonincreasing. Combining Eqs. (14) and (15) we have

$$\begin{aligned} & F(\mathbf{x}((\ell+1)B)) - F(\mathbf{x}^*) \\ & \leq F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*) - \sum_{i=1}^n \frac{(f'_i(x_i(\ell B)) - f'_i(x_i^*))^2}{8Ln^2} \\ & \leq F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*) - \frac{(F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*))^2}{8Ln^2 R_0^2}. \end{aligned} \quad (16)$$

We now show the last inequality implies Eq. (12) via some standard equation manipulations. Letting $\Delta(k) = F(\mathbf{x}(k)) - F(\mathbf{x}^*)$, note that $\Delta(k)$ is nonincreasing by Lemma 3. We have just shown

$$\Delta((\ell+1)B) \leq \Delta(\ell B) - \frac{\Delta^2(\ell B)}{8LR_0^2 n^2}.$$

Dividing both sides of this by $\Delta((\ell+1)B)\Delta(\ell B)$ and rearranging, we obtain

$$\begin{aligned} \frac{1}{\Delta((\ell+1)B)} & \geq \frac{1}{\Delta(\ell B)} + \frac{1}{8LR_0^2 n^2} \frac{\Delta(\ell B)}{\Delta((\ell+1)B)} \\ & \geq \frac{1}{\Delta(\ell B)} + \frac{1}{8LR_0^2 n^2}, \end{aligned}$$

where we used the monotonicity of $\Delta(k)$. Summing up this inequality up over $\ell = 0, \dots, t-1$, we obtain $F(\mathbf{x}(tB)) - F(\mathbf{x}^*) \leq \frac{8LR_0^2 n^2}{t}$ and using the monotonicity of $F(\mathbf{x}(k))$ we obtain Eq. (12).

Turning now to Eq. (13), let us define as before $g_i(x) = f_i(x) - qx$ where $q = f'_1(x_1^*)$, and further let $G(\mathbf{x}) = \sum_{i=1}^n g_i(x_i)$. Observe that $G(\mathbf{x})$ is μ -strongly convex and has global minimizer at $\mathbf{x}^*$. Consequently if $\mathbf{x} \in \mathcal{S}$,

$$\begin{aligned} \sum_{i=1}^n (f'_i(x_i) - f'_i(x_i^*))^2 & = \|\nabla G(\mathbf{x})\|_2^2 \geq 2\mu(G(\mathbf{x}) - G(\mathbf{x}^*)) \\ & = 2\mu(F(\mathbf{x}) - F(\mathbf{x}^*)), \end{aligned}$$

where the final equality used the fact that the sum of the entries of $\mathbf{x}$ and $\mathbf{x}^*$ is the same since both are in $\mathcal{S}$. Thus from Eq. (16),

$$\begin{aligned} & F(\mathbf{x}((\ell+1)B)) - F(\mathbf{x}^*) \\ & \leq F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*) - \sum_{i=1}^n \frac{2\mu(F(\mathbf{x}(\ell B)) - F(\mathbf{x}^*))}{8Ln^2} \end{aligned}$$

which immediately implies Eq. (13). $\square$

Remark 1. Note that although Eq. (1) does not have constraints on the variables x_i , for certain functions $f_i(x_i)$ our algorithm automatically solves a constrained version of the problem. For example, if the initial conditions $x_i(0)$ are all nonnegative and $f'_i(0) = f'_j(0)$ for all i, j , then by Lemma 1 the constraint $x_i \geq 0$ will automatically be satisfied throughout the execution of the gradient balancing

Figure 2: Convergence time for gradient balancing as a function of the number of nodes for the line graph.

Figure 3: Convergence time for gradient balancing as a function of the number of nodes for the lollipop graph.

method. In other words, the constraints $x_i \geq 0$ can be added “for free.” The condition on the functions $f_i(x)$ is somewhat restrictive, but admissible f_i include, for example, all polynomials with non-negative coefficients whose linear coefficient is zero.

5. A simulation

We now describe a simulation of the gradient balancing protocol on some particular graphs. We consider local objective functions $f_i(x_i) = w_i(x_i - a_i)^4$ where the nonnegative coefficient w_i and the coefficient a_i are chosen uniformly on $[0, 1]$. We set $K = 0$. We show simulations of the line and lollipop graphs in Figure 2 and Figure 3, respectively, where we plot the first time $F(\mathbf{x}(k)) - F(\mathbf{x}^*) < 1/100$ on the y-axis. The figures appear to be broadly consistent with the quadratic bound of Theorem 1.

6. Concluding remarks

We have proposed a new deterministic, distributed method for the network resource allocation problem with convergence time that scales quadratically in n on time-varying undirected graphs. The

main open questions left by this work include improving the convergence time and obtaining similar convergence rate guarantees in the distributed setting with local constraints at each node.

References

- [1] K. J. Arrow and L. Hurwicz. Decentralization and computation in resource allocation. In *Essays in Economics and Econometrics*, pages 34–104. University of North Carolina Press, 1960.
- [2] A. Beck. The 2-coordinate descent method for solving double-sided simplex constrained minimization problems. *Journal of Optimization Theory and Applications*, 162(3):892–919, 2014.
- [3] A. Beck, A. Nedic, A. Ozdaglar, and M. Teboulle. An $O(1/k)$ gradient method for network resource allocation problems. *IEEE Transactions on Control of Network Systems*, 1(1):2325–5870, 2014.
- [4] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation*. Athena Press, 1989.
- [5] G. Cavarero, R. Carli, and S. Zampieri. A distributed control algorithm for the minimization of the power generation cost in smart micro-grid. Technical report, 2014.
- [6] A. Cherukuri and J. Cortes. Initialization-free distributed coordination for economic dispatch under varying loads and generator commitment. Technical report, 2014.
- [7] A. Cherukuri and J. Cortes. Distributed generator coordination for initialization and anytime optimization in economic dispatch. *IEEE Transactions on Control of Network Systems*, 2015.
- [8] Y. C. Ho, L. D. Servi, and R. Suri. A class of center-free resource allocation algorithms. *Large Scale Systems*, 1:51–62, 1980.
- [9] M. Ismail, A. Abdrabou, and W. Zhang. Cooperative decentralized resource allocation in heterogeneous wireless access medium. *IEEE Transactions on Wireless Communications*, 12(2):714–724, 2012.
- [10] J.F. Kurose and R.A. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Trans. on Computers*, 38:705–717, 1989.
- [11] H. Lakshmanan and D.P. de Farias. Decentralized resource allocation in dynamic networks of agents. *SIAM Journal on Optimization*, 19(2):911–940, 2008.
- [12] I. Necoara. Random coordinate descent algorithms for multi-agent convex optimization over networks. *IEEE Trans. on Automatic Control*, 58(8):2001–2012, 2013.
- [13] I. Necoara. Distributed and parallel random coordinate descent methods for convex programming over big data networks. Technical report, 2015.
- [14] I. Necoara, Y. Nesterov, and F. Glineur. A random coordinate descent method on large-scale optimization problems with linear constraints. In *Proceedings of IC-COPT 2013*, pages 1–25, 2011.
- [15] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quantization effect. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.
- [16] L. Xiao and S. Boyd. Optimal scaling of a gradient method for distributed resource allocation. *J. of Optimization Theory and Appl.*, 129(3):469–488, 2006.