

UC Office of the President

Recent Work

Title

Resilient Distributed Optimization Algorithms for Resource Allocation

Permalink

<https://escholarship.org/uc/item/5b0455mw>

Authors

Uribe, Cesar

Wai, Hoi-To

Alizadeh, Mahnoosh

Publication Date

2019-09-01

Peer reviewed

Resilient Distributed Optimization Algorithms for Resource Allocation

César A. Uribe[†], Hoi-To Wai[†], Mahnoosh Alizadeh^{*}

September 11, 2019

Abstract

Distributed algorithms provide flexibility over centralized algorithms for resource allocation problems, e.g., cyber-physical systems. However, the distributed nature of these algorithms often makes the systems susceptible to man-in-the-middle attacks, especially when messages are transmitted between price-taking agents and a central coordinator. We propose a resilient strategy for distributed algorithms under the framework of primal-dual distributed optimization. We formulate a robust optimization model that accounts for Byzantine attacks on the communication channels between agents and coordinator. We propose a resilient primal-dual algorithm using state-of-the-art robust statistics methods. The proposed algorithm is shown to converge to a neighborhood of the robust optimization model, where the neighborhood's radius is proportional to the fraction of attacked channels.

1 Introduction

Consider the following multi-agent optimization problem involving the average of parameters in the constraints:

$$\begin{aligned} \min_{\boldsymbol{\theta}_i \in \mathbb{R}^d, \forall i} \quad & U(\boldsymbol{\theta}) := \frac{1}{N} \sum_{i=1}^N U_i(\boldsymbol{\theta}_i) \\ \text{s.t.} \quad & g_t \left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i \right) \leq 0, \quad t = 1, \dots, T, \\ & \boldsymbol{\theta}_i \in \mathcal{C}_i, \quad i = 1, \dots, N, \end{aligned} \tag{1}$$

where both $U_i : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g_t : \mathbb{R}^d \rightarrow \mathbb{R}$ are continuously differentiable, convex functions, and \mathcal{C}_i is a compact convex set in \mathbb{R}^d . We let $\mathbf{0} \in \mathcal{C}_i$ and

$$\max_{\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathcal{C}_i} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\| \leq R, \quad i = 1, \dots, N, \tag{2}$$

such that R is an upper bound on the diameters of \mathcal{C}_i .

^{*}CAU and HTW have contributed equally. CAU is with LIDS, MIT, Cambridge, MA, USA. HTW is with Dept. of SEEM, CUHK, Shatin, Hong Kong. MA is with Dept. of ECE, UCSB, Santa Barbara, CA, USA. This work is partially supported by UCOP Grant LFR-18-548175 and CUHK Direct Grant #4055113. E-mails: cauribe@mit.edu, htwai@se.cuhk.edu.hk, alizadeh@ucsb.edu

Problem (1) arises in many resource allocation problems with a set of potentially *nonlinear* constraints on the amount of allowable resources, see Section 1.1 for a detailed exploration.

We consider a system where there exists a central coordinator and N agents. In this context, the function $U_i(\boldsymbol{\theta}_i)$ and parameter $\boldsymbol{\theta}_i$ are the utility of the i th agent and the resource controlled by agent i , respectively. As the agents work independently, it is desirable to design algorithms that allow the N agents to solve (1) cooperatively through communication with the central coordinator. Among others, the primal-dual optimization methods [1] have been advocated as they naturally give rise to decomposable algorithms that favor distributed implementation [2]. In addition to their practical success, these methods are supported by strong theoretical guarantees where fast convergence to an optimal solution of (1) is well established. However, the distributed nature of these methods also exposes the system to vulnerabilities not faced by traditional centralized systems. Precisely, existing algorithms assume the agents, and the communication links between central server and agents, to be *completely trustworthy*. However, an attacker can take over a sub-system operated by the agents, and deliberately edit the messages in these communication links, *i.e.*, a Byzantine attack. This might result in an unstable system with possible damages to hardware and the system overall.

In this paper, we propose strategies for securing primal-dual distributed algorithms, e.g., in [1], tailored to solving a relaxed version of the resource allocation problem (1). A key observation is that the existing algorithms depend on reliably computing the average of a set of parameter vectors, $\{\boldsymbol{\theta}_i\}_{i=1}^N$, transmitted by the agents. As a remedy, we apply robust statistics techniques as a subroutine, therefore proposing a *resilient* distributed algorithm that is proven to converge to a neighborhood of the optimal solution of a robust version of (1).

Vulnerabilities of various types of distributed algorithms have been identified and addressed in a number of recent studies. Relevant examples are [3–7] which study secure decentralized algorithms on a general network topology but consider consensus-based optimization models. Moreover, [8–10] consider a similar optimization architecture as this paper, yet they focus on securing distributed algorithms for machine learning tasks which assumes i.i.d. functions, a fundamentally different setting from the current paper. Our work is also related to the literature on robust statistics [11, 12], and particularly, with the recently rekindled research efforts on high dimensional robust statistics [13–15]. These works will be the working horse for our attack resilient algorithm.

Our contributions and organization are as follows. First, we derive a formal model for attack resilient resource allocation via a conservative approximation for the robust optimization problem [cf. Section 3]. Second, we apply and derive new robust estimation results to secure distributed resource allocation algorithms [cf. Section 4]. Third, we provide a non-asymptotic convergence guarantee of the proposed attack resilient algorithm [cf. Section 4.1]. In particular, our algorithm is shown to converge to a $\mathcal{O}(\alpha^2)$ neighborhood to the optimal solution of (1), where $\alpha \in [0, \frac{1}{2})$ is the fraction of attacked links.

Notations. Unless otherwise specified, $\|\cdot\|$ denotes the standard Euclidean norm. For any $N \in \mathbb{N}$, $[N]$ denotes the finite set $\{1, \dots, N\}$.

1.1 Motivating Examples

Our set-up here can be employed in a wide range of optimization problems for resource allocation and networked control in multi-agent systems, e.g., in the pioneering example of congestion control in data networks [16, 17]; in determining the optimal price of electricity and enabling more efficient demand supply balancing (a.k.a. demand response) in smart power distribution systems [18, 19]; in managing user transmit powers and data rates in wireless cellular networks [20]; in determining optimal caching policies by content delivery networks [21]; in optimizing power consumption in wireless sensor networks with energy-restricted batteries [22, 23]; and in designing congestion control systems in urban traffic networks [24]. These examples would have different utility functions and constraint sets that can be handled through our general formulation in (1). For example, in the power/rate control problem in data networks, the cost functions are usually logarithmic functions associated with rate θ_i , e.g., $U_i(\theta_i) = -\beta_i \log(\theta_i)$. In demand response applications in power distribution systems, the utilities capture the users' benefits from operating their electric appliances under different settings. For example, we can capture the cost function of temperature θ_i controlled by a price-responsive air conditioner as $U_i(\theta_i) = b_i(\theta_i - \theta_{\text{comf}})^2 - c_i$ [19]. In terms of constraints, our general nonlinear constraint formulation can not only capture common linear resource constraints such as link capacity in data networks [16, 17], but can also handle important non-linear constraints arising in many different applications. For example, in radial power distribution systems, nonlinear convexified power flow constraints can be included for distributed demand response optimization (to see a description of distribution system power flow constraints, see, e.g., [25, 26]). This can enable our algorithm to perform demand supply balancing in power distribution systems in a *distributed* and *resilient* fashion.

2 Primal-dual Algorithm for Resource Allocation

This section reviews the basic primal-dual algorithm for resource allocation. Let $\boldsymbol{\lambda} \in \mathbb{R}_+^T$ be the dual variable. We consider the Lagrangian function of (1):

$$\mathcal{L}(\{\boldsymbol{\theta}_i\}_{i=1}^N; \boldsymbol{\lambda}) := \frac{1}{N} \sum_{i=1}^N U_i(\boldsymbol{\theta}_i) + \sum_{t=1}^T \lambda_t g_t\left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i\right).$$

Assuming strong duality holds (e.g., under the Slater's condition), solving problem (1) is equivalent to solving its dual problem:

$$\max_{\boldsymbol{\lambda} \in \mathbb{R}_+^T} \min_{\boldsymbol{\theta}_i \in \mathcal{C}_i, \forall i} \mathcal{L}(\{\boldsymbol{\theta}_i\}_{i=1}^N; \boldsymbol{\lambda}). \quad (\text{P})$$

For a given $\boldsymbol{\lambda}$, the inner minimization of (P) is known as the Lagrangian relaxation of (1), which can be interpreted as a *penalized* resource allocation problem [19].

In a distributed setting, the goal is to solve (1) where the agents only observe a *pricing signal* received from the central coordinator, and this pricing signal is to be updated iteratively at the central coordinator. As suggested in [1], we apply the primal-dual algorithm (PDA) to a regularized version of (P). Let us define

$$\begin{aligned} \mathcal{L}_v(\{\boldsymbol{\theta}_i\}_{i=1}^N; \boldsymbol{\lambda}) := \\ \mathcal{L}(\{\boldsymbol{\theta}_i\}_{i=1}^N; \boldsymbol{\lambda}) + \frac{v}{2N} \sum_{i=1}^N \|\boldsymbol{\theta}_i\|^2 - \frac{v}{2} \|\boldsymbol{\lambda}\|^2, \end{aligned} \quad (3)$$

Algorithm 1 PD-DRA Procedure.

1: **for** $k = 1, 2, \dots$ **do**

2: (*Message exchanges stage*):

(a) Central coordinator receives $\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N$ from agents and computes $\bar{\boldsymbol{\theta}}^{(k)}, \{\nabla_{\boldsymbol{\theta}} g_t(\bar{\boldsymbol{\theta}}^{(k)})\}_{t=1}^T$.

(b) Central coordinator broadcasts the vectors $\bar{\boldsymbol{\theta}}^{(k)}, \bar{\boldsymbol{g}}^{(k)} := \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\boldsymbol{\theta}} g_t(\bar{\boldsymbol{\theta}}^{(k)})$ to agents.

3: (*Computation stage*):

(a) Agent i computes the update for $\boldsymbol{\theta}_i^{(k+1)}$ according to (4a) using the received $\bar{\boldsymbol{\theta}}^{(k)}$.

(b) The central coordinator computes the update for $\boldsymbol{\lambda}^{(k+1)}$ according to (4b).

4: **end for**

such that $\mathcal{L}_v(\cdot)$ is v -strongly convex and v -strongly concave in $\{\boldsymbol{\theta}_i\}_{i=1}^N$ and $\boldsymbol{\lambda}$, respectively. Let $k \in \mathbb{Z}_+$ be the iteration index, $\gamma > 0$ be the step sizes, the PDA recursion is described by:

$$\boldsymbol{\theta}_i^{(k+1)} = \tag{4a}$$

$$\mathcal{P}_{\mathcal{C}_i}(\boldsymbol{\theta}_i^{(k)} - \gamma \nabla_{\boldsymbol{\theta}_i} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N; \boldsymbol{\lambda}^{(k)})), \forall i \in [N]$$

$$\boldsymbol{\lambda}^{(k+1)} = [\boldsymbol{\lambda}^{(k)} + \gamma \nabla_{\boldsymbol{\lambda}} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N; \boldsymbol{\lambda}^{(k)})]_+ \tag{4b}$$

where $\mathcal{P}_{\mathcal{C}_i}(\cdot)$ is the Euclidean projection operator, $[\cdot]_+$ denotes $\max\{0, \cdot\}$, and the gradients are:

$$\nabla_{\boldsymbol{\theta}_i} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N; \boldsymbol{\lambda}^{(k)}) = \frac{1}{N} \left(\nabla_{\boldsymbol{\theta}_i} U_i(\boldsymbol{\theta}_i^{(k)}) + v \boldsymbol{\theta}_i^{(k)} \right) \tag{5}$$

$$[\nabla_{\boldsymbol{\lambda}} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N; \boldsymbol{\lambda}^{(k)})]_t = g_t \left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i^{(k)} \right) - v \lambda_t^{(k)}, \tag{6}$$

for all i, t . We denoted $[\boldsymbol{x}]_t$ as the t th element of $\boldsymbol{x} \in \mathbb{R}^T$. In particular, observe that (4) performs a projected gradient descent/ascent on the primal/dual variables.

From the above, both gradients with respect to (w.r.t.) $\boldsymbol{\theta}_i$ and λ_t depend only on the average parameter $\bar{\boldsymbol{\theta}}^{(k)} := \frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i^{(k)}$. We summarize the primal dual distributed resource allocation (PD-DRA) procedure in Algorithm 1. In addition to solving the general problem (1), Algorithm 1 also serves as a general solution method to popular resource allocation problems [19].

As the regularized primal-dual problem is strongly convex/concave in primal/dual variables, Algorithm 1 converges linearly to an optimal solution [1]. To study this, let us denote $\boldsymbol{z}^{(k)} = (\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N, \boldsymbol{\lambda}^{(k)})$ as the primal-dual variable at the k th iteration,

$$\boldsymbol{\Phi}(\boldsymbol{z}^{(k)}) := \begin{pmatrix} \nabla_{\boldsymbol{\theta}} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N, \boldsymbol{\lambda}^{(k)}) \\ \nabla_{\boldsymbol{\lambda}} \mathcal{L}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i=1}^N, \boldsymbol{\lambda}^{(k)}) \end{pmatrix}. \tag{7}$$

Fact 1. [1, Theorem 3.5] Assume that the map $\boldsymbol{\Phi}(\boldsymbol{z}^{(k)})$ is $L_{\boldsymbol{\Phi}}$ Lipschitz continuous. For all $k \geq 1$, we have

$$\|\boldsymbol{z}^{(k+1)} - \boldsymbol{z}^*\|^2 \leq (1 - 2\gamma v + \gamma^2 L_{\boldsymbol{\Phi}}^2) \|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2, \tag{8}$$




Figure 1: Illustrating the PD-DRA algorithm under attack. The uplink for agent j is compromised such that the correct $\theta_j^{(k)}$ is not transmitted to the central node. The up/downlink for agent i are operating properly.

where \mathbf{z}^* is a saddle point to the regularized version of (P). Setting $\gamma = v/L_{\Phi}^2$ gives $\|\mathbf{z}^{(k+1)} - \mathbf{z}^*\|^2 \leq (1 - v^2/L_{\Phi}^2)\|\mathbf{z}^{(k)} - \mathbf{z}^*\|^2, \forall k \geq 1$.

3 Problem Formulation

Despite the simplicity and the strong theoretical guarantee, the PD-DRA method is susceptible to attacks on the channels between the central coordinator and the agents, as described below.

Attack Model. We consider a situation when *uplink* channels between agents and the central coordinator are compromised [see Fig. 1]. Let $\mathcal{A} \subset [N]$ be the set of *compromised uplink channels*, whose identities are unknown to the central coordinator. We define $\mathcal{H} := [N] \setminus \mathcal{A}$ as the set of trustworthy channels. At iteration k , instead of receiving $\theta_i^{(k)}$ from each agent $i \in [N]$ [cf. Step 2(a)], the central coordinator receives the following messages:

$$\mathbf{r}_i^{(k)} = \begin{cases} \theta_i^{(k)}, & \text{if } i \in \mathcal{H}, \\ \mathbf{b}_i^{(k)}, & \text{if } i \in \mathcal{A}. \end{cases} \quad (9)$$

We focus on a Byzantine attack scenario such that the messages, $\mathbf{b}_i^{(k)}$, communicated on the attacked channels can be arbitrary. Under such scenario, if the central coordinator forms the naive average $\hat{\boldsymbol{\theta}}^{(k)} = 1/N \sum_{i=1}^N \mathbf{r}_i^{(k)}$ and computes the gradients $\nabla g_t(\hat{\boldsymbol{\theta}}^{(k)})$ accordingly, this may result in uncontrollable error since the deviation $\hat{\boldsymbol{\theta}}^{(k)} - (1/N) \sum_{i=1}^N \theta_i^{(k)}$ can be arbitrarily large. It is anticipated that the PD-DRA method would not provide a solution to the regularized version of (P).

Robust Optimization Model. In light of the Byzantine attack, it is impossible to optimize the original problem (P) since the contribution from $U_i(\cdot) : i \in \mathcal{A}$ becomes unknown to the central coordinator. As a compromise, we focus on optimizing the cost function of agents with trustworthy uplinks and the following robust optimization problem as our target model:

$$\min_{\boldsymbol{\theta}_i \in \mathcal{C}_i, i \in \mathcal{H}} \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} U_i(\boldsymbol{\theta}_i) \quad (10a)$$

$$\text{s.t.} \quad \max_{\boldsymbol{\theta}_j \in \mathcal{C}_j, j \in \mathcal{A}} g_t \left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i \right) \leq 0, \quad \forall t, \quad (10b)$$

note that $\{\boldsymbol{\theta}_j\}_{j \in \mathcal{A}}$ is taken away from the decision variables and we have included (10b) to account for the *worst case* scenario for the resource usage of the agents with compromised uplinks. This is to ensure that the physical operation limit of the system will not be violated under attack. Consider the following assumption which will be assumed throughout the paper:

H1. For all $\boldsymbol{\theta} \in \mathbb{R}^d$, the gradient of g_t is bounded with $\|\nabla g_t(\boldsymbol{\theta})\| \leq B$ and is L -Lipschitz continuous.

We define

$$\bar{g}_t(\boldsymbol{\theta}) := g_t(\boldsymbol{\theta}) + \frac{|\mathcal{A}|}{N}(RB + \frac{1}{2}LR^2), \quad (11)$$

Lemma 1. Under H1. The following problem yields a conservative approximation of (10), i.e., its feasible set is a subset of the feasible set of (10):

$$\begin{aligned} \min_{\boldsymbol{\theta}_i \in \mathcal{C}_i, i \in \mathcal{H}} \quad & \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} U_i(\boldsymbol{\theta}_i) \\ \text{s.t.} \quad & \bar{g}_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) \leq 0, \quad \forall t \in [T], \end{aligned} \quad (12)$$

Similar to PD-DRA, we define the regularized Lagrangian function of (12) as:

$$\begin{aligned} \bar{\mathcal{L}}_v(\{\boldsymbol{\theta}_i\}_{i \in \mathcal{H}}; \boldsymbol{\lambda}; \mathcal{H}) \\ := \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} U_i(\boldsymbol{\theta}_i) + \sum_{t=1}^T \lambda_t \bar{g}_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) \\ + \frac{v}{2|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|\boldsymbol{\theta}_i\|^2 - \frac{v}{2} \|\boldsymbol{\lambda}\|^2. \end{aligned} \quad (13)$$

Again, the regularized Lagrangian function is v -strongly convex and concave in $\boldsymbol{\theta}$ and $\boldsymbol{\lambda}$, respectively.

Our main task is to tackle the following modified problem of (P) under Byzantine attack on (some of) the uplinks:

$$\max_{\boldsymbol{\lambda} \in \mathbb{R}_+^T} \min_{\boldsymbol{\theta}_i \in \mathcal{C}_i, \forall i \in \mathcal{H}} \bar{\mathcal{L}}_v(\{\boldsymbol{\theta}_i\}_{i \in \mathcal{H}}; \boldsymbol{\lambda}; \mathcal{H}), \quad (\text{P}')$$

and we let $\hat{\boldsymbol{z}}^* = (\hat{\boldsymbol{\theta}}^*, \hat{\boldsymbol{\lambda}}^*)$ be the optimal solution to (P'). Notice that (P') bears a similar form as (P) and thus one may apply the PD-DRA method to the former naturally. However, such application requires the central coordinator to compute the sample average

$$\bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} := \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i^{(k)}, \quad (14)$$

at each iteration. However, the above might not be computationally feasible under the attack model, since the central coordinator is oblivious to the identity of \mathcal{H} . This is the main objective in the design of our scheme.

4 Robust Distributed Resource Allocation

In this section, we describe two estimators for approximating $\bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ [cf. (14)] from the received messages (9) without knowing the identity of links in \mathcal{H} . To simplify notations, we define $\alpha \geq |\mathcal{A}|/N$ as a known upper bound to the fraction of compromised channels and assume $\alpha < 1/2$ where less than half of the channels are compromised.

As discussed after (14), the problem at hand is *robust mean estimation*, whose applications to robust distributed optimization has been considered in the machine learning literature [9, 10, 14] under the assumption that the ‘trustworthy’ signals are drawn i.i.d. from a Gaussian distribution. Our setting is different since the signals $\theta_i^{(k)}$, $i \in \mathcal{H}$ are variables from the previous iteration whose distribution is non-Gaussian in general. Our analysis will be developed without such assumption on the distribution.

We first consider a simple median-based estimator applied to each coordinate $j = 1, \dots, d$. First, define the coordinate-wise median as:

$$[\theta_{\text{med}}^{(k)}]_j = \text{med}(\{[r_i^{(k)}]_j\}_{i=1}^N), \quad (15)$$

where $\text{med}(\cdot)$ computes the median of the operand. Then, our estimator is computed as the mean of the nearest $(1 - \alpha)N$ neighbors of $[\theta_{\text{med}}^{(k)}]_j$. To formally describe this, let us define:

$$\mathcal{N}_j^{(k)} = \{i \in [N] : |[r_i^{(k)}]_j - [\theta_{\text{med}}^{(k)}]_j| \leq r_j^{(k)}\}, \quad (16)$$

where $r_j^{(k)}$ is chosen as $|\mathcal{N}_j^{(k)}| = (1 - \alpha)N$. Our estimator is:

$$[\hat{\theta}_{\mathcal{H}}^{(k)}]_j = \frac{1}{(1-\alpha)N} \sum_{i \in \mathcal{N}_j^{(k)}} [r_i^{(k)}]_j. \quad (17)$$

The following bounds the performance of (17).

Proposition 1. *Suppose that $\max_{i \in \mathcal{H}} \|\theta_i^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\|_{\infty} \leq r$, then for any $\alpha \in (0, \frac{1}{2})$, it holds that*

$$\|\hat{\theta}_{\mathcal{H}}^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\| \leq \frac{\alpha}{1-\alpha} \left(2 + \sqrt{\frac{(1-\alpha)^2}{1-2\alpha}}\right) r \sqrt{d}. \quad (18)$$

Under mild assumptions, the condition $\max_{i \in \mathcal{H}} \|\theta_i^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\|_{\infty} \leq r$ can be satisfied with $r = \Theta(R)$, as implied by the compactness of \mathcal{C}_i [cf. (2)]. Moreover, for sufficiently small α , the right hand side on (18) can be approximated by $\mathcal{O}(\alpha R \sqrt{d})$. However, this median-based estimator may perform poorly for large α (especially when $\alpha \rightarrow 1/2$) or dimension d . For these situations, a more sophisticated estimator is required, as detailed next.

To derive the second estimator, we apply an auxiliary result from [15] which provides an algorithm for estimating $\bar{\theta}_{\mathcal{H}}^{(k)}$, as summarized in Algorithm 2. We observe:

Proposition 2. [15, Proposition 16] *Suppose that $\lambda_{\max}(\frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} (\theta_i^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)})(\theta_i^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)})^{\top}) \leq \sigma^2$. For any $\alpha \in [0, \frac{1}{4})$, Algorithm 2 produces an output such that $\|\bar{\theta}_{\mathcal{H}}^{(k)} - \hat{\theta}_{\mathcal{H}}^{(k)}\| = \mathcal{O}(\sigma \sqrt{\alpha})$.*

Again, similar to Proposition 1, the required condition above can be satisfied with $\sigma = \Theta(R)$ under mild conditions. Thus, Proposition 2 states that Algorithm 2 recovers $\bar{\theta}_{\mathcal{H}}^{(k)}$ up to an error of $\mathcal{O}(\sqrt{\alpha}R)$. Note that this bound is dimension free unlike the median estimator analyzed in Proposition 1.

The idea behind Algorithm 2 is to sequentially identify and remove the subset of points that cannot be re-constructed from the mean of the data points. The solution of the optimization problem in

Algorithm 2 Recovering the mean of a set [15].

- 1: **Input:** α , $\boldsymbol{\theta}_i^{(k)}$, $c_i = 1$ for all $i = 1, \dots, N$, and $\mathcal{B} = \{1, \dots, N\}$.
- 2: Set $X_{\mathcal{B}} = [\dots \boldsymbol{\theta}_j^{(k)} \dots]^\top$ for $j \in \mathcal{B}$ as the concatenated data matrix.
- 3: Let $Y \in \mathbb{R}^{d \times d}$ and $W \in \mathbb{R}^{\mathcal{B} \times \mathcal{B}}$ be the maximizer/minimizer of the saddle point problem

$$\max_{\substack{Y \succeq 0, \\ \text{tr}(Y) \leq 1}} \min_{\substack{0 \leq W_{ij}, \\ W_{ij} \leq \frac{4-\alpha}{\alpha(2+\alpha)n}, \\ \sum_j W_{ji} = 1}} \sum_{i \in \mathcal{B}} c_i (\boldsymbol{\theta}_i^{(k)} - X_{\mathcal{B}} w_i)^\top Y (\boldsymbol{\theta}_i^{(k)} - X_{\mathcal{B}} w_i)$$

- 4: Let $\tau_i^* = (\boldsymbol{\theta}_i^{(k)} - X_{\mathcal{B}} w_i)^\top Y (\boldsymbol{\theta}_i^{(k)} - X_{\mathcal{B}} w_i)$.
 - 5: **if** $\sum_{i \in \mathcal{B}} c_i \tau_i^* > 4n\sigma^2$ **then**
 - 6: For $i \in \mathcal{B}$, replace c_i with $(1 - \frac{\tau_i^*}{\max_{j \in \mathcal{B}} \tau_j^*}) c_i$.
 - 7: For all i with $c_i < \frac{1}{2}$, remove from \mathcal{B} .
 - 8: Go back to Line 3.
 - 9: **end if**
 - 10: Set W_1 as the result of zeroing out all singular values of W that are greater than 0.9.
 - 11: Set $Z = X_{\mathcal{B}} W_0$ where $W_0 = (W - W_1)(I - W_1)^{-1}$.
 - 12: **if** $\text{rank}(Z) = 1$ **then**
 - 13: **Output:** $\hat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ as average of the columns of $X_{\mathcal{B}}$.
 - 14: **else**
 - 15: **Output:** $\hat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ as a column of Z at random.
 - 16: **end if**
-

Line 3 measures how well can we recover the data points as an average of the other $|\mathcal{H}|$ points. The bounded sample variance assumption guarantees that one can re-construct any element in the set \mathcal{H} from its mean, thus, all such points that introduce a large error, as quantified by c_i can be safely removed. Line 5 quantifies the magnitude of the optimal point of Line 3, and if such value is large, such points that introduce a large error are down-weighted. The process is repeated until the optimal solution of Line 3 is small enough and a low rank approximation of the optimal W can be used to return the sample mean estimate.

Attack Resilient PD-DRA method. The above section provides the enabling tool for developing the resilient PD-DRA method, which we summarize in Algorithm 3. The algorithm behaves similarly as Algorithm 1 applied to (\mathbf{P}') , with the exception that the central coordinator is oblivious to \mathcal{H} , and it uses a robust mean estimator to find an approximate average for the signals sent through the trustworthy links. This approximate value is used to compute the new price signals, and sent back to agents. In particular, the primal-dual updates are described by

$$\boldsymbol{\theta}_i^{(k+1)} = \mathcal{P}_{\mathcal{C}_i}(\boldsymbol{\theta}_i^{(k)} - \frac{\gamma}{N}(\hat{\mathbf{g}}_{\mathcal{H}}^{(k)} + \nabla U_i(\boldsymbol{\theta}_i^{(k)}) + v\boldsymbol{\theta}_i^{(k)})), \quad (19)$$

$$\lambda_i^{(k+1)} = [\lambda_i^{(k)} + \gamma(\bar{g}_t(\frac{|\mathcal{H}|}{N}\hat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}) - v\lambda_t^{(k)})]_+. \quad (20)$$

Lemma 2. *Algorithm 3 is a primal-dual algorithm [1] for (\mathbf{P}') with perturbed gradients:*

$$\hat{\mathbf{g}}_{\boldsymbol{\theta}}^{(k)} = \nabla_{\boldsymbol{\theta}} \bar{\mathcal{L}}_v(\boldsymbol{\theta}^{(k)}; \boldsymbol{\lambda}^{(k)}; \mathcal{H}) + \mathbf{e}_{\boldsymbol{\theta}}^{(k)}, \quad (21a)$$

$$\hat{\mathbf{g}}_{\boldsymbol{\lambda}}^{(k)} = \nabla_{\boldsymbol{\lambda}} \bar{\mathcal{L}}_v(\boldsymbol{\theta}^{(k)}; \boldsymbol{\lambda}^{(k)}; \mathcal{H}) + \mathbf{e}_{\boldsymbol{\lambda}}^{(k)}, \quad (21b)$$

Algorithm 3 Resilient PD-DRA

- 1: **Input:** Each agent has initial state $\boldsymbol{\theta}_i^{(0)}$.
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: *(At the Central Coordinator):*
 - (a) Receives $\{\mathbf{r}_i^{(k)}\}_{i=1}^N$, see (9), from agents.
 - (b) Computes robust mean $\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ using the estimator (17) or Algorithm 2.
 - (c) Broadcasts the vectors $\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ and $\widehat{\mathbf{g}}_{\mathcal{H}}^{(k)} := \sum_{t=1}^T \lambda_t \nabla_{\boldsymbol{\theta}} \overline{g}_t(\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)})$ to agents.
 - (d) Computes the update for $\boldsymbol{\lambda}^{(k+1)}$ with (20).
 - 4: *(At each agent i):*
 - (a) Agent receives $\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}$ and $\widehat{\mathbf{g}}_{\mathcal{H}}^{(k)}$.
 - (b) Agent computes update for $\boldsymbol{\theta}_i^{(k+1)}$ with (19).
 - 5: **end for**
-

where we have used concatenated variable as $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N)$ and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_T)$. Under H1 and assuming that $\lambda_t^{(k)} \leq \bar{\lambda}$ for all k , we have:

$$\|\mathbf{e}_{\boldsymbol{\theta}}^{(k)}\| \leq \bar{\lambda} L T \|\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}\|, \quad (22)$$

$$\|\mathbf{e}_{\boldsymbol{\lambda}}^{(k)}\| \leq B T \|\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}\|. \quad (23)$$

The assumption $\lambda_t^{(k)} \leq \bar{\lambda}$ can be guaranteed since $\overline{g}_t(\frac{|\mathcal{H}|}{N} \widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)})$ is bounded.

4.1 Convergence Analysis

Finally, based on Lemma 2, we can analyze the convergence of Algorithm 3. Let $\widehat{\mathbf{z}}^* = (\widehat{\boldsymbol{\theta}}^*, \widehat{\boldsymbol{\lambda}}^*)$ be a saddle point of (P') and define

$$\overline{\boldsymbol{\Phi}}(\mathbf{z}^{(k)}) := \begin{pmatrix} \nabla_{\boldsymbol{\theta}} \overline{\mathcal{L}}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i \in \mathcal{H}}, \boldsymbol{\lambda}^{(k)}; \mathcal{H}) \\ -\nabla_{\boldsymbol{\lambda}} \overline{\mathcal{L}}_v(\{\boldsymbol{\theta}_i^{(k)}\}_{i \in \mathcal{H}}, \boldsymbol{\lambda}^{(k)}; \mathcal{H}) \end{pmatrix}, \quad (24)$$

Theorem 1. *Assume the map $\overline{\boldsymbol{\Phi}}(\mathbf{z}^{(k)})$ is L_{Φ} -Lipschitz continuous. For Algorithm 3, for all $k \geq 0$ it holds*

$$\begin{aligned} \|\mathbf{z}^{(k+1)} - \widehat{\mathbf{z}}^*\|^2 &\leq (1 - \gamma v + 2\gamma^2 L_{\Phi}^2) \|\mathbf{z}^{(k)} - \widehat{\mathbf{z}}^*\|^2 \\ &+ \left(\frac{4\gamma}{v} + 2\gamma^2\right) E_k. \end{aligned} \quad (25)$$

where $E_k := \|\mathbf{e}_{\boldsymbol{\theta}}^{(k)}\|^2 + \|\mathbf{e}_{\boldsymbol{\lambda}}^{(k)}\|^2$ is the total perturbation at iteration k . Moreover, if we choose $\gamma < v/2L_{\Phi}^2$ and E_k is upper bounded by \overline{E} for all k , then

$$\limsup_{k \rightarrow \infty} \|\mathbf{z}^{(k)} - \widehat{\mathbf{z}}^*\|^2 \leq \frac{\frac{4}{v} + 2\gamma}{v - 2\gamma L_{\Phi}^2} \overline{E}. \quad (26)$$

Combining the results from the last subsection, the theorem shows the desired result that the resilient PD-DRA method converges to a $\mathcal{O}(\alpha^2 R^2 d)$ neighborhood of the saddle point of (\mathbf{P}') , if the median-based estimator (17) is used [or $\mathcal{O}(\alpha R^2)$ if Algorithm 2 is used], where α is the fraction of attacked uplink channels. Moreover, it shows that the convergence rate to the neighborhood is linear, which is similar to the classical PDA analysis [1].

Interestingly, Theorem 1 illustrates a trade-off in the choice of the step size γ between convergence speed and accuracy. In specific, (25) shows that the rate of convergence factor $1 - \gamma v + 2\gamma^2 L_{\Phi}^2$ can be minimized by setting $\gamma = v/(4L_{\Phi}^2)$. However, in the meantime, the asymptotic upper bound in (26) is increasing with γ and it can be minimized by setting $\gamma \rightarrow 0$. This will be a design criterion to be explored in practical implementations.

5 Conclusions

In this paper, we studied the strategies for securing a primal-dual algorithm for distributed resource allocation. Particularly, we propose a resilient distributed algorithm based on primal-dual optimization and robust statistics. We derive bounds for the performance of the studied algorithm and show that it converges to a neighborhood of a robustified resource allocation problem when the number of attacked channels is small.

Acknowledgement

The authors would like to thank the anonymous reviewers for feedback, and Mr. Berkay Turan (UCSB) for pointing out typos in the original submission of this paper.

A Proof of Lemma 1

Since g_t is L_t -smooth, the following holds

$$g_t\left(\frac{1}{N} \sum_{i=1}^N \boldsymbol{\theta}_i\right) \leq g_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) + \frac{1}{N} \sum_{j \in \mathcal{A}} \left\langle \boldsymbol{\theta}_j, \nabla g_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) \right\rangle + \frac{L_t}{2N^2} \left\| \sum_{j \in \mathcal{A}} \boldsymbol{\theta}_j \right\|^2 \quad (27)$$

Furthermore, observe that the gradient of g_t is uniformly bounded by B and the diameter of \mathcal{C}_j is R , then the right hand side of (27) can be upper bounded by

$$g_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) + \frac{1}{N} \sum_{j \in \mathcal{A}} (RB + \frac{1}{2}LR^2) \quad (28)$$

As such, defining $c_t := \frac{|\mathcal{A}|}{N} (RB + \frac{1}{2}LR^2)$, it can be seen that

$$g_t\left(\frac{1}{N} \sum_{i \in \mathcal{H}} \boldsymbol{\theta}_i\right) + c_t \leq 0, \quad t = 1, \dots, T \quad (29)$$

implies the desired constraint in (10).

B Proof of Proposition 1

Fix any $j \in [d]$. The assumption implies that for all $i \in \mathcal{H}$, one has

$$|[\boldsymbol{\theta}_i^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j| \leq r. \quad (30)$$

We observe that $|\mathcal{H}| \geq (1-\alpha)N$. Applying [8, Lemma 1] shows that the median estimator¹ satisfies

$$|[\boldsymbol{\theta}_{\text{med}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j| \leq (1-\alpha)\sqrt{\frac{1}{1-2\alpha}}r. \quad (31)$$

The above implies that for all $i \in \mathcal{H}$, we have

$$|[\boldsymbol{\theta}_i^{(k)} - \boldsymbol{\theta}_{\text{med}}^{(k)}]_j| \leq \left(1 + \sqrt{\frac{(1-\alpha)^2}{1-2\alpha}}\right)r. \quad (32)$$

This implies that $r_j^{(k)} \leq \left(1 + \sqrt{\frac{(1-\alpha)^2}{1-2\alpha}}\right)r$ since $|\mathcal{H}| \geq (1-\alpha)N$. We then bound the performance of $\widehat{\boldsymbol{\theta}}^{(k)}$:

$$\begin{aligned} (1-\alpha)N[\widehat{\boldsymbol{\theta}}^{(k)}]_j &= \sum_{i \in \mathcal{N}_j^{(k)}} [\mathbf{r}_i^{(k)}]_j \\ &= \sum_{i \in \mathcal{H}} [\mathbf{r}_i^{(k)}]_j - \sum_{i \in \mathcal{H} \setminus \mathcal{N}_j^{(k)}} [\mathbf{r}_i^{(k)}]_j + \sum_{i \in \mathcal{A} \cap \mathcal{N}_j^{(k)}} [\mathbf{r}_i^{(k)}]_j, \end{aligned} \quad (33)$$

thus

$$\begin{aligned} (1-\alpha)N[\widehat{\boldsymbol{\theta}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j &= - \sum_{i \in \mathcal{H} \setminus \mathcal{N}_j^{(k)}} [\mathbf{r}_i^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j + \sum_{i \in \mathcal{A} \cap \mathcal{N}_j^{(k)}} [\mathbf{r}_i^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j. \end{aligned}$$

Notice that $|\mathcal{A} \cap \mathcal{N}_j^{(k)}| \leq \alpha N$ and thus $|\mathcal{H} \setminus \mathcal{N}_j^{(k)}| \leq \alpha N$. Gathering terms shows

$$|[\widehat{\boldsymbol{\theta}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}]_j| \leq \frac{\alpha N}{(1-\alpha)N} \left(2 + \sqrt{\frac{(1-\alpha)^2}{1-2\alpha}}\right)r. \quad (34)$$

The above holds for all $j \in [d]$. Applying the norm equivalence shows the desired bound.

C Proof of Lemma 2

Comparing the equations in (21) with (19), (20), we identify that

$$[\mathbf{e}_{\boldsymbol{\theta}}^{(k)}]_i = \frac{1}{N} \sum_{t=1}^T \lambda_t^{(k)} \left(\nabla_{\boldsymbol{\theta}} \bar{g}_t \left(\frac{|\mathcal{H}|}{N} \widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} \right) - \nabla_{\boldsymbol{\theta}} \bar{g}_t \left(\frac{|\mathcal{H}|}{N} \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} \right) \right) \quad (35)$$

¹At each coordinate, the median is the geometric median estimator of one dimension in [8].

$$[\mathbf{e}_\lambda^{(k)}]_t = \bar{g}_t(\frac{|\mathcal{H}|}{N}\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}) - \bar{g}_t(\frac{|\mathcal{H}|}{N}\bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}) \quad (36)$$

where $[\mathbf{e}_\theta^{(k)}]_i$ denotes the i th block of $\mathbf{e}_\theta^{(k)}$. Using H1 and the said assumptions, we immediately see that

$$\|[\mathbf{e}_\theta^{(k)}]_i\| \leq \bar{\lambda} \frac{LT}{N} \|\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}\| \quad (37)$$

which then implies (22). H1 implies that \bar{g}_t is B -Lipschitz continuous, therefore

$$|[\mathbf{e}_\lambda^{(k)}]_t| \leq B \|\widehat{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)} - \bar{\boldsymbol{\theta}}_{\mathcal{H}}^{(k)}\|, \quad (38)$$

which implies (23).

D Proof of Theorem 1

Based on Proposition 2, our idea is to perform a perturbation analysis on the PDA algorithm. Without loss of generality, we assume $N = 1$ and denote $\boldsymbol{\theta} = \boldsymbol{\theta}_1$. To simplify notations, we also drop the subscript, denote the modified and regularized Lagrangian function as $\mathcal{L} = \bar{\mathcal{L}}_v$. Furthermore, we denote the saddle point to (P') as $\mathbf{z}^* = (\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)$.

Using the fact that $\boldsymbol{\theta}^* = \mathcal{P}_C(\boldsymbol{\theta}^*) = \mathcal{P}_C(\boldsymbol{\theta}^* - \gamma \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*))$, we observe that in the primal update:

$$\begin{aligned} & \|\boldsymbol{\theta}^{(k+1)} - \boldsymbol{\theta}^*\|^2 \\ & \stackrel{(a)}{\leq} \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 - 2\gamma \langle \widehat{\mathbf{g}}_{\boldsymbol{\theta}}^{(k)} - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*), \boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^* \rangle \\ & \quad + \gamma^2 \|\widehat{\mathbf{g}}_{\boldsymbol{\theta}}^{(k)} - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)\|^2 \end{aligned}$$

where (a) is due to the projection inequality $\|\mathcal{P}_C(\mathbf{x} - \mathbf{y})\| \leq \|\mathbf{x} - \mathbf{y}\|$. Furthermore, using the Young's inequality, for any $c_0, c_1 > 0$, we have

$$\begin{aligned} & \|\boldsymbol{\theta}^{(k+1)} - \boldsymbol{\theta}^*\|^2 \\ & \leq \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 \\ & \quad - 2\gamma \langle \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*), \boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^* \rangle \\ & \quad + \gamma^2 (1 + c_0) \|\nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)\|^2 \\ & \quad - 2\gamma \langle \mathbf{e}_{\boldsymbol{\theta}}^{(k)}, \boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^* \rangle + \gamma^2 (1 + \frac{1}{c_0}) \|\mathbf{e}_{\boldsymbol{\theta}}^{(k)}\|^2 \\ & \leq (1 + 2c_1\gamma) \|\boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^*\|^2 \\ & \quad - 2\gamma \langle \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*), \boldsymbol{\theta}^{(k)} - \boldsymbol{\theta}^* \rangle \\ & \quad + \gamma^2 (1 + c_0) \|\nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)\|^2 \\ & \quad + \left(\frac{2\gamma}{c_1} + \gamma^2 + \frac{\gamma^2}{c_0} \right) \|\mathbf{e}_{\boldsymbol{\theta}}^{(k)}\|^2. \end{aligned}$$

Similarly, in the dual update we get,

$$\begin{aligned}
& \|\boldsymbol{\lambda}^{(k+1)} - \boldsymbol{\lambda}^*\|^2 \\
& \leq \|\boldsymbol{\lambda}^{(k)} - \boldsymbol{\lambda}^*\|^2 + \gamma^2 \|\widehat{\boldsymbol{g}}_{\boldsymbol{\lambda}}^{(k)} - \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)\|^2 \\
& \quad + 2\gamma \langle \widehat{\boldsymbol{g}}_{\boldsymbol{\lambda}}^{(k)} - \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*), \boldsymbol{\lambda}^{(k)} - \boldsymbol{\lambda}^* \rangle \\
& \leq (1 + 2c_1\gamma) \|\boldsymbol{\lambda}^{(k)} - \boldsymbol{\lambda}^*\|^2 \\
& \quad + 2\gamma \langle \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*), \boldsymbol{\lambda}^{(k)} - \boldsymbol{\lambda}^* \rangle \\
& \quad + \gamma^2(1 + c_0) \|\nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^{(k)}, \boldsymbol{\lambda}^{(k)}) - \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\theta}^*, \boldsymbol{\lambda}^*)\|^2 \\
& \quad + \left(\frac{2\gamma}{c_1} + \gamma^2 + \frac{\gamma^2}{c_0} \right) \|\boldsymbol{e}_{\boldsymbol{\lambda}}^{(k)}\|^2.
\end{aligned}$$

Summing up the two inequalities gives:

$$\begin{aligned}
& \|\boldsymbol{z}^{(k+1)} - \boldsymbol{z}^*\|^2 \\
& \leq (1 + 2c_1\gamma) \|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2 + \left(\frac{2\gamma}{c_1} + \gamma^2 + \frac{\gamma^2}{c_0} \right) E_k \\
& \quad - 2\gamma \langle \boldsymbol{\Phi}(\boldsymbol{z}^{(k)}) - \boldsymbol{\Phi}(\boldsymbol{z}^*), \boldsymbol{z}^{(k)} - \boldsymbol{z}^* \rangle \\
& \quad + \gamma^2(1 + c_0) \|\boldsymbol{\Phi}(\boldsymbol{z}^{(k)}) - \boldsymbol{\Phi}(\boldsymbol{z}^*)\|^2 \\
& \stackrel{(a)}{\leq} \left(1 + 2\gamma(c_1 - v) + \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2 \right) \|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2 \\
& \quad + \left(\frac{2\gamma}{c_1} + \gamma^2 + \frac{\gamma^2}{c_0} \right) E_k,
\end{aligned}$$

where (a) uses the strong monotonicity and smoothness of the map $\boldsymbol{\Phi}$. Setting $c_1 = v/2$ yields

$$\begin{aligned}
& \|\boldsymbol{z}^{(k+1)} - \boldsymbol{z}^*\|^2 \\
& \leq \left(1 - \gamma v + \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2 \right) \|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2 \\
& \quad + \left(\frac{4\gamma}{v} + \gamma^2 + \frac{\gamma^2}{c_0} \right) E_k.
\end{aligned} \tag{39}$$

Observe that we can choose γ such that $1 - \gamma v + \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2 < 1$. Moreover, the above inequality implies that $\|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2$ evaluates to

$$\begin{aligned}
& \|\boldsymbol{z}^{(k+1)} - \boldsymbol{z}^*\|^2 \\
& \leq (1 - \gamma v + \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2)^k \|\boldsymbol{z}^{(0)} - \boldsymbol{z}^*\|^2 + \\
& \quad \sum_{\ell=1}^k (1 - \gamma v + \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2)^{k-\ell} \left(\frac{4\gamma}{v} + \gamma^2 + \frac{\gamma^2}{c_0} \right) E_{\ell}
\end{aligned}$$

If $E_k \leq \overline{E}$ for all k , then $\boldsymbol{z}^{(k)}$ converges to a neighborhood of \boldsymbol{z}^* of radius

$$\limsup_{k \rightarrow \infty} \|\boldsymbol{z}^{(k)} - \boldsymbol{z}^*\|^2 \leq \frac{\frac{4\gamma}{v} + \gamma^2 + \frac{\gamma^2}{c_0}}{\gamma v - \gamma^2(1 + c_0)L_{\boldsymbol{\Phi}}^2} \overline{E} \tag{40}$$

Setting $c_0 = 1$ concludes the proof.

References

- [1] J. Koshal, A. Nedić, and U. V. Shanbhag, “Multiuser optimization: Distributed algorithms and error analysis,” *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 1046–1081, 2011.
- [2] D. P. Palomar and M. Chiang, “A tutorial on decomposition methods for network utility maximization,” *IEEE JSAC*, vol. 24, no. 8, pp. 1439–1451, 2006.
- [3] S. Sundaram and C. N. Hadjicostis, “Distributed function calculation via linear iterative strategies in the presence of malicious agents,” *IEEE TAC*, vol. 56, no. 7, pp. 1495–1508, 2011.
- [4] F. Pasqualetti, A. Bicchi, and F. Bullo, “Consensus computation in unreliable networks: A system theoretic approach,” *IEEE TAC*, vol. 57, no. 1, pp. 90–104, 2012.
- [5] R. Gentz, S. X. Wu, H.-T. Wai, A. Scaglione, and A. Leshem, “Data injection attacks in randomized gossiping,” *IEEE TSIPN*, vol. 2, no. 4, pp. 523–538, 2016.
- [6] S. Sundaram and B. Ghahesifard, “Distributed optimization under adversarial nodes,” *IEEE TAC*, 2018.
- [7] Y. Chen, S. Kar, and J. Moura, “Resilient distributed estimation: Sensor attacks,” *IEEE Transactions on Automatic Control*, 2018.
- [8] J. Feng, H. Xu, and S. Mannor, “Distributed robust learning,” *arXiv preprint arXiv:1409.5937*, 2014.
- [9] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, “Byzantine-robust distributed learning: Towards optimal statistical rates,” *arXiv preprint arXiv:1803.01498*, 2018.
- [10] D. Alistarh, Z. Allen-Zhu, and J. Li, “Byzantine stochastic gradient descent,” in *NeurIPS*, pp. 4618–4628, 2018.
- [11] D. L. Donoho and P. J. Huber, “The notion of breakdown point,” *A festschrift for Erich L. Lehmann*, vol. 157184, 1983.
- [12] P. J. Huber, *Robust statistics*. Springer, 2011.
- [13] S. Minsker *et al.*, “Geometric median and robust estimation in banach spaces,” *Bernoulli*, vol. 21, no. 4, pp. 2308–2335, 2015.
- [14] I. Diakonikolas, G. Kamath, D. M. Kane, J. Li, A. Moitra, and A. Stewart, “Robust estimators in high dimensions without the computational intractability,” in *IEEE FOCS*, pp. 655–664, 2016.
- [15] J. Steinhardt, M. Charikar, and G. Valiant, “Resilience: A criterion for learning in the presence of arbitrary outliers,” *arXiv preprint arXiv:1703.04940*, 2017.
- [16] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, “Rate control for communication networks: shadow prices, proportional fairness and stability,” *Journal of the Operational Research Society*, vol. 49, pp. 237–252, Mar 1998.

- [17] S. H. Low and D. E. Lapsley, “Optimization flow control. i. basic algorithm and convergence,” *IEEE/ACM Transactions on Networking*, vol. 7, pp. 861–874, Dec 1999.
- [18] P. Samadi, A. Mohsenian-Rad, R. Schober, V. W. S. Wong, and J. Jatskevich, “Optimal real-time pricing algorithm based on utility maximization for smart grid,” in *IEEE Smart-GridComm*, pp. 415–420, Oct 2010.
- [19] N. Li, L. Chen, and S. H. Low, “Optimal demand response based on utility maximization in power networks,” in *2011 IEEE Power and Energy Society General Meeting*, pp. 1–8, July 2011.
- [20] M. Chiang and J. Bell, “Balancing supply and demand of bandwidth in wireless cellular networks: utility maximization over powers and rates,” in *IEEE INFOCOM 2004*, vol. 4, pp. 2800–2811 vol.4, March 2004.
- [21] M. Dehghan, L. Massoulie, D. Towsley, D. Menasche, and Y. C. Tay, “A utility optimization approach to network cache design,” in *IEEE INFOCOM*, pp. 1–9, April 2016.
- [22] M. Zhao, J. Li, and Y. Yang, “Joint mobile energy replenishment and data gathering in wireless rechargeable sensor networks,” in *Proceedings of the 23rd International Teletraffic Congress, ITC '11*, pp. 238–245, International Teletraffic Congress, 2011.
- [23] R. Deng, Y. Zhang, S. He, J. Chen, and X. Shen, “Maximizing network utility of rechargeable sensor networks with spatiotemporally coupled constraints,” *IEEE JSAC*, vol. 34, pp. 1307–1319, May 2016.
- [24] N. Mehr, J. Lioris, R. Horowitz, and R. Pedarsani, “Joint perimeter and signal control of urban traffic via network utility maximization,” in *IEEE ITSC*, pp. 1–6, Oct 2017.
- [25] J. Lavaei, D. Tse, and B. Zhang, “Geometry of power flows and optimization in distribution networks,” *IEEE Transactions on Power Systems*, vol. 29, no. 2, pp. 572–583, 2014.
- [26] J. A. Taylor, *Convex optimization of power systems*. Cambridge University Press, 2015.