

Resilient Primal-Dual Optimization Algorithms for Distributed Resource Allocation

Berkay Turan[†] César A. Uribe[†] Hoi-To Wai Mahnoosh Alizadeh

Abstract—Distributed algorithms for multi-agent resource allocation can provide privacy and scalability over centralized algorithms in many cyber-physical systems. However, the distributed nature of these algorithms can render these systems vulnerable to man-in-the-middle attacks that can lead to non-convergence and infeasibility of resource allocation schemes. In this paper, we propose attack-resilient distributed algorithms based on primal-dual optimization when Byzantine attackers are present in the system. In particular, we design attack-resilient primal-dual algorithms for static and dynamic impersonation attacks by means of robust statistics. For static impersonation attacks, we formulate a robustified optimization model and show that our algorithm guarantees convergence to a neighborhood of the optimal solution of the robustified problem. On the other hand, a robust optimization model is not required for the dynamic impersonation attack scenario and we are able to design an algorithm that is shown to converge to a near-optimal solution of the original problem. We analyze the performances of our algorithms through both theoretical and computational studies.

I. INTRODUCTION

A number of multi-agent optimization problems arise in a wide range of resource allocation systems that fall under the general umbrella of *Network Utility Maximization* problems: in the pioneering example of congestion control in data networks [1], [2]; in determining the optimal price of electricity and enabling more efficient demand supply balancing in smart power distribution systems [3], [4]; in managing user transmit powers and data rates in wireless cellular networks [5]; in determining optimal caching policies by content delivery networks [6]; in optimizing power consumption in wireless sensor networks with energy-restricted batteries [7], [8]; and in designing congestion control systems in urban traffic networks [9]. The shared goal among the above-mentioned problems is to minimize the sum of N user-specific cost functions, subject to a set of coupling constraints that depend on users’ decisions.

In these resource allocation problems, the user-specific cost functions and the set of coupling constraints are considered private information to the users and to a central coordinator, respectively. Consequently, it is necessary to solve these problems in a distributed fashion allowing the agents to cooperate through communication with a central coordinator. Among others, primal-dual optimization methods [10] have been advocated as they naturally give rise to decomposable

algorithms that favor distributed implementation [11]. In addition to their practical success, these methods are supported by strong theoretical guarantees where fast convergence to a near-optimal solution is well established [10].

However, the distributed nature of these methods also exposes the system to vulnerabilities not faced by their traditional centralized counterpart. Many of the existing algorithms assume the agents, and the communication channels between the central coordinator and the agents, to be *completely trustworthy*. In this paper, we consider the setting where these communications are susceptible to adversarial attacks. An attacker can take over network sub-systems, and deliberately edit the messages communicated to the central coordinator to any arbitrary value, i.e., a Byzantine attack. As we will demonstrate, this might result in an unstable system with possible damages to hardware and the system overall.

Our goal is to design attack-resilient primal-dual algorithms in order to solve multi-agent resource allocation problems in presence of Byzantine attackers. If a communication channel is attacked and becomes compromised, the attacker can modify messages and/or inject fresh messages into the network on the agents’ behalf. We consider two scenarios with different attacker capabilities. A static impersonation attack scenario considers the set of agents communicating through compromised channels to be the same for the duration of the algorithm, whereas a dynamic impersonation attack scenario considers the case where all agents are susceptible to attacks and hence communicate through compromised channels for a *limited fraction* of the algorithm’s runtime. Our main contributions are as follows:

- We propose resilient distributed resource allocation algorithms under the two aforementioned attack scenarios that rely on robust mean estimation.
- We provide convergence guarantees of the proposed algorithms. We show that our algorithm for the dynamic impersonation attack scenario converges to the optimal solution of the regularized problem, while our algorithm for the static impersonation attack scenario converges to an $\mathcal{O}(\alpha_1^2)$ neighborhood of the optimal solution of a robustified and regularized optimization model, where $\alpha_1 \in [0, \frac{1}{2})$ is a known upper bound on fraction of attacked channels.
- We provide empirical evidence that supports our theoretical results on convergence and preventing constraint violation. We do so via computational simulations on electric vehicle charging and power distribution applications.

Related work: Vulnerabilities of various types of distributed algorithms have been identified and addressed in a number of recent studies. Relevant examples can be found in [12]–

[†]Authors have equal contribution. B. Turan and M. Alizadeh are with Dept. of ECE, UCSB, Santa Barbara, CA, USA. C. A. Uribe is with LIDS, MIT, Cambridge, MA, USA. H. T. Wai is with Dept. of SEEM, CUHK, Shatin, Hong Kong. This work is partially supported by UCOP Grant LFR-18-548175, NSF grant #1847096, CUHK Direct Grant #4055113, and the Yahoo! Research Faculty Engagement Program. E-mails: bturan@ucsb.edu, cauribe@mit.edu, htwai@se.cuhk.edu.hk, alizadeh@ucsb.edu

[20] which study secure decentralized algorithms on a general network topology but consider consensus-based optimization models. There are two fundamental differences between distributed resource allocation and consensus problems that make these algorithms inapplicable in our case:

- In resource allocation problems, each agent is solving for their own optimal level of resource consumption, i.e., each agent is solving for their own parameter, whereas consensus problems focus on all agents solving for a shared (global) parameter.
- Unlike resilient consensus algorithms, in resource allocation problems pertaining to access to critical infrastructure systems such as power or transportation networks, one cannot simply block a set of users' access to the network even if they are deemed likely to be attackers.

A recently popular line of works in [21]–[25] focuses on building resilient algorithms for distributed statistical learning. A crucial difference from this work is that they assume identical functions across the agents. In fact, we employ robust statistics [26], [27] to develop our resilient algorithms, and particularly, we develop novel results for robust mean estimation, a topic that is recently rekindled in [28]–[30].

The present paper is a revised and extended version of the preliminary conference report [31]. This paper expands [31] into multiple attack scenarios and includes numerical studies. **Paper Organization:** The remainder of the paper is organized as follows. In Section II, we provide an overview of the basic primal-dual algorithm for resource allocation. In Section III, we formally define two Byzantine attack models and demonstrate how Byzantine attacks can alter the primal-dual optimization procedure. In Section IV, we present two attack-resilient primal-dual algorithms corresponding to the different attack scenarios along with their convergence analysis. In Section V, we provide numerical results for our algorithms.

Notations. Unless otherwise specified, $\|\cdot\|$ denotes the standard Euclidean norm. For any $N \in \mathbb{N}$, $[N]$ denotes the finite set $\{1, \dots, N\}$. Given θ , θ_i indicates the i 'th block/entry of θ that corresponds to the parameter of agent i . $\theta_{i,j}$ denotes the j 'th element of vector θ_i .

II. OVERVIEW OF PRIMAL-DUAL ALGORITHM FOR RESOURCE ALLOCATION

We consider the following multi-agent optimization problem with an objective to minimize the average cost incurred by the agents, subject to a set of constraints that are functions of the average of the agents' parameters:

$$\begin{aligned} \min_{\theta_i \in \mathbb{R}^d, \forall i} \quad & f(\theta) := \frac{1}{N} \sum_{i=1}^N f_i(\theta_i) \\ \text{subject to} \quad & g_t \left(\frac{1}{N} \sum_{i=1}^N \theta_i \right) \leq 0, \quad t = 1, \dots, T, \\ & \theta_i \in \mathcal{C}_i, \quad i = 1, \dots, N, \end{aligned} \quad (1)$$

where $f_i(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ is the continuously differentiable and convex cost function of agent i and $g_t(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ are continuously differentiable and convex set of constraints. The parameter θ_i of agent i is constrained to be in a compact convex set $\mathcal{C}_i \in \mathbb{R}^d$.

Algorithm 1: PD-DRA Procedure.

- 1: **for** $k = 1, 2, \dots$ **do**
- 2: (*Communication stage*):
 - (a) Central coordinator receives $\{\theta_i^{(k)}\}_{i=1}^N$ from agents and computes $\bar{\theta}^{(k)} := \frac{1}{N} \sum_{i=1}^N \theta_i^{(k)}$, $\{\nabla_{\theta} g_t(\bar{\theta}^{(k)})\}_{t=1}^T$.
 - (b) Central coordinator broadcasts the vector $\bar{g}^{(k)} := \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\theta} g_t(\bar{\theta}^{(k)})$ to agents.
- 3: (*Computation stage*):
 - (a) Agent i computes the update for $\theta_i^{(k+1)}$ according to (4a) using the received $\bar{g}^{(k)}$.
 - (b) The central coordinator computes the update for $\lambda^{(k+1)}$ according to (4b).
- 4: **end for**

Running Example (Resource Allocation Problem). *Throughout the paper, we use the following toy example as a running example to clarify the concepts and the methods: We consider an EV charging example with 5 agents. The cost function $f_i(\cdot)$ is monotone decreasing and is the same for all agents. As an example, we set $f_i(\theta_i) = (\theta_i - 10)^2$ as the quadratic cost function which is monotonically decreasing for $0 \leq \theta_i \leq 10$. There is a charging station with 5 EV charging points, three of which have a maximum charging rate of 7kW, and two have a rate of 10kW. The total rate at which the charging station is able to deliver electricity is determined by the grid, and let it be upper bounded by 25kW (hence, the average rate is upper bounded by $\frac{25}{5} = 5$ kW). Accordingly, the constraints of this system are stated as:*

$$\begin{aligned} g((1/5) \sum_{i=1}^5 \theta_i) &:= (1/5) \sum_{i=1}^5 \theta_i - 5 \leq 0, \\ 0 \leq \theta_i &\leq 7, \quad i = 1, 2, 3, \\ 0 \leq \theta_i &\leq 10, \quad i = 4, 5. \end{aligned}$$

Note that θ is a real number, hence dimension $d = 1$. The optimal solution in this example is to deliver electricity at a rate of 5kW to all agents due to symmetry.

The optimization problem in (1) can not be solved centrally, because the utility functions $f_i(\cdot)$ are private to the agents, and furthermore the coupling constraints on the resources are only known by a central coordinator. Accordingly, the goal of the primal-dual distributed resource allocation (PD-DRA) procedure in Algorithm 1 is to solve (1) in a distributed manner, where the agents observe a pricing signal received from the central coordinator and communicate their parameters to the central coordinator [10]. Consequent to this information exchange, the pricing signal and the agents' parameters are updated by the central coordinator and by the individual agents, respectively.

In order to derive the update rules used by Algorithm 1, we first consider the Lagrangian function of (1):

$$\mathcal{L}(\{\theta_i\}_{i=1}^N; \lambda) := \frac{1}{N} \sum_{i=1}^N f_i(\theta_i) + \sum_{t=1}^T \lambda_t g_t \left(\frac{1}{N} \sum_{i=1}^N \theta_i \right), \quad (2)$$

where $\lambda_t \geq 0$ is the dual variable associated with constraint $g_t(\cdot)$ and $\lambda = [\lambda_1 \dots \lambda_T]^T \in \mathbb{R}_+^T$ is the vector of the dual variables. Under strong duality (e.g., when the Slater's

condition holds), solving problem (1) is equivalent to solving its dual problem:

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

As suggested in [10], we consider a regularized version of (P). Let us define

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

such that $\mathcal{L}_v(\cdot)$ is v -strongly convex and v -strongly concave in $\{\theta_i\}_{i=1}^N$ and λ , respectively.

Remark 1. Adding regularization terms is a typical technique used in optimization, called dual smoothing [32]. We add the regularization terms for the purposes of convergence analysis used in this paper, which can be applied on strongly convex/concave functions. Indeed adding the regularization terms might change the solution of the original optimization problem. However, as explained in [33, Proposition 5.2], by an appropriate selection of the regularization parameters, we can recover an optimality gap guarantee for the original problem based on the solution to the regularized problem.

We define the regularized problem as:

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

Let $\gamma > 0$ be the step size and $k \in \mathbb{Z}_+$ be the iteration index. The primal-dual recursion performs projected gradient descent/ascent on the primal/dual variables as follows:

$$\theta_i^{(k+1)} = \quad (4a)$$

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

$$\lambda^{(k+1)} = [\lambda^{(k)} + \gamma \nabla_{\lambda} \mathcal{L}_v(\{\theta_i^{(k)}\}_{i=1}^N; \lambda^{(k)})]_+, \quad (4b)$$

where $\mathcal{P}_{\mathcal{C}_i}(\cdot)$ is the Euclidean projection operator to set \mathcal{C}_i and $[\cdot]_+$ denotes $\max\{\cdot, 0\}$ operator. According to (3), the gradients with respect to (w.r.t.) the primal and the dual variables are given respectively by:

$$\begin{aligned} \nabla_{\theta_i} \mathcal{L}_v(\{\theta_i^{(k)}\}_{i=1}^N; \lambda^{(k)}) = \frac{1}{N} \left(\nabla_{\theta_i} f_i(\theta_i^{(k)}) + v \theta_i^{(k)} \right. \\ \left. + \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\theta} g_t(\theta) \Big|_{\theta = \frac{1}{N} \sum_{i=1}^N \theta_i^{(k)}} \right), \end{aligned} \quad (5a)$$

$$[\nabla_{\lambda} \mathcal{L}_v(\{\theta_i^{(k)}\}_{i=1}^N; \lambda^{(k)})]_t = g_t \left(\frac{1}{N} \sum_{i=1}^N \theta_i^{(k)} \right) - v \lambda_t^{(k)}, \quad (5b)$$

for all i, t . It is worthwhile to highlight that both gradients depend on the average parameter $\bar{\theta}^{(k)} := \frac{1}{N} \sum_{i=1}^N \theta_i^{(k)}$. From the above equations (5a) and (5b), we can determine which variables should be communicated between the central coordinator and the agents so that the gradients can be computed locally, see Algorithm 1.

Since the regularized primal-dual problem is strongly convex/concave in primal/dual variables, Algorithm 1 converges linearly to the optimal solution of (P_v) [10]. To study this, let us concatenate the primal and the dual variables and denote $\mathbf{z}^{(k)} := (\{\theta_i^{(k)}\}_{i=1}^N, \lambda^{(k)})$ as the primal-dual variable at the k th iteration and define the mapping $\Phi(\mathbf{z}^{(k)})$ as:

$$\Phi(\mathbf{z}^{(k)}) := \begin{pmatrix} \nabla_{\theta} \mathcal{L}_v(\{\theta_i^{(k)}\}_{i=1}^N, \lambda^{(k)}) \\ -\nabla_{\lambda} \mathcal{L}_v(\{\theta_i^{(k)}\}_{i=1}^N, \lambda^{(k)}) \end{pmatrix}. \quad (6)$$

Proposition 1. [10, Theorem 3.5] Assume that the map $\Phi(\mathbf{z}^{(k)})$ is L_{Φ} Lipschitz continuous. For all $k \geq 1$, we have

$$\|\mathbf{z}^{(k+1)} - \mathbf{z}^*\|^2 \leq (1 - 2\gamma v + \gamma^2 L_{\Phi}^2) \|\mathbf{z}^{(k)} - \mathbf{z}^*\|^2, \quad (7)$$

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

III. PROBLEM FORMULATION

Even though the PD-DRA provides strong theoretical convergence guarantee, it relies on error-free communication between the central coordinator and the agents, and is not robust to attacks on the channels between the agents and the central coordinator, as described below.

We study a situation when the uplink communication channels between some of the agents and the central coordinator are compromised.¹ Let $\mathcal{A}^{(k)} \subset [N]$ be the set of agents communicating through compromised uplink channels at iteration k , whose identities are unknown to the central coordinator, and let $\mathcal{H}^{(k)} := [N] \setminus \mathcal{A}^{(k)}$ be the set of agents communicating through trustworthy uplink channels at iteration k . Instead of receiving $\theta_i^{(k)}$ from each agent $i \in [N]$ at iteration k (Algorithm 1 Step 2(a)), the central coordinator receives the following messages:

$$\mathbf{r}_i^{(k)} = \begin{cases} \theta_i^{(k)}, & \text{if } i \in \mathcal{H}^{(k)}, \\ \mathbf{b}_i^{(k)}, & \text{if } i \in \mathcal{A}^{(k)}. \end{cases} \quad (8)$$

We consider a Byzantine attack scenario, under which the messages sent through the compromised channels, $\mathbf{b}_i^{(k)}$, can be chosen arbitrarily by an adversary. This also encompasses faulty messages due to erroneous inputs or erroneous channels, since we set no restrictions on $\mathbf{b}_i^{(k)}$. The adversary's goal is to harm the system and cause suboptimality. When the messages are erroneous or chosen adversarially, the central coordinator computes the gradients and therefore the pricing signal using these erroneous messages. The agents then update their parameters based on this erroneous pricing signal, which can lead to an overall suboptimal resource allocation. Moreover, the choice of the compromised channels $\mathcal{A}^{(k)}$ affects the impact of the attack and the precautions to be taken in order to defend against the attack. As such, we study two Byzantine attack scenarios that differ in the set of the compromised channels as illustrated in Figure 1.

Running Example (Byzantine Attack). Let agent 1 be communicating through a compromised channel at all iterations, i.e., $\mathcal{A}^{(k)} = \{1\}, \forall k$. The compromised message sent to the central coordinator is $\mathbf{b}_1^{(k)} = 1kW, \forall k$. This means that irrespective of $\theta_1^{(k)}$, the central coordinator receives a message indicating agent 1 is willing to charge at rate of $1kW$.

¹This paper studies the case where only uplink channels are compromised. However, the case of downlink corruption can also be addressed. Since the downlink channel is a broadcast channel, a compromised downlink channel results in no agent receiving a trustworthy pricing signal. In that case, there is no optimization method based solution to that problem since there is no communication. If we assume however that all the downlink channels are point-to-point between the central coordinator and each agent, the methods developed in this paper can be applied in a similar fashion.

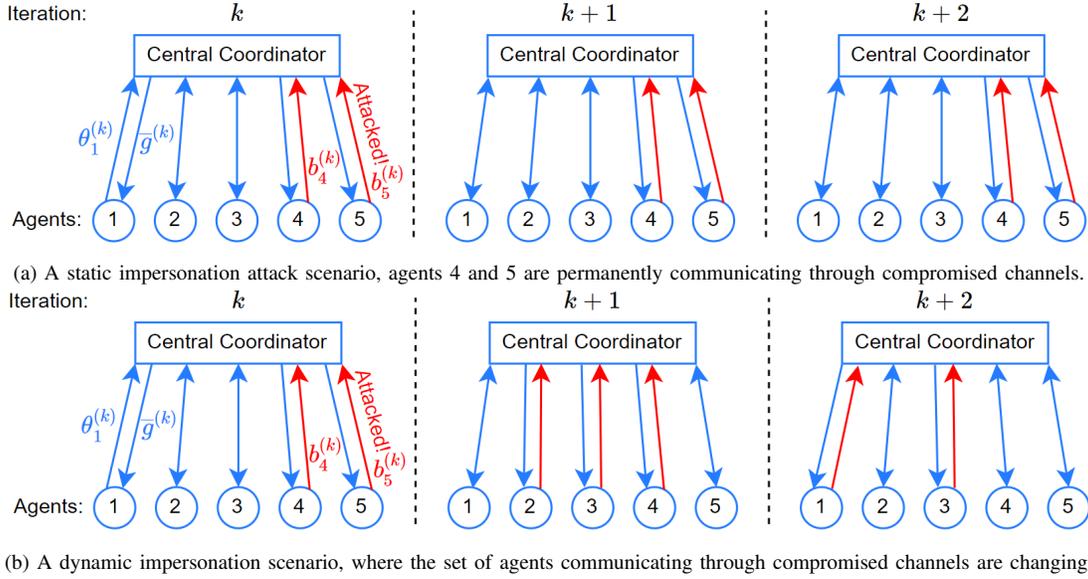


Fig. 1: Illustration of (a) static impersonation attack, and (b) dynamic impersonation attack. Blue arrows represent trustworthy channels, whereas red arrows represent compromised channels.

A. Attack scenarios

1) A *static impersonation attack*, where an adversary takes over a subset of uplink channels permanently and the set of agents communicating through compromised channels is fixed (i.e., $\mathcal{A}^{(k)} = \mathcal{A}$, $\forall k$). Consequently, the central coordinator is never able to communicate reliably with agents $i \in \mathcal{A}$. In this case, it is not feasible to optimize the original problem (P) since the contribution from $f(\theta_i) : i \in \mathcal{A}$ becomes unknown to the central coordinator. Yet, we assume that it is also not possible to deny access to resources to agents who are suspected of potentially being under attack. As a compromise, we formulate the following optimization problem:

$$\begin{aligned} \min_{\theta_i \in \mathcal{C}_i, i \in \mathcal{H}} \quad & f(\theta) := \frac{1}{N} \sum_{i \in \mathcal{H}} f_i(\theta_i) \\ \text{subject to} \quad & \max_{\theta_j \in \mathcal{C}_j, j \in \mathcal{A}} g_t \left(\frac{1}{N} \sum_{i=1}^N \theta_i \right) \leq 0, \quad \forall t \in [T]. \end{aligned} \quad (9)$$

The objective of (9) is to minimize the cost of the agents with trustworthy channels subject to a *robust* set of constraints that consider the *worst case* scenario, in which the parameters of the agents with compromised channels are assumed to be maximizing the constraints (e.g., those agents are assumed to be consuming the maximum amount of resources). It is critical to mention that during a primal-dual algorithm scheme, the messages received through the compromised channels can be anything. The robust approach is to however ignore those messages, and assume that the parameters of the agents communicating through those channels are maximizing the constraints so that the operation of the system is feasible under any circumstance. Our goal is to develop an attack-resilient PD-DRA to solve the robust optimization problem (9).

Running Example (Robust Optimization Model). *Since agent 1 is sending a compromised message of 1kW and their true parameter can be anything, the worst-case approach is to*

assume that they are charging at the maximum rate, which is 7kW for that agent. Hence, the robust constraint is:

$$\begin{aligned} \max_{\theta_j \in \mathcal{C}_j, j \in \mathcal{A}} g \left(\frac{1}{5} \sum_{i=1}^5 \theta_i \right) &= \max_{\theta_j \in \mathcal{C}_j, j \in \mathcal{A}} \frac{1}{5} \sum_{i=1}^5 \theta_i - 5 \\ &= \frac{1}{5} \sum_{i \in \mathcal{H}} \theta_i + \max_{\theta_j \in \mathcal{C}_j, j \in \mathcal{A}} \frac{1}{5} \sum_{j \in \mathcal{A}} \theta_j - 5 = \frac{4}{5} \bar{\theta}_{\mathcal{H}} - 3.6, \end{aligned}$$

where we used $|\mathcal{H}| = 4$ and the notation $\bar{\theta}_{\mathcal{H}} = \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \theta_i$. The robust constraint states that:

$$\frac{4}{5} \bar{\theta}_{\mathcal{H}} - 3.6 \leq 0 \Rightarrow \bar{\theta}_{\mathcal{H}} \leq 4.5.$$

The optimal solution in this case is to deliver electricity at a rate of 4.5kW to the trustworthy agents. Since the compromised agent has the same cost function, their true charging rate will also be 4.5kW, even though the message sent is 1kW and the central coordinator assumes their charging rate is 7kW.

2) A *dynamic impersonation attack*, where all the agents might be affected by the adversarial attacks but only for a *limited* fraction of time and hence, the set of agents communicating through compromised channels $\mathcal{A}^{(k)}$ has to dynamically change with iteration k . As opposed to the static case, this scenario considers the case where the central coordinator is able to communicate reliably with all the agents at some iterations. Due to this distinction, it is necessary to mention that the static attack is not a special case of the dynamic attack and both scenarios are distinguishable from each other. The dynamic scenario could be applicable when agents do not have dedicated communication channels to the central coordinator and instead communicate over random access systems which are more appropriate for distributed deployments. Hence, each user periodically accesses authenticated network devices/subsystems that are controlled by Byzantine adversaries and can alter the user's message. Our goal is to develop an attack-resilient PD-DRA algorithm that can still solve the original regularized problem (P_v) in this environment.

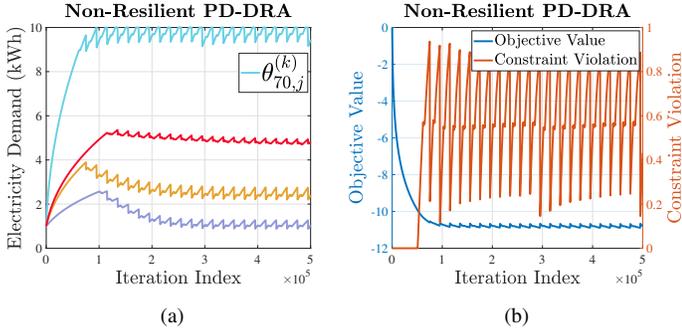


Fig. 2: Illustration of basic PD-DRA algorithm failure under static impersonation attack. (a) The agents' parameters do not converge, (b) the objective function does not converge and moreover there is constraint violation. We only display one constraint for brevity.

B. Limitations of the Basic PD-DRA Algorithm

Applying the basic PD-DRA algorithm under a Byzantine attack scenario can lead to undesirable outcomes. Recall that the gradients in (5) depend on the average parameter $\bar{\theta}^{(k)}$. Under a Byzantine attack scenario, if the central coordinator forms the naive average $\tilde{\theta}^{(k)} = (1/N) \sum_{i=1}^N r_i^{(k)}$ and computes the gradients $\nabla g_t(\tilde{\theta}^{(k)})$ accordingly, this may result in large error since the deviation $\tilde{\theta}^{(k)} - (1/N) \sum_{i=1}^N \theta_i^{(k)}$ can be large (proportional to the maximum diameter of \mathcal{C}_i 's). This in turn can obstruct convergence and also overload the system by causing constraint violations.

Running Example (Basic PD-DRA Failure). *If the central coordinator believes all the agents are sending trustworthy information, then the optimal solution will occur when one agent is demanding 1kW and the others are demanding 6kW (so that the average is 5kW). But since the 1kW message is compromised and all the agents have same cost function, the compromised agent's true electricity demand is also at a rate of 6kW. Hence, the solution delivers electricity at an average rate of 6kW, which is infeasible.*

We preview our numerical result of applying the basic PD-DRA method under a static impersonation attack scenario for an optimal electric vehicle charging application in Figure 2. For constraint $g_t(\cdot)$, we define constraint violation as $\max\{0, g_t(\bar{\theta}^{(k)})\}$. Observe that the PD-DRA method does not provide convergence and the first constraint is being violated. From resource allocation perspective, this means that the agents are asking to consume more resources than the available amount in the system, which is infeasible. For details regarding the experimental setup, please see Section V.

IV. RESILIENT PD-DRA ALGORITHMS

Motivated by the failure of the basic PD-DRA procedure under Byzantine attack scenarios, resilient PD-DRA algorithms are necessary to optimize multi-agent systems in a distributed manner when the system is susceptible to attacks. We hold the following assumption to be true throughout the rest of the paper and propose two different attack resilient PD-DRA algorithms corresponding to the different attack scenarios outlined in Section III.

Assumption 1. *For all $\theta \in \mathbb{R}^d$ and for all t , the gradient of g_t is bounded with $\|\nabla g_t(\theta)\| \leq B$ and is L -Lipschitz continuous. Moreover, since maximum resource that can be consumed by an agent is bounded due to limited amount of resources, we let $\mathbf{0} \in \mathcal{C}_i$ and upper bound the diameters of \mathcal{C}_i by R :*

$$\max_{\theta, \theta' \in \mathcal{C}_i} \|\theta - \theta'\| \leq R, \quad i = 1, \dots, N. \quad (10)$$

Running Example (Assumptions). *The constraint in our running example satisfies $\nabla g(\theta) = 1$, which is bounded by $B = 1$ and is $L = 0$ -Lipschitz continuous. Since the maximum charging rate is upper bounded by 7kW for three of the agents and by 10kW for two of the agents, $R = 10$.*

A. Static Impersonation Attack

Under this attack scenario, given the complete lack of any credible information on the resource consumption parameters of the agents that permanently communicate through compromised channels, the central coordinator can only hope to solve the robust optimization model defined in (9) instead. This formulation considers a worst-case scenario on how much resources the compromised agents will consume, which ensures constraint satisfaction in all cases. However, the constraints in (9) require the knowledge of the set \mathcal{A} and the sets $\mathcal{C}_j, \forall j \in \mathcal{A}$, yet the central coordinator lacks this information.

Hence, in order to develop a robust optimization model that can handle the worst-case scenario without the knowledge of \mathcal{A} , we let $\alpha_1 \geq |\mathcal{A}|/N$ as a known upper bound to the fraction of agents communicating through compromised channels and assume $\alpha_1 < 1/2$, where less than half of the agents are communicating through compromised channels.² Let $\bar{\theta}_{\mathcal{H}} := \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \theta_i$ be the mean of the agent's parameters that are sent through trustworthy channels. We then define the following set of constraints

$$\bar{g}_t(\theta) := g_t(\theta) + \alpha_1 (RB + \frac{1}{2}LR^2), \quad (11)$$

and formulate a conservative approximation of (9):

Lemma 1. *Under Assumption 1 the following problem yields a conservative approximation of (9), i.e., its feasible set is a subset of the feasible set of (9):*

$$\begin{aligned} \min_{\theta_i \in \mathcal{C}_i, i \in \mathcal{H}} \quad & \frac{1}{N} \sum_{i \in \mathcal{H}} f_i(\theta_i) \\ \text{subject to} \quad & \bar{g}_t((1 - \alpha_1)\bar{\theta}_{\mathcal{H}}) \leq 0, \quad \forall t \in [T], \end{aligned} \quad (12)$$

The proof can be found in Appendix A, which is provided in the online version [34]. All proofs and Appendices can be found in the online version [34].

Remark 2. *The proof of Lemma 1 is done by upper bounding constraints of (9) using Assumption 1 and the fact that $\alpha_1 \geq |\mathcal{A}|/N$. The looser these upper bounds compared to the true values, the more conservative is (12). This approach*

²If more than half of the agents communicate through compromised channels, then the adversary controls the majority and therefore the median, which will be used to estimate the average parameter later in the paper. In that case, there is no optimization based solution the central coordinator can implement in order to securely run the system.

potentially leaves less resources available to the agents communicating through trustworthy channels by assuming more than $|\mathcal{A}|$ number of agents having maximum possible impact on the constraints, irrespective of their set \mathcal{C}_i or the true value/gradient of the constraints.

Running Example (Conservative Approximation). With $B = 1$, $L = 0$, and $R = 10$, the conservative approximation of the running example has the following constraint:

$$\begin{aligned} \bar{g}((1 - \alpha_1)\bar{\theta}_{\mathcal{H}}) &= g((1 - \alpha_1)\bar{\theta}_{\mathcal{H}}) + \alpha_1(RB + \frac{1}{2}LR^2) \\ &= (1 - \alpha_1)\bar{\theta}_{\mathcal{H}} - 5 + 10\alpha_1 \end{aligned}$$

If $\alpha_1 = |\mathcal{A}|/N = 0.2$, then the upper bound is the fraction of compromised channels. In that case, the constraint is:

$$0.8\bar{\theta}_{\mathcal{H}} - 3 \leq 0 \Leftrightarrow \bar{\theta}_{\mathcal{H}} \leq 3.75,$$

which is more conservative compared to the constraint of the robust optimization model (which was $\bar{\theta}_{\mathcal{H}} \leq 4.5$). The optimal solution in this case is to deliver electricity at a rate of 3.75kW to the agents. The conservatism arises due to the difference between agent-specific maximum charging rate 7kW and the absolute maximum charging rate 10kW. Since the constraint is linear, the gradient is constant. Hence, the smoothness and Lipschitz bounds hold with equality without causing additional conservatism.

If however $\alpha_1 = 0.4$, then the central coordinator assumes two agents communicating through compromised channels. In this case the conservative approximation has the constraint as

$$0.6\bar{\theta}_{\mathcal{H}} - 1 \leq 0 \Leftrightarrow \bar{\theta}_{\mathcal{H}} \leq \frac{5}{3},$$

which results in charging at an even slower rate since the central coordinator has to be robust against two agents charging at the maximum rate of 10kW.

To develop an attack resilient PD-DRA algorithm, we again define the regularized Lagrangian function of (12):

$$\begin{aligned} \bar{\mathcal{L}}_v(\{\theta_i\}_{i \in \mathcal{H}}; \lambda; \mathcal{H}) \\ := \frac{1}{N} \sum_{i \in \mathcal{H}} f_i(\theta_i) + \sum_{t=1}^T \lambda_t \bar{g}_t((1 - \alpha_1)\bar{\theta}_{\mathcal{H}}) \\ + \frac{v}{2N} \sum_{i \in \mathcal{H}} \|\theta_i\|^2 - \frac{v}{2} \|\lambda\|^2. \end{aligned} \quad (13)$$

The above function is $(1 - \alpha_1)v$ -strongly convex and concave in θ and λ , respectively (since $(1 - \alpha_1) \leq \frac{|\mathcal{H}|}{N} \leq 1$). Our main task is to tackle the following modified problem of (P) under Byzantine attack on (some of) the uplinks:

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

Notice that (P'_v) bears a similar form as (P) and thus one may apply the PD-DRA method to the former. The gradients with respect to primal/dual variables are given by:

$$\begin{aligned} \nabla_{\theta_i} \bar{\mathcal{L}}_v(\{\theta_i^{(k)}\}_{i \in \mathcal{H}}; \lambda^{(k)}; \mathcal{H}) &= \frac{1}{N} \left(\nabla_{\theta_i} f_i(\theta_i^{(k)}) + v \theta_i^{(k)} \right. \\ &\quad \left. + \frac{(1 - \alpha_1)N}{|\mathcal{H}|} \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\theta} \bar{g}_t(\theta) \Big|_{\theta = (1 - \alpha_1)\bar{\theta}_{\mathcal{H}}^{(k)}} \right), \forall i \in \mathcal{H}, \end{aligned} \quad (14a)$$

$$[\nabla_{\lambda} \bar{\mathcal{L}}_v(\{\theta_i^{(k)}\}_{i \in \mathcal{H}}; \lambda^{(k)}; \mathcal{H})]_t = \bar{g}_t((1 - \alpha_1)\bar{\theta}_{\mathcal{H}}^{(k)}) - v \lambda_t^{(k)}. \quad (14b)$$

Algorithm 2: Robust PD-DRA Algorithm

- 1: **Input:** Each agent has initial state $\theta_i^{(0)}$.
 - 2: **for** $k = 1, 2, \dots$ **do**
 - 3: (At the Central Coordinator):
 - (a) Receives $\{r_i^{(k)}\}_{i=1}^N$, see (8), from agents.
 - (b) Computes robust mean $\hat{\theta}_{\mathcal{H}}^{(k)}$ using the estimator (16).
 - (c) Broadcasts the vector $\hat{g}_{\mathcal{H}}^{(k)} := \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\theta} \bar{g}_t((1 - \alpha_1)\hat{\theta}_{\mathcal{H}}^{(k)})$ to agents.
 - (d) Computes the update for $\lambda^{(k+1)}$ with (18b).
 - 4: (At each agent $i \in \mathcal{H}$):
 - (a) Agent receives $\hat{g}_{\mathcal{H}}^{(k)}$.
 - (b) Agent computes update for $\theta_i^{(k+1)}$ with (18a).
 - 5: **end for**
-

However, such application requires the central coordinator to compute the sample average

$$\bar{\theta}_{\mathcal{H}}^{(k)} = \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \theta_i^{(k)}, \quad (15)$$

at each iteration. The above might not be computationally feasible under the attack model, since the central coordinator is oblivious to the identity of \mathcal{H} . As a solution, the central coordinator computes the robust mean $\hat{\theta}_{\mathcal{H}}^{(k)}$ of the received parameters $\{r_i^{(k)}\}_{i \in [N]}$ using a median-based mean estimator described next.

1) *Overview of Median-Based Mean Estimation:* Consider a set of N vectors $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^N$, among which at least $(1 - \alpha_1)N$ are trustworthy ($\mathbf{x}_i \in \mathcal{H}$) and at most $\alpha_1 N$ are compromised ($\mathbf{x}_i \in \mathcal{A}$). We consider a simple median-based estimator applied to each coordinate $j = 1, \dots, d$. First, define the coordinate-wise median as:

$$[\mathbf{x}_{\text{med}}]_j = \text{med}(\{[\mathbf{x}_i]_j\}_{i=1}^N),$$

where $\text{med}(\cdot)$ computes the median of the operand. Then, our estimator is computed as the mean of the nearest $(1 - \alpha_1)N$ neighbors of $[\mathbf{x}_{\text{med}}]_j$. Our estimator is:

$$[\hat{\mathbf{x}}_{\mathcal{H}}]_j = \frac{1}{(1 - \alpha_1)N} \sum_{i \in \mathcal{N}_j} [\mathbf{x}_i]_j, \quad (16)$$

where we have defined the set with $|\mathcal{N}_j| = (1 - \alpha_1)N$ as:

$$\mathcal{N}_j = \{i \in [N] : |[\mathbf{x}_i - \mathbf{x}_{\text{med}}]_j| \leq r_j\},$$

such that r_j is chosen to satisfy $|\mathcal{N}_j| = (1 - \alpha_1)N$.

The following bounds the performance of (16):

Proposition 2. *Let $\bar{\mathbf{x}}_{\mathcal{H}}$ be the mean of the trustworthy vectors. Suppose that $\max_{i \in \mathcal{H}} \|\mathbf{x}_i - \bar{\mathbf{x}}_{\mathcal{H}}\|_{\infty} \leq r$, then for any $\alpha_1 \in (0, \frac{1}{2})$, it holds that:*

$$\|\hat{\mathbf{x}}_{\mathcal{H}} - \bar{\mathbf{x}}_{\mathcal{H}}\| \leq \frac{2\alpha_1}{1 - \alpha_1} \left(1 + \sqrt{\frac{(1 - \alpha_1)^2}{1 - 2\alpha_1}} \right) r \sqrt{d}. \quad (17)$$

The proof can be found in Appendix B of the online version [34]. We note that for sufficiently small α_1 , the right hand side on (17) can be approximated by $\mathcal{O}(\alpha_1 r \sqrt{d})$. Using this median-based mean estimator, we propose the robust PD-DRA algorithm as follows.

2) *Robust PD-DRA Algorithm*: We summarize the static impersonation attack resilient PD-DRA method in Algorithm 2. The algorithm behaves similarly as Algorithm 1 applied to (\mathbf{P}'_v) , 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, as illustrated in Figure 3. This approximate value is used to compute the new price signals, and sent back to agents. In particular, the primal-dual updates are:

$$\theta_i^{(k+1)} = \mathcal{P}_{C_i}(\theta_i^{(k)} - \frac{\gamma}{N}(\widehat{\mathbf{g}}_{\mathcal{H}}^{(k)} + \nabla_{\theta_i} f_i(\theta_i^{(k)}) + v\lambda_t^{(k)})), \quad (18a)$$

$$\lambda_t^{(k+1)} = [\lambda_t^{(k)} + \gamma(\widehat{g}_t((1 - \alpha_1)\widehat{\theta}_{\mathcal{H}}^{(k)} - v\lambda_t^{(k)}))]_+. \quad (18b)$$

We note that the update rule in (18a) is valid for agents in set \mathcal{H} , because the gradients of the Lagrangian are defined only for those agents in (14a). The agents in set \mathcal{A} may or may not use the same update rule, however, this does not have any impact on the algorithm as they can never communicate their true parameters to the central coordinator.

Lemma 2. *Algorithm 2 is a primal-dual algorithm [10] for (\mathbf{P}'_v) with perturbed gradients:*

$$\widehat{\mathbf{g}}_{\theta}^{(k)} = \nabla_{\theta} \overline{\mathcal{L}}_v(\theta^{(k)}; \lambda^{(k)}; \mathcal{H}) + \mathbf{e}_{\theta}^{(k)}, \quad (19a)$$

$$\widehat{\mathbf{g}}_{\lambda}^{(k)} = \nabla_{\lambda} \overline{\mathcal{L}}_v(\theta^{(k)}; \lambda^{(k)}; \mathcal{H}) + \mathbf{e}_{\lambda}^{(k)}, \quad (19b)$$

where we have used concatenated variable as $\theta = (\{\theta_i\}_{i \in \mathcal{H}})$. Under Assumption 1 and assuming that $\lambda_t^{(k)} \leq \bar{\lambda}$ for all k , we have:

$$\|\mathbf{e}_{\theta}^{(k)}\| \leq (1 - \alpha_1)\bar{\lambda}LT\|\widehat{\theta}_{\mathcal{H}}^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\| + \frac{|\mathcal{H}| - (1 - \alpha_1)N}{|\mathcal{H}|}\bar{\lambda}BT, \quad (20)$$

$$\|\mathbf{e}_{\lambda}^{(k)}\| \leq (1 - \alpha_1)BT\|\widehat{\theta}_{\mathcal{H}}^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\|. \quad (21)$$

The proof can be found in Appendix C of the online version [34]. The assumption $\lambda_t^{(k)} \leq \bar{\lambda}$ can be guaranteed since $\widehat{g}_t((1 - \alpha_1)\widehat{\theta}_{\mathcal{H}}^{(k)})$ is bounded, which is proven in Appendix H of the online version [34]. Furthermore, the performance analysis for the median based estimator shows that

$$\|\widehat{\theta}_{\mathcal{H}}^{(k)} - \bar{\theta}_{\mathcal{H}}^{(k)}\| = \mathcal{O}(\alpha_1 R\sqrt{d}) \quad (22)$$

when α_1 is small. Finally, based on Lemma 2, we can analyze the convergence of Algorithm 2. Let $\widehat{\mathbf{z}}^* = (\widehat{\theta}^*, \widehat{\lambda}^*)$ be a saddle point of (\mathbf{P}'_v) and define

$$\overline{\Phi}(\mathbf{z}^{(k)}) := \begin{pmatrix} \nabla_{\theta} \overline{\mathcal{L}}_v(\{\theta_i^{(k)}\}_{i \in \mathcal{H}}, \lambda^{(k)}; \mathcal{H}) \\ -\nabla_{\lambda} \overline{\mathcal{L}}_v(\{\theta_i^{(k)}\}_{i \in \mathcal{H}}, \lambda^{(k)}; \mathcal{H}) \end{pmatrix}. \quad (23)$$

We are ready to present our main result for static attacks.

Theorem 1. *Assume the map $\overline{\Phi}(\mathbf{z}^{(k)})$ is L_{Φ} -Lipschitz continuous. For Algorithm 2, for all $k \geq 0$ it holds:*

$$\|\mathbf{z}^{(k+1)} - \widehat{\mathbf{z}}^*\|^2 \leq (1 - \gamma v' + 2\gamma^2 L_{\Phi}^2)\|\mathbf{z}^{(k)} - \widehat{\mathbf{z}}^*\|^2 + \left(\frac{4\gamma}{v'} + 2\gamma^2\right)E_k. \quad (24)$$

where $v' := (1 - \alpha_1)v$ and $E_k := \|\mathbf{e}_{\theta}^{(k)}\|^2 + \|\mathbf{e}_{\lambda}^{(k)}\|^2$ is the total perturbation at iteration k . Moreover, if we choose $\gamma < v'/2L_{\Phi}^2$ and E_k is upper bounded by \bar{E} for all k , then

$$\limsup_{k \rightarrow \infty} \|\mathbf{z}^{(k)} - \widehat{\mathbf{z}}^*\|^2 \leq \frac{4}{v'} + 2\gamma \bar{E}. \quad (25)$$

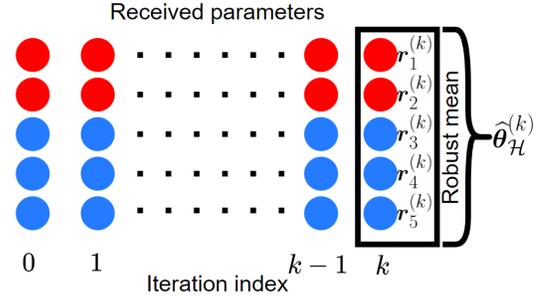


Fig. 3: Robust mean estimation under static impersonation attack. Red/blue circles correspond to parameters received through compromised/trustworthy channels, respectively. In this example, there are $N = 5$ agents and agents 1 and 2 are always communicating through compromised channels. At iteration k , the central coordinator computes the robust mean $\widehat{\theta}_{\mathcal{H}}^{(k)}$ of the received parameters $\{r_i^{(k)}\}_{i \in [N]}$.

The proof can be found in Appendix D of the online version [34]. Combining with (22) shows that the resilient PD-DRA method converges to a $\mathcal{O}(\alpha_1^2 R^2 d)$ neighborhood of the saddle point of (\mathbf{P}'_v) . Moreover, it shows that the convergence rate to the neighborhood is linear, which is similar to the classical analysis in [10].

B. Dynamic Impersonation Attack

Under this attack scenario, the set of agents communicating through compromised channels is dynamically changing with iterations. We make the following assumption on how frequently each agent's communications are compromised:

Assumption 2. *Let m be a fixed window size and $\alpha_2 < 0.5$ be a known upper bound on how frequent an agent communicates through a compromised channel. Then, for all $k \geq m - 1$ and for all agents $i \in [N]$, among the received parameters $\{r_i^{(k-\ell)}\}_{\ell=0}^{m-1}$ at most $\alpha_2 m$ are sent through compromised channels.*

It is important to recall that the dynamic attack scenario does not generalize the static attack scenario and there is a significant distinction between the two. The static attack scenario assumes that a fixed set of agents' communications are *permanently* compromised. It may occur when the attacker compromises set of communication channels and those channels are assigned to the agents via a static channel allocation scheme.

On the contrary, for the dynamic attack scenario, each user's communications are vulnerable to attacks for *at most* a given α_2 fraction of iterations over a window of size m under Assumption 2, and hence each agent is able to communicate reliably with the central coordinator at some iterations. This scheme may occur when the attacker compromises a fixed set of communication channels (same as the static scenario), however, the channels are assigned to the agents via a dynamic channel allocation scheme (e.g., do a round-robin channel allocation. If there are m communication channels out of which $\alpha_2 m$ are compromised, assigning channels dynamically in a cyclic way to the agents ensures that over a window of m , every agent has sent $\alpha_2 m$ compromised messages). Although the attacker behaves the same way, we can simulate both scenarios by static/dynamic channel allocation. In cyber-physical

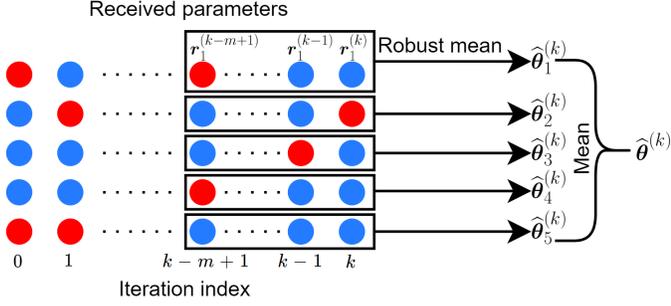


Fig. 4: Robust mean estimation under dynamic impersonation attacks. Red/blue circles correspond to parameters received through compromised/trustworthy channels, respectively. In this example, there are $N = 5$ agents and the set of agents communicating through compromised channels is changing at every iteration. At iteration k , the central coordinator computes the robust mean $\hat{\theta}_i^{(k)}$ of the received parameters $\{r_i^{(k-\ell)}\}_{\ell=0}^{m-1}$ for all agents $i \in [N]$. Then, computes the naive average of $\{\hat{\theta}_i^{(k)}\}_{i=1}^N$ to get the average parameter $\hat{\theta}^{(k)}$.

systems, such dynamic allocation schemes are commonly used (e.g., dynamic IP assignment to be protected from hackers).

Interestingly, it is possible to develop an algorithm that converges to the optimal solution of Problem (P_v). The intuition behind is that the received parameters over a long period of time contain a fraction of trustworthy information that can be extracted by the algorithm to perform faithful computations.

Our algorithm is similar in nature to an *averaging gradient* scheme where the primal-dual updates utilize the averages of time delayed gradients. Furthermore, the scheme is combined with the *robust mean* estimator developed in Sec. IV-A1 to approximate the averages of outdated gradients, as illustrated in Figure 4. Specifically, the central coordinator chooses a window size of m . For any iteration $k \geq m - 1$, instead of using $r_i^{(k)}$ for computing the average parameter $\bar{\theta}^{(k)}$ and the gradients, the central coordinator computes the robust mean $\hat{\theta}_i^{(k)}$ from the received parameters $\{r_i^{(k-\ell)}\}_{\ell=0}^{m-1}$ using the median-based mean estimator (16) for all agents $i \in [N]$, applied on the sequence of historical received parameters. Note that we have replaced α_1 by α_2 , N by m in this application. It then uses $\hat{\theta}^{(k)} := \frac{1}{N} \sum_{i=1}^N \hat{\theta}_i^{(k)}$ for computation of the primal-dual updates.

We summarize our robust averaging PD-DRA method in Algorithm 3. The primal-dual updates are described by:

$$\theta_i^{(k+1)} = \mathcal{P}_{C_i}(\theta_i^{(k)} - \frac{\gamma}{N}(\hat{g}^{(k)} + \nabla_{\theta_i} f_i(\theta_i^{(k)}) + v\theta_i^{(k)})), \quad (26a)$$

$$\lambda_t^{(k+1)} = [\lambda_t^{(k)} + \gamma(g_t(\hat{\theta}^{(k)}) - v\lambda_t^{(k)})]_+. \quad (26b)$$

Lemma 3. *Algorithm 3 is a primal-dual algorithm for (P_v) with perturbed gradients:*

$$\hat{g}_{\theta}^{(k)} = \nabla_{\theta} \mathcal{L}_v(\theta^{(k)}; \lambda^{(k)}) + e_{\theta}^{(k)}, \quad (27a)$$

$$\hat{g}_{\lambda}^{(k)} = \nabla_{\lambda} \mathcal{L}_v(\theta^{(k)}; \lambda^{(k)}) + e_{\lambda}^{(k)}, \quad (27b)$$

where we have used concatenated variable as $\theta = (\{\theta_i\}_{i \in N})$. Under Assumption 1 and assuming that $\lambda_t^{(k)} \leq \bar{\lambda}$ for all k , we have:

$$\|e_{\theta}^{(k)}\| \leq \frac{\bar{\lambda}LT}{N} \sum_{i=1}^N \|\theta_i^{(k)} - \hat{\theta}_i^{(k)}\|, \quad (28a)$$

$$\|e_{\lambda}^{(k)}\| \leq \frac{BT}{N} \sum_{i=1}^N \|\theta_i^{(k)} - \hat{\theta}_i^{(k)}\|. \quad (28b)$$

Algorithm 3: Averaging PD-DRA Algorithm

- 1: **Input:** Each agent has initial state $\theta_i^{(0)}$.
- 2: **for** $k = 0, 1, \dots, m - 2$ **do**
- 3: Apply basic PD-DRA (Run Algorithm 1).
- 4: **end for**
- 5: **for** $k = m - 1, m, \dots$ **do**
- 6: *(At the Central Coordinator):*
 - (a) Receives $\{r_i^{(k)}\}_{i=1}^N$, see (8), from agents.
 - (b) For all agents $i = 1, \dots, N$, computes robust mean $\hat{\theta}_i^{(k)}$ of $\{r_i^{(k-\ell)}\}_{\ell=0}^{m-1}$ using the estimator (16) with parameters $\alpha_1 \rightarrow \alpha_2, N \rightarrow m$.
 - (c) Computes $\hat{\theta}^{(k)} := \frac{1}{N} \sum_{i=1}^N \hat{\theta}_i^{(k)}$.
 - (d) Broadcasts the vector $\hat{g}^{(k)} := \sum_{t=1}^T \lambda_t^{(k)} \nabla_{\theta} g_t(\hat{\theta}^{(k)})$ to agents.
 - (e) Computes the update for $\lambda^{(k+1)}$ with (26b).
- 7: *(At each agent i):*
 - (a) Agent receives $\hat{g}^{(k)}$.
 - (b) Agent computes update for $\theta_i^{(k+1)}$ with (26a).
- 8: **end for**

The proof can be found in Appendix E of the online version [34]. The assumption $\lambda_t^{(k)} \leq \bar{\lambda}$ can be guaranteed since $g_t(\hat{\theta}^{(k)})$ is bounded, which is proven in Appendix H of the online version [34]. Let $z^{(k)} := (\{\theta_i^{(k)}\}_{i=1}^N, \lambda^{(k)})$ be the primal-dual variable at the k th iteration and define the mapping $\Phi(z^{(k)})$ as in (6). We observe that the algorithm's behavior is similar to the incremental aggregated gradient method in [35]–[37]. The following Lemma, which is inspired by [35]–[37], upper bounds the perturbation in the gradients in (28) by the maximum optimality gap in a finite window of size $2m - 1$:

Lemma 4. *Assume the map $\Phi(z^{(k)})$ is L_{Φ} -Lipschitz continuous. Let $E_k := \|e_{\theta}^{(k)}\|^2 + \|e_{\lambda}^{(k)}\|^2$. Then, for all $k \geq 2(m - 1)$ we have:*

$$E_k \leq \gamma^2 \bar{C} \max_{0 \leq \ell \leq 2(m-1)} \|z^{(k-\ell)} - z^*\|^2, \quad (29)$$

where

$$\bar{C} = \left(\frac{T^2(\bar{\lambda}^2 L^2 + B^2)}{N} \right) \times \left(\frac{1}{L_{\Phi}} + (1 + \sqrt{d})\bar{\lambda}LT \right)^2 \times \left(\frac{1 + C_{\alpha}}{1 - \alpha_2} + C_{\alpha} \right)^2 \times (m - 1)^2, \quad (30)$$

and

$$C_{\alpha} = \frac{2\alpha_2}{1 - \alpha_2} \left(1 + \sqrt{\frac{(1 - \alpha_2)^2}{1 - 2\alpha_2}} \right) \sqrt{d}.$$

The proof can be found in Appendix F of the online version [34]. Using on Lemmas 3 and 4, we can analyze the converge of Algorithm 3:

Theorem 2. *Assume the map $\Phi(z^{(k)})$ is L_{Φ} -Lipschitz continuous. For Algorithm 3, for all $k \geq 2(m - 1)$ it holds that:*

$$\|z^{(k+1)} - z^*\|^2 \leq (1 - \gamma v + 2\gamma^2 L_{\Phi}^2) \|z^{(k)} - z^*\|^2 + \left(\frac{4\gamma}{v} + 2\gamma^2 \right) \gamma^2 \bar{C} \max_{0 \leq \ell \leq 2(m-1)} \|z^{(k-\ell)} - z^*\|^2. \quad (31)$$

Moreover, if we choose γ sufficiently small such that it satisfies

$$v - 2\gamma L_{\Phi}^2 - \frac{4\bar{C}\gamma^2}{v} - 2\bar{C}\gamma^3 > 0,$$

then:

$$\|z^{(k)} - z^*\|^2 \leq \rho^{k-2(m-1)} \|z^{(2(m-1))} - z^*\|^2, \quad (32)$$

and

$$\lim_{k \rightarrow \infty} \|z^{(k)} - z^*\|^2 = 0, \quad (33)$$

where $\rho = (1 - \gamma v + 2\gamma^2 L_{\Phi}^2 + \frac{4\bar{C}\gamma^3}{v} + 2\bar{C}\gamma^4)^{\frac{1}{1+2(m-1)}}$.

The proof can be found in Appendix G of the online version [34]. Theorem 2 shows that the robust averaging PD-DRA method converges *geometrically* to the optimal solution of (P_v) under said assumptions.

C. Remarks

A few remarks highlighting design criteria to be explored in practical implementations are in order:

- Theorem 1 illustrates a trade-off in the choice of the step size γ between convergence speed and accuracy. In particular, (24) shows that the rate of convergence factor $1 - \gamma v + 2\gamma^2 L_{\Phi}^2$ can be minimized by setting $\gamma = v/(4L_{\Phi}^2)$. Meanwhile, the asymptotic upper bound in (25) is increasing with γ and it can be minimized by setting $\gamma \rightarrow 0$.
- Theorem 2 illustrates a trade-off between the window size m and the convergence rate. Observe that increasing the window size m decreases the rate of convergence by increasing ρ (Equations (30) and (32)). On the other hand, the likelihood that Assumption 2 holds true in a stochastic setting (e.g., channels being compromised with some probability) increases with a larger window size m .
- Under the dynamic impersonation attack scenario, the choice of α_2 does not affect convergence accuracy to the saddle point of (P_v) , but only changes the convergence rate. As such, choosing the largest α_2 such that $\alpha_2 m = \lfloor \frac{m-1}{2} \rfloor$ (i.e., assuming maximum possible number of iterates received through compromised channels) makes the algorithm robustly applicable to all dynamic impersonation attack scenarios regardless of the frequency of the attack.
- In case the central coordinator can not identify the attack scenario as static or dynamic impersonation (or the attack can be a mixture of both), a mixture of both Algorithms 2 and 3 can be applied. In particular, this can be done by adding Step 6(b) of Algorithm 3 before Step 3(b) of Algorithm 2, and applying the rest of the Algorithm 2 as it is. The central coordinator first computes robust parameters $\hat{\theta}_i^{(k)}$ by computing the robust mean of $\{r_i^{(k-\ell)}\}_{\ell=0}^{m-1}$ for all agents, and then computes the robust mean of $\{\hat{\theta}_i^{(k)}\}_{i=1}^N$. This effectively makes Algorithm 2 robust to possible dynamic impersonation attacks on uplink channels that are thought to be trustworthy for all iterations.

V. NUMERICAL STUDY

In this section, we demonstrate the performance of our methods and verify our theoretical claims by applying our

algorithms for: 1) an electric vehicle (EV) charging coordinator under static impersonation attack, 2) an electric vehicle charging coordinator under dynamic impersonation attack, and 3) a power distribution network with flexible demand under dynamic impersonation attack. The EV charging coordinator problem resembles classic network utility maximization problems such as those studied in communication networks whereas the power distribution network problem has more nuisances that we will discuss next. To solve the convex optimization problems in order to get the optimal solutions, we used CVX, a package for specifying and solving convex programs [38].

A. Electric Vehicle (EV) Charging Facility

In this study, the aim is to optimize EV charging demand over time. We consider multiple EVs receiving charge under the same local feeder/transformer. Each agent (or EV owner) has different utility of charging at different times. Hence, at a given time period, it is desired to charge those EVs who have a higher utility (or less cost) for that time period. This problem falls into the broad category of network utility maximization problems, which can be formulated as:

$$\min_{\theta_i \in \mathbb{R}_+^d, \forall i} f(\theta) = \frac{1}{N} \sum_{i=1}^N f_i(\theta_i) \quad (34a)$$

$$\text{subject to} \quad \frac{1}{N} \sum_{i=1}^N \theta_i \preceq \bar{e}, \quad (34b)$$

$$\theta_i^{\min} \preceq \theta_i \preceq \theta_i^{\max}, \quad \forall i, \quad (34c)$$

$$\Theta_i^{\min} \leq \mathbf{1}^T \theta_i \leq \Theta_i^{\max}, \quad \forall i, \quad (34d)$$

where $N \times \bar{e} \in \mathbb{R}^d$ is the vector of maximum available transformer capacity in all time periods and \preceq denotes component-wise inequality between the vectors. The available capacity changes with time of day as exogenous load on the transformer varies with time as well. The elements $\{\theta_{i,j}\}_{j=1}^d$ of the vector θ_i correspond to the electricity demand of the EV i at time slots $j = 1 \dots d$. The constraint (34c) restricts the amount an EV can charge at each time slot, whereas the constraint (34d) bounds the total amount an EV can charge. For this study, we set the cost function to be:

$$f_i(\theta) = - \sum_{j=1}^d \beta_{i,j} \log \theta_{i,j}, \quad (35)$$

where $\beta_{i,j}$ are generated randomly from a uniform distribution in $[0, 1]$. We study this problem under both attack scenarios for $N = 100$ EVs.

1) *Static impersonation attack:* We simulated various static impersonation attack scenarios and ran Algorithm 2. The results are displayed in Figure 5.

In Figure 5a, we plot agent 70's electricity demand for some time periods, with $|\mathcal{A}|/N = 0.2$ and $\alpha_1 = 0.3$. Each different color corresponds to a different dimension of the parameter vector (i.e., electricity demand for different time periods). A colored solid line corresponds to a dimension of the parameter vector iterates generated by the algorithm. A dashed line with the same marker and color as a solid line is

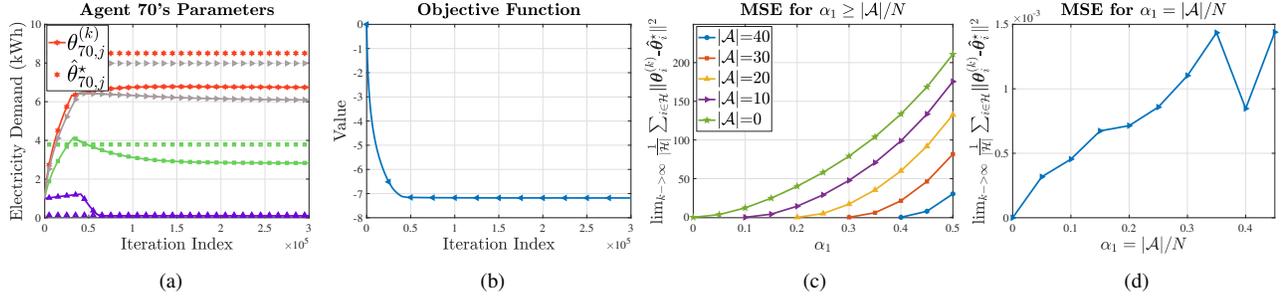
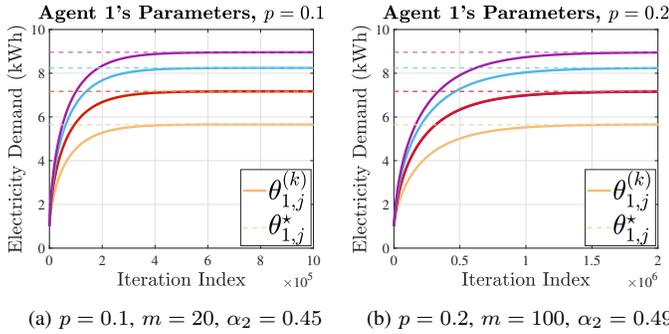


Fig. 5: Numerical study results for optimal electric vehicle charging under static impersonation attack. (a) Optimal parameter of agent 70 converges to a neighborhood of the optimal solution of the robust optimization problem for $|\mathcal{A}|/N = 0.2$ and $\alpha_1 = 0.3$, (b) The algorithm provides convergence of the objective function value, (c) Mean squared error for different number of compromised channels and different choices of upper bound α_1 , (d) Mean squared error when $\alpha_1 = |\mathcal{A}|/N$.



(a) $p = 0.1$, $m = 20$, $\alpha_2 = 0.45$ (b) $p = 0.2$, $m = 100$, $\alpha_2 = 0.49$

Fig. 6: Numerical study results demonstrating convergence of Algorithm 3 for optimal electric vehicle charging under two dynamic impersonation attack scenarios: (a) $p = 0.1$, (b) $p = 0.2$. Observe that the number of iterations it takes to converge for (b) is much larger than for (a).

the optimal value corresponding to that dimension of the parameter vector, which is the solution of the regularized robust optimization problem (formulated as (P'_v)) of (34). Observe that Algorithm 2 successfully provides convergence to a close neighborhood of the optimal solution of the regularized robust optimization problem. Furthermore, in Figure 5b we show that the objective function value converges, as opposed to a non-resilient PD-DRA method that is shown to oscillate and violate the constraint in Figure 2. Our robust optimization model on the other hand ensures there is no constraint violation.

In Figure 5c, we plot the mean squared error (MSE) in primal variables θ_i for different number of compromised channels $|\mathcal{A}|$ and different choices of α_1 , which is the upper bound on fraction of compromised links known by the central coordinator. The MSE is calculated by:

$$\text{MSE} = \lim_{k \rightarrow \infty} \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|\theta_i^{(k)} - \hat{\theta}_i^*\|^2, \quad (36)$$

where $\hat{\theta}_i^*$ is the solution to (P'_v) with $\alpha_1 = |\mathcal{A}|/N$, i.e., the solution to the regularized and robustified problem with the knowledge of the compromised channels. Naturally, the looser the upper bound, the larger the error, since it increases the amount of conservatism. Hence, having an accurate upper bound on fraction of compromised channels significantly improves the performance.

Finally, in Figure 5d we exhibit the efficacy of our approach with median-based mean estimation. We plot the mean squared error in primal variables, when the upper bound on α_1 is tight, i.e., $\alpha_1 = |\mathcal{A}|/N$. The error tends to increase with $|\mathcal{A}|/N$, however, considering the magnitude, the error is negligible

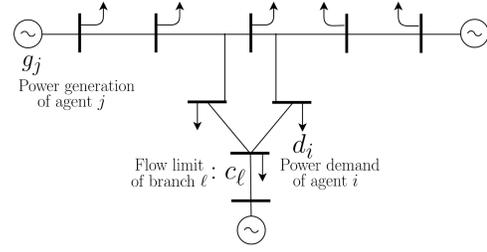


Fig. 7: IEEE 9 bus system with 3 generators (supplies) represented by sources and 8 loads (demands) represented by arrows.

and we can conclude that the median-based mean estimator performs well.

2) *Dynamic impersonation attack*: We simulated a dynamic impersonation attack scenario and ran Algorithm 3. To simulate a dynamic impersonation attack, we assigned a probability p for an uplink to be compromised at each iteration³. For $p = 0.1$, we picked a window size $m = 20$ and $\alpha_2 = 0.45$, whereas for $p = 0.2$, we picked a window size $m = 100$ and $\alpha_2 = 0.49$. The results are displayed in Figure 6. Each different color corresponds to a different dimension of the parameter vector. A colored solid line corresponds to a dimension of the parameter vector iterates generated by the algorithm. A dashed line with the same color as a solid line is the optimal value corresponding to that dimension of the parameter vector, which is the solution of the regularized optimization problem (formulated as (P_v)) of (34).

In both scenarios, Algorithm 3 successfully provides convergence to the optimal solution of the regularized problem. Observe that for $p = 0.2$, we chose a larger window size and a larger α_2 in order to meet Assumption 2. However, this restricts us to choose a smaller step size γ as dictated by Theorem 2 and in turn slower convergence. This highlights an important trade-off between robustness and convergence rate, where a larger window size m and larger α_2 makes the algorithm more robust while decreasing the convergence rate.

B. Power Distribution Network

We consider the IEEE $N = 9$ bus system with $N_g = 3$ generators and $N_\ell = 8$ loads as shown in Figure 7. The power

³Although a probabilistic scenario does not guarantee that Assumption 2 holds, with sufficiently large window size m and α_2 , it holds with high probability at each iteration. Even though we do not study this scenario theoretically, our algorithm still performs well.

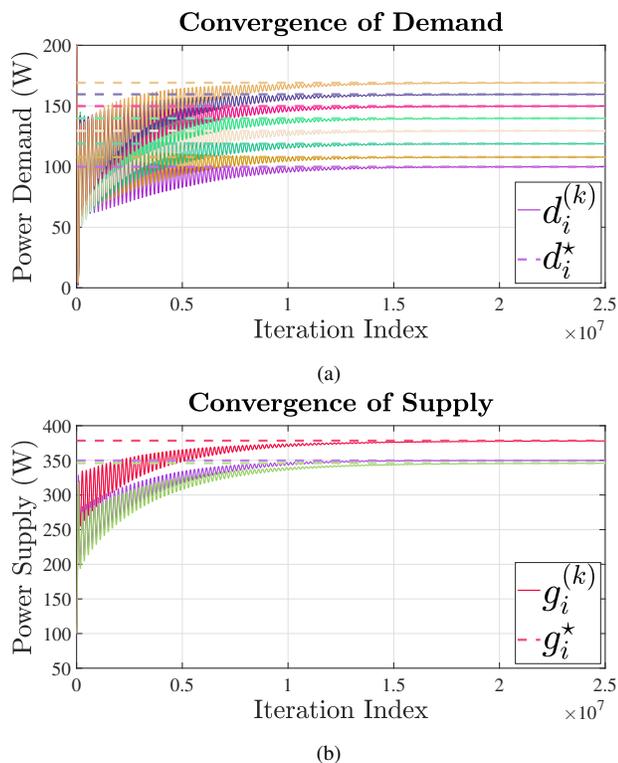


Fig. 8: Numerical study results for power network under dynamic impersonation attack: (a)/(b) displaying convergence of the demand/supply, respectively.

network cost minimization problem can be stated as:

$$\min_{d_i, g_i \in \mathbb{R}^+} f(\mathbf{d}, \mathbf{g}) = - \sum_{i=1}^{N_l} U_i(d_i) + \sum_{i=1}^{N_g} C_i(g_i) \quad (37a)$$

$$\text{subject to } \mathbf{1}^T(\mathbf{d} - \mathbf{g}) = 0, \quad (37b)$$

$$\mathbf{H}(\mathbf{d} - \mathbf{g}) \leq \mathbf{c}, \quad (37c)$$

where $\mathbf{d} = [d_1 \dots d_N]^T$ and $\mathbf{g} = [g_1 \dots g_N]^T$ are the vectors of load and generation at each node, respectively ($d_i = 0$ for nodes without load and $g_j = 0$ for nodes without generators). The first constraint (37b) ensures the power supply is equal to the demand, and the second constraint (37c) is the power flow constraint limiting the power flow on each branch.

Observe that the formulation in (37) does not directly match with our general formulation in (1) mainly due to the presence of equality constraint (37b), which prevents the application of the robustified formulation in (9) and hence the robust PD-DRA algorithm for static impersonation attacks. Nevertheless, our algorithm for dynamic impersonation attacks can still be applied since it does not require any robustified constraints (which cannot be done for equality constraints).

We have chosen the utility function for load i to be $U_i(d_i) = \beta_i \log d_i$ and randomly generated β_i from a uniform distribution in $[500, 1000]$. For generators, we set the cost function $C_i(g_i) = e^{c_i g_i}$, where $c_1 = 0.01$, $c_2 = 0.011$, $c_3 = 0.012$. We obtained the Power Transfer Distribution Factor (PTDF) matrix \mathbf{H} and the vector of flow limits \mathbf{c} from MATPOWER [39]. To simulate a dynamic impersonation attack scenario, we assigned a probability p for an uplink to be compromised at each iteration. We ran Algorithm 3 for $p = 0.15$, $m = 75$ and $\alpha_2 = 0.49$.

The results are shown in Figure 8. In both Figures 8a and 8b, each different color corresponds to a different agent. A colored solid line corresponds to an agent's parameter iterates generated by the algorithm. A dashed line with the same color as a solid line is the to the optimal value of that agent's parameter, which is the solution of the regularized optimization problem (formulated as (P_v)) of (37). Our algorithm successfully generates sequences that convergence to the optimal solution of the regularized problem for both power supplying and power demanding agents.

VI. CONCLUSION

In this paper, we studied two strategies for establishing primal-dual algorithms for resource allocation in presence of Byzantine attackers. Specifically, we consider static and dynamic impersonation attack scenarios and propose an attack-resilient primal-dual algorithm for each scenario based on robust mean estimation techniques. We derive bounds for the performance (in terms of distance to optimality) of the proposed algorithms and show that our algorithm for static impersonation attack converges to a neighborhood of the optimal solution of the regularized and robustified resource allocation problem, whereas our algorithm for dynamic impersonation attack converges to the optimal solution of the original regularized problem. We verify our theoretical results via computational simulations for network utility maximization problems involving optimal distributed resource allocation, such as power distribution networks.

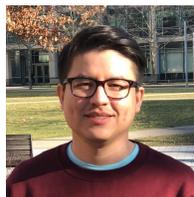
REFERENCES

- [1] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *J OPER RES SOC*, vol. 49, no. 3, pp. 237–252, Mar 1998.
- [2] S. H. Low and D. E. Lapsley, "Optimization flow control. i. basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, Dec 1999.
- [3] 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 SmartGridComm*, Oct 2010, pp. 415–420.
- [4] 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*, July 2011, pp. 1–8.
- [5] 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, March 2004, pp. 2800–2811 vol.4.
- [6] M. Dehghan, L. Massoulie, D. Towsley, D. Menasche, and Y. C. Tay, "A utility optimization approach to network cache design," in *IEEE INFOCOM*, April 2016, pp. 1–9.
- [7] 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*, ser. ITC '11. International Teletraffic Congress, 2011, pp. 238–245.
- [8] 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, no. 5, pp. 1307–1319, May 2016.
- [9] N. Mehr, J. Lioris, R. Horowitz, and R. Pedarsani, "Joint perimeter and signal control of urban traffic via network utility maximization," in *IEEE ITSC*, Oct 2017, pp. 1–6.
- [10] 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.
- [11] 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.
- [12] 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.

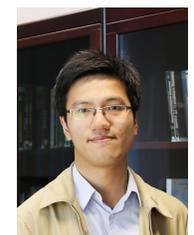
- [13] F. Pasqualetti, A. Bicchi, and F. Bullo, "Consensus computation in unreliable networks: A system theoretic approach," *IEEE TAC*, vol. 57, no. 1, pp. 90–104, 2012.
- [14] R. Gentz, S. X. Wu, H.-T. Wai, A. Scaglione, and A. Leshem, "Data injection attacks in randomized gossiping," *IEEE TSIPN*, vol. 2, no. 4, pp. 523–538, 2016.
- [15] S. Sundaram and B. Gharesifard, "Distributed optimization under adversarial nodes," *IEEE TAC*, 2018.
- [16] Y. Chen, S. Kar, and J. Moura, "Resilient distributed estimation: Sensor attacks," *IEEE Transactions on Automatic Control*, 2018.
- [17] W. Ben-Ameur, P. Bianchi, and J. Jakubowicz, "Robust distributed consensus using total variation," *IEEE Transactions on Automatic Control*, vol. 61, no. 6, pp. 1550–1564, 2016.
- [18] H. J. LeBlanc, H. Zhang, X. Koutsoukos, and S. Sundaram, "Resilient asymptotic consensus in robust networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 4, pp. 766–781, April 2013.
- [19] J. S. Baras and X. Liu, "Trust is the cure to distributed consensus with adversaries," in *2019 27th Mediterranean Conference on Control and Automation (MED)*, July 2019, pp. 195–202.
- [20] N. Ravi, A. Scaglione, and A. Nedić, "A case of distributed optimization in adversarial environment," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2019, pp. 5252–5256.
- [21] J. Feng, H. Xu, and S. Mannor, "Distributed robust learning," *arXiv preprint arXiv:1409.5937*, 2014.
- [22] D. Yin, Y. Chen, K. Ramchandran, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," *arXiv preprint arXiv:1803.01498*, 2018.
- [23] D. Alistarh, Z. Allen-Zhu, and J. Li, "Byzantine stochastic gradient descent," in *NeurIPS*, 2018, pp. 4618–4628.
- [24] Y. Chen, L. Su, and J. Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 2, pp. 44:1–44:25, 2017.
- [25] D. Data, L. Song, and S. Diggavi, "Data encoding methods for byzantine-resilient distributed optimization," in *2019 IEEE International Symposium on Information Theory (ISIT)*, July 2019, pp. 2719–2723.
- [26] D. L. Donoho and P. J. Huber, "The notion of breakdown point," *A festschrift for Erich L. Lehmann*, vol. 157184, 1983.
- [27] P. J. Huber, *Robust statistics*. Springer, 2011.
- [28] S. Minsker *et al.*, "Geometric median and robust estimation in banach spaces," *Bernoulli*, vol. 21, no. 4, pp. 2308–2335, 2015.
- [29] 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*, 2016, pp. 655–664.
- [30] J. Steinhardt, M. Charikar, and G. Valiant, "Resilience: A criterion for learning in the presence of arbitrary outliers," *arXiv preprint arXiv:1703.04940*, 2017.
- [31] C. A. Uribe, H. Wai, and M. Alizadeh, "Resilient distributed optimization algorithms for resource allocation," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 8341–8346.
- [32] Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical programming*, vol. 103, no. 1, pp. 127–152, 2005.
- [33] C. A. Uribe, S. Lee, A. Gasnikov, and A. Nedić, "A dual approach for optimal algorithms in distributed optimization over networks," *arXiv preprint arXiv:1809.00710*, 2018.
- [34] B. Turan, C. A. Uribe, H.-T. Wai, and M. Alizadeh, "Resilient primal-dual optimization algorithms for distributed resource allocation," *arXiv preprint arXiv:2001.00612*, 2020.
- [35] D. Blatt, A. O. Hero, and H. Gauchman, "A convergent incremental gradient method with a constant step size," *SIAM Journal on Optimization*, vol. 18, no. 1, pp. 29–51, 2007.
- [36] M. Gürbüzbalaban, A. Ozdaglar, and P. A. Parrilo, "On the convergence rate of incremental aggregated gradient algorithms," *SIAM Journal on Optimization*, vol. 27, no. 2, pp. 1035–1048, 2017.
- [37] P. Tseng and S. Yun, "Incrementally updated gradient methods for constrained and regularized optimization," *Journal of Optimization Theory and Applications*, vol. 160, no. 3, pp. 832–853, Mar 2014.
- [38] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," <http://cvxr.com/cvx>, Mar. 2014.
- [39] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "Matpower: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Transactions on Power Systems*, vol. 26, no. 1, pp. 12–19, Feb 2011.
- [40] H. R. Feyzmahdavian, A. Aytakin, and M. Johansson, "A delayed proximal gradient method with linear convergence rate," in *2014 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, Sep. 2014, pp. 1–6.



BERKAY TURAN is pursuing the Ph.D. degree in Electrical and Computer Engineering at the University of California, Santa Barbara. He received the B.Sc. degree in Electrical and Electronics Engineering as well as the B.Sc. degree in Physics degree from Boğaziçi University, Istanbul, Turkey, in 2018. His research interests include optimization and reinforcement learning for the design, control, and analysis of smart infrastructure systems such as the power grid and transportation systems.



CÉSAR A. URIBE received the M.Sc. degrees in systems and control from Delft University of Technology, in The Netherlands, and in applied mathematics from the University of Illinois at Urbana-Champaign, in 2013 and 2016, respectively. He also received the PhD degree in electrical and computer engineering at the University of Illinois at Urbana-Champaign in 2018. He is currently a Postdoctoral Associate in the Laboratory for Information and Decision Systems-LIDS at the Massachusetts Institute of Technology-MIT. His research interests include distributed learning and optimization, decentralized control, algorithm analysis, and computational optimal transport.



Hoi-To Wai (S'11–M'18) received his PhD degree from Arizona State University (ASU) in Electrical Engineering in Fall 2017, B. Eng. (with First Class Honor) and M. Phil. degrees in Electronic Engineering from The Chinese University of Hong Kong (CUHK) in 2010 and 2012, respectively. He is an Assistant Professor in the Department of Systems Engineering & Engineering Management at CUHK. He has held research positions at ASU, UC Davis, Telecom ParisTech, Ecole Polytechnique, LIDS, MIT. Hoi-To's research interests are in the broad area of signal processing, machine learning and distributed optimization, with a focus of their applications to network science. His dissertation has received the 2017's Dean's Dissertation Award from the Ira A. Fulton Schools of Engineering of ASU and he is a recipient of a Best Student Paper Award at ICASSP 2018.



MAHNOOSH ALIZADEH is an assistant professor of Electrical and Computer Engineering at the University of California Santa Barbara. She received the B.Sc. degree ('09) in Electrical Engineering from Sharif University of Technology and the M.Sc. ('13) and Ph.D. ('14) degrees in Electrical and Computer Engineering from the University of California Davis, where she was the recipient of the Richard C. Dorf award for outstanding research accomplishment. From 2014 to 2016, she was a postdoctoral scholar at Stanford University. Her research interests are focused on designing network control, optimization, and learning frameworks to promote efficiency and resiliency in societal-scale cyber-physical systems. Dr. Alizadeh is a recipient of the NSF CAREER award and the best paper award from HICSS-53 power systems track.