

Distributed Learning in Non-Convex Environments

– Part I: Agreement at a Linear Rate

Stefan Vlaski, *Student Member, IEEE*, and Ali H. Sayed, *Fellow, IEEE*

Abstract—Driven by the need to solve increasingly complex optimization problems in signal processing and machine learning, there has been increasing interest in understanding the behavior of gradient-descent algorithms in non-convex environments. Most available works on distributed non-convex optimization problems focus on the deterministic setting where exact gradients are available at each agent. In this work and its Part II, we consider stochastic cost functions, where exact gradients are replaced by stochastic approximations and the resulting gradient noise persistently seeps into the dynamics of the algorithm. We establish that the diffusion learning strategy continues to yield meaningful estimates non-convex scenarios in the sense that the iterates by the individual agents will cluster in a small region around the network centroid. We use this insight to motivate a short-term model for network evolution over a finite-horizon. In Part II [3] of this work, we leverage this model to establish descent of the diffusion strategy through saddle points in $O(1/\mu)$ steps, where μ denotes the step-size, and the return of approximately second-order stationary points in a polynomial number of iterations.

Index Terms—Stochastic optimization, adaptation, non-convex cost, gradient noise, stationary points, distributed optimization, diffusion learning.

I. INTRODUCTION

THE broad objective of distributed adaptation and learning is the solution of global, stochastic optimization problems by networked agents through localized interactions and in the absence of information about the statistical properties of the data. When constant, rather than diminishing, step-sizes are employed, the resulting algorithms are adaptive in nature and are able to adapt to drifts in the data statistics. In this work, we consider a collection of N agents, where each agent k is equipped with a stochastic risk of the form $J_k(w) = \mathbb{E}_{x_k} Q_k(w; x_k)$ with $Q_k(w; x_k)$ referring to the loss function, $w \in \mathbb{R}^M$ denoting a parameter vector, and x_k referring to the stochastic data. The expectation is over the probability distribution of the data. The objective of the network is to seek the Pareto solution:

$$\min_w J(w), \quad \text{where } J(w) \triangleq \sum_{k=1}^N p_k J_k(w) \quad (1)$$

where the p_k are positive weights that are normalized to add up to one and will be specified further below; in particular,

The authors are with the Institute of Electrical Engineering, École Polytechnique Fédérale de Lausanne. S. Vlaski is also with the Department of Electrical Engineering, University of California, Los Angeles. This work was supported in part by NSF grant CCF-1524250. Emails: {stefan.vlaski, ali.sayed}@epfl.ch. A limited short version of this work appears in the conference publication [1], and a summary of some of the results without proofs appears in [2].

in the special case when the $\{p_k\}$ are identical, they can be removed from (1). Algorithms for the solution of (1) have been studied extensively over recent years both with inexact [4]–[7] and exact [8]–[10] gradients. Here, we focus on the following diffusion strategy, which has been shown in previous works to provide enhanced performance and stability guarantees under constant step-size learning and adaptive scenarios [5], [11]:

$$\phi_{k,i} = w_{k,i-1} - \mu \widehat{\nabla} J_k(w_{k,i-1}) \quad (2a)$$

$$w_{k,i} = \sum_{\ell=1}^N a_{\ell k} \phi_{\ell,i} \quad (2b)$$

where $\widehat{\nabla} J_k(\cdot)$ denotes a stochastic approximation for the true local gradient $\nabla J_k(\cdot)$ and i denotes the iteration index. The intermediate estimate $\phi_{k,i}$ is obtained at agent k by taking a stochastic gradient update relative to the local cost $J_k(\cdot)$. The intermediate estimates are then fused across local neighborhoods where $a_{\ell k}$ are convex combination weights satisfying:

$$a_{\ell k} \geq 0, \quad \sum_{\ell \in \mathcal{N}_k} a_{\ell k} = 1, \quad a_{\ell k} = 0 \text{ if } \ell \notin \mathcal{N}_k \quad (3)$$

The symbol $\mathcal{N}_k$ denotes the set of neighbors of agent k .

Assumption 1 (Strongly-connected graph). We shall assume that the graph described by the weighted combination matrix $A = [a_{\ell k}]$ is strongly-connected [5]. This means that there exists a path with nonzero weights between any two agents in the network and, moreover, at least one agent has a nontrivial self-loop, $a_{kk} > 0$. ■

It then follows from the Perron-Frobenius theorem [5], [12], [13] that A has a single eigenvalue at one while all other eigenvalues are strictly inside the unit circle, so that $\rho(A) = 1$. Moreover, if we let p denote the right-eigenvector of A that is associated with the eigenvalue at one, and if we normalize the entries of p to add up to one, then it also holds that all entries of p are strictly positive, i.e.,

$$Ap = p, \quad \mathbf{1}^T p = 1, \quad p_k > 0 \quad (4)$$

where the $\{p_k\}$ denote the individual entries of the Perron vector, p .

Notation: We denote random variables using boldface letters, while deterministic quantities are set using regular font. In the absence of further specification, $\|\cdot\|$ denotes the Euclidean norm. For two matrices A and B , we denote $A \geq B$, if $A - B$ is positive semi-definite. We say that $f(\mu) \in O(g(\mu))$ if $\lim_{\mu \rightarrow 0} \frac{|f(\mu)|}{g(\mu)} \leq C$ for some constant independent of μ ,

and $f(\mu) \in o(g(\mu))$ if $\lim_{\mu \rightarrow 0} \frac{f(\mu)}{g(\mu)} = 0$. We use $O(g(\mu))$ and $o(g(\mu))$ as placeholders for terms $f(\mu) \in O(g(\mu))$ and $f(\mu) \in o(g(\mu))$, respectively. For example, $u(\mu) \leq v(\mu) + o(g(\mu))$ means that there exists $f(\mu) \in o(g(\mu))$, such that $u(\mu) \leq v(\mu) + f(\mu)$.

A. Related Works

The performance of the diffusion algorithm (2a)–(2b) has been studied extensively in differentiable settings [6], [11], with extensions to multi-task [14], constrained [15], and non-differentiable [16] environments. A common assumption in these works, along with others studying the behavior of distributed optimization algorithms in general, is that of *convexity* (or strong-convexity) of the aggregate risk $J(w)$. While many problems of interest such as least-squares estimation [5], logistic regression [5], and support vector machines [17] are convex, there has been recent heightened interest in the optimization of *non-convex* cost functions. Such problems arise naturally in the design of robust estimators [18] and the training of complex machine learning architectures such as those involving dictionary learning [19] and artificial neural networks [20].

Motivated by these applications, recent works have pursued the study of optimization algorithms for non-convex problems, both in the centralized and distributed settings [21]–[39]. While some works focus on establishing convergence to a stationary point [31]–[33], there has been growing interest in examining the ability of gradient descent implementations to escape from saddle points, since such points represent bottlenecks to the underlying learning problem [20]. We defer a detailed discussion on the plethora of related works on second-order guarantees [21]–[30], [35], [36], [40] to Part II [3], where we will be establishing the ability of the diffusion strategy (2a)–(2b) to escape strict-saddle points efficiently. For ease of reference, the modeling conditions and results from this and related works are summarized in Table I.

The key contributions of Parts I and II this work are three-fold. To the best of our knowledge, we present the first analysis establishing *efficient* (i.e., polynomial) escape from strict-saddle points in the *distributed* setting. Second, we establish that the gradient noise process is sufficient to ensure efficient escape without the need to alter it by adding artificial forms of perturbations, interlacing steps with small and large step-sizes, or imposing a dispersive noise assumption. Third, relative to the existing literature on *centralized* non-convex optimization, where the focus is mostly on deterministic or *finite-sum* optimization, our modeling conditions are specifically tailored to the scenario of learning from stochastic *streaming* data. In particular, we only impose bounds on the gradient noise variance in expectation, rather than assume a bound with probability one [25], [29] or a sub-Gaussian distribution [30]. Furthermore, we assume that any Lipschitz conditions only hold on the *expected* stochastic gradient approximation, rather than for every realization, with probability one [26]–[28].

B. Preview of Results

We first establish that in non-convex environments, as was already shown earlier in [6] for convex environments, the

evolution of the individual iterates $w_{k,i}$ at the agents continues to be well-described by the evolution of the weighted centroid vector $\sum_{k=1}^N p_k w_{k,i}$ in the sense that the iterates from across the network will cluster around this centroid after sufficient iterations (Theorem 1). We subsequently consider two cases separately and establish descent in both of them. The first case corresponds to the region where the gradient at the network centroid is large and establish that descent can occur in one iteration (Theorem 2). The second and more challenging case occurs when the gradient norm is small, but there is a sufficiently negative eigenvalue in the Hessian matrix. In this regime, we introduce a short-term model, which allows for tractable analysis over several iterations, and establish that it tracks the true dynamics of the diffusion algorithm well enough. We then leverage this result to establish in Part II [3] that the recursion will continue to descend along the aggregate cost at a rate that is linear in μ per $O(1/\mu)$ iterations. Combined with the first result, this descent relation allows us to provide guarantees about the second-order optimality of the returned iterates.

The flow of the argument is summarized in Fig. 1. We decompose $\mathbb{R}^M$ into the set of approximate first-order stationary points, i.e., those with $\|\nabla J(w)\|^2 \leq O(\mu)$ and the complement, i.e., the large-gradient regime. For the large-gradient regime, descent is established in Theorem 2. We proceed to further decompose the set of approximate first-order stationary points into those that are τ -strict-saddle, i.e., those that have a Hessian with significant negative eigenvalue $\lambda_{\min}(\nabla^2 J(w)) \leq -\tau$, and the complement, which are approximately second-order stationary points. For τ -strict-saddle points we establish descent in Part II [3, Theorem 1]. Finally, in Part II [3, Theorem 2], we conclude that the centroid will reach an approximately second-order stationary point in a *polynomial* number of iterations.

II. EVOLUTION ANALYSIS

We shall perform the analysis under the following common assumptions on the gradients and their approximations.

Assumption 2 (Lipschitz gradients). *For each k , the gradient $\nabla J_k(\cdot)$ is Lipschitz, namely, for any $x, y \in \mathbb{R}^M$:*

$$\|\nabla J_k(x) - \nabla J_k(y)\| \leq \delta \|x - y\| \quad (5)$$

In light of (1) and Jensen's inequality, this implies for the aggregate cost:

$$\|\nabla J(x) - \nabla J(y)\| \leq \delta \|x - y\| \quad (6)$$

The Lipschitz gradient conditions (5) and (6) imply bounds on the both the function value and the Hessian matrix (when it exists), which will be used regularly throughout the derivations. In particular, we have for the function values:

$$J(y) \leq J(x) + \nabla J(x)^\top (y - x) + \frac{\delta}{2} \|x - y\|^2 \quad (7)$$

For the Hessian matrix we have [5]:

$$-\delta I \leq \nabla^2 J(x) \leq \delta I \quad (8)$$

	Modeling conditions					Results	
	Gradient	Hessian	Initialization	Perturbations	Step-size	Stationary	Saddle
Centralized							
[21]	Lipschitz	—	—	SGD + Annealing	diminishing	✓	asymptotic [†]
[22]	Lipschitz & bounded*	Lipschitz	—	i.i.d. and bounded w.p. 1	constant	✓	polynomial
[23]	Lipschitz	—	Random	—	constant	✓	asymptotic
[24]	Lipschitz	Lipschitz	—	Selective & bounded w.p. 1	constant	✓	polynomial
[25]	Lipschitz	Lipschitz	—	SGD, bounded w.p. 1	alternating	✓	polynomial
[26]	Lipschitz	Lipschitz	—	Bounded variance, Lipschitz w.p. 1	constant	✓	polynomial
[27]	Lipschitz	Lipschitz	—	Bounded variance, Lipschitz w.p. 1	constant	✓	polynomial
[28]	Lipschitz	Lipschitz	—	Bounded variance, Lipschitz w.p. 1	constant	✓	polynomial
[29]	Lipschitz	Lipschitz	—	SGD, bounded w.p. 1	constant	✓	polynomial
[30]	Lipschitz	Lipschitz	—	SGD + Gaussian	constant	✓	polynomial
Decentralized							
[31]	Lipschitz & bounded	—	—	—	constant	✓	—
[32]	Lipschitz	—	—	SGD	constant	✓	—
[33]	Lipschitz	—	—	—	constant	✓	—
[34]	Lipschitz & bounded	—	—	i.i.d.	diminishing	✓	—
[35]	Lipschitz	Exists	Random	—	constant	✓	asymptotic
[36]	Bounded disagreement	—	—	SGD + Annealing	diminishing	✓	asymptotic [†]
This work	Bounded disagreement	Lipschitz	—	Bounded moments	constant	✓	polynomial

TABLE I

COMPARISON OF MODELING ASSUMPTIONS AND RESULTS FOR GRADIENT-BASED METHODS. STATEMENTS MARKED WITH * ARE NOT EXPLICITLY STATED BUT ARE IMPLIED BY OTHER CONDITIONS. THE WORKS MARKED WITH [†] ESTABLISH GLOBAL (ASYMPTOTIC) CONVERGENCE, WHICH OF COURSE IMPLIES ESCAPE FROM SADDLE-POINTS.




Fig. 1. Classification of approximately stationary points. Theorem 2 in this work establishes descent in the green branch. The red branch is treated in Part II [3, Theorem 1]. The two results are combined in [3, Theorem 2] to establish the return of an approximately second-order stationary point with high probability. Constants are defined in Definition 2 and the referenced theorems.

Assumption 3 (Bounded gradient disagreement). For each pair of agents k and ℓ , the gradient disagreement is bounded, namely, for any $x \in \mathbb{R}^M$:

$$\|\nabla J_k(x) - \nabla J_\ell(x)\| \leq G \quad (9)$$

This assumption is similar to the one used in [36] under a diminishing step-size with annealing. Note that condition (9) is weaker than the more common assumption of bounded gradients. Condition (9) is automatically satisfied in cases where the expected risks $J_k(\cdot)$ are common (though agents still may see different realizations of data), or in the case of centralized stochastic gradient descent where the number

of agents is one. This condition is also satisfied whenever agent-specific risks with bounded gradients are regularized by common regularizers with potentially unbounded gradients, as is common in many machine learning applications. Observe that (9) implies a similar condition on the deviation from the centralized gradient via Jensen's inequality:

$$\begin{aligned} \|\nabla J_k(x) - \nabla J(x)\| &= \left\| \sum_{\ell=1}^N p_\ell (\nabla J_k(x) - \nabla J_\ell(x)) \right\| \\ &\leq \sum_{\ell=1}^N p_\ell \|\nabla J_k(x) - \nabla J_\ell(x)\| \leq G \end{aligned} \quad (10)$$

Definition 1 (Filtration). We denote by $\mathcal{F}_i$ the filtration generated by the random processes $\mathbf{w}_{k,j}$ for all k and $j \leq i$:

$$\mathcal{F}_i \triangleq \{\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_i\} \quad (11)$$

where $\mathbf{w}_j \triangleq \text{col}\{\mathbf{w}_{1,j}, \dots, \mathbf{w}_{N,j}\}$ contains the iterates across the network at time j . Informally, $\mathcal{F}_i$ captures all information that is available about the stochastic processes $\mathbf{w}_{k,j}$ across the network up to time i . ■

Assumption 4 (Gradient noise process). For each k , the gradient noise process is defined as

$$\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1}) = \widehat{\nabla} J_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1}) \quad (12)$$

and satisfies

$$\mathbb{E}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1}) | \mathcal{F}_{i-1}\} = 0 \quad (13a)$$

$$\mathbb{E}\{\|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^4 | \mathcal{F}_{i-1}\} \leq \sigma_k^4 \quad (13b)$$

for some non-negative constants σ_k^4 . We also assume that the gradient noise processes are pairwise uncorrelated over the space conditioned on $\mathcal{F}_{i-1}$, i.e.:

$$\mathbb{E}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1}) \mathbf{s}_{\ell,i}(\mathbf{w}_{\ell,i-1})^\top | \mathcal{F}_{i-1}\} = 0 \quad (14)$$

Property (13a) means that the gradient noise construction is unbiased on average. Property (13b) means that the fourth-moment of the gradient noise is bounded. These properties are automatically satisfied for several costs of interest [5], [11], and weaker than imposing a uniform bound to hold with probability one, or the perturbations to be sub-Gaussian. Note, that the bound on the fourth-order moment, in light of Jensen's inequality, immediately implies:

$$\begin{aligned} & \mathbb{E}\{\|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^2 | \mathcal{F}_{i-1}\} \\ & \leq \sqrt{\mathbb{E}\{\|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^4 | \mathcal{F}_{i-1}\}} \stackrel{(13b)}{\leq} \sigma_k^2 \end{aligned} \quad (15)$$

While our primary interest is in the development of algorithms that allow for learning from *streaming* data, we remark briefly that the results obtained in this work are equally applicable to empirical risk minimization via stochastic gradient descent, by assuming that the streaming data is selected according to a particular distribution.

Example 1 (Empirical Risk Minimization). Suppose the costs $J_k(\cdot)$ are empirical based on locally collected data $\{x_{k,s}\}_{s=1}^S$ and take the form:

$$J_k(w) = \frac{1}{S} \sum_{s=1}^S Q(w, x_{k,s}) \quad (16)$$

In empirical risk minimization (ERM) problems, we are interested in finding a vector w° that minimizes the following empirical risk over the data across the *entire* network:

$$w^\circ \triangleq \arg \min_w \frac{1}{N} \sum_{k=1}^N \left(\frac{1}{S} \sum_{s=1}^S Q(w, x_{k,s}) \right) \quad (17)$$

If we introduce the uniformly-distributed random variable $\mathbf{x}_k = x_{k,s}$ with probability $\frac{1}{S}$ for all s , then the cost (17) is equivalent to solving:

$$w^\circ = \arg \min_w \frac{1}{N} \sum_{k=1}^N \mathbb{E}_{\mathbf{x}_k} Q(w, \mathbf{x}_k) \quad (18)$$

which is of the same form as (1) with $p_k = \frac{1}{N}$. The resulting gradient noise process satisfies the assumptions imposed in this work under appropriate conditions on the risk $Q(\cdot, \cdot)$. This observation has been leveraged to accurately quantify the performance of stochastic gradient descent, as well as mini-batch and importance sampling generalizations, for empirical minimization of convex risks in [8]. ■

A. Network basis transformation

In analyzing the dynamics of the distributed algorithm (2a)–(2b), it is useful to introduce the following extended quantities by collecting variables from across the network:

$$\mathbf{w}_i \triangleq \text{col}\{\mathbf{w}_{1,i}, \dots, \mathbf{w}_{N,i}\} \quad (19)$$

$$\mathcal{A} \triangleq A \otimes I_N \quad (20)$$

$$\widehat{\mathbf{g}}(\mathbf{w}_i) \triangleq \text{col}\{\widehat{\nabla} J_1(\mathbf{w}_{1,i}), \dots, \widehat{\nabla} J_N(\mathbf{w}_{N,i})\} \quad (21)$$

where $\otimes$ denotes the Kronecker product operation and $\text{col}\{\cdot\}$ is obtained by stacking the arguments columnwise. We can then write the diffusion recursion (2a)–(2b) compactly as:

$$\mathbf{w}_i = \mathcal{A}^\top (\mathbf{w}_{i-1} - \mu \widehat{\mathbf{g}}(\mathbf{w}_{i-1})) \quad (22)$$

By construction, the combination matrix A is left-stochastic and primitive and hence admits a Jordan decomposition of the form $A = V_\epsilon J V_\epsilon^{-1}$ with [5], [6]:

$$V_\epsilon = \begin{bmatrix} p & V_R \end{bmatrix}, \quad J = \begin{bmatrix} 1 & 0 \\ 0 & J_\epsilon \end{bmatrix}, \quad V_\epsilon^{-1} = \begin{bmatrix} \mathbf{1}^\top \\ V_L^\top \end{bmatrix} \quad (23)$$

where J_ϵ is a block Jordan matrix with the eigenvalues $\lambda_2(A)$ through $\lambda_N(A)$ on the diagonal and ϵ on the first lower sub-diagonal. The extended matrix $\mathcal{A}$ then satisfies $\mathcal{A} = V_\epsilon \mathcal{J} V_\epsilon^{-1}$ with $V_\epsilon = V_\epsilon \otimes I_N$, $\mathcal{J} = J \otimes I_N$, $V_\epsilon^{-1} = V_\epsilon^{-1} \otimes I_N$. The spectral properties of A and its corresponding eigendecomposition have been exploited extensively in the study of the diffusion learning strategy in the *convex* setting [5], [6], and will continue to be useful in *non-convex* scenarios.

Multiplying both sides of (22) by $(p^\top \otimes I)$ from the left, we obtain in light of (4):

$$(p^\top \otimes I) \mathbf{w}_i = (p^\top \otimes I) \mathbf{w}_{i-1} - \mu (p^\top \otimes I) \widehat{\mathbf{g}}(\mathbf{w}_{i-1}) \quad (24)$$

Letting $\mathbf{w}_{c,i} \triangleq \sum_{k=1}^K p_k \mathbf{w}_{k,i} = (p^\top \otimes I) \mathbf{w}_i$ and exploiting the block-structure of the gradient term, we find:

$$\mathbf{w}_{c,i} = \mathbf{w}_{c,i-1} - \mu \sum_{k=1}^N p_k \widehat{\nabla} J_k(\mathbf{w}_{k,i-1}) \quad (25)$$

Note that $\mathbf{w}_{c,i}$ is a convex combination of iterates across the network and can be viewed as a weighted centroid. The recursion for $\mathbf{w}_{c,i}$ is reminiscent of a stochastic gradient step associated with the aggregate cost $\sum_{k=1}^N p_k J_k(w)$ with the exact gradients $\nabla J_k(\cdot)$ replaced by stochastic approximations

$\widehat{\nabla J}_k(\cdot)$ and with the stochastic gradients evaluated at $\mathbf{w}_{k,i-1}$, rather than $\mathbf{w}_{c,i-1}$. In fact, we can write:

$$\mathbf{w}_{c,i} = \mathbf{w}_{c,i-1} - \mu \sum_{k=1}^N p_k \nabla J_k(\mathbf{w}_{c,i-1}) - \mu \mathbf{d}_{i-1} - \mu \mathbf{s}_i \quad (26)$$

where we defined the perturbation terms:

$$\mathbf{d}_{i-1} \triangleq \sum_{k=1}^N p_k (\nabla J_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{c,i-1})) \quad (27)$$

$$\mathbf{s}_i \triangleq \sum_{k=1}^N p_k (\widehat{\nabla J}_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1})) \quad (28)$$

We use the subscript $i-1$ for $\mathbf{d}_{i-1}$ to emphasize that it depends on data up to time $i-1$, in contrast to $\mathbf{s}_i$ which is also dependent on the most recent data from time i . Observe that $\mathbf{d}_{i-1}$ arises from the disagreement within the network, and in particular that if each $\mathbf{w}_{k,i-1}$ remains close to the network centroid $\mathbf{w}_{c,i-1}$, this perturbation will be small in light of the Lipschitz condition (5) on the gradients. The second perturbation term $\mathbf{s}_i$ arises from the noise introduced by stochastic gradient approximations at each agent. We now establish that recursion (26) will continue to exhibit some of the desired properties of (centralized) gradient descent, despite the presence of persistent and coupled perturbation terms.

B. Network disagreement

To begin with, we study more closely the evolution of the individual estimates $\mathbf{w}_{k,i}$ relative to the network centroid $\mathbf{w}_{c,i}$. Multiplying (22) by $\mathcal{V}_R^\top \triangleq (V_R^\top \otimes I)$ from the left yields in light of (23):

$$\begin{aligned} \mathcal{V}_R^\top \mathbf{w}_i &= \mathcal{V}_R^\top \mathcal{A}^\top \mathbf{w}_{i-1} - \mu \mathcal{V}_R^\top \mathcal{A}^\top \widehat{\mathbf{g}}(\mathbf{w}_{i-1}) \\ &= J_\epsilon^\top \mathcal{V}_R^\top \mathbf{w}_{i-1} - \mu J_\epsilon^\top \mathcal{V}_R^\top \widehat{\mathbf{g}}(\mathbf{w}_{i-1}) \end{aligned} \quad (29)$$

Then, for the deviation from the network centroid:

$$\begin{aligned} &\mathbf{w}_i - \mathbf{w}_{c,i} \\ &= (I - (\mathbb{1} p^\top \otimes I)) \mathbf{w}_i \\ &= ((V_\epsilon^{-1} \otimes I)^\top (V_\epsilon \otimes I)^\top - (\mathbb{1} p^\top \otimes I)) \mathbf{w}_i \\ &\stackrel{(23)}{=} \mathcal{V}_L \mathcal{V}_R^\top \mathbf{w}_i \end{aligned} \quad (30)$$

where $\mathcal{V}_L \triangleq (V_L \otimes I)$, so that the deviation from the centroid can be easily recovered from $\mathcal{V}_R^\top \mathbf{w}_i$ in (29). Proceeding with (29), we find:

$$\begin{aligned} &\|\mathcal{V}_R^\top \mathbf{w}_i\|^4 \\ &\stackrel{(a)}{\leq} \|J_\epsilon^\top\|^4 \|\mathcal{V}_R^\top \mathbf{w}_{i-1} - \mu \mathcal{V}_R^\top \widehat{\mathbf{g}}(\mathbf{w}_{i-1})\|^4 \\ &\stackrel{(b)}{\leq} \|J_\epsilon^\top\|^4 \|\mathcal{V}_R^\top \mathbf{w}_{i-1}\|^4 + \mu^4 \frac{\|J_\epsilon^\top\|^4}{(1 - \|J_\epsilon^\top\|)^3} \|\mathcal{V}_R^\top \widehat{\mathbf{g}}(\mathbf{w}_{i-1})\|^4 \end{aligned} \quad (31)$$

where (a) follows from the sub-multiplicative property of norms, and (b) follows from Jensen's inequality $\|a+b\|^4 \leq \frac{1}{\alpha^3} \|a\|^4 + \frac{1}{(1-\alpha)^3} \|b\|^4$ with

$$\alpha = \|J_\epsilon^\top\| \triangleq \sqrt{\rho(J_\epsilon J_\epsilon^\top)} \leq \sqrt{\|J_\epsilon J_\epsilon^\top\|_1} \leq \sqrt{\lambda_2^2 + \epsilon} < 1 \quad (32)$$

for sufficiently small ϵ due to Assumption 1, where $\lambda_2 \triangleq \rho(A - \mathbb{1} p^\top)$. We observe that the term $\|\mathcal{V}_R^\top \mathbf{w}_i\|^4$ contracts at an exponential rate given by $\|J_\epsilon^\top\| \approx \lambda_2$ for small ϵ , also known as the mixing rate of the graph. Iterating this relation and applying Assumptions 1–4, we obtain the following result.

Theorem 1 (Network disagreement (4th order)). *Under assumptions 1–4, the network disagreement is bounded after sufficient iterations $i \geq i_o$ by:*

$$\begin{aligned} &\mathbb{E} \|\mathbf{w}_i - (\mathbb{1} p^\top \otimes I) \mathbf{w}_i\|^4 \\ &\leq \mu^4 \|\mathcal{V}_L\|^4 \frac{16 \|J_\epsilon^\top\|^4}{(1 - \|J_\epsilon^\top\|)^4} \|\mathcal{V}_R^\top\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \end{aligned} \quad (33)$$

where $i_o = \frac{4 \log(\mu)}{\log(\|J_\epsilon^\top\|)} + O(1) = o(\mu^{-1})$ and $o(\mu^{-1})$ denotes a term that is higher in order than μ^{-1} . Convergence to agreement occurs at a linear rate given by $\|J_\epsilon^\top\|$.

Proof: Appendix B. ■

Note again, that Jensen's inequality immediately implies for the second-order moment:

$$\begin{aligned} &\mathbb{E} \|\mathbf{w}_i - (\mathbb{1} p^\top \otimes I) \mathbf{w}_i\|^2 \leq \sqrt{\mathbb{E} \|\mathbf{w}_i - (\mathbb{1} p^\top \otimes I) \mathbf{w}_i\|^4} \\ &\stackrel{(a)}{\leq} \mu^2 \|\mathcal{V}_L\|^2 \frac{4 \|J_\epsilon^\top\|^2 \|\mathcal{V}_R^\top\|^2 N}{(1 - \|J_\epsilon^\top\|)^2} \left(G^2 + \max_k \sigma_k^2 \right) \end{aligned} \quad (34)$$

where (a) follows from (33) and sub-additivity of the square root, i.e. $\sqrt{x+y} \leq \sqrt{x} + \sqrt{y}$. This result establishes that, for every agent k , we have after sufficient iterations $i \geq i_o$:

$$\mathbb{E} \|\mathbf{w}_{k,i} - \mathbf{w}_{c,i}\|^2 \leq O(\mu^2) \quad (35)$$

or, by Markov's inequality [41]:

$$\Pr \left\{ \|\mathbf{w}_{k,i} - \mathbf{w}_{c,i}\|^2 \geq C\mu \right\} \leq O(\mu) \quad (36)$$

for a C independent of μ . Hence $\mathbf{w}_{k,i}$ will be arbitrarily close to $\mathbf{w}_{c,i}$ with arbitrarily high probability for all agents. This result has two implications. First, it allows us to use the network centroid $\mathbf{w}_{c,i}$ as a proxy for all iterates $\mathbf{w}_{k,i}$ in the network, since all agents will cluster around the network centroid after sufficient iterations. Second, it allows us to bound the perturbation terms encountered in (26).

Lemma 1 (Perturbation bounds (2nd and 4th order)). *Under assumptions 1–4 and for sufficiently small step-sizes μ , the perturbation terms are bounded as:*

$$\left(\mathbb{E} \|\mathbf{d}_{i-1}\|^2 \right)^2 \leq \mathbb{E} \|\mathbf{d}_{i-1}\|^4 \leq O(\mu^4) \quad (37)$$

$$\mathbb{E} \left\{ \|\mathbf{s}_i\|^2 | \mathcal{F}_{i-1} \right\} \leq \sum_{k=1}^N p_k^2 \sigma_k^2 \triangleq \sigma^2 \quad (38)$$

$$\mathbb{E} \left\{ \|\mathbf{s}_i\|^4 | \mathcal{F}_{i-1} \right\} \leq \max_k \sigma_k^4 \quad (39)$$

after sufficient iterations $i \geq i_o$ where we defined $\sigma^2 \triangleq \sum_{k=1}^N p_k^2 \sigma_k^2$.

Proof: Appendix C. ■

Definition 2 (Sets). To simplify the notation in the sequel, we introduce following sets:

$$\mathcal{G} \triangleq \left\{ w : \|\nabla J(w)\|^2 \geq \mu \frac{c_2}{c_1} \left(1 + \frac{1}{\pi}\right) \right\} \quad (40)$$

$$\mathcal{H} \triangleq \{ w : w \in \mathcal{G}^C, \lambda_{\min}(\nabla^2 J(w)) \leq -\tau \} \quad (41)$$

$$\mathcal{M} \triangleq \{ w : w \in \mathcal{G}^C, \lambda_{\min}(\nabla^2 J(w)) > -\tau \} \quad (42)$$

where τ is a small positive parameter, c_1 and c_2 are constants:

$$c_1 \triangleq \frac{1}{2} (1 - 2\mu\delta) = O(1) \quad (43)$$

$$c_2 \triangleq \delta\sigma^2/2 = O(1) \quad (44)$$

and $0 < \pi < 1$ is a parameter to be chosen. Note that $\mathcal{G}^C = \mathcal{H} \cup \mathcal{M}$. We also define the probabilities $\pi_i^{\mathcal{G}} \triangleq \Pr\{\mathbf{w}_{c,i} \in \mathcal{G}\}$, $\pi_i^{\mathcal{H}} \triangleq \Pr\{\mathbf{w}_{c,i} \in \mathcal{H}\}$ and $\pi_i^{\mathcal{M}} \triangleq \Pr\{\mathbf{w}_{c,i} \in \mathcal{M}\}$. Then for all i , we have $\pi_i^{\mathcal{G}} + \pi_i^{\mathcal{H}} + \pi_i^{\mathcal{M}} = 1$. ■

The definitions (40)–(42) decompose the parameter-space $\mathbb{R}^M$ into two disjoint sets $\mathcal{G}$ and $\mathcal{G}^C$, and further sub-divides $\mathcal{G}^C$ into $\mathcal{H}$ and $\mathcal{M}$. The set $\mathcal{G}$ denotes the set all points w where the norm of the gradient is large, while $\mathcal{G}^C = \mathcal{H} \cup \mathcal{M}$ denotes the set of all points where the norm of the gradient is small, i.e., approximately first-order stationary points. In a manner similar to related works on the escape from strict-saddle points, we further decompose the set $\mathcal{G}^C$ of approximate first-order stationary points into those points $w \in \mathcal{H}$ that do have a significant negative eigenvalue, and those in $\mathcal{M}$ that do not [22], [24]. Points in the parameter space that have a small gradient norm and *no* significant negative eigenvalue are referred to as approximately *second-order* stationary points, while points in $\mathcal{H}$ are known as *strict* saddle-points due to the presence of a strictly negative eigenvalue in the Hessian matrix. In the sequel, we will establish descent for centroids in $\mathcal{G}$ in Theorem 2 and centroids in $\mathcal{H}$ in Part II [3, Theorem 1], and hence the approach of a point in $\mathcal{M}$ with high probability after a polynomial number of iterations in Part II [3, Theorem 2]. Second-order stationary points are generally more likely to be “good” minimizers than first-order stationary points, which could even correspond to local maxima. Furthermore, for a certain class of cost functions, known as “strict-saddle” functions, second-order stationary points always correspond to local minima for sufficiently small τ [22].

C. Evolution of the network centroid

Having established in (34), that after sufficient iterations, all agents in the network will have contracted around the centroid in a small cluster for small step-sizes, we can now leverage $\mathbf{w}_{c,i}$ as a proxy for all $\mathbf{w}_{k,i}$. From Assumption 2 and (7), we have the following bound:

$$J(\mathbf{w}_{c,i}) \leq J(\mathbf{w}_{c,i-1}) + \nabla J(\mathbf{w}_{c,i-1})^\top (\mathbf{w}_{c,i} - \mathbf{w}_{c,i-1}) + \frac{\delta}{2} \|\mathbf{w}_{c,i} - \mathbf{w}_{c,i-1}\|^2 \quad (45)$$

From (26), we then obtain:

$$\begin{aligned} J(\mathbf{w}_{c,i}) &\leq J(\mathbf{w}_{c,i-1}) - \mu \|\nabla J(\mathbf{w}_{c,i-1})\|^2 \\ &\quad - \mu \nabla J(\mathbf{w}_{c,i-1})^\top (\mathbf{d}_{i-1} + \mathbf{s}_i) \\ &\quad + \mu^2 \frac{\delta}{2} \|\nabla J(\mathbf{w}_{c,i-1}) + \mathbf{d}_{i-1} + \mathbf{s}_i\|^2 \end{aligned} \quad (46)$$

This relation, along with (26) and the results from Lemma 1, allow us to establish the following theorem.

Theorem 2 (Descent relation). Beginning at $\mathbf{w}_{c,i-1}$ in the large gradient regime $\mathcal{G}$, we can bound after $i \geq i_o$:

$$\begin{aligned} &\mathbb{E}\{J(\mathbf{w}_{c,i}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} \\ &\leq \mathbb{E}\{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} - \mu^2 \frac{c_2}{\pi} + \frac{O(\mu^3)}{\pi_{i-1}^{\mathcal{G}}} \end{aligned} \quad (47)$$

as long as $\pi_{i-1}^{\mathcal{G}} = \Pr\{\mathbf{w}_{c,i-1} \in \mathcal{G}\} \neq 0$ where c_2 and π are defined in Definition 2. On the other hand, beginning at $\mathbf{w}_{c,i-1} \in \mathcal{M}$, we can bound:

$$\begin{aligned} &\mathbb{E}\{J(\mathbf{w}_{c,i}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} \\ &\leq \mathbb{E}\{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} + \mu^2 c_2 + \frac{O(\mu^3)}{\pi_{i-1}^{\mathcal{M}}} \end{aligned} \quad (48)$$

as long as $\pi_{i-1}^{\mathcal{M}} = \Pr\{\mathbf{w}_{c,i-1} \in \mathcal{M}\} \neq 0$. Finally, we bound the difference between the initial function value and the function value when consensus is achieved (see Lemma 1) as:

$$\begin{aligned} &\mathbb{E}J(\mathbf{w}_{c,i_o}) \\ &\leq J(\mathbf{w}_{c,0}) + \frac{\mu}{2} (1 + 2\mu\delta) \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \frac{\|\mathcal{V}_R^\top \mathbf{w}_0\|^2}{1 - \|J_\epsilon^\top\|^{1/2}} + o(\mu) \end{aligned} \quad (49)$$

Proof: Appendix E. ■

Relation (47) guarantees a lower bound on the expected improvement when the gradient norm at the current iterate is sufficiently large, i.e. $\mathbf{w}_{c,i-1} \in \mathcal{G}$ is not an approximately first-order stationary point. On the other hand, when $\mathbf{w}_{c,i-1} \in \mathcal{M}$, inequality (48) establishes an upper bound on the expected ascent. The respective bounds can be balanced by appropriately choosing π , which will be leveraged in Part II [3]. Relation (49) ensures that any deterioration in performance until agreement according to Theorem 1 is reached is negligible and will also be used in Part II [3]. We are left to treat the third possibility, namely $\mathbf{w}_{c,i-1} \in \mathcal{H}$. In this case, since the norm of the gradient is small, it is no longer possible to guarantee descent in a single iteration. We shall study the dynamics in more detail in the sequel.

D. Behavior around stationary points

In the vicinity of saddle-points, the norm of the gradient is not sufficiently large to guarantee descent at every iteration as indicated by (47). Instead, we will study the cumulative effect of the gradient, as well as perturbations, over several iterations. For this purpose, we introduce the following second-order condition on the cost functions, which is common in the literature [5], [22], [24].

Assumption 5 (Lipschitz Hessians). Each $J_k(\cdot)$ is twice-differentiable with Hessian $\nabla^2 J_k(\cdot)$ and, there exists $\rho \geq 0$ such that:

$$\|\nabla^2 J_k(x) - \nabla^2 J_k(y)\| \leq \rho \|x - y\| \quad (50)$$

By Jensen's inequality, this implies that $J(\cdot) = \sum_{k=1}^N p_k J_k(\cdot)$ also satisfies:

$$\|\nabla^2 J(x) - \nabla^2 J(y)\| \leq \rho \|x - y\| \quad (51)$$

■

Let i^* denote an arbitrary point in time. We use i^* in order to emphasize approximately first-order stationary points, where the norm of the gradient is small. Such approximately first-order stationary points $\mathbf{w}_{c,i^*} \in \mathcal{G}^C$ could either be in the set of approximately second-order stationary points $\mathcal{M}$ or in the set of strict-saddle points $\mathcal{H}$. Our objective is to show that when $\mathbf{w}_{c,i^*} \in \mathcal{H}$, we can guarantee descent after several iterations. To this end, starting at i^* , we have for $i \geq 0$:

$$\mathbf{w}_{c,i^*+i+1} = \mathbf{w}_{c,i^*+i} - \mu \nabla J(\mathbf{w}_{c,i^*+i}) - \mu \mathbf{d}_{i^*+i} - \mu \mathbf{s}_{i^*+i+1} \quad (52)$$

Subsequent analysis will rely on an auxiliary model, referred to as a short-term model. It will be seen that this model is more tractable and evolves “close” to the true recursion under the second-order smoothness condition on the Hessian matrix (51) and as long as the iterates remain close to a stationary point. A similar approach has been introduced and used to great advantage in the form of a “long-term model” to derive accurate mean-square deviation performance expressions for strongly-convex costs in [5], [11], [42], [43]. The approach was also used to provide a “quadratic approximation” to establish the ability of stochastic gradient based algorithms to escape from strict saddle-points in the single-agent case under i.i.d. perturbations in [22].

For the driving gradient term in (52), we have from the mean-value theorem [5]:

$$\nabla J(\mathbf{w}_{c,i^*+i}) - \nabla J(\mathbf{w}_{c,i^*}) = \mathbf{H}_{i^*+i}(\mathbf{w}_{c,i^*+i} - \mathbf{w}_{c,i^*}) \quad (53)$$

where

$$\mathbf{H}_{i^*+i} \triangleq \int_0^1 \nabla^2 J((1-t)\mathbf{w}_{c,i^*+i} + t\mathbf{w}_{c,i^*}) dt \quad (54)$$

Subtracting (52) from $\mathbf{w}_{c,i^*}$, we obtain:

$$\begin{aligned} & \mathbf{w}_{c,i^*} - \mathbf{w}_{c,i^*+i+1} \\ &= \mathbf{w}_{c,i^*} - \mathbf{w}_{c,i^*+i} + \mu \nabla J(\mathbf{w}_{c,i^*+i}) + \mu \mathbf{d}_{i^*+i} + \mu \mathbf{s}_{i^*+i+1} \\ &= (I - \mu \mathbf{H}_{i^*+i})(\mathbf{w}_{c,i^*} - \mathbf{w}_{c,i^*+i}) + \mu \nabla J(\mathbf{w}_{c,i^*}) \\ & \quad + \mu \mathbf{d}_{i^*+i} + \mu \mathbf{s}_{i^*+i+1} \end{aligned} \quad (55)$$

We introduce short-hand notation for the deviation:

$$\tilde{\mathbf{w}}_i^{i^*} \triangleq \mathbf{w}_{c,i^*} - \mathbf{w}_{c,i^*+i} \quad (56)$$

Note that $\tilde{\mathbf{w}}_i^{i^*}$ denotes the deviation of the network centroid $\mathbf{w}_{c,i^*+i}$ at time $i^* + i$ from the initial, approximately first-order stationary point $\mathbf{w}_{c,i^*}$. Establishing escape from saddle-points is equivalent to establishing the growth of $\tilde{\mathbf{w}}_i^{i^*}$ whenever $\mathbf{w}_{c,i^*} \in \mathcal{H}$. We hence expect the deviation to grow over

time, but would like to establish that $\mathbf{w}_{c,i^*+i}$ moves away from $\mathbf{w}_{c,i^*}$ in a direction of descent. We can then write more compactly:

$$\begin{aligned} \tilde{\mathbf{w}}_{i+1}^{i^*} &= (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i^*} + \mu \nabla J(\mathbf{w}_{c,i^*}) \\ & \quad + \mu \mathbf{d}_{i^*+i} + \mu \mathbf{s}_{i^*+i+1} \end{aligned} \quad (57)$$

The time-varying nature of $\mathbf{H}_{i^*+i}$ makes this recursion difficult to study. Additionally, while $\mathbf{d}_{i^*+i}$ was shown in Theorem 1 to be small in the mean-fourth sense, in contrast to $\mathbf{s}_{i^*+i+1}$, it need not be conditionally zero-mean. We hence introduce the following auxiliary recursion, initialized at $\mathbf{w}'_{c,i^*} = \mathbf{w}_{c,i^*}$, where $\mathbf{H}_{i^*+i}$ is replaced by $\nabla^2 J(\mathbf{w}_{c,i^*})$ and the perturbation term $\mu \mathbf{d}_{i^*+i}$ is omitted:

$$\begin{aligned} \mathbf{w}_{c,i^*} - \mathbf{w}'_{c,i^*+i+1} &= (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) (\mathbf{w}_{c,i^*} - \mathbf{w}'_{c,i^*+i}) \\ & \quad + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{s}_{i^*+i+1} \end{aligned} \quad (58)$$

or, more compactly, with $\tilde{\mathbf{w}}_i'^{i^*} \triangleq \mathbf{w}_{c,i^*} - \mathbf{w}'_{c,i^*+i}$:

$$\tilde{\mathbf{w}}_{i+1}'^{i^*} = (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i'^{i^*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{s}_{i^*+i+1} \quad (59)$$

Of course, this second model is only useful in studying the behavior of the original recursion (52) if the iterates generated by both models remain close to each other, which we shall prove to be true. Results along this line have been established in the centralized and distributed contexts for strongly-convex costs [5], [11] and in the centralized setting for strict saddle points [22]. We show here that this conclusion holds more generally in the vicinity of $O(\mu)$ -first-order stationary points. Before establishing deviation bounds, we establish a short lemma which will be used repeatedly.

Lemma 2 (A limiting result). For $T, \mu, \delta > 0$ and $k \in \mathbb{Z}_+$ with $\mu < \frac{1}{\delta}$, we have:

$$\lim_{\mu \rightarrow 0} \left(\frac{(1 + \mu\delta)^k}{(1 - \mu\delta)^{k-1}} \right)^{\frac{T}{\mu}} = e^{-T\delta + 2kT\delta} = O(1) \quad (60)$$

Proof: Appendix D. ■

Lemma 3 (Deviation bounds). Suppose $\Pr\{\mathbf{w}_{c,i^*} \in \mathcal{H}\} \neq 0$. Then, for $i^* \geq i_o$ the following quantities are conditionally bounded:

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i^*} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu) + \frac{O(\mu^2)}{\pi_{i^*}^{\mathcal{H}}} \quad (61)$$

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i^*} \right\|^3 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu^{3/2}) + \frac{O(\mu^3)}{\pi_{i^*}^{\mathcal{H}}} \quad (62)$$

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i^*} \right\|^4 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu^2) + \frac{O(\mu^4)}{\pi_{i^*}^{\mathcal{H}}} \quad (63)$$

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i^*} - \tilde{\mathbf{w}}_i'^{i^*} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu^2) + \frac{O(\mu^2)}{\pi_{i^*}^{\mathcal{H}}} \quad (64)$$

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i'^{i^*} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu) + \frac{O(\mu^2)}{\pi_{i^*}^{\mathcal{H}}} \quad (65)$$

$$\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i'^{i^*} \right\|^3 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq O(\mu^{3/2}) + \frac{O(\mu^3)}{\pi_{i^*}^{\mathcal{H}}} \quad (66)$$

for $i \leq \frac{T}{\mu}$, where T denotes an arbitrary constant that is independent of the step-size μ .

Proof: Appendix F. ■

These deviation bounds establish that, beginning at a strict-saddle point w_{c,i^*} at time i^* the iterates will remain close to w_{c,i^*} for the next $O(1/\mu)$ iterations. Consequently, the short-term model will be sufficiently accurate for the next $O(1/\mu)$ iterations. We will establish formally in Part II [3] that the small-deviation bounds in Lemma 3 ensure descent of the true recursion can be inferred by studying only the evolution of the short-term model, which is significantly more tractable.

III. APPLICATION: ROBUST REGRESSION

Consider a scenario where each agent k in the network observes streaming realizations $\{\gamma(k, i), \mathbf{h}_{k,i}\}$ from the linear model $\gamma(k) = \mathbf{h}_k^\top w^o + v(k)$ where $\gamma(k)$ denotes scalar observations and $v(k)$ denotes measurement noise. One common approach for estimating w^o in a distributed setting is via least-mean-square error estimation, resulting in the local cost functions:

$$J_k^{\text{MSE}}(w) = \mathbb{E} \left\| \gamma(k) - \mathbf{h}_k^\top w \right\|^2 \quad (67)$$

The resulting problem is convex and has been studied extensively in the literature. While effective under the assumption of Gaussian noise, and similar well-behaved noise conditions, this approach is susceptible to outliers caused by heavy-tailed distributions for $v(k)$ [18]. This is caused by the fact that the quadratic risk penalizes errors proportionally to their squared norm, and as such has a tendency to over-correct outliers, even if they are rare. Several alternative robust cost functions have been suggested in the literature. We consider two in particular in order to illustrate the advantages of allowing for non-convex costs in the context of robust estimation, namely the Huber loss $Q_k^H(w; \mathbf{x}_k)$ and Tukey's biweight loss $Q_k^B(w; \mathbf{x}_k)$ [18]. For ease of notation, let $e(w) \triangleq \gamma(k) - \mathbf{h}_k^\top w$. Then:

$$Q_k^H(w; \mathbf{x}_k) = \begin{cases} \frac{1}{2}|e(w)|^2, & \text{for } |e(w)| \leq c_H \\ c_H|e(w)| - \frac{1}{2}c_H^2, & \text{for } |e(w)| > c_H. \end{cases} \quad (68)$$

$$Q_k^B(w; \mathbf{x}_k) = \begin{cases} \frac{c_B^2}{6} \left(1 - \left(1 - \frac{|e(w)|^2}{c_B^2} \right)^3 \right), & \text{for } |e(w)| \leq c_B \\ \frac{c_B^2}{6}, & \text{otherwise} \end{cases} \quad (69)$$

where c_H, c_B are tuning constants. The Huber cost is merely convex (and not strongly-convex), while the Tukey loss is non-convex. Both losses satisfy assumptions 1–4 imposed in this work. In particular, since the Huber risk $J_k^H(w)$ has a unique, local minimum, which also happens to be locally strongly-convex, we can conclude that despite the absence of strong-convexity, the algorithm will converge to within $O(\mu)$ of the global minimum. The Tukey loss on the other hand, is non-convex, and is therefore a more challenging problem. The setting for the simulation results is shown in Figure 2.

Performance is illustrated in Fig. 3. We first show the performance of each cost in the nominal scenario, where $v(k) \sim \mathcal{N}(0, \sigma_v^2)$. We observe that the distributed strategies outperform the non-cooperative ones, and that despite differences in the rate of convergence, there is negligible difference in the performance of the mean-square-error, Huber and Tukey variations. In the presence of outliers, modeled as a bimodal




Fig. 2. Graph with $N = 20$ nodes (left) and regressor power $\text{Tr}(R_{h,k})$ at each agent (right).




Fig. 3. Average performance across agents in the nominal (top) and corrupted case (bottom). The non-cooperative solution consists of a local stochastic gradient step of the form (2a) without a combination step of the form (2b).

distribution with $v(k) \sim (1 - \epsilon)\mathcal{N}(0, \sigma_v^2) + \epsilon\mathcal{N}(10, \sigma_v^2)$ and $\epsilon = 0.1$, the performance of the mean-square-error solution dramatically deteriorates, as is to be expected in the presence of deviations from the nominal model. We illustrate the evolution of average disagreement in Fig. 4 and observe that disagreement, in all cases, is several decibel smaller than the mean-square deviation.

APPENDIX A A LEMMA ON CONDITIONING

Throughout the derivations, we will frequently rely on appropriate conditionings to make the analysis tractable. A frequent theme will be the exchange of conditioning on filtrations by




Fig. 4. Average disagreement across agents in the nominal (top) and corrupted case (bottom).

conditioning on events. To this end, the following lemma will be used repeatedly.

Lemma 4 (Conditioning). Suppose $\mathbf{w} \in \mathbb{R}^M$ is a random variable measurable by $\mathcal{F}$. In other words, $\mathbf{w}$ is deterministic conditioned on $\mathcal{F}$ and

$$\mathbb{E}\{\mathbf{w} | \mathcal{F}\} = \mathbf{w}. \quad (70)$$

Then,

$$\mathbb{E}\{\mathbb{E}\{\mathbf{x} | \mathcal{F}\} | \mathbf{w} \in \mathcal{S}\} = \mathbb{E}\{\mathbf{x} | \mathbf{w} \in \mathcal{S}\} \quad (71)$$

for any deterministic set $\mathcal{S} \subseteq \mathbb{R}^M$ with $\Pr\{\mathbf{w} \in \mathcal{S}\} \neq 0$ and random $\mathbf{x} \in \mathbb{R}^M$ with $\mathbb{E}\{\mathbf{x} | \mathbf{w} \in \mathcal{S}\} < \infty$.

Proof. The result is a direct consequence of the law of total expectation. $\square$

APPENDIX B PROOF OF THEOREM 1

Starting from (29), taking norms of both sides and computing the fourth power, we find:

$$\begin{aligned} & \|\mathcal{V}_R^T \mathbf{w}_i\|^4 \\ & \leq \|J_\epsilon^T\|^4 \|\mathcal{V}_R^T \mathbf{w}_{i-1} + \mu \mathcal{V}_R^T \hat{\mathbf{g}}(\mathbf{w}_{i-1})\|^4 \\ & \stackrel{(a)}{\leq} \|J_\epsilon^T\|^4 \|\mathcal{V}_R^T \mathbf{w}_{i-1}\|^4 + \mu^4 \frac{\|J_\epsilon^T\|^4}{(1 - \|J_\epsilon^T\|)^3} \|\mathcal{V}_R^T \hat{\mathbf{g}}(\mathbf{w}_{i-1})\|^4 \end{aligned} \quad (72)$$

where step (a) follows from convexity of $\|\cdot\|^4$ and Jensen's inequality, i.e. $\|a + b\|^4 = \frac{1}{\alpha^3} \|a\|^4 + \frac{1}{(1-\alpha)^3} \|b\|^4$. To begin with, we study the stochastic gradient term in some greater detail. We have:

$$\begin{aligned} & \|\mathcal{V}_R^T \hat{\mathbf{g}}(\mathbf{w}_{i-1})\|^4 \\ & = \|\mathcal{V}_R^T g(\mathbf{w}_{i-1}) + \mathcal{V}_R^T \text{col}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\}\|^4 \\ & \leq 8 \|\mathcal{V}_R^T g(\mathbf{w}_{i-1})\|^4 + 8 \|\mathcal{V}_R^T \text{col}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\}\|^4 \end{aligned} \quad (73)$$

For the first term we have:

$$\begin{aligned} & 8 \|\mathcal{V}_R^T g(\mathbf{w}_{i-1})\|^4 \\ & \stackrel{(a)}{=} 8 \|\mathcal{V}_R^T g(\mathbf{w}_{i-1}) - \mathcal{V}_R^T (\mathbb{1} p^T \otimes I) g(\mathbf{w}_{i-1})\|^4 \\ & \stackrel{(b)}{\leq} 8 \|\mathcal{V}_R^T\|^4 \|g(\mathbf{w}_{i-1}) - (\mathbb{1} p^T \otimes I) g(\mathbf{w}_{i-1})\|^4 \\ & \stackrel{(c)}{=} 8 \|\mathcal{V}_R^T\|^4 \left(\sum_{k=1}^N \|\nabla J_k(\mathbf{w}_{k,i-1}) - \nabla J(\mathbf{w}_{k,i-1})\|^2 \right)^2 \\ & \stackrel{(9)}{\leq} 8 \|\mathcal{V}_R^T\|^4 \left(\sum_{k=1}^N G^2 \right)^2 \leq 8 \|\mathcal{V}_R^T\|^4 N^2 G^4 \end{aligned} \quad (74)$$

where (a) follows from the fact that (23) implies $\mathcal{V}_R^T \mathbb{1} = 0$, (b) follows from the sub-multiplicity of norms and (c) expands $\|\cdot\|^2$. For the gradient noise term we find under expectation:

$$\begin{aligned} & 8 \mathbb{E} \|\mathcal{V}_R^T \text{col}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\}\|^4 \\ & = 8 \|\mathcal{V}_R^T\|^4 \mathbb{E} \|\text{col}\{\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\}\|^4 \\ & = 8 \|\mathcal{V}_R^T\|^4 \mathbb{E} \left(\sum_{k=1}^N \|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^2 \right)^2 \\ & \stackrel{(a)}{\leq} 8 \|\mathcal{V}_R^T\|^4 N \sum_{k=1}^N \mathbb{E} \|\mathbf{s}_{k,i}(\mathbf{w}_{k,i-1})\|^4 \\ & \stackrel{(13b)}{\leq} 8 \|\mathcal{V}_R^T\|^4 N \sum_{k=1}^N \sigma_k^4 = 8 \|\mathcal{V}_R^T\|^4 N^2 \max_k \sigma_k^4 \end{aligned} \quad (75)$$

where (a) follows from Cauchy-Schwarz, which implies $\left(\sum_{k=1}^N x_k \right)^2 \leq N \sum_{k=1}^N x_k^2$. Plugging these relations back into (72), we obtain:

$$\begin{aligned} \mathbb{E} \|\mathcal{V}_R^T \mathbf{w}_i\|^4 & \leq \|J_\epsilon^T\| \mathbb{E} \|\mathcal{V}_R^T \mathbf{w}_{i-1}\|^4 \\ & \quad + \mu^4 \frac{8 \|J_\epsilon^T\|^4}{(1 - \|J_\epsilon^T\|)^3} \|\mathcal{V}_R^T\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \end{aligned} \quad (76)$$

We can iterate, starting from $i = 0$, to obtain:

$$\begin{aligned} & \mathbb{E} \|\mathcal{V}_R^T \mathbf{w}_i\|^4 \\ & \leq \|J_\epsilon^T\|^i \mathbb{E} \|\mathcal{V}_R^T \mathbf{w}_0\|^4 \\ & \quad + \mu^4 \frac{8 \|J_\epsilon^T\|^4}{(1 - \|J_\epsilon^T\|)^3} \|\mathcal{V}_R^T\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \sum_{n=1}^i \|J_\epsilon^T\|^{n-1} \\ & \stackrel{(a)}{\leq} \|J_\epsilon^T\|^i \mathbb{E} \|\mathcal{V}_R^T \mathbf{w}_0\|^4 \\ & \quad + \mu^4 \frac{8 \|J_\epsilon^T\|^4}{(1 - \|J_\epsilon^T\|)^4} \|\mathcal{V}_R^T\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \\ & \stackrel{(b)}{\leq} \mu^4 \frac{16 \|J_\epsilon^T\|^4}{(1 - \|J_\epsilon^T\|)^4} \|\mathcal{V}_R^T\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \end{aligned} \quad (77)$$

where (a) follows from $\sum_{n=1}^i \|J_\epsilon^\top\|^{n-1} \leq \sum_{n=1}^\infty \|J_\epsilon^\top\|^{n-1} = (1 - \|J_\epsilon^\top\|)^{-1}$, and (b) holds whenever:

$$\begin{aligned} & \|J_\epsilon^\top\|^i \mathbb{E} \|\mathcal{V}_R^\top \mathcal{W}_0\|^4 \\ & \leq \mu^4 \frac{8 \|J_\epsilon^\top\|^4}{(1 - \|J_\epsilon^\top\|)^4} \|\mathcal{V}_R^\top\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \\ \iff & \|J_\epsilon^\top\|^i \leq \mu^4 \frac{8 \|J_\epsilon^\top\|^4}{(1 - \|J_\epsilon^\top\|)^4} \frac{\|\mathcal{V}_R^\top\|^4}{\mathbb{E} \|\mathcal{V}_R^\top \mathcal{W}_0\|^4} N^2 \left(G^4 + \max_k \sigma_k^4 \right) \\ \iff & i \log(\|J_\epsilon^\top\|) \leq 4 \log(\mu) + O(1) \\ \iff & i \geq \frac{4 \log(\mu)}{\log(\|J_\epsilon^\top\|)} + O(1) \end{aligned} \quad (78)$$

Finally, we have from (30) under (78):

$$\begin{aligned} & \mathbb{E} \|\mathbf{w}_i - (\mathbf{1} p^\top \otimes I) \mathbf{w}_i\|^4 = \mathbb{E} \|\mathcal{V}_L \mathcal{V}_R^\top \mathbf{w}_i\|^4 \\ & \stackrel{(77)}{\leq} \mu^4 \|\mathcal{V}_L\|^4 \frac{16 \|J_\epsilon^\top\|^4}{(1 - \|J_\epsilon^\top\|)^4} \|\mathcal{V}_R^\top\|^4 N^2 \left(G^4 + \max_k \sigma_k^4 \right) \end{aligned} \quad (79)$$

We conclude that all agents in the network will contract around the centroid vector $(\mathbf{1} p^\top \otimes I) \mathbf{w}_i$ after sufficient iterations i_o , where:

$$i_o = \frac{4 \log(\mu)}{\log(\|J_\epsilon^\top\|)} + O(1) \stackrel{(a)}{=} o(\mu^{-1}) \quad (80)$$

where (a) follows since $\lim_{\mu \rightarrow 0} \mu \log(\mu) = 0$.

APPENDIX C PROOF OF LEMMA 1

We begin by studying the perturbation term $\mathbf{s}_i$. We have:

$$\begin{aligned} & \mathbb{E} \left\{ \|\mathbf{s}_i\|^2 | \mathcal{F}_{i-1} \right\} \\ & = \mathbb{E} \left\{ \left\| \sum_{k=1}^N p_k \left(\widehat{\nabla J}_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1}) \right) \right\|^2 | \mathcal{F}_{i-1} \right\} \\ & \stackrel{(a)}{=} \sum_{k=1}^N p_k^2 \mathbb{E} \left\{ \left\| \widehat{\nabla J}_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1}) \right\|^2 | \mathcal{F}_{i-1} \right\} \\ & \stackrel{(b)}{\leq} \sum_{k=1}^N p_k^2 \sigma_k^2 \end{aligned} \quad (81)$$

where (a) follows from the uncorrelatedness assumption (14) and (b) follows from the bound on the second-order moment (15). For the fourth-moment, we have:

$$\begin{aligned} & \mathbb{E} \left\{ \|\mathbf{s}_i\|^4 | \mathcal{F}_{i-1} \right\} \\ & = \mathbb{E} \left\{ \left\| \sum_{k=1}^N p_k \left(\widehat{\nabla J}_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1}) \right) \right\|^4 | \mathcal{F}_{i-1} \right\} \\ & \stackrel{(a)}{\leq} \sum_{k=1}^N p_k \mathbb{E} \left\{ \left\| \widehat{\nabla J}_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{k,i-1}) \right\|^4 | \mathcal{F}_{i-1} \right\} \\ & \stackrel{(b)}{\leq} \sum_{k=1}^N p_k \sigma_k^4 \leq \max_k \sigma_k^4 \end{aligned} \quad (82)$$

where (a) follows from $\sum_{k=1}^N p_k = 1$ and Jensen's inequality and (b) follows from the fourth-order moment condition in Assumption 4. For the second perturbation term, we have:

$$\begin{aligned} \|\mathbf{d}_{i-1}\|^4 & = \left\| \sum_{k=1}^N p_k (\nabla J_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{c,i-1})) \right\|^4 \\ & \stackrel{(a)}{\leq} \sum_{k=1}^N p_k \|\nabla J_k(\mathbf{w}_{k,i-1}) - \nabla J_k(\mathbf{w}_{c,i-1})\|^4 \\ & \stackrel{(b)}{\leq} \delta^4 \sum_{k=1}^N p_k \|\mathbf{w}_{k,i-1} - \mathbf{w}_{c,i-1}\|^4 \\ & \leq \delta^4 p_{\max} \sum_{k=1}^N \|\mathbf{w}_{k,i-1} - \mathbf{w}_{c,i-1}\|^4 \\ & \leq \delta^4 p_{\max} \left(\sum_{k=1}^N \|\mathbf{w}_{k,i-1} - \mathbf{w}_{c,i-1}\|^2 \right)^2 \\ & = \delta^4 p_{\max} \|\mathbf{w}_{i-1} - \mathbf{w}_{c,i-1}\|^4 \end{aligned} \quad (83)$$

where (a) again follows from Jensen's inequality, (b) follows from the Lipschitz gradient condition in Assumption 2, and we introduced $\mathbf{w}_{c,i-1} \triangleq \mathbf{1} \otimes \mathbf{w}_{c,i-1}$. Result (37) follows by applying (79) to (83).

APPENDIX D PROOF OF LEMMA 2

Since the logarithm is continuous over $\mathbb{R}_+$, we have:

$$\begin{aligned} & \log \left(\lim_{\mu \rightarrow 0} \left(\frac{(1 + \mu\delta)^k}{(1 - \mu\delta)^{k-1}} \right)^{\frac{T}{\mu}} \right) \\ & = \lim_{\mu \rightarrow 0} \log \left(\left(\frac{(1 + \mu\delta)^k}{(1 - \mu\delta)^{k-1}} \right)^{\frac{T}{\mu}} \right) \\ & = kT \lim_{\mu \rightarrow 0} \frac{\log(1 + \mu\delta)}{\mu} - (k-1)T \lim_{\mu \rightarrow 0} \frac{\log(1 - \mu\delta)}{\mu} \end{aligned} \quad (84)$$

We examine the fraction inside the limit more closely. Since both the numerator and denominator of the fraction approach zero as $\mu \rightarrow 0$, we apply L'Hôpital's rule:

$$\lim_{\mu \rightarrow 0} \frac{\log(1 \pm \mu\delta)}{\mu} = \lim_{\mu \rightarrow 0} \frac{\pm\delta}{1 \pm \mu\delta} = \pm\delta \quad (85)$$

Hence, we find:

$$\lim_{\mu \rightarrow 0} \left(\frac{(1 + \mu\delta)^k}{(1 - \mu\delta)^{k-1}} \right)^{\frac{T}{\mu}} = e^{kT\delta + (k-1)T\delta} = e^{-T\delta + 2kT\delta} \quad (86)$$

APPENDIX E PROOF OF THEOREM 2

We begin with (46) and take expectations conditioned on $\mathbf{w}_{i-1}$ to obtain:

$$\begin{aligned} & \mathbb{E} \{ J(\mathbf{w}_{c,i}) | \mathbf{w}_{i-1} \} \\ & \stackrel{(a)}{\leq} J(\mathbf{w}_{c,i-1}) - \mu \|\nabla J(\mathbf{w}_{c,i-1})\|^2 - \mu \nabla J(\mathbf{w}_{c,i-1})^\top \mathbf{d}_{i-1} \\ & \quad + \mu^2 \frac{\delta}{2} \|\nabla J(\mathbf{w}_{c,i-1}) + \mathbf{d}_{i-1}\|^2 + \mu^2 \frac{\delta}{2} \mathbb{E} \left\{ \|\mathbf{s}_i\|^2 | \mathbf{w}_{i-1} \right\} \\ & \stackrel{(b)}{\leq} J(\mathbf{w}_{c,i-1}) - \mu \|\nabla J(\mathbf{w}_{c,i-1})\|^2 + \frac{\mu}{2} \|\nabla J(\mathbf{w}_{c,i-1})\|^2 \\ & \quad + \frac{\mu}{2} \|\mathbf{d}_{i-1}\|^2 + \mu^2 \delta \|\nabla J(\mathbf{w}_{c,i-1})\|^2 + \mu^2 \delta \|\mathbf{d}_{i-1}\|^2 \\ & \quad + \mu^2 \frac{\delta}{2} \mathbb{E} \left\{ \|\mathbf{s}_i\|^2 | \mathbf{w}_{i-1} \right\} \end{aligned}$$

$$\begin{aligned}
&\stackrel{(c)}{\leq} J(\mathbf{w}_{c,i-1}) - \frac{\mu}{2} (1 - 2\mu\delta) \|\nabla J(\mathbf{w}_{c,i-1})\|^2 \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \|\mathbf{d}_{i-1}\|^2 + \mu^2 \frac{\delta}{2} \sigma^2
\end{aligned} \tag{87}$$

where cross-terms were removed in (a) due to the conditional zero-mean condition (13a), (b) follows from $\|a+b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ and from $-2a^\top b \leq \|a\|^2 + \|b\|^2$ and (c) is a result of grouping terms and Lemma 1.

Note that (87) continues to be random due to the conditioning on $\mathcal{W}_{i-1}$, but that it holds for every choice of $\mathcal{W}_{i-1}$ with probability 1. Furthermore, since $\mathbf{w}_{c,i-1} = \sum_{k=1}^N p_k \mathbf{w}_{k,i-1}$, the centroid $\mathbf{w}_{c,i-1}$ is deterministic conditioned on $\mathcal{W}_{i-1}$. As such, the event $\mathbf{w}_{c,i-1} \in \mathcal{G}$ is deterministic conditioned on $\mathcal{W}_{i-1}$, and (87) holds for every $\mathbf{w}_{c,i-1} \in \mathcal{G}$. We can hence take expectations over $\mathbf{w}_{c,i-1} \in \mathcal{G}$ and apply Lemma 4 to find:

$$\begin{aligned}
&\mathbb{E} \{J(\mathbf{w}_{c,i}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} \\
&\leq \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} \\
&\quad - \frac{\mu}{2} (1 - 2\mu\delta) \mathbb{E} \left\{ \|\nabla J(\mathbf{w}_{c,i-1})\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} + \mu^2 \frac{\delta}{2} \sigma^2 \\
&\stackrel{(a)}{\leq} \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} - \mu^2 c_1 \frac{c_2}{c_1} \left(1 + \frac{1}{\pi}\right) \\
&\quad + O(\mu) \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} + \mu^2 c_2 \\
&\stackrel{(b)}{\leq} \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} - \mu^2 \frac{c_2}{\pi} \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\}
\end{aligned} \tag{88}$$

In step (a) we applied Definition 2, and in particular, that from (40) $\|\nabla J(\mathbf{w}_{c,i-1})\|^2 \geq \mu \frac{c_2}{c_1} (1 + \frac{1}{\pi})$ whenever $\mathbf{w}_{c,i-1} \in \mathcal{G}$, which implies:

$$\mathbb{E} \left\{ \|\nabla J(\mathbf{w}_{c,i-1})\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} \geq \mu \frac{c_2}{c_1} \left(1 + \frac{1}{\pi}\right) \tag{89}$$

We also collected constants into c_1 and c_2 defined in (43)–(44) for brevity. Step (b) is obtained by grouping terms. Note that from Lemma 1, we have a bound on $\mathbb{E} \|\mathbf{d}_{i-1}\|^2$, but not on the partial expectation conditioned over $\mathbf{w}_{c,i-1} \in \mathcal{G}$. We can decompose the full expectation:

$$\begin{aligned}
&\mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 \right\} \\
&= \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} \cdot \pi_{i-1}^{\mathcal{G}} \\
&\quad + \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G}^C \right\} \cdot \pi_{i-1}^{\mathcal{G}^C} \stackrel{(37)}{\leq} O(\mu^2)
\end{aligned} \tag{90}$$

which implies

$$\mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{G} \right\} \leq \frac{O(\mu^2)}{\pi_{i-1}^{\mathcal{G}}} \tag{91}$$

so that we obtain for (88):

$$\begin{aligned}
&\mathbb{E} \{J(\mathbf{w}_{c,i}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} \\
&\leq \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{G}\} - \mu^2 \frac{c_2}{\pi} + \frac{O(\mu^3)}{\pi_{i-1}^{\mathcal{G}}}
\end{aligned} \tag{92}$$

Similarly:

$$\begin{aligned}
&\mathbb{E} \{J(\mathbf{w}_{c,i}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} \\
&\leq \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} \\
&\quad - \frac{\mu}{2} (1 - 2\mu\delta) \mathbb{E} \left\{ \|\nabla J(\mathbf{w}_{c,i-1})\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{M} \right\} \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{M} \right\} + \mu^2 \frac{\delta}{2} \sigma^2 \\
&\stackrel{(a)}{\leq} \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} + \mu^2 c_2 \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \left\{ \|\mathbf{d}_{i-1}\|^2 | \mathbf{w}_{c,i-1} \in \mathcal{M} \right\} \\
&\stackrel{(b)}{\leq} \mathbb{E} \{J(\mathbf{w}_{c,i-1}) | \mathbf{w}_{c,i-1} \in \mathcal{M}\} + \mu^2 c_2 + \frac{O(\mu^3)}{\pi_{i-1}^{\mathcal{M}}}
\end{aligned} \tag{93}$$

where (a) follows from the fact that $\|\nabla J(\mathbf{w}_{c,i-1})\|^2 \geq 0$ with probability 1 and (b) made use of the same argument that led to (92).

To bound the expected function value at time i_o , we take full expectation of (87) and obtain:

$$\begin{aligned}
&\mathbb{E} J(\mathbf{w}_{c,i}) \\
&\leq \mathbb{E} J(\mathbf{w}_{c,i-1}) - \frac{\mu}{2} (1 - 2\mu\delta) \mathbb{E} \|\nabla J(\mathbf{w}_{c,i-1})\|^2 \\
&\quad + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \|\mathbf{d}_{i-1}\|^2 + \mu^2 \frac{\delta}{2} \sigma^2
\end{aligned} \tag{94}$$

$$\stackrel{(a)}{\leq} \mathbb{E} J(\mathbf{w}_{c,i-1}) + \frac{\mu}{2} (1 + 2\mu\delta) \mathbb{E} \|\mathbf{d}_{i-1}\|^2 + \mu^2 \frac{\delta}{2} \sigma^2 \tag{95}$$

where (a) follows since $\frac{\mu}{2} (1 - 2\mu\delta) \mathbb{E} \|\nabla J(\mathbf{w}_{c,i-1})\|^2 \geq 0$ for $\mu \leq \frac{1}{2\delta}$. Iterating, starting at $i = 0$, we find:

$$\mathbb{E} J(\mathbf{w}_{c,i_o}) \leq J(\mathbf{w}_{c,0}) + \frac{\mu}{2} (1 + 2\mu\delta) \sum_{i=0}^{i_o-1} \mathbb{E} \|\mathbf{d}_i\|^2 + i_o \mu^2 \frac{\delta}{2} \sigma^2 \tag{96}$$

We bound the term involving the aggregate disagreement using Eq. (77). After taking square-roots and applying sub-additivity of the square root, we find:

$$\begin{aligned}
&\sqrt{\mathbb{E} \|\mathcal{V}_R^\top \mathbf{w}_i\|^4} \\
&\stackrel{(77)}{\leq} \|J_\epsilon^\top\|^{i/2} \|\mathcal{V}_R^\top \mathbf{w}_0\|^2 \\
&\quad + \mu^2 \frac{\sqrt{8} \|J_\epsilon^\top\|^2}{(1 - \|J_\epsilon^\top\|)^2} \|\mathcal{V}_R^\top\|^2 N \left(G^2 + \max_k \sigma_k^2 \right)
\end{aligned} \tag{97}$$

so that we have:

$$\begin{aligned}
&\sum_{i=0}^{i_o-1} \mathbb{E} \|\mathbf{d}_i\|^2 \\
&\stackrel{(a)}{\leq} \sum_{i=0}^{i_o-1} \sqrt{\mathbb{E} \|\mathbf{d}_i\|^4} \\
&\stackrel{(83)}{\leq} \sum_{i=0}^{i_o-1} \sqrt{\mathbb{E} \delta^4 p_{\max} \|\mathbf{w}_i - \mathbf{w}_{c,i}\|^4} \\
&\stackrel{(30)}{=} \sum_{i=0}^{i_o-1} \sqrt{\mathbb{E} \delta^4 p_{\max} \|\mathcal{V}_L \mathcal{V}_R^\top \mathbf{w}_i\|^4} \\
&\leq \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \sum_{i=0}^{i_o-1} \sqrt{\mathbb{E} \|\mathcal{V}_R^\top \mathbf{w}_i\|^4}
\end{aligned}$$

$$\begin{aligned}
&\leq \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \left(\sum_{i=0}^{i_o-1} \|J_\epsilon^\top\|^{i/2} \|\mathcal{V}_R^\top \mathcal{W}_0\|^2 \right) \\
&\quad + i_o \mu^2 \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \frac{\sqrt{8} \|J_\epsilon^\top\|^2}{(1 - \|J_\epsilon^\top\|)^2} \|\mathcal{V}_R^\top\|^2 N \left(G^2 + \max_k \sigma_k^2 \right) \\
&\leq \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \frac{\|\mathcal{V}_R^\top \mathcal{W}_0\|^2}{1 - \|J_\epsilon^\top\|^{1/2}} + i_o \cdot O(\mu^2) \quad (98)
\end{aligned}$$

where (a) follows from Jensen's inequality. Replacing the bounds in (96), we then find:

$$\begin{aligned}
&\mathbb{E} J(\mathbf{w}_{c,i_o}) \\
&\leq J(\mathbf{w}_{c,0}) + \frac{\mu}{2} (1 + 2\mu\delta) \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \frac{\|\mathcal{V}_R^\top \mathcal{W}_0\|^2}{1 - \|J_\epsilon^\top\|^{1/2}} \\
&\quad + i_o \cdot O(\mu^3) + i_o \cdot O(\mu^2) \\
&\stackrel{(a)}{\leq} J(\mathbf{w}_{c,0}) + \frac{\mu}{2} (1 + 2\mu\delta) \delta^2 \sqrt{p_{\max}} \|\mathcal{V}_L\|^2 \frac{\|\mathcal{V}_R^\top \mathcal{W}_0\|^2}{1 - \|J_\epsilon^\top\|^{1/2}} + o(\mu) \quad (99)
\end{aligned}$$

where (a) follows from Lemma 1, since $i_o = o(\mu^{-1})$.

APPENDIX F PROOF OF LEMMA 3

We refer to (57). Suppose $i \leq \frac{T}{\mu}$, where T is an arbitrary constant independent of μ . We then have for $i \geq 0$:

$$\begin{aligned}
&\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_{i+1}^{i*} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \\
&\stackrel{(57)}{=} \mathbb{E} \left\{ \left\| (I - \mu \mathbf{H}_{i*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i*}) \right. \right. \\
&\quad \left. \left. + \mu \mathbf{d}_{i*+i} + \mu \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \\
&\stackrel{(a)}{=} \left\| (I - \mu \mathbf{H}_{i*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i*}) + \mu \mathbf{d}_{i*+i} \right\|^2 \\
&\quad + \mu^2 \mathbb{E} \left\{ \left\| \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \\
&\stackrel{(b)}{=} \frac{1}{1 - \mu\delta} \left\| (I - \mu \mathbf{H}_{i*+i}) \tilde{\mathbf{w}}_i^{i*} \right\|^2 + \frac{\mu}{\delta} \left\| \nabla J(\mathbf{w}_{c,i*}) + \mathbf{d}_{i*+i} \right\|^2 \\
&\quad + \mu^2 \mathbb{E} \left\{ \left\| \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \\
&\stackrel{(c)}{=} \frac{1}{1 - \mu\delta} \left\| (I - \mu \mathbf{H}_{i*+i}) \tilde{\mathbf{w}}_i^{i*} \right\|^2 + 2 \frac{\mu}{\delta} \left\| \nabla J(\mathbf{w}_{c,i*}) \right\|^2 \\
&\quad + 2 \frac{\mu}{\delta} \left\| \mathbf{d}_{i*+i} \right\|^2 + \mu^2 \mathbb{E} \left\{ \left\| \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \\
&\stackrel{(d)}{\leq} \frac{(1 + \mu\delta)^2}{1 - \mu\delta} \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^2 + 2 \frac{\mu}{\delta} \left\| \nabla J(\mathbf{w}_{c,i*}) \right\|^2 \\
&\quad + 2 \frac{\mu}{\delta} \left\| \mathbf{d}_{i*+i} \right\|^2 + \mu^2 \mathbb{E} \left\{ \left\| \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathcal{F}_{i*+i} \right\} \quad (100)
\end{aligned}$$

where (a) follows from the conditional zero-mean property of the gradient noise term in Assumption 4, (b) follows from Jensen's inequality

$$\|a + b\|^2 \leq \frac{1}{\alpha} \|a\|^2 + \frac{1}{1 - \alpha} \|b\|^2 \quad (101)$$

with $\alpha = \mu\delta < 1$ and (c) follows from the same inequality with $\alpha = \frac{1}{2}$. Step (d) follows from the sub-multiplicative property of norms along with $-\delta I \leq \nabla^2 J(\mathbf{w}_{c,i*}) \leq \delta I$, which

follows from the Lipschitz gradient condition in Assumption 2. Since $\mathbf{w}_{c,i*}$ is deterministic conditioned on $\mathcal{F}_{i*+i}$ we can now take expectations over $\mathbf{w}_{c,i*} \in \mathcal{H}$ to obtain:

$$\begin{aligned}
&\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_{i+1}^{i*} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\leq \frac{(1 + \mu\delta)^2}{1 - \mu\delta} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\quad + 2 \frac{\mu}{\delta} \mathbb{E} \left\{ \left\| \mathbf{d}_{i*+i} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\quad + 2 \frac{\mu}{\delta} \mathbb{E} \left\{ \left\| \nabla J(\mathbf{w}_{c,i*}) \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\quad + \mu^2 \mathbb{E} \left\{ \left\| \mathbf{s}_{i*+i+1} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\stackrel{(a)}{\leq} \frac{(1 + \mu\delta)^2}{1 - \mu\delta} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} + 2 \frac{\mu}{\delta} \cdot \frac{O(\mu^2)}{\pi_{i*}^{\mathcal{H}}} \\
&\quad + 2 \frac{\mu}{\delta} \cdot O(\mu) + O(\mu^2) \\
&\leq \frac{(1 + \mu\delta)^2}{1 - \mu\delta} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} + O(\mu^2) + \frac{O(\mu^3)}{\pi_{i*}^{\mathcal{H}}} \quad (102)
\end{aligned}$$

where (a) follows from the perturbation bounds in Lemma 1 and the starting assumption that $\mathbf{w}_{c,i*}$ is an $O(\mu)$ -square stationary point. Note that, at time $i = 0$, we have:

$$\tilde{\mathbf{w}}_0^{i*} = \mathbf{w}_{c,i*} - \mathbf{w}_{c,i*+0} = 0 \quad (103)$$

and hence the initial deviation is zero, by definition. Iterating, starting at $i = 0$ yields:

$$\begin{aligned}
&\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^2 \middle| \mathbf{w}_{c,i*} \in \mathcal{H} \right\} \\
&\leq \left(\sum_{n=0}^{i-1} \left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^n \right) \left(O(\mu^2) + \frac{O(\mu^3)}{\pi_{i*}^{\mathcal{H}}} \right) \\
&= \frac{1 - \left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^i}{1 - \frac{(1 + \mu\delta)^2}{1 - \mu\delta}} \left(O(\mu^2) + \frac{O(\mu^3)}{\pi_{i*}^{\mathcal{H}}} \right) \\
&= \frac{\left(\left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^i - 1 \right) (1 - \mu\delta)}{3\delta + \mu\delta^2} \left(O(\mu) + \frac{O(\mu^2)}{\pi_{i*}^{\mathcal{H}}} \right) \\
&\leq \frac{\left(\left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^{\frac{T}{\mu}} - 1 \right) (1 - \mu\delta)}{3\delta + \mu\delta^2} \left(O(\mu) + \frac{O(\mu^2)}{\pi_{i*}^{\mathcal{H}}} \right) \\
&= O(\mu) + \frac{O(\mu^2)}{\pi_{i*}^{\mathcal{H}}} \quad (104)
\end{aligned}$$

where the last line follows from Lemma 2 after noting that:

$$\frac{\left(\left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^{\frac{T}{\mu}} - 1 \right) (1 - \mu\delta)}{3\delta + \mu\delta^2} \leq \frac{\left(\frac{(1 + \mu\delta)^2}{1 - \mu\delta} \right)^{\frac{T}{\mu}} - 1}{3\delta} \quad (105)$$

This establishes (61). We proceed to establish a bound on the fourth-order moment. Using the inequality [5]:

$$\|a + b\|^4 \leq \|a\|^4 + 3\|b\|^4 + 8\|a\|^2\|b\|^2 + 4\|a\|^2(a^\top b) \quad (106)$$

we have:

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_{i+1}^{i*} \right\|^4 \middle| \mathcal{F}_{i^*+i} \right\} \\
& \stackrel{(a)}{=} \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{d}_{i^*+i} \right\|^4 \\
& \quad + 3\mu^4 \mathbb{E} \left\{ \left\| \mathbf{s}_{i^*+i+1} \right\|^4 \middle| \mathcal{F}_{i^*+i} \right\} \\
& \quad + 8\mu^2 \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{d}_{i^*+i} \right\|^2 \\
& \quad \times \mathbb{E} \left\{ \left\| \mathbf{s}_{i^*+i+1} \right\|^2 \middle| \mathcal{F}_{i^*+i} \right\} \\
& \stackrel{(b)}{=} \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{d}_{i^*+i} \right\|^4 + O(\mu^4) \\
& \quad + \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{d}_{i^*+i} \right\|^2 O(\mu^2) \\
& \stackrel{(c)}{=} \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{d}_{i^*+i} \right\|^4 + O(\mu^4) \\
& \quad + \left(\left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} \right\|^2 \right. \\
& \quad \left. + \mu^2 \left\| \nabla J(\mathbf{w}_{c,i^*}) \right\|^2 + \mu^2 \left\| \mathbf{d}_{i^*+i} \right\|^2 \right) O(\mu^2) \quad (107)
\end{aligned}$$

where in step (a) we dropped cross-terms due to the conditional zero-mean property of the gradient noise in Assumption 4, step (b) follows from the fourth-order conditions on the gradient noise in Assumption 4 along with the perturbation bounds in Lemma 1, and (c) follows from Jensen's inequality, i.e. $\|a + b + c\|^2 \leq 3\|a\|^2 + 3\|b\|^2 + 3\|c\|^2$. Taking expectations over $\mathbf{w}_{c,i^*} \in \mathcal{H}$ on both sides and collecting constant factors along with μ in appropriate $O(\cdot)$ terms:

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_{i+1}^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \leq \mathbb{E} \left\{ \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) \right. \right. \\
& \quad \left. \left. + \mu \mathbf{d}_{i^*+i} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} + O(\mu^3) + \frac{O(\mu^6)}{\pi_i^{\mathcal{H}}} \quad (108)
\end{aligned}$$

Finally, from Jensen's inequality, we find for $0 < \alpha < 1$:

$$\|a + b\|^4 = \frac{1}{\alpha^3} \|a\|^4 + \frac{1}{(1-\alpha)^3} \|b\|^4 \quad (109)$$

and hence for $\alpha = 1 - \mu\delta$ and $0 < \mu < \frac{1}{\delta}$:

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) \right. \right. \\
& \quad \left. \left. + \mu \mathbf{d}_{i^*+i} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \stackrel{(109)}{\leq} \frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \quad + \frac{\mu^4}{\mu^3 \delta^3} \mathbb{E} \left\{ \left\| \nabla J(\mathbf{w}_{c,i^*}) + \mathbf{d}_{i^*+i} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \leq \frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} + O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \quad (110)
\end{aligned}$$

where the last line follows after applying (109) to $\|\nabla J(\mathbf{w}_{c,i^*}) + \mathbf{d}_{i^*+i}\|^4$. Hence,

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_{i+1}^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \leq \frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} + O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \quad (111)
\end{aligned}$$

Recall again that $\tilde{\mathbf{w}}_0^{i*} = 0$ and therefore iterating yields:

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\
& \leq \left(\sum_{n=0}^{i-1} \left(\frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \right)^n \right) \left(O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \right) \\
& = \frac{1 - \left(\frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \right)^i}{1 - \frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3}} \left(O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \right) \\
& \leq \frac{\left(\frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \right)^i - 1}{(1 + \mu\delta)^4 - (1 - \mu\delta)^3} \left(O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \right) \\
& \stackrel{(a)}{\leq} \frac{\left(\frac{(1 + \mu\delta)^4}{(1 - \mu\delta)^3} \right)^i - 1}{O(\mu)} \left(O(\mu^3) + \frac{O(\mu^5)}{\pi_i^{\mathcal{H}}} \right) \\
& \leq O(\mu^2) + \frac{O(\mu^4)}{\pi_i^{\mathcal{H}}} \quad (112)
\end{aligned}$$

where in (a) we expanded $(1 + \mu\delta)^4 - (1 - \mu\delta)^3 = O(\mu)$ and the last step follows from Lemma 2. This establishes (63). Eq. (62) then follows from Jensen's inequality via:

$$\begin{aligned}
& \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^3 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \leq \left(\mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \middle| \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \right)^{3/4} \\
& \leq \left(O(\mu^2) + \frac{O(\mu^4)}{\pi_i^{\mathcal{H}}} \right)^{3/4} \leq O(\mu^{3/2}) + \frac{O(\mu^3)}{\pi_i^{\mathcal{H}}} \quad (113)
\end{aligned}$$

We now study the difference between the short-term model (59) and the true recursion (57). We have:

$$\begin{aligned}
& \mathbf{w}_{c,i^*+i+1} - \mathbf{w}'_{c,i^*+i+1} \\
& = - (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} - \mu \nabla J(\mathbf{w}_{c,i^*}) - \mu \mathbf{d}_{i^*+i} - \mu \mathbf{s}_{i^*+i+1} \\
& \quad + (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} + \mu \nabla J(\mathbf{w}_{c,i^*}) + \mu \mathbf{s}_{i^*+i+1} \\
& = - (I - \mu \mathbf{H}_{i^*+i}) \tilde{\mathbf{w}}_i^{i*} - \mu \mathbf{d}_{i^*+i} + (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} \\
& = (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) (\mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i}) - \mu \mathbf{d}_{i^*+i} \\
& \quad + \mu (\mathbf{H}_{i^*+i} - \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} \quad (114)
\end{aligned}$$

Before proceeding, note that the difference between the Hessians in the driving term can be bounded as:

$$\begin{aligned}
& \left\| \nabla^2 J(\mathbf{w}_{c,i^*}) - \mathbf{H}_{i^*+i} \right\| \\
& = \left\| \int_0^1 (\nabla^2 J(\mathbf{w}_{c,i^*}) - \nabla^2 J((1-t)\mathbf{w}_{c,i^*+i} + t\mathbf{w}_{c,i^*})) dt \right\| \\
& \stackrel{(a)}{\leq} \int_0^1 \left\| \nabla^2 J(\mathbf{w}_{c,i^*}) - \nabla^2 J((1-t)\mathbf{w}_{c,i^*+i} + t\mathbf{w}_{c,i^*}) \right\| dt \\
& \stackrel{(b)}{\leq} \rho \int_0^1 \left\| (1-t)\mathbf{w}_{c,i^*} - (1-t)\mathbf{w}_{c,i^*+i} \right\| dt
\end{aligned}$$

$$= \rho \left\| \tilde{\mathbf{w}}_i^{i*} \right\| \int_0^1 (1-t) dt = \frac{\rho}{2} \left\| \tilde{\mathbf{w}}_i^{i*} \right\| \quad (115)$$

where (a) follows Jensen's inequality and (b) follows from the Lipschitz Hessian assumption 5. Returning to (114) and taking norms yields:

$$\begin{aligned} & \left\| \mathbf{w}_{c,i^*+i+1} - \mathbf{w}'_{c,i^*+i+1} \right\|^2 \\ &= \left\| (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) (\mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i}) \right. \\ & \quad \left. - \mu \mathbf{d}_{i^*+i} + \mu (\mathbf{H}_{i^*+i} - \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} \right\|^2 \\ &\stackrel{(a)}{\leq} \frac{1}{1-\mu\delta} \left\| (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) (\mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i}) \right\|^2 \\ & \quad + \frac{\mu^2}{\mu\delta} \left\| \mathbf{d}_{i^*+i} + (\mathbf{H}_{i^*+i} - \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} \right\|^2 \\ &\stackrel{(b)}{\leq} \frac{1}{1-\mu\delta} \left\| (I - \mu \nabla^2 J(\mathbf{w}_{c,i^*})) (\mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i}) \right\|^2 \\ & \quad + 2 \frac{\mu}{\delta} \left(\left\| \mathbf{d}_{i^*+i} \right\|^2 + \left\| (\mathbf{H}_{i^*+i} - \nabla^2 J(\mathbf{w}_{c,i^*})) \tilde{\mathbf{w}}_i^{i*} \right\|^2 \right) \\ &\stackrel{(115)}{\leq} \frac{(1+\mu\delta)^2}{1-\mu\delta} \left\| \mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i} \right\|^2 \\ & \quad + 2 \frac{\mu}{\delta} \left(\left\| \mathbf{d}_{i^*+i} \right\|^2 + \frac{\rho}{2} \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \right) \end{aligned} \quad (116)$$

where (a) again follows from Jensen's inequality (101) with $\alpha = 1 - \mu\delta$ and (b) follows from the same inequality with $\alpha = \frac{1}{2}$. Taking expectations over $\mathbf{w}_{c,i^*} \in \mathcal{H}$ yields:

$$\begin{aligned} & \mathbb{E} \left\{ \left\| \mathbf{w}_{c,i^*+i+1} - \mathbf{w}'_{c,i^*+i+1} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\ &\leq \frac{(1+\mu\delta)^2}{1-\mu\delta} \mathbb{E} \left\{ \left\| \mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\ & \quad + 2 \frac{\mu}{\delta} \mathbb{E} \left\{ \left\| \mathbf{d}_{i^*+i} \right\|^2 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\ & \quad + \frac{\rho\mu}{\delta} \mathbb{E} \left\{ \left\| \tilde{\mathbf{w}}_i^{i*} \right\|^4 \mid \mathbf{w}_{c,i^*} \in \mathcal{H} \right\} \\ &\stackrel{(a)}{\leq} \frac{(1+\mu\delta)^2}{1-\mu\delta} \mathbb{E} \left\| \mathbf{w}_{c,i^*+i} - \mathbf{w}'_{c,i^*+i} \right\|^2 + O(\mu^3) + \frac{O(\mu^3)}{\pi_{i^*}^{\mathcal{H}}} \end{aligned} \quad (117)$$

where (a) follows from the bound on the network disagreement in Lemma 3.

Since both the true and the short-term model are initialized at $\mathbf{w}_{c,i^*}$, we have $\mathbf{w}_{c,i^*+0} - \mathbf{w}'_{c,i^*+0} = 0$. Iterating and applying the same argument as above leads to:

$$\mathbb{E} \left\| \mathbf{w}_{c,i^*+i+1} - \mathbf{w}'_{c,i^*+i+1} \right\|^2 \leq O(\mu^2) + \frac{O(\mu^2)}{\pi_{i^*}^{\mathcal{H}}} \quad (118)$$

which is (64). Eq. (65) and (66) follow using arguments analogous to (61) and (62).

REFERENCES

- [1] S. Vlaski and A. H. Sayed, "Diffusion learning in non-convex environments," in *Proc. of IEEE ICASSP*, Brighton, UK, May 2019, pp. 5262–5266.
- [2] S. Vlaski and A. H. Sayed, "Second-order guarantees in centralized, federated and decentralized non-convex optimization," *Communications in Information Systems*, vol. 20, pp. 353 – 388, 2020, also available as arXiv:2003.14366.
- [3] S. Vlaski and A. H. Sayed, "Distributed learning in non-convex environments – Part II: Polynomial escape from saddle-points," *submitted for publication, see also arXiv version*, July 2019.
- [4] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Trans. Automatic Control*, vol. 54, no. 1, pp. 48–61, Jan 2009.
- [5] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, July 2014.
- [6] J. Chen and A. H. Sayed, "On the learning behavior of adaptive networks - Part I: Transient analysis," *IEEE Transactions on Information Theory*, vol. 61, no. 6, pp. 3487–3517, June 2015.
- [7] R. Xin, A. K. Sahu, U. A. Khan, S. Kar, "Distributed stochastic optimization with gradient tracking over strongly-connected networks," *available as arXiv:1903.07266*, March 2019.
- [8] K. Yuan, B. Ying, S. Vlaski, and A. H. Sayed, "Stochastic gradient descent with finite samples sizes," in *Proc. of IEEE MLSP*, Vietri sul Mare, Italy, Sep. 2016, pp. 1–6.
- [9] W. Shi, Q. Ling, G. Wu, and W. Yin, "Extra: An exact first-order algorithm for decentralized consensus optimization," *SIAM Journal on Optimization*, vol. 25, no. 2, pp. 944–966, 2015.
- [10] K. Yuan, B. Ying, X. Zhao, and A. H. Sayed, "Exact diffusion for distributed optimization and learning – Part II: Convergence analysis," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 724–739, Feb 2019.
- [11] A. H. Sayed, "Adaptive networks," *Proceedings of the IEEE*, vol. 102, no. 4, pp. 460–497, April 2014.
- [12] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 2003.
- [13] S. U. Pillai, T. Suel, and S. Cha, "The Perron-Frobenius theorem: Some of its applications," *IEEE Signal Processing Magazine*, vol. 22, no. 2, pp. 62–75, March 2005.
- [14] R. Nassif, C. Richard, A. Ferrari, and A. H. Sayed, "Proximal multitask learning over networks with sparsity-inducing coregularization," *IEEE Transactions on Signal Processing*, vol. 64, no. 23, pp. 6329–6344, Dec 2016.
- [15] Z. J. Towfic and A. H. Sayed, "Adaptive penalty-based distributed stochastic convex optimization," *IEEE Trans. on Signal Process.*, vol. 62, no. 15, pp. 3924–3938, Aug. 2014.
- [16] B. Ying and A. H. Sayed, "Performance limits of stochastic sub-gradient learning, Part II: Multi-agent case," *Signal Processing*, vol. 144, pp. 253 – 264, 2018.
- [17] M. A. Hearst, S. T. Dumais, E. Osuna, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intelligent Systems and their Applications*, vol. 13, no. 4, pp. 18–28, July 1998.
- [18] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*, Cambridge University Press, 2018.
- [19] I. Tosic and P. Frossard, "Dictionary learning," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27–38, March 2011.
- [20] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun, "The Loss Surfaces of Multilayer Networks," in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, San Diego, May 2015, pp. 192–204.
- [21] S. Gelfand and S. Mitter, "Recursive stochastic algorithms for global optimization in $\mathbb{R}^d$," *SIAM Journal on Control and Optimization*, vol. 29, no. 5, pp. 999–1018, 1991.
- [22] R. Ge, F. Huang, C. Jin, and Y. Yuan, "Escaping from saddle points—online stochastic gradient for tensor decomposition," in *Proc. of Conference on Learning Theory*, Paris, France, 2015, pp. 797–842.
- [23] J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht, "Gradient descent only converges to minimizers," in *29th Annual Conference on Learning Theory*, New York, 2016, pp. 1246–1257.
- [24] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, "How to escape saddle points efficiently," in *Proc. of ICML*, Sydney, Australia, Aug. 2017, pp. 1724–1732.
- [25] H. Daneshmand, J. Kohler, A. Lucchi and T. Hofmann, "Escaping saddles with stochastic gradients," *available as arXiv:1803.05999*, March 2018.
- [26] C. Fang, C. J. Li, Z. Lin, and T. Zhang, "SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator," in *Proc. of NIPS*, pp. 689–699, Montreal, Canada, 2018.
- [27] Z. Allen-Zhu, "Natasha 2: Faster non-convex optimization than SGD," in *Proc. of NIPS*, pp. 2675–2686, Montreal, Canada, Dec. 2018.
- [28] Z. Allen-Zhu and Y. Li, "NEON2: Finding local minima via first-order oracles," in *Proc. of NIPS*, pp. 3716–3726, Montreal, Canada, Dec. 2018.

- [29] C. Fang, Z. Lin and T. Zhang, "Sharp analysis for nonconvex sgd escaping from saddle points," *available as arXiv:1902.00247*, Feb. 2019.
- [30] C. Jin, P. Netrapalli, R. Ge, S. M. Kakade and M. I. Jordan, "Stochastic gradient descent escapes saddle points efficiently," *available as arXiv:1902.04811*, Feb. 2019.
- [31] P. Di Lorenzo and G. Scutari, "NEXT: in-network nonconvex optimization," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 2, no. 2, pp. 120–136, June 2016.
- [32] X. Lian, C. Zhang, H. Zhang, C.-J. Hsieh, W. Zhang, and J. Liu, "Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., pp. 5330–5340. Curran Associates, Inc., 2017.
- [33] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *Journal of Scientific Computing*, vol. 78, no. 1, pp. 29–63, Jan. 2019.
- [34] T. Tatarenko and B. Touri, "Non-convex distributed optimization," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3744–3757, Aug. 2017.
- [35] A. Daneshmand, G. Scutari and V. Kungurtsev, "Second-order guarantees of distributed gradient algorithms," *available as arXiv:1809.08694*, Sep. 2018.
- [36] B. Swenson, S. Kar, H. V. Poor and J. M. F. Moura, "Annealing for distributed global optimization," *available as arXiv:1903.07258*, March 2019.
- [37] P. Jain and P. Kar, "Non-convex optimization for machine learning," *Foundations and Trends in Machine Learning*, vol. 10, no. 3-4, pp. 142–336, 2017.
- [38] S. J. Reddi, A. Hefny, S. Sra, B. Póczos, and A. Smola, "Stochastic variance reduction for nonconvex optimization," in *Proc. of ICML*, New York, NY, USA, 2016, pp. 314–323.
- [39] R. Ge, Z. Li, W. Wang and X. Wang, "Stabilized SVRG: Simple variance reduction for nonconvex optimization," *available as arXiv:1905.00529*, May 2019.
- [40] Y. Nesterov and B.T. Polyak, "Cubic regularization of newton method and its global performance," *Mathematical Programming*, vol. 108, no. 1, pp. 177–205, Aug 2006.
- [41] A. Klenke, *Probability Theory: A Comprehensive Course*, Springer, 2013.
- [42] J. Chen and A. H. Sayed, "On the learning behavior of adaptive networks – Part II: Performance analysis," *IEEE Transactions on Information Theory*, vol. 61, no. 6, pp. 3518–3548, June 2015.
- [43] X. Zhao and A. H. Sayed, "Asynchronous adaptation and learning over networks – Part II: Performance analysis," *IEEE Transactions on Signal Processing*, vol. 63, no. 4, pp. 827–842, Feb 2015.



Highly Cited Researcher, and is a member of the US National Academy of Engineering.

Ali H. Sayed (S'90-M'92-SM'99-F'01) is Dean of Engineering at EPFL, Switzerland, and principal investigator of the Adaptive Systems Laboratory. He has served as distinguished professor and former chairman of electrical engineering at UCLA. An author of over 550 scholarly publications and six books, his research involves several areas including adaptation and learning theories, data and network sciences, statistical inference, multi-agent systems, and optimization. His work has been recognized with several major awards. He is recognized as a



Stefan Vlaski (S'13) received the B.Sc. degree in Electrical Engineering from Technical University Darmstadt, Germany in 2013 and the M.S. in Electrical Engineering as well as Ph.D. in Electrical and Computer Engineering from the University of California, Los Angeles in 2014 and 2019 respectively. He is currently a postdoctoral researcher at the Adaptive Systems Laboratory, EPFL, Switzerland. His research interests are in machine learning, signal processing, and optimization. His current focus is on the development and study of learning algorithms

with a particular emphasis on adaptive and decentralized solutions. He was a recipient of the German National Scholarship at TU Darmstadt and the Graduate Division Fellowship at UCLA.