

Decentralized Sparse Linear Regression via Gradient-Tracking

Marie Maros*

MMAROS@TAMU.EDU

*Wm Michael Barnes '64 Department of Industrial & Systems Engineering
Texas A&M University
College Station, TX 77843, USA*

Gesualdo Scutari**

GSCUTARI@PURDUE.EDU

*School of Industrial Engineering
School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907, USA*

Ying Sun*

YBS5190@PSU.EDU

*School of Electrical Engineering and Computer Science
The Pennsylvania State University
University Park, PA 16802, USA*

Guang Cheng*

GUANGCHENG@UCLA.EDU

*Department of Statistics
University of California Los Angeles
Los Angeles, USA.*

* *The order of the first three authors is alphabetic.*

** *Corresponding author.*

Editor: Pradeep Ravikumar

Abstract

We study sparse linear regression over a network of agents, modeled as an undirected graph without a center node. The estimation of the s -sparse parameter is formulated as a constrained LASSO problem wherein each agent owns a subset of the N total observations. We analyze the convergence rate and statistical guarantees of a distributed projected gradient tracking-based algorithm under high-dimensional scaling, allowing the ambient dimension d to grow with (and possibly exceed) the sample size N . Our theory shows that, under standard notions of restricted strong convexity and smoothness of the average loss functions, suitable conditions on the network connectivity and algorithm tuning, the distributed algorithm converges globally at a *linear* rate to an estimate that is within the centralized *statistical precision* of the model, $O(s \log d/N)$. When $s \log d/N = o(1)$, a condition necessary for statistical consistency, an ε -optimal solution is attained after $\mathcal{O}(\kappa \log(1/\varepsilon))$ gradient computations and $\mathcal{O}(\kappa/(1-\rho) \log(1/\varepsilon))$ communication rounds, where κ is the restricted condition number of the loss function and ρ measures the network connectivity. The computation cost matches that of the centralized projected gradient algorithm despite having data distributed; whereas the communication rounds reduce as the network connectivity improves. Overall, our study reveals interesting connections between statistical efficiency, network connectivity & topology, and convergence rate in the high dimensional setting.




Figure 1: Two network architectures: master-worker and mesh networks.

Keywords: High-dimensional estimation, distributed convex optimization, linear convergence, sparse linear regression, gradient-tracking¹

1. Introduction

Datasets with massive sample size and high-dimensionality are ubiquitous in modern science and engineering; examples include genomics and biomedical data, social media, financial time-series, and e-commerce data, just to name a few. The archetypal scenario is the one wherein data are produced, collected, or stored at different times and locations. The sheer volume and spatial/temporal disparity of data render centralized processing and storage on standalone machines prohibitively expensive or outright impossible. This has motivated in recent years a surge of interest in developing distributed methods that enable analytics and computations across multiple machines (henceforth referred to as agents), connected by a communication network. Broadly speaking, two architectures have captured most of the interest of the research community, namely: (i) the *master-worker* topology and (ii) a general, connected, multiagent architecture (a.k.a. *mesh network*). They are depicted in Fig. 1 and briefly commented next.

In a *master-worker* architecture, there is typically one or more master nodes connected to all other nodes, termed workers. The workers store part of the data, execute (in parallel) intermediate computations, and communicate (iteratively) suitable outcomes to the master node, which maintains and updates the authoritative copy of the optimization variables, producing eventually the final output. These federated architectures have been adopted in several applications to parallelize and decompose a variety of learning and optimization tasks; see, e.g., (Li et al., 2020; McMahan et al., 2017; Boyd et al., 2011). The statistical community is best acquainted with *divide-and-conquer* (D&C) methods whereby the master node combines the estimators produced by the workers using their local datasets. The idea has been widely applied in statistical estimation and inference such as M-estimation (Zhang et al., 2013b; Chen and Xie, 2014; Shi et al., 2018), sparse high dimensional models Battay et al. (2018); Lee et al. (2017), PCA Fan et al. (2019), matrix factorization Mackey et al. (2011), quantile regression Volgushev et al. (2019); Chen et al. (2019) and non- and semi-parametric methods including (Neiswanger et al., 2014; Shang and Cheng, 2017; Zhao et al.,

1. This work was conducted when Sun, Maros, and Cheng were at Purdue University.

2016; Zhang et al., 2013a; Kleiner et al., 2014), just to name a few.

However, such architectures may be impractical or undesirable in various scenarios. For example, in large-scale systems, the master node can become a systemic bottleneck: (i) its failure can disrupt the entire network, (ii) the node’s communication resources may become insufficient as the network size grows (owing to restricted bandwidth), and (iii) the implementation of some algorithms may require prohibitive computation efforts on standalone machines (Hendriks et al., 2020), particularly when processing high-dimensional and large datasets. Furthermore, there are scenarios where establishing a star topology (or a spanning tree) is simply not feasible or too resource-expensive; this is the case, e.g., of wireless networks with low-power devices, which can communicate only with nodes in their physical proximity. To circumventing these impediments, a natural approach is to eliminate master nodes, transitioning to peer-to-peer architectures wherein each node is connected exclusively to a subset of other nodes. These are the systems considered in this work.

1.1 Notation

Throughout the paper vectors are boldfaced and matrices are boldfaced and capitalized. We denote by $\mathbf{1}$ the vector of all ones, and by $\mathbf{J} = \mathbf{1}\mathbf{1}^\top$ the projection onto the consensus subspace; dimensions of these quantities will be clear from the context. Let $\|\cdot\|_p$ denote ℓ_p norm, for $p \geq 1$. When p is not specified. With a slight abuse of notation, we also denote by $\|\cdot\|_p$ the operator norm when applied to matrices, induced by the vector ℓ_p norm; $\|\cdot\|_F$ denotes the Frobenius norm.

1.2 Sparse linear regression over mesh networks

We study sparse linear regression problem over a mesh network of m agents, modeled as a general connected, undirected graph. Each agent i takes n linear measures of an s -sparse signal $\boldsymbol{\theta}^* \in \mathbb{R}^d$:

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\theta}^* + \mathbf{n}_i, \quad (1)$$

where $\mathbf{y}_i \in \mathbb{R}^n$ is the vector of measurements, $\mathbf{X}_i \in \mathbb{R}^{n \times d}$ is the design matrix, and $\mathbf{n}_i \in \mathbb{R}^n$ is the observation noise. For simplicity we assume that all agents acquire the same number n of measurements, accumulating to a total of $N = m \cdot n$ over the network. Our focus is on the high-dimensional setting where both the ambient dimension d and the sample size N grow, with d faster than N .

A standard approach to estimate $\boldsymbol{\theta}^*$ from $\{(\mathbf{y}_i, \mathbf{X}_i)\}_{i=1}^m$ is solving the LASSO problem, whose constrained form reads

$$\hat{\boldsymbol{\theta}} \in \operatorname{argmin}_{\|\boldsymbol{\theta}\|_1 \leq r} \left\{ \mathcal{L}(\boldsymbol{\theta}) \triangleq \frac{1}{m} \sum_{i=1}^m \mathcal{L}_i(\boldsymbol{\theta}) \right\}, \quad \text{where} \quad \mathcal{L}_i(\boldsymbol{\theta}) \triangleq \frac{1}{2n} \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\theta}\|^2. \quad (2)$$

In (2), each \mathcal{L}_i is the least squares loss, local to agent i , and the ℓ_1 -norm constraint (with r known to all the agents) aims at promoting sparsity on the solution $\hat{\boldsymbol{\theta}}$.

As an instance of the empirical risk minimization problem with nonsmooth regularization, there is a large body of distributed first order algorithms applicable for (2). Among these methods, the gradient tracking based ones have been demonstrated successful in improving the algorithmic efficiency both empirically and theoretically (Xu et al., 2018; Di Lorenzo and Scutari, 2016; Nedić et al., 2016). As such, this technique has been widely

applied to design distributed algorithms in various problem setups. Based on decentralizing the proximal gradient algorithm (PGD), the DGT algorithm for problem (2) decouples the optimization by letting each agent i locally maintain an estimator $\theta_i \in \mathbb{R}^d$ of the common variable θ . At every iteration t , each agent $i \in [m]$ performs in parallel the following update:

$$\text{comm. step: } \theta_i^t = \sum_{j=1}^m w_{ij} \theta_j^{t-\frac{1}{2}}, \quad \mathbf{g}_i^t = \sum_{j=1}^m w_{ij} \left(\mathbf{g}_j^{t-1} + \nabla \mathcal{L}_j(\theta_j^t) - \nabla \mathcal{L}_j(\theta_j^{t-1}) \right), \quad (3a)$$

$$\text{comp. step: } \theta_i^{t+\frac{1}{2}} = \prod_{\|\theta_i\|_1 \leq r} (\theta_i^t - \gamma^{-1} \mathbf{g}_i^t). \quad (3b)$$

In the computation step, each agent performs a proximal (projected) gradient update, where \mathbf{g}_i^t plays the role of estimating the centralized gradient $\nabla \mathcal{L}(\theta_i^t)$; $\prod_{\|\theta_i\|_1 \leq r}$ denotes the Euclidean projection onto the ℓ_1 ball $\{\theta_i \in \mathbb{R}^d : \|\theta_i\|_1 \leq r\}$. In the communication step, each agent updates the local variable $\theta_i^{t-\frac{1}{2}}$ and gradient estimator \mathbf{g}_i^{t-1} based on local averaging. The update of \mathbf{g}_i^t (gradient tracking step) can be explained as follows: At each iteration t , agent i subtracts the outdated gradient $\nabla \mathcal{L}_i(\theta_i^{t-1})$ at the previous iteration and adds the new one $\nabla \mathcal{L}_i(\theta_i^t)$; this refreshes the “memory” of \mathbf{g}_i and ensures that the sum of the \mathbf{g}_i^t 's is equal to the sum gradient $\sum_{i=1}^m \nabla \mathcal{L}_i(\theta_i^t)$. Then, agents perform a local averaging by computing a weighted sum using w_{ij} , enforcing thus consensus among the \mathbf{g}_i 's. If both θ_i 's and \mathbf{g}_i 's are asymptotically consensual, then

$$\mathbf{g}_i^t \rightarrow \frac{1}{m} \sum_{j=1}^m \mathbf{g}_j^t = \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\theta_j^t) \rightarrow \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\theta_j^t) \quad \text{as } t \rightarrow \infty. \quad (4)$$

That is, \mathbf{g}_i^t converges to $\nabla \mathcal{L}(\theta_i^t)$, as desired.

Despite the vast literature on decentralized algorithms, existing results of DGT are of pure optimization type; hence they lack guarantees in the high-dimensional setting, even for the sparse linear regression problem (2). In contrast, the PGD algorithm as its centralized counterpart has been proven to converge at a fast linear rate to statistically optimal solutions (Agarwal et al., 2012). For a detailed review of the literature, we refer readers to Section 1.4. Given this disparity, it naturally raises the question of whether DGT can enjoy the same fast rate and statistical optimality of PGD in the decentralized setting, and how the network configuration affects its performance. This paper addresses this open question.

1.3 Major contributions

We provide the statistical and convergence analysis of the DGT algorithm for solving the sparse linear regression problem (2). Our major contributions are the following.

- (i) **Statistical-computational guarantees:** By exploiting the Restricted Strong Convexity (RSC) and Smoothness (RSM) properties of \mathcal{L} (Agarwal et al., 2012), we establish regularity conditions under which the sequence generated by DGT provably convergences at a *linear rate* to an estimate whose distance to $\hat{\theta}$ is *smaller than the statistical precision* $\|\hat{\theta} - \theta^*\|$. Specifically, to enter an ε -neighborhood of such precision, it takes

$$O(\kappa \log(1/\varepsilon)) \quad (5)$$

iterations (gradient evaluations) as long as $\rho \leq \text{poly}(m, \kappa)^{-1}$, where κ is the restricted condition number of \mathcal{L} and ρ measures the network connectivity. The condition on ρ , when not met by the given network, can be enforced via multiple rounds of communication per iteration, resulting in

$$O\left(\frac{\kappa}{1-\rho} \log(m\kappa) \log(1/\varepsilon)\right) \quad (6)$$

overall communication rounds. As an implication, our result shows statistically optimal estimators can be computed at a linear rate by the DGT algorithm in the limit $N, d \rightarrow \infty$ and $s \log d/N = o(1)$. Furthermore, (5) matches the iteration complexity of the *centralized* proximal gradient descent (PGD) (Agarwal et al., 2012).

(ii) **Sample/network scaling:** The communication complexity (6) remains unchanged, for $s \log d/N$ being constant and for a fixed network. We also study the scalability of (6) with respect to the network size m , revealing tradeoffs between communication rounds and communication costs for different graph topologies, as summarized in Table 1. This sheds light on which network architectures are more favorable for high-dimensional estimation. For instance, among the commonly used network architectures, DGT running on the Erdős-Rényi graph with edge connecting probability $p = \log m/m$ achieves the lowest communication cost at the busiest node, and ties with the implementation on the star topology in the amount of total channel use.

(iii) **Technical novelties:** Our analysis represents a shift from both centralized approaches in high-dimension (e.g., (Agarwal et al., 2012)) and the prevailing techniques used for studying gradient-tracking algorithms in decentralized optimization (e.g., (Nedić et al., 2016; Sun et al., 2022c; Xu et al., 2021)), across several aspects.

(a) *Inexact descent under RSC/RSM:* Linear convergence of gradient-tracking methods has been proved under the assumption of *global* strong convexity of \mathcal{L} and global smoothness of \mathcal{L}_i 's—conditions that fall short in the high-dimensional setting where $d \gg N$. In contrast, our analysis employs the more nuanced RSC and RSM conditions—better suited for high-dimensional data and known to hold with high probability across various data generation models (Wainwright, 2019). Our first result is an inexact descent property of the estimation error (detailed in Proposition 11), achieving linear shrinkage subject to consensus and gradient-tracking errors. This can be viewed as a generalization of (Agarwal et al., 2012) to more complex scenarios posed by the mentioned errors (not present in (Agarwal et al., 2012)) and decentralized contexts.

(b) *A new analysis of gradient tracking errors in high-dimensions:* Traditional analyses of gradient tracking errors leverage the Lipschitz continuity of each $\nabla \mathcal{L}_i$ to establish *uniform* bounds on the tracking errors $\mathbf{g}_i^t - \nabla \mathcal{L}(\boldsymbol{\theta}_i^t)$, in terms of consensus and optimization errors (see Sec. 3.2.2). However, this approach proves inadequate in high-dimensional settings, where in the asymptotics $d/N \rightarrow \infty$, the Lipschitz constants of $\nabla \mathcal{L}_i$'s scale as $\mathcal{O}(d)$, for various predictor models (Wainwright, 2019), thereby diverging as $d \uparrow \infty$. Fundamentally, this issue arises because the gradients $\nabla \mathcal{L}_i(\boldsymbol{\theta}_i)$ remain dense vectors even when all $\boldsymbol{\theta}_i$'s are sparse. This would translate in convergence rates scaling unfavorably with the ambient dimension. Our new analysis posits that, under

the RSC and RSM conditions, it is sufficient to manage the tracking errors only along *sparse directions*. Within these directions, \mathbf{g}_i^t is proved to serve as a sufficiently accurate estimate of $\nabla \mathcal{L}(\boldsymbol{\theta}_i^t)$ (see Proposition 12). This crucial adjustment enables us to achieve convergence rates for DGT that are *independent of the ambient dimensions*.

(c) *Combining error dynamics via the z -transform*: The refined treatment of gradient-tracking errors, as discussed in (b), significantly complicates the convergence analysis of DGT. The error dynamics at any given iteration now depend on the consensus/tracking errors and optimization residuals from *all past iterates*, starting from initialization. This extended dependency chain precludes the use of simpler, traditional analytical techniques, such as those in (Xu et al., 2018; Sun et al., 2022c; Nedić et al., 2016) based on the small gain theorem, typically applied only to dynamics with *single-hop* time dependencies. To effectively manage this historical information, we employ the z -transform for finite-length sequences and obtain a sufficient condition for the linear convergence of DGT. This approach represents a technique of independent interest, offering potential applications to other problems of decentralized optimization.




Figure 2: Some commonly used graphs \mathcal{G} for the communication network. Dashed line in the Erdős-Rényi graph stands for the existence of an edge with some probability.

	line	2-d grid	star	complete	push-pull (star)	Erdős-Rényi	Erdős-Rényi
$(1 - \rho)^{-1}$	$O(m^2)$	$O(m \log m)$	$O(m^2)$	1	1	$O(1) [p = \log m/m]$	$O(1) [p = O(1)]$
max degree	2	4	m	m	m	$O(\log m)$	$O(m)$
channel use/comm.	$O(m)$	$O(m)$	$O(m)$	$O(m^2)$	$O(m)$	$O(m \log m)$	$O(m^2)$
max channel use	$\tilde{O}(m^2)$	$\tilde{O}(m)$	$\tilde{O}(m^3)$	$\tilde{O}(m)$	$\tilde{O}(m)$	$\tilde{O}(1)$	$\tilde{O}(m)$
total channel use	$\tilde{O}(m^3)$	$\tilde{O}(m^2)$	$\tilde{O}(m^3)$	$\tilde{O}(m^2)$	$\tilde{O}(m)$	$\tilde{O}(m)$	$\tilde{O}(m^2)$

Table 1: Scalability of the communication complexity of DGT with the network size m , for the network topology listed in Fig. 2. **Max degree**: maximum degree of the nodes; **channel use/comm.**: total channel use per communication round; **max channel use**: channel use of the busiest node to reach ε optimization precision; **total channel use**: total channel uses to reach ε optimization precision. \tilde{O} hides the logarithmic factors in communication complexity. Push-pull(star) refers to the centralized PGD algorithm implemented over star-networks, we refer the readers to the description around (11) for more details.

1.4 Related works

Distributed methods over master-worker systems. The last decade has witnessed a significant surge in research on statistical methods within master-worker architectures, as highlighted in Sec. 1. Much of this research has focused on D&C strategies (Zhang et al., 2013b; Chen and Xie, 2014; Shi et al., 2018; Battey et al., 2018; Lee et al., 2017; Fan et al., 2019; Mackey et al., 2011; Volgushev et al., 2019; Chen et al., 2019; Neiswanger et al., 2014; Shang and Cheng, 2017; Zhao et al., 2016; Zhang et al., 2013a; Kleiner et al., 2014). These methods typically involve a single round of communication from clients to the server and achieve statistical optimality, but *only with a limited number of client nodes*, restricting their use to small networks. To address these limitations, iterative approaches have been developed that allow for larger networks by removing the cap on the number of workers; examples include (Jordan et al., 2018; Fan et al., 2021; Wang et al., 2017). At the cost of multiple communication rounds, they effectively remove restrictions on the number of workers, at the expense of requiring a *minimum sample size* at each client to ensure the statistical optimality of the final estimator at the server.

All these algorithms heavily rely on the presence of a server node that, at each iteration, aggregates local gradients from clients into a centralized gradient and distributes updates. This setup ensures that the iterates are generated using *exact* gradient information, mimicking a fully centralized setting. Without such a centralized node, this mechanism vanishes, rendering these algorithms inapplicable on mesh networks. Moreover, simply replacing the exact gradient with approximations obtained by averaging neighboring local gradients is inadequate. This approach fails to achieve statistically optimal estimates or guarantee convergence, as it does not align with the assumptions and analytical frameworks of centralized methods, which do not account for gradient inaccuracies.

Distributed methods over mesh networks. Early decentralized methods over mesh networks are distributed (sub-)gradient descent (DGD) type methods including (Nedić and Ozdaglar, 2009; Nedić et al., 2010; Chen and Ozdaglar, 2012; Chen and Sayed, 2012; Nedić and Olshevsky, 2014). These algorithms are not directly applicable to (2) because they cannot handle constraints. Subsequent schemes handling constraints include (Sayed, 2014; Shi et al., 2015; Di Lorenzo and Scutari, 2016; Sun et al., 2022c,b; Xu et al., 2021).

Based on the order of local computation and communication, the DGD algorithm has the following two variations:

$$\begin{aligned}\boldsymbol{\theta}_i^{t+1} &= \prod_{\|\boldsymbol{\theta}_i\|_1 \leq r} \left(\sum_{j=1}^m w_{ij} \boldsymbol{\theta}_j^t - \gamma^{-1} \nabla \mathcal{L}_i(\boldsymbol{\theta}_i^t) \right), & \text{DGD-CTA} \\ \boldsymbol{\theta}_i^{t+1} &= \prod_{\|\boldsymbol{\theta}_i\|_1 \leq r} \left(\sum_{j=1}^m w_{ij} \left(\boldsymbol{\theta}_j^t - \gamma^{-1} \nabla \mathcal{L}_j(\boldsymbol{\theta}_j^t) \right) \right). & \text{DGD-ATC}\end{aligned}$$

In DGD-CTA, each agent at iteration t first averages the neighboring variables through a consensus step and then performs a gradient step using its local gradient. Conversely, in the DGD-ATC, each agent first executes the local gradient descent step before carrying out the consensus averaging. It is well established that DGD algorithms, even when applied to smooth and strongly convex loss functions \mathcal{L} , converge at a sublinear rate. To overcome

this drawback, the DGT [cf. eq. (3)] (Xu et al., 2018; Di Lorenzo and Scutari, 2016; Sun et al., 2022c; Nedić et al., 2016) and various primal-dual algorithms (Shi et al., 2015; Yuan et al., 2018; Xu et al., 2021) utilize an improved estimate of the centralized gradient at each agent’s location. Consequently, these algorithms achieve a linear convergence rate for loss functions \mathcal{L} that are both strongly convex and smooth. For a detailed overview of these methods, please see (Nedic, 2020; Nedić et al., 2018).

Despite intensive study, existing convergence results for these algorithms, derived purely from an optimization perspective, lead to pessimistic conclusions when applied to the high-dimensional problem (2). Specifically, since $N < d$, the average loss function \mathcal{L} lacks strong convexity, resulting in sublinear convergence rates for any *fixed* ambient dimension d . Moreover, when d increases at a rate faster than N , which is the typical regime in high-dimension (Wainwright, 2019), the Lipschitz constant of $\nabla\mathcal{L}$ will diverge as well, under standard choices of the design matrices $\{\mathbf{X}\}_{i=1}^m$. This requires progressively smaller step sizes, ultimately hindering the convergence of the algorithms.

First order algorithms in the high-dimensional setting. It is well-recognized that convergence analysis and parameter tuning of *centralized* first-order methods for (2) must exploit the landscape properties of \mathcal{L} as determined by the underlying statistical model. Notably, linear convergence to statistical optimal estimates is established in (Agarwal et al., 2012; Wang et al., 2014; Lin and Xiao, 2014; Xiao and Zhang, 2013) for the proximal gradient algorithm applied to (2), under the RSC and RSM properties of \mathcal{L} , which hold with high probability across various random data generation models and noise distribution (Wainwright, 2019). These conditions, coupled with a proper algorithm tuning, ensure that the algorithmic trajectory remains within a confined region where \mathcal{L} is approximately strongly convex and smooth, thereby securing linear convergence up to statistical optimality. An intriguing question is whether *distributed* first-order algorithms can similarly leverage this landscape to achieve rapid convergence toward statistically optimal estimates, particularly when local agents lack sufficient data samples to ensure statistical consistency locally.

In contrast to their centralized counterparts, the performance of distributed algorithms in high-dimension is less clear. While research is still evolving, existing studies (concurrent to this work (Sun et al., 2022a)) suggest that not all categories of distributed algorithms consistently benefit from these favorable landscape properties, indicating a more complex interaction among algorithmic design, distributed nature of data, and network topology, to be revealed. For instance, the analysis of a DGD-CTA type algorithm for the Lagrangian formulation of the LASSO problem, developed in the companion paper (Ji et al., 2023a), proves that the iterates generated by the algorithm enter an ε -neighborhood of a statistically optimal estimate of $\boldsymbol{\theta}^*$ after

$$\mathcal{O}\left(\frac{\kappa}{1-\rho} \cdot d \cdot m \log m \cdot \log \frac{1}{\varepsilon}\right) \quad (7)$$

number of iterations, where $\rho \in [0, 1)$ is a measure of the connectivity of the network (the smaller ρ , the more connected the graph, see Assumption 4 in Sec. 2 for a formal definition); and κ is the restricted condition number of \mathcal{L} . This result shows that DGD-CTA converges at a linear rate, but the number of iteration towards statistically optimal estimates grows as $\mathcal{O}(d)$, even in the asymptotic regime $N, d \uparrow \infty$ where the statistical error $\mathcal{O}((s \log d)/N)$ remains unchanged. In contrast, the DGD-ATC algorithm achieve statistical

optimal estimates at a converges at rate (Ji et al., 2023b)

$$\tilde{\mathcal{O}}\left(\frac{\kappa}{1-\rho} \cdot \log(dm) \cdot \log \frac{1}{\varepsilon}\right), \quad (8)$$

which has the more favorable $\mathcal{O}(\log d)$ scalability with d , compared to the CTA version.

The key difference between the two DGD variants lies in their local optimization approaches. In the CTA version, each agent uses its local gradient evaluated at a point that is *inconsistent* with the iterate obtained after the communication step, necessitating a step size of $\mathcal{O}(d^{-1})$ regardless of network connectivity. Conversely, by swapping the order of the communication and optimization steps, the ATC version allows for a *constant* step size in well-connected networks, enhancing scalability with the ambient dimension. Building on the insights from these analyses, in a paper (Maros and Scutari, 2022) subsequent to this work, we introduce a modified version of DGD (termed DGD²) that achieves statistical optimality at comparable communication and computational cost of DGT.

1.5 Paper organization

The remainder of the paper is structured as follows: Sec. 2 outlines the primary assumptions and details the statistical-computational guarantees—the proofs are reported in Sec. 3, with intermediate results available in the appendix. Numerical simulations corroborating our theoretical findings are discussed in Sec. 4. Finally, Sec. 5 draws some conclusions.

2. Convergence Analysis

We start by introducing the assumptions related to the LASSO problem (2) and the network topology, used for analyzing the convergence of DGT, given in (3).

2.1 Problem Setup and Main Assumptions

Assumption 1 (Global RSC/RSM) *Given $\mathbf{X} = [\mathbf{X}_1^\top, \dots, \mathbf{X}_m^\top]^\top \in \mathbb{R}^{N \times d}$, with $N = n \cdot m$, the following conditions hold:*

$$\begin{aligned} \text{(Global RSC/RSM)} \quad & \frac{\|\mathbf{X}\mathbf{u}\|^2}{N} \geq \mu_\Sigma \|\mathbf{u}\|^2 - \tau_\mu \|\mathbf{u}\|_1^2, \\ & \frac{\|\mathbf{X}\mathbf{u}\|^2}{N} \leq L_\Sigma \|\mathbf{u}\|^2 + \tau_g \|\mathbf{u}\|_1^2, \end{aligned} \quad \forall \mathbf{u} \in \mathbb{R}^d, \quad (9)$$

where (μ_Σ, τ_μ) and (L_Σ, τ_g) are positive RSC/RSM parameters.

Assumption 1 is a standard requirement for ensuring linear convergence of the PGD algorithm (Agarwal et al., 2012) to statistically optimal estimates, even in centralized settings. Additionally, we impose the following local restricted smoothness condition on the agents' losses, which is instrumental to ensure that the gradient tracking procedure estimates the centralized gradient accurately.

Assumption 2 (Local RSM) *For each local design matrix \mathbf{X}_i , the following holds:*

$$\frac{\|\mathbf{X}_i \mathbf{u}\|^2}{n} \leq \ell_\Sigma \|\mathbf{u}\|^2 + \tau_\ell \|\mathbf{u}\|_1^2, \quad \forall \mathbf{u} \in \mathbb{R}^d, \quad (10)$$

where (ℓ_Σ, τ_ℓ) are positive RSM parameters.

The mesh network is modeled as an undirected graph $\mathcal{G} \triangleq (\mathcal{V}, \mathcal{E})$, with $\mathcal{V} = \{1, \dots, m\}$ being the set of m nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ being the set of edges corresponding to the communication links: $(i, j) \in \mathcal{E}$ iff there exists a communication link between agent i and j . We denote by $\mathcal{N}_i = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\} \cup \{i\}$ the set of neighbors of agent i (including the agent itself). We make the following standard Assumption 3 on the graph connectivity, which is necessary to achieve consensus over the network.

Assumption 3 (Connected network) *The graph \mathcal{G} is connected.*

To achieve consensus and ensure implementation over the communication network \mathcal{G} , standard conditions in the literature require the following constraints on the weights w_{ij} used in the DGT algorithm (3), where \mathcal{P}_K denotes the set of polynomials with degree no larger than K .

Assumption 4 *[On the gossip matrices] The matrix $\mathbf{W} = (w_{ij})_{i,j=1}^m$ satisfies the following:*
(a) $\mathbf{W} = P^K(\bar{\mathbf{W}})$, where $P_K \in \mathcal{P}_K$ with $P_K(1) = 1$, and $\bar{\mathbf{W}} \triangleq (\bar{w}_{ij})_{i,j=1}^m$ has a sparsity pattern compliant with \mathcal{G} , that is

i) $\bar{w}_{ii} > 0$, for all $i = 1, \dots, m$;

ii) $\bar{w}_{ij} > 0$, if $(i, j) \in \mathcal{E}$; and $\bar{w}_{ij} = 0$ otherwise.

Furthermore, $\bar{\mathbf{W}}$ is doubly stochastic, that is, $\mathbf{1}^\top \bar{\mathbf{W}} = \mathbf{1}^\top$ and $\bar{\mathbf{W}} \mathbf{1} = \mathbf{1}$.

(b) Let $\rho \triangleq \|\mathbf{W} - \mathbf{J}\|_2$; then, it holds $\|(\mathbf{W} - \mathbf{J})^t\|_\infty \leq c_m \rho^t$, for all $t = 1, 2, \dots$, and some $c_m \geq 1$.

Remark 5 *Assumption 3 and Assumption 4(a) imply $\rho < 1$ and $\|(\mathbf{W} - \mathbf{J})^t\|_2 = \rho^t$. Using the norm bound $\|(\mathbf{W} - \mathbf{J})^t\|_\infty \leq \sqrt{m} \|(\mathbf{W} - \mathbf{J})^t\|_2$, we infer that Assumption 4(b) holds with $c_m \leq \sqrt{m}$.*

Assumption 4(a) requires $\bar{\mathbf{W}}$ being a consensus-forcing matrix with sparsity pattern matching the topology of communication network graph \mathcal{G} . When implementing the DGT algorithm (3) with $w_{ij} = \bar{w}_{ij}$, step (3a) can be executed via one-hop communication with each agent's immediate neighbors. Several rules of choosing $\bar{\mathbf{W}}$ have been proposed in the literature satisfying Assumption 4, such as the Laplacian, the Metropolis-Hasting, and the maximum-degree weights rules; see, e.g., (Nedić et al., 2018) and references therein. When the connectivity of \mathcal{G} is relatively weak, it may be beneficial to run multiple communication rounds to accelerate information propagation. This is realized by letting $\mathbf{W} = P_K(\bar{\mathbf{W}})$. When $K > 1$, K rounds of communications per iteration t are employed. For instance, by letting $\mathbf{W} = \bar{\mathbf{W}}^K$, the connectivity parameter can be improved from $\bar{\rho} = \|\bar{\mathbf{W}} - \mathbf{J}\|_2$ to $\|\bar{\mathbf{W}}^K - \mathbf{J}\|_2 = \bar{\rho}^K$. Faster information mixing can be obtained using suitably designed polynomials $P_K(\bar{\mathbf{W}})$, such as Chebyshev (Wien, 2011; Scaman et al., 2017) or orthogonal (a.k.a. Jacobi) (Berthier et al., 2020) polynomials.

It is worth noting that the DGT algorithm contains as a special instance, the proximal gradient descent when the graph \mathcal{G} is a star-network. In this setting, by electing the first node as the master node (star center), the centralized PGD iterate

$$\boldsymbol{\theta}^{t+1} = \prod_{\|\boldsymbol{\theta}\|_1 \leq r} (\boldsymbol{\theta}^t - \gamma^{-1} \nabla \mathcal{L}(\boldsymbol{\theta}^t)) \quad (11)$$

can be naturally implemented on the star-network. Specifically, at each iteration t , the master node broadcasts $\boldsymbol{\theta}^t$. Then each client node j evaluates $\nabla \mathcal{L}_j(\boldsymbol{\theta}^t)$ and sends it to the master. The master node averages the gradients $\nabla \mathcal{L}(\boldsymbol{\theta}^t) = \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\boldsymbol{\theta}^t)$ based on which the optimization step (11) is computed. The new value $\boldsymbol{\theta}^{t+1}$ is broadcast back to the clients. It is not difficult to check that the above procedure can be encompassed by the DGT algorithm (3) using the mixing matrix $\mathbf{W} = \mathbf{W}'\mathbf{W}'^\top$, with $\mathbf{W}' = 1/\sqrt{m}[\mathbf{1}_m, \mathbf{0}_{m \times m-1}]$. Note that the resulting \mathbf{W} is the same as that of a complete graph, but implemented in a more communication efficient way.

We are ready to present the main convergence properties of DGT and their implications. Our statements consists of two parts. In Sec. 2.2, we establish conditions under which the network average optimization error $\frac{1}{m} \sum_{i=1}^m \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2$ shrinks linearly up to a tolerance $o(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2)$, and provide an explicit expression of the rate. In Sec. 2.3, we show that for the statistical model subsumed in (1) with random design matrix, this result holds with high probability. We also discuss the impact of the network topology on the algorithm's performance. The proofs of the theorems are deferred to Sec. 3.

2.2 Convergence under RSC/RSM

Under Assumption 1-4, define

$$\Delta_{\text{stat}} \triangleq \frac{\rho}{2L_\Sigma} \left(\frac{\ell_\Sigma}{\mu_\Sigma} \frac{5C_2^2 c_m^2}{(1-\rho)^2} \right) \tau_\ell \nu^2 + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2, \quad \text{with} \quad \nu \triangleq 2\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 + 2\sqrt{s}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|, \quad (12)$$

and

$$\lambda \triangleq \frac{1 - (2\kappa)^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - 2C_1 s\tau_g/L_\Sigma}, \quad \text{with} \quad \kappa \triangleq \frac{L_\Sigma}{\mu_\Sigma}, \quad (13)$$

where $C_1 > 0$, $C_2 > 1$ are universal constants.

The convergence rate of DGT is given by the following theorem.

Theorem 6 *Given the linear model (1) with $\boldsymbol{\theta}^*$ being s -sparse, consider the LASSO problem (2) over a network \mathcal{G} satisfying Assumption 3. Suppose that the global \mathcal{L} and local \mathcal{L}_i losses satisfy Assumption 1 and Assumption 2, respectively. Let $\{(\boldsymbol{\theta}_i^t)_{i=1}^m\}$ be the sequence generated by the DGT algorithm with step size γ chosen as*

$$\gamma = L_\Sigma + 4 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}, \quad (14)$$

and the weight matrix \mathbf{W} satisfying Assumption 4. Further, assume that

$$\mu_\Sigma > 36C_1 s(\tau_\mu + \tau_g) \quad (15)$$

and

$$\rho \leq \left\{ 2 \left(75c_m^2 C_2^2 \frac{\ell_\Sigma^2}{\mu_\Sigma^2} + \frac{\ell_\Sigma}{\mu_\Sigma} \cdot 6C_2^2 c_m^2 s\tau_\ell \right) \right\}^{-2} = \text{poly} \left(m, \frac{s\tau_\ell}{\mu_\Sigma}, \frac{\ell_\Sigma}{\mu_\Sigma} \right)^{-1}. \quad (16)$$

Then, for any optimal solution $\hat{\boldsymbol{\theta}}$ of the LASSO problem satisfying $\|\hat{\boldsymbol{\theta}}\|_1 = r$, we have

$$\frac{1}{m} \sum_{i=1}^m \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 \leq B \cdot \lambda^t + \frac{\Delta_{\text{stat}}}{1 - \lambda}, \quad \forall t = 1, 2, \dots, \quad (17)$$

for $\lambda \in (0, 1)$, where B is a constant depending on the initialization [c.f. (69)].

The residual error Δ_{stat} depends on the parameters of the RSC/RSM conditions, the network connectivity ρ , and the statistical error $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2$. The next corollary shows that these parameters can be chosen such that $\Delta_{\text{stat}} = o(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2)$.

Corollary 7 *Instate assumptions of Theorem 6; and suppose $r \leq \|\boldsymbol{\theta}^*\|_1$ and*

$$s(\tau_\mu + \tau_g) = o(1) \quad \text{and} \quad s \cdot \rho c_m^2 \frac{\ell_\Sigma}{\mu_\Sigma} \tau_\ell = o(1). \quad (18)$$

Then, we have

$$\frac{1}{m} \sum_{i=1}^m \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 \leq \lambda^t \cdot B + o(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2). \quad (19)$$

• **Iteration complexity.** The contraction factor λ determining the linear decay of the optimization error depends on the restricted condition number κ and the RSC/RSM parameters. The lack of strong convexity and smoothness deteriorates the convergence rate by introducing the terms related to the tolerance parameters τ_μ, τ_g . Nevertheless, it can be verified that under condition (15), $\lambda \leq 1 - (4\kappa)^{-1}$, and thus such degradation is limited. This matches the convergence rate of the centralized PGD up to a constant factor and improves on existing analyses of distributed algorithms whose convergence to a solution of (2) is certified only at *sublinear* rate (see Sec. 1.4).

When special settings are considered, the convergence rate reduces to well-known expressions. In particular, for strongly convex and smooth losses, i.e., $\tau_\mu = \tau_g = \tau_\ell = 0$, we have $\Delta_{\text{stat}} = 0$, implying that DGT converges to an ε -solution of (2) in $O(\kappa \log(1/\varepsilon))$ iterations. This recovers the result in (Sun et al., 2022c). When the network is fully connected, the first term in Δ_{stat} vanishes (as $\rho = 0$), implying that DGT converges linearly to a neighborhood of $\hat{\boldsymbol{\theta}}$ of size $O((\tau_\nu + \tau_g)\nu^2/L_\Sigma)$. This rate matches (up to constant factors) that of the centralized PGD under RSC/RSM established in (Agarwal et al., 2012).

• **Interplay between optimization and communication.** Convergence of DGT is established, in particular, under condition (16). This reveals an interesting interplay between communication cost and hardness of the optimization problem. The quantity ℓ_Σ/μ_Σ can be viewed as the restricted condition number of the local losses \mathcal{L}_i (cf. Assumption 2); larger ℓ_Σ/μ_Σ and $s\tau_\ell/\mu_\Sigma$ correspond to more ill conditioned functions \mathcal{L}_i . Therefore, the more ill conditioned \mathcal{L}_i , the more connected the network must be to achieve convergence the rate of centralized PGD. When the network topology is given, the associated ρ might not satisfy (16). If so, one can still enforce the desired threshold on ρ by employing multiple rounds of communications. This will be discussed in detail in the next subsection.

2.3 Guarantees for sparse linear regression under random designs

We present some implications of Theorem 6 and Corollary 7 for the model (1), under the following assumption on the design matrix and noise.

Assumption 8 *Consider the linear model (1). The design matrices $\{\mathbf{X}_i\}_{i=1}^m$ are i.i.d. random matrices drawn from the Σ -Gaussian ensemble; and the elements of $\{\mathbf{n}_i\}_{i=1}^m$ are i.i.d. σ^2 -sub-Gaussian random variables, independent of the \mathbf{X}_i 's. Define*

$$\mu_\Sigma \triangleq \frac{1}{2}\sigma_{\min}(\Sigma), \quad L_\Sigma \triangleq 2\sigma_{\max}(\Sigma), \quad \text{and} \quad \zeta \triangleq \max_j \Sigma_{jj}.$$

Theorem 9 *Given the linear model (1), with $\boldsymbol{\theta}^*$ being s -sparse, consider the LASSO problem (2) over a network \mathcal{G} satisfying Assumption 3, and design matrices $\{\mathbf{X}_i\}_{i=1}^m$ and the measurement noise satisfying Assumption 8. Suppose*

$$\|\hat{\boldsymbol{\theta}}\|_1 = r \leq \|\boldsymbol{\theta}^*\|_1, \quad s \log d/N < c_5 \cdot \frac{\mu_\Sigma}{\zeta}, \quad \text{and} \quad \rho < (c_6 m^8 \kappa^4)^{-1} = \text{poly}(m, \kappa)^{-1}, \quad (20)$$

for some constants $c_5, c_6 > 0$. Let $\{(\boldsymbol{\theta}_i^t)_{i=1}^m\}$ be the sequence generated by the DGT algorithm with step size γ chosen according to (14) and the weight matrix \mathbf{W} satisfying Assumption 4. Then, the following holds:

$$\frac{1}{m} \sum_{i=1}^m \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 \leq B \cdot \left(\frac{1 - (2\kappa)^{-1} + C' s \log d/N}{1 - C' s \log d/N} \right)^t + c_7 \left(\frac{\zeta}{L_\Sigma} \frac{s \log d}{N} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2, \quad (21)$$

with probability at least $1 - c_8 \exp(-c_9 \log d)$, for some constants $C', c_7, c_8, c_9 > 0$. The expression of $B > 0$ is given in (109).

Theorem 9 reveals several interesting properties of DGT, as discussed next.

- **Scalability with respect to the problem dimension.** For a fixed network with connectivity ρ satisfying (20), the dependency of the convergence rate on the ambient dimension d , the total sample size N , and sparsity level s is only through the ratio $s \log d/N$. This implies that the convergence rate (21) is invariant under the asymptotics $s, N, d \rightarrow \infty$ and $s \log d/N = C$ (C is an universal constant)—resulting in the global sample scaling $N = O(s \log d)$.

- **Near optimal sample complexity.** When $s \log d/N = o(1)$, the residual error in (21) is of smaller order than the statistical precision $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2$. Therefore, DGT takes $O(\kappa \log(1/\varepsilon))$ iterations to reach an ε -neighborhood of a statistically optimal solution; this rate is of the same order of the centralized PGD under RSC/RSM (Agarwal et al., 2012). Note also that the condition $s \log d/N = o(1)$ (almost) matches the optimal sample complexity $N = \Omega(s \log(d/s))$ for model (1) (Wainwright, 2019), which is necessary for any centralized method to consistently estimate the d -dimensional, s -sparse $\boldsymbol{\theta}^*$ from N samples.

Since $N = n \cdot m$, the condition $s \log d/N = o(1)$ reveals an interesting interplay between statistical error, convergence rate, and network connectivity/communication cost, which is peculiar of the distributed setting: when agents do not have enough local samples n for statistical consistency but the overall sample size N across the network suffices—a situation that happens, e.g., when m is large (large-scale network)—DGT still achieves centralized

statistical accuracy at linear rate, provided that the network is “sufficiently” connected, i.e., $\rho \leq \text{poly}(m, \kappa)^{-1}$. Therefore, the insufficiency of the local sample size is compensated by higher communication costs. The impact of network parameters and topology on the convergence of DGT is further elaborated next.

• **Logarithmic scalability of the communication cost with the network size.** When the graph \mathcal{G} is not part of the design but given a-priori along with $\bar{\mathbf{W}}$ (satisfying Assumption 4), the condition $\rho \leq \text{poly}(m, \kappa)^{-1}$ can be satisfied by running multiple rounds of consensus steps. More precisely, since $\bar{\rho} = \|\bar{\mathbf{W}} - \mathbf{J}\| < 1$ (Assumption 4), one can construct $\mathbf{W} = \bar{\mathbf{W}}^K$ such that $\rho = \bar{\rho}^K \leq \text{poly}(m, \kappa)^{-1}$. This corresponds to running K consensus steps over the graph \mathcal{G} using each time the matrix $\bar{\mathbf{W}}$ in (3a). Note that for fixed m (graph), such K is always finite. In fact, one can see that taking $K = O(\log(m\kappa) \cdot (1 - \rho)^{-1})$ fulfills the requirement. If Chebyshev acceleration is used to employ multiple rounds of communications, coupled with a symmetric matrix $\bar{\mathbf{W}}$ (satisfying Assumption (4)), one can show that the number of communication steps K reduces to $O(\log(m\kappa) \cdot (1 - \rho)^{-1/2})$. Combining this fact with the iteration complexity (21), one can conclude that, in the setting of Theorem 9, DGT based on $\mathbf{W} = \bar{\mathbf{W}}^K$ drives $\frac{1}{m} \sum_{i=1}^m \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2$ within an ε -neighborhood of a statistical optimal solution in the following number of communication steps

$$O\left(\frac{\kappa}{1 - \rho} \log(m\kappa) \log(1/\varepsilon)\right). \quad (22)$$

This approach significantly improves the communication complexity compared to DGD-like algorithms, such as those discussed in (Ji et al., 2023a,b). Unlike (7) and (8), (22) does not depend on the ambient dimension d ; therefore, DGT avoids the speed-accuracy trade-off inherent in DGD algorithms. This marks the first instance of a distributed algorithm demonstrating such a desirable property in high-dimensional settings.

• **Impact of the network topology.** The network topology affects the convergence rate of DGT through the quantities m and ρ . A larger m or ρ increases communication complexity, as indicated by (22). In addition, for networks of the same size m , the efficiency of information propagation depends on the graph’s topology, which in (22) is captured by the dependency of ρ on m (Nedić et al., 2018). Table 1 provides, for some commonly used network topologies (see Fig. 2 for some examples), the dependency of ρ on m , when the lazy Metropolis rule is adopted; for other rules, see (Charron-Bost, 2020). While $(1 - \rho)^{-1}$ affects the total number of communications, it does not capture the overall cost of communications, which also depends on the edge-density of the graph. Hence, counting as one channel use per communication, with each edge shared between two nodes, Table 1 reports also the following quantities, capturing different communications costs: the total channel use per communication round, the channel use of the busiest node, and the total channel uses to reach an ε neighborhood of a statistically optimal solution.

From the first row of the table and (22), we infer that DGT running over a graph with stronger connectivity (smaller ρ) requires less communication rounds to converge. On the other hand, as a trade-off, better connectivity generally results from denser graphs, and thus yields higher communications costs (heavier channel uses), as reported in the table. We deduce that DGT running on the Erdős-Rényi graph with edge connecting probability $p = \log m/m$ achieves the lowest communication cost at the busiest node, and ties with

the implementation on the star (“push-pull”) in the amount of total channel use². For the path graph and 2-d grid graph, even though the channel use per communication round is low, DGT takes more iterations to converge, and overall a higher total communication cost. Comparing the two instances of the Erdős-Rényi graphs (the last two columns in the table), one may prefer a sparser graph at the price of more communication rounds to a denser graph and less communication rounds. Finally, even though the Erdős-Rényi is relatively efficient for DGT, in practice one may not be able to construct such a topology, e.g., due to geographic constraints. In such scenarios, a path/grid graph is still a valuable choice.

3. Proof of Main Results

3.1 Preliminaries

We begin rewriting the iterates (3a)-(3b) in matrix form and introducing some notation, which is convenient for our developments. Given $(\boldsymbol{\theta}_i^t)_{i=1}^m$, $(\mathbf{g}_i^t)_{i=1}^m$, and $(\nabla \mathcal{L}(\boldsymbol{\theta}_i^t))_{i=1}^m$, we introduce the following associated matrices:

$$\begin{aligned} \boldsymbol{\Theta}^t &\triangleq [\boldsymbol{\theta}_1^t, \dots, \boldsymbol{\theta}_m^t]^\top, \quad \Delta \boldsymbol{\Theta}^t = [\Delta \boldsymbol{\theta}_1^t, \dots, \Delta \boldsymbol{\theta}_m^t]^\top \triangleq \boldsymbol{\Theta}^{t+\frac{1}{2}} - \boldsymbol{\Theta}^t, \\ \mathbf{G}^t &\triangleq [\mathbf{g}_1^t, \dots, \mathbf{g}_m^t]^\top, \quad \text{and} \quad \nabla \mathcal{L}(\boldsymbol{\Theta}^t) \triangleq [\nabla \mathcal{L}_1(\boldsymbol{\theta}_1^t), \dots, \nabla \mathcal{L}_m(\boldsymbol{\theta}_m^t)]^\top. \end{aligned} \quad (23)$$

Furthermore, denote $\widehat{\boldsymbol{\Theta}} = \underbrace{[\widehat{\boldsymbol{\theta}}, \dots, \widehat{\boldsymbol{\theta}}]^\top}_{m \text{ times}}$. Using (23), (3a) can be rewritten in compact form:

$$\boldsymbol{\Theta}^t = \mathbf{W} (\boldsymbol{\Theta}^{t-1} + \Delta \boldsymbol{\Theta}^{t-1}), \quad (24a)$$

$$\mathbf{G}^t = \mathbf{W} (\mathbf{G}^{t-1} + \nabla \mathcal{L}(\boldsymbol{\Theta}^t) - \nabla \mathcal{L}(\boldsymbol{\Theta}^{t-1})). \quad (24b)$$

Introducing the average quantities

$$\bar{\boldsymbol{\theta}}^t \triangleq \frac{1}{m} \sum_{i=1}^m \boldsymbol{\theta}_i^t \quad \text{and} \quad \bar{\mathbf{g}}^t \triangleq \frac{1}{m} \sum_{i=1}^m \mathbf{g}_i^t, \quad (25)$$

we define the consensus errors on the $\boldsymbol{\theta}$, \mathbf{g} -vectors as

$$\boldsymbol{\theta}_{i,\perp}^t \triangleq \boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}^t \quad \text{and} \quad \mathbf{g}_{i,\perp}^t \triangleq \mathbf{g}_i^t - \bar{\mathbf{g}}^t, \quad i = 1, \dots, m. \quad (26)$$

Note that in our setting, $\boldsymbol{\theta}_{i,\perp}^0 = 0$, for all $i = 1, \dots, m$. The matrix counterparts of (26) are

$$\boldsymbol{\Theta}_\perp^t \triangleq [\boldsymbol{\theta}_{1,\perp}^t, \dots, \boldsymbol{\theta}_{m,\perp}^t]^\top = (\mathbf{I} - \mathbf{J}) \boldsymbol{\Theta}^t \quad \text{and} \quad \mathbf{G}_\perp^t \triangleq [\mathbf{g}_{1,\perp}^t, \dots, \mathbf{g}_{m,\perp}^t]^\top = (\mathbf{I} - \mathbf{J}) \mathbf{G}^t, \quad (27)$$

where $\mathbf{J} \triangleq \mathbf{1}\mathbf{1}^\top/m$.

Using (24) and the fact that $(\mathbf{I} - \mathbf{J})\mathbf{W} = \mathbf{W} - \mathbf{J}$, the dynamic of (27) reads

$$\boldsymbol{\Theta}_\perp^t = (\mathbf{W} - \mathbf{J})(\boldsymbol{\Theta}_\perp^{t-1} + \Delta \boldsymbol{\Theta}^{t-1}) \quad (28a)$$

$$\mathbf{G}_\perp^t = (\mathbf{W} - \mathbf{J})(\mathbf{G}_\perp^{t-1} + \nabla \mathcal{L}(\boldsymbol{\Theta}^t) - \nabla \mathcal{L}(\boldsymbol{\Theta}^{t-1})). \quad (28b)$$

The ℓ_1 -norm constraint plays a key role in enforcing sparsity of feasible points. Specifically, for any $\boldsymbol{\theta} \in \mathbb{R}^d$, with $\|\boldsymbol{\theta}\|_1 \leq r$, $\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}$ is approximately s -sparse, as quantified next.

2. We note that the performance on “push-pull” (star) are better than those on the Erdős-Rényi when the log factor in \tilde{O} is explicitly accounted.

Lemma 10 (Agarwal et al. (2012)) *Let $\hat{\boldsymbol{\theta}}$ be any optimal solution of Problem (2) such that $\|\hat{\boldsymbol{\theta}}\|_1 = r$. Then for any $\|\boldsymbol{\theta}\|_1 \leq r$, there holds*

$$\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_1 \leq 2\sqrt{s}\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\| + \underbrace{2\|\Delta^*\|_1 + 2\sqrt{s}\|\Delta^*\|}_{\nu}, \quad (29)$$

where $\Delta^* \triangleq \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*$ is the statistical error.

3.2 Analysis of the error dynamics

Due to the double stochasticity of \mathbf{W} , the iterates $\boldsymbol{\theta}_i^{t+\frac{1}{2}}$ and $\boldsymbol{\theta}_i^t$ generated by DGT must be feasible. As a result of Lemma 10, the directions $\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}$ and $\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}$ satisfy the approximate sparsity property defined by (29). This adherence is a fundamental property leveraged in our proof to establish the desired convergence results, marking a notable shift from traditional proof techniques found in the literature on decentralized optimization. Based on this key property, we structure our analysis into the following steps.

- (i) **Inexact descent property (Sec. 3.2.1):** We demonstrate that the optimization residual $\|\boldsymbol{\theta}^t - \hat{\boldsymbol{\theta}}\|^2$ decreases linearly, up to a tracking error term arising from the discrepancy between $\nabla\mathcal{L}(\boldsymbol{\theta}_i^t)$ and its estimator \mathbf{g}_i^t , and a tolerance term related to ν due to the approximate sparsity of the direction $\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}$.
- (ii) **Bounding the tracking error (Sec. 3.2.2):** We establish that the tracking error can be effectively bounded by an average of all historical optimization residuals, weighted by the network connectivity ρ and a ν -dependent tolerance term, akin to the one described in (i). Aiming at exploiting the local restricted smoothness of \mathcal{L}_i and the approximate sparsity of the direction $\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}$, our approach introduces a novel line of analysis that diverges from traditional proof techniques used to study convergence of gradient-tracking algorithms.
- (iii) **Combining error dynamics via the z -transform (Sec. 3.2.3):** We combine the two error dynamics in (i) and (ii) employing the z -transform for finite-length sequences, to effectively deal with *historical* optimization residuals, and obtain a sufficient condition for the algorithm's convergence. Our method deviates from conventional approaches commonly found in the literature on gradient-tracking methods, such as (Xu et al., 2018; Sun et al., 2022c; Nedić et al., 2016), which often rely on the small gain theorem to analyze dynamics with only single-hop time dependencies. By employing the z -transform, we broaden the analytical framework to incorporate more complex time dependencies.

3.2.1 INEXACT DESCENT

We begin with the analysis of the optimization step (3b). Since each \mathbf{g}_i aims to estimate the local gradient $\nabla\mathcal{L}(\boldsymbol{\theta}_i)$, step (3b) can be regarded as an inexact version of the proximal gradient update:

$$\boldsymbol{\theta}_i^{t+\frac{1}{2}} = \underset{\|\boldsymbol{\theta}\|_1 \leq r}{\operatorname{argmin}} \left\{ \nabla\mathcal{L}(\boldsymbol{\theta}_i^t)^\top (\boldsymbol{\theta}_i - \boldsymbol{\theta}_i^t) + \frac{\gamma}{2} \|\boldsymbol{\theta}_i - \boldsymbol{\theta}_i^t\|^2 \right\}. \quad (30)$$

Therefore, we base our proof on the analysis of (30) while taking into account the gradient estimation error $\nabla\mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t$, due to the use in (30) of the gradient tracking vector \mathbf{g}_i^t rather than $\nabla\mathcal{L}(\boldsymbol{\theta}_i^t)$. The following proposition establishes the one-step descent of the optimization error, up to additive errors.

Proposition 11 *Consider the LASSO problem (2) over the network \mathcal{G} , under Assumptions 3 and 1; and $\hat{\boldsymbol{\theta}}$ such that $\|\hat{\boldsymbol{\theta}}\| = r$. Let $\{\boldsymbol{\theta}_i^t\}$ be the sequence generated by Algorithm 3 under Assumption 4. Then, the following holds:*

$$\begin{aligned} \left(1 - \frac{C_1 s \tau_g}{\gamma}\right) \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \hat{\boldsymbol{\Theta}}\|_F^2 &\leq \left(1 - \frac{\mu_\Sigma}{\gamma} + \frac{C_1 s(\tau_\mu + \tau_g)}{\gamma}\right) \|\boldsymbol{\Theta}^t - \hat{\boldsymbol{\Theta}}\|_F^2 \\ &\quad - \left(1 - \frac{L_\Sigma}{\gamma}\right) \|\Delta\boldsymbol{\Theta}^t\|_F^2 + \underbrace{\frac{2}{\gamma} \sum_{i=1}^m (\nabla\mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}})}_{\triangleq \delta^t} + \frac{C_1(\tau_\mu + \tau_g)}{\gamma} m \nu^2, \end{aligned} \quad (31)$$

for some universal constant $C_1 > 0$.

Proof See Appendix A. ■

Note that in a fully connected network, $\nabla\mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t = \mathbf{0}$ and $\boldsymbol{\theta}_i^t = \boldsymbol{\theta}_j^t$, for all $i, j = 1, \dots, m$, and $t = 1, \dots$. Setting $\gamma = L_\Sigma$, (31) recovers the convergence result of the centralized PGD algorithm, as in (Agarwal et al., 2012, Thm. 1).

The subsequent analysis focuses on the dynamics of the gradient tracking error δ^t .

3.2.2 BOUNDING THE TRACKING ERROR IN HIGH DIMENSION

To study the evolution of the tracking error δ^t , we resort to the dynamics of \mathbf{g}_i^t . Recalling the definition of $\bar{\mathbf{g}}^t$ [cf. (25)] and the tracking dynamics (3a), we have

$$\bar{\mathbf{g}}^t = \bar{\mathbf{g}}^{t-1} + \frac{1}{m} \sum_{i=1}^m (\nabla\mathcal{L}_i(\boldsymbol{\theta}_i^t) - \nabla\mathcal{L}_i(\boldsymbol{\theta}_i^{t-1})) \quad \text{and hence} \quad \bar{\mathbf{g}}^t = \frac{1}{m} \sum_{i=1}^m \nabla\mathcal{L}_i(\boldsymbol{\theta}_i^t),$$

where we used the initialization $\bar{\mathbf{g}}^0 = \frac{1}{m} \sum_{i=1}^m (\nabla\mathcal{L}_i(\boldsymbol{\theta}_i^0))$.

This average preserving property of the sum-gradient suggests the following decomposition of the tracking error δ^t :

$$\begin{aligned} \delta^t &= \sum_{i=1}^m \left(\nabla\mathcal{L}(\boldsymbol{\theta}_i^t) - \frac{1}{m} \sum_{j=1}^m \nabla\mathcal{L}_j(\boldsymbol{\theta}_j^t) + \bar{\mathbf{g}}^t - \mathbf{g}_i^t \right)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) \\ &= \sum_{i=1}^m \left(\frac{1}{m} \sum_{j=1}^m \nabla\mathcal{L}_j(\boldsymbol{\theta}_i^t) - \frac{1}{m} \sum_{j=1}^m \nabla\mathcal{L}_j(\boldsymbol{\theta}_j^t) \right)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) + \sum_{i=1}^m (\bar{\mathbf{g}}^t - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}). \end{aligned} \quad (32)$$

Limitation of current proof techniques: The subsequent analysis, aiming at bounding the gradient-tracking error δ^t , necessitates a departure from traditional techniques developed in the literature of gradient-tracking algorithms, such as (Xu et al., 2018; Sun et al., 2022c;

Nedić et al., 2016). Traditional approaches often handle the error δ^t in (32) using the Cauchy-Schwarz inequality, which yields the following type of bound:

$$\delta^t \leq C' \cdot \ell_{\max} \cdot \|\Theta_{\perp}^t\|_F^2 + C'' \|\Theta^{t+\frac{1}{2}} - \hat{\Theta}\|_F^2 + C''' \|\mathbf{G}_{\perp}^t\|_F^2, \quad (33)$$

where $\ell_{\max} = \max_i \|(1/n)\mathbf{X}_i^{\top} \mathbf{X}_i\|_2$ is the largest Lipschitz constants of the local gradients $\nabla \mathcal{L}_i$. The analysis would proceed absorbing the second term $\|\Theta^{t+\frac{1}{2}} - \hat{\Theta}\|_F^2$ in (33) into the LHS of (31), and using the dynamics (28a)-(28b) along with the Lipschitz continuity of $\nabla \mathcal{L}_i$'s to obtain the following bounds on the remaining consensus and tracking errors $\|\Theta_{\perp}^t\|_F^2$ and $\|\mathbf{G}_{\perp}^t\|_F^2$ in (33) (Sun et al., 2022c):

$$\|\Theta_{\perp}^t\|_F \leq \rho \|\Theta_{\perp}^{t-1}\|_F + \rho \|\Delta \Theta^{t-1}\|_F, \quad (34a)$$

$$\|\mathbf{G}_{\perp}^t\|_F \leq \rho \|\mathbf{G}_{\perp}^{t-1}\|_F + \rho \ell_{\max} \|\Theta_{\perp}^{t-1}\|_F + \rho \ell_{\max} \|\Delta \Theta^{t-1}\|_F. \quad (34b)$$

The rest of the proof in (Sun et al., 2022c) would chain the inequalities (31), (33), and (34) using small-gain arguments, yielding linear convergence rates, under a suitable choice of the stepsize γ (and the network connectivity ρ), *provided that strong convexity of \mathcal{L} is assumed*—see (Sun et al., 2022c, Propositions 3.6 & 3.8).

This line of analysis fails in high dimensional scenarios, for the following reasons. **(i)** The centralized objective functions \mathcal{L} is RSC and not strongly convex globally, and the local objective functions \mathcal{L}_i are only convex. **(ii)** In the typical asymptotics of high-dimensional problems— $d/N \rightarrow \infty$ and $s \log d/N = O(1)$ —the Lipschitz constants ℓ_{\max} of $\nabla \mathcal{L}_i$'s scale as $\mathcal{O}(d)$, for various predictor x'_i models (Wainwright, 2019), thereby diverging as $d \uparrow \infty$. Even postulating to be able to relax the strong convexity requirement of \mathcal{L} to the RSC, establishing linear convergence along the path of the classical analysis will require the network connectivity ρ (or the step-size γ) to counteract the unfavorable $1/d$ scaling of ℓ_{\max} , ensuring $\rho \ell_{\max} = \mathcal{O}(1)$, if optimal centralized statistical precision is to be achieved. This would translate in the unsatisfactory scaling of the total number of communications as $\log d$ (or d). Notice that, this issue remains even under more restrictive assumptions on the sample size, such as $s \log d/n = \mathcal{O}(1)$. At the high level, this arises because the gradients $\nabla \mathcal{L}_i(\theta_i)$'s remain non-sparse vectors even when θ_i 's are all sparse.

A new line of analysis: The challenges outlined above underscore the need of a new mechanism to control the gradient error δ^t in high-dimensional decentralized settings. The key observation is that the undesirable scaling of the tracking error δ^t with the ambient dimension d , steaming from ℓ_{\max} , is unavoidable if one pursues the statistically agnostic bound of $\|\mathbf{G}_{\perp}^t\|_F$ as given in (34b), which *globally* relates the gradient tracking error $\|\mathbf{G}_{\perp}^t\|_F$ to the consensus error. Our proposed approach capitalizes on the observation that δ^t primarily depends on the component of the gradient error projected along the directions $\theta_i^{t+\frac{1}{2}} - \hat{\theta}$, which are *approximately sparse* vectors. This contrasts with (34b), which considers the entirety of the gradient tracking error. More precisely, our proof leverages the fact that, for any set of row-wise feasible points $\Theta_1, \Theta_2, \Theta_3$, and Θ_4 , the inner products

$$(\Theta_1 - \Theta_2)^{\top} (\nabla \mathcal{L}_j(\Theta_3) - \nabla \mathcal{L}_j(\Theta_4)), \forall j \in [m] \quad (35a)$$

$$(\Theta_1 - \Theta_2)^{\top} (\nabla \mathcal{L}(\Theta_3) - \nabla \mathcal{L}(\Theta_4)) \quad (35b)$$

can be controlled by quantities that do scale logarithmically with the ambient dimension at most. Therefore, in place of breaking the inner products using the Cauchy-Schwartz inequality, yielding to one-hop bounds as (34b), our approach necessitates working directly with inner product relationships. The final result is a bound of δ^t that depends on the entire algorithm trajectory from the initialization. This is formalized in the next proposition.

Proposition 12 *Under the setting of Proposition 11 and the local restricted smoothness Assumption 2, the gradient tracking error δ^t ($t = 1, \dots$) can be bounded as*

$$\begin{aligned} \delta^t \leq & \frac{C_2}{2} \left\{ \frac{\rho \cdot c_m \epsilon}{1 - \rho} (\ell_\Sigma + s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 + \frac{c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta_\perp^s\|_F^2 \right. \\ & + c_m s \tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m s \tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s-\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\ & \left. + m c_m \rho^t (A_3 \cdot \epsilon^{-1} + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) + \frac{\rho \cdot m c_m}{1 - \rho} \tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) \right\} \end{aligned} \quad (36)$$

for some $C_2 > 0$; where $\epsilon > 0$, $\Theta^{-\frac{1}{2}} \triangleq \Theta^0$ and

$$A_3 \triangleq \ell_\Sigma \|\theta^*\|^2 + \tau_\ell \|\theta^*\|_1^2, \quad c_g \triangleq \|\nabla \mathcal{L}_j(\theta^*)\|_\infty + \|\nabla \mathcal{L}(\theta^*)\|_\infty. \quad (37)$$

Proof See Appendix B. ■

Proposition 12 reveals that the gradient tracking error δ^t can be bounded by the discounted cumulative sum of historical optimization errors $\{\|\Theta^{s-1/2} - \widehat{\Theta}\|\}_{s=0}^{t+1}$, historical step-length $\{\|\Delta \Theta^s\|\}_{s=0}^{t-1}$, and consensus errors $\{\|\Theta_\perp^s\|\}_{s=0}^{t-1}$. Furthermore, one can check that δ^t decreases if ρ , m , or the right-hand-side tolerance parameter τ_ℓ decrease, which is a consequence of a more connected network, and larger local sample size n , respectively.

We can now combine the statements obtained in Propositions 11 and 12, and write

$$\begin{aligned} & \left(1 - \frac{1}{\gamma} \left(C_1 s \tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_\Sigma + s\tau_\ell) \right) \right) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\ \leq & \left(1 - \frac{\mu_\Sigma}{\gamma} + \frac{C_1 s (\tau_\mu + \tau_g)}{\gamma} \right) \|\Theta^t - \widehat{\Theta}\|_F^2 - \left(1 - \frac{L_\Sigma}{\gamma} \right) \|\Delta \Theta^t\|_F^2 + \frac{C_1 (\tau_\mu + \tau_g)}{\gamma} m \nu^2 \\ & + \frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 + \frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta_\perp^s\|_F^2 \\ & + \frac{C_2}{\gamma} c_m s \tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{C_2}{\gamma} c_m s \tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s-\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\ & + \frac{C_2}{\gamma} m c_m \rho^t (A_3 \cdot \epsilon^{-1} + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) + \frac{C_2}{\gamma} \frac{\rho \cdot m c_m}{1 - \rho} \tau_\ell \nu^2 (\epsilon + \epsilon^{-1}). \end{aligned} \quad (38)$$

Our last step bounds the consensus error $\|\Theta_\perp^t\|$ using $\{\|\Delta \Theta^s\|\}_{s=0}^{t-1}$, as given next. The bound deviates from (34a), as it is a function of the entirety sequence $\{\|\Delta \Theta^s\|\}_{s \geq 0}$. This will allow us to exploit the term $-\|\Delta \Theta^t\|^2$ in (38) to control consensus errors.

Proposition 13 *The consensus error Θ_{\perp}^t can be bounded as*

$$\|\Theta_{\perp}^t\|_F^2 \leq 2\rho^{2t} \underbrace{\|\Theta_{\perp}^0\|_F^2}_{=0} + \frac{2\rho}{1-\rho} \cdot \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta\Theta^s\|_F^2. \quad (39)$$

Proof The proof is standard and is provided in Appendix D for completeness. \blacksquare

3.2.3 COMBINING THE ERROR DYNAMICS

To establish linear convergence of DGT from (38)-(39), we will employ the z -transform for finite length sequences to deal with the convolution-sum terms. Notice that small gain arguments typically used in (Nedić et al., 2016; Sun et al., 2022c) are not directly applicable to chain the inequalities (38)-(39) in our scenario. This is because, unlike (Nedić et al., 2016; Sun et al., 2022c), each sequence in the above inequalities is influenced by the complete historical dependencies of the others, rather than a simpler single-hop dependence. This motivates the introduction of the z -transform to handle an arbitrarily long memory.

We begin briefly reviewing the relevant properties of the z -transform and then apply it to the system (38)-(39).

• THE z -TRANSFORM AND PROPERTIES

Let $\{a(t)\}$ be a sequence of nonnegative numbers. Define the length K z -transform of $\{a(t)\}$ as

$$A^K(z) \triangleq \sum_{t=1}^K a(t)z^{-t}, \quad z \in \mathbb{R}_{++}. \quad (40)$$

Some properties of the z -transform instrumental to our developments are summarized next, whose proofs are provided in Appendix M.

Lemma 14 *Let $A^K(z)$ be the length K z -transform of a nonnegative sequence $\{a(t)\}$, and $\rho \in (0, 1)$. Then, the following hold:*

- (i) $\sum_{t=1}^K a(t+1)z^{-t} \geq zA^K(z) - a(1), \quad \forall z \in \mathbb{R}_{++};$
- (ii) $\sum_{t=1}^K \left(\sum_{s=0}^{t-1} \rho^{t-s} a(s) \right) z^{-t} \leq \frac{\rho}{z-\rho} \cdot (A^K(z) + a(0)), \quad \forall z \in (\rho, 1);$
- (iii) $\sum_{t=1}^K \left(\sum_{s=0}^{t-1} \rho^{t-s} a(s+1) \right) z^{-t} \leq z \cdot \frac{\rho}{z-\rho} A^K(z), \quad \forall z \in (\rho, 1).$

Lemma 15 *Suppose a nonnegative sequence $\{a(t)\}$ satisfies*

$$\sum_{t=1}^K a(t)z^{-t} \leq B + c \cdot \sum_{t=1}^K z^{-t}, \quad \text{for all } K \geq 1 \quad \text{and } z \in (\bar{z}, 1), \quad (41)$$

and some $B, c > 0$. Then

$$a(t) \leq B \cdot \bar{z}^t + \frac{c}{1-\bar{z}}. \quad (42)$$

When $B < +\infty$, $\{a(t)\}$ converges linearly at rate \bar{z} up to a constant.

• THE Z-TRANSFORM SYSTEM

Introduce the following quantities:

$$D_{\Theta}^K(z) = \sum_{t=1}^K z^{-t} \|\Theta^{t-\frac{1}{2}} - \hat{\Theta}\|^2, \quad \Delta_{\Theta}^K(z) = \sum_{t=1}^K z^{-t} \|\Delta\Theta^t\|^2, \quad \Theta_{\perp}^K(z) = \sum_{t=1}^K z^{-t} \|\Theta_{\perp}^t\|^2. \quad (43)$$

To apply (43) to (38) we first bound the term $\|\Theta^t - \hat{\Theta}\|^2$ on the RHS there as $\|\Theta^t - \hat{\Theta}\|^2 \leq \|\Theta^{t-\frac{1}{2}} - \hat{\Theta}\|^2$ under condition

$$\gamma \geq \mu_{\Sigma} - C_1 s(\tau_{\mu} + \tau_g). \quad (44)$$

Further assuming

$$\gamma \geq C_1 s\tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_{\Sigma} + s\tau_{\ell}), \quad (45)$$

we can multiply z^{-t} ($z > \rho$) to both sides of (38) and (39) while summing up from $t = 1$ to $t = K$ and applying Lemma 14 (recall that $\Delta\Theta^0 = \mathbf{0}$), leading to

$$\begin{aligned} & \left(1 - \frac{1}{\gamma} \left\{ C_1 s\tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_{\Sigma} + s\tau_{\ell}) + \frac{C_2 \rho \cdot c_m \epsilon^{-1} s\tau_{\ell}}{z - \rho} \right\}\right) z D_{\Theta}^K(z) \\ & - \left(1 - \frac{\mu_{\Sigma}}{\gamma} + \frac{C_1 s(\tau_{\mu} + \tau_g)}{\gamma} + \frac{C_2}{\gamma} c_m s\tau_{\ell} \epsilon^{-1} \frac{\rho}{z - \rho}\right) D_{\Theta}^K(z) \\ & \leq - \left(1 - \frac{L_{\Sigma}}{\gamma} - \frac{C_2}{\gamma} \frac{c_m \ell_{\Sigma}}{\epsilon} \frac{\rho}{z - \rho} - \frac{C_2}{\gamma} \frac{c_m \ell_{\Sigma}}{\epsilon} \frac{2\rho}{1 - \rho} \cdot \left(\frac{\rho}{z - \rho}\right)^2\right) \Delta_{\Theta}^K(z) \\ & + \bar{\Delta}_{\text{stat}} \cdot \sum_{t=1}^K z^{-t} + B(z), \end{aligned} \quad (46)$$

with

$$\begin{aligned} B(z) & \triangleq \frac{C_2}{\gamma} c_m s\tau_{\ell} \epsilon^{-1} \frac{\rho}{z - \rho} \|\Theta^0 - \hat{\Theta}\|^2 + \frac{C_2}{\gamma} m c_m (A_3 \cdot \epsilon^{-1} + s c_g^2 \ell_{\Sigma}^{-1} \epsilon^{-1} + c_g \nu) \frac{\rho}{z - \rho} \\ & + \left(1 - \frac{1}{\gamma} \left\{ C_1 s\tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_{\Sigma} + s\tau_{\ell}) \right\}\right) \|\Theta^{1/2} - \hat{\Theta}\|^2 \end{aligned} \quad (47)$$

and

$$\bar{\Delta}_{\text{stat}} = \frac{C_2}{\gamma} \frac{\rho \cdot c_m}{1 - \rho} \tau_{\ell} \nu^2 (\epsilon + \epsilon^{-1}) + \frac{C_1 (\tau_{\mu} + \tau_g)}{\gamma} \nu^2. \quad (48)$$

The following conditions are sufficient to invoke Lemma 15 and prove the convergence of sequence $\{(1/m) \|\Theta^{t+\frac{1}{2}} - \hat{\Theta}\|_F^2\}$:

- i) The step size condition (44) and (45) hold.

ii) There exists some $z \in (\rho, 1)$ such that

$$\begin{aligned} & \left[1 - \frac{1}{\gamma} \left(C_1 s \tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_\Sigma + s \tau_\ell) + \frac{C_2 \rho \cdot c_m \epsilon^{-1} s \tau_\ell}{z - \rho} \right) \right] z \\ & > 1 - \frac{\mu_\Sigma}{\gamma} + \frac{C_1 s (\tau_\mu + \tau_g)}{\gamma} + \frac{C_2}{\gamma} c_m s \tau_\ell \epsilon^{-1} \frac{\rho}{z - \rho}; \end{aligned} \quad (49a)$$

$$1 - \frac{L_\Sigma}{\gamma} - \frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \frac{\rho}{z - \rho} - \frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \frac{2\rho}{1 - \rho} \cdot \left(\frac{\rho}{z - \rho} \right)^2 > 0. \quad (49b)$$

3.3 Linear convergence under RSC/RSM (Proof of Theorem 6)

The proof is organized into the following steps:

- In Sec. 3.3.1 we provide choices of ϵ and γ for the existence of $z > \rho$ satisfying (49).
- Under such choices, we bound $B(z)$ [cf. (47)] and $\bar{\Delta}_{\text{stat}}$ [cf. (48)] in Sec. 3.3.2.
- In Sec. 3.3.3 we establish conditions for the linear convergence of $\left\{ \frac{1}{m} \|\Theta^{t+\frac{1}{2}} - \hat{\Theta}\|_F^2 \right\}_t$ up to a residual error depending on the statistical precision $\|\hat{\Theta} - \Theta^*\|^2$.

3.3.1 CHOICE OF ϵ AND γ

Lemma 16 and 17 below provide conditions and the range of $z > \rho$ for (49b) and (49a) to hold, respectively. For convenience we set ϵ as

$$\epsilon = \frac{\mu_\Sigma}{\ell_\Sigma} \frac{1 - \rho}{2C_2 c_m}. \quad (50)$$

Lemma 16 *Given $\gamma > L_\Sigma$, condition (49b) is satisfied for all $z > z_\rho$, with*

$$z_\rho \triangleq \rho + \frac{4\rho^2}{1 - \rho} \left(\sqrt{1 + (\gamma - L_\Sigma) \left(\frac{\mu_\Sigma}{\ell_\Sigma^2} \frac{2\rho}{C_2^2 c_m^2} \right)} - 1 \right)^{-1}. \quad (51)$$

Proof See Appendix E. ■

We focus now on condition (49a). Denote for simplicity

$$\begin{aligned} A & \triangleq \frac{1 - \frac{\mu_\Sigma}{\gamma} + \frac{C_1 s (\tau_\mu + \tau_g)}{\gamma}}{\frac{C_2}{\gamma} c_m s \tau_\ell \epsilon^{-1}} = \frac{\gamma - \mu_\Sigma + C_1 s (\tau_\mu + \tau_g)}{\frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2^2 c_m^2}{1 - \rho} s \tau_\ell}, \\ D & \triangleq \frac{1 - \frac{1}{\gamma} \left\{ C_1 s \tau_g + \frac{C_2 \rho \cdot c_m \epsilon}{1 - \rho} (\ell_\Sigma + s \tau_\ell) \right\}}{\frac{C_2}{\gamma} c_m s \tau_\ell \epsilon^{-1}} = \frac{\gamma - \left\{ C_1 s \tau_g + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} (\ell_\Sigma + s \tau_\ell) \right\}}{\frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2^2 c_m^2}{1 - \rho} s \tau_\ell}, \end{aligned} \quad (52)$$

and

$$\sigma \triangleq \frac{A}{D} = \frac{\gamma - \mu_\Sigma + C_1 s (\tau_\mu + \tau_g)}{\gamma - \left\{ C_1 s \tau_g + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} (\ell_\Sigma + s \tau_\ell) \right\}}. \quad (53)$$

Note that, under (44) and (45) strictly satisfied, $A, D > 0$. Then (49a) can be rewritten succinctly as

$$z \geq \frac{A + G_{\text{net}}(z)}{D - G_{\text{net}}(z)}, \quad G_{\text{net}}(z) \triangleq \frac{\rho}{z - \rho}. \quad (54)$$

Lemma 17 *Given γ satisfying (44) and (45), condition (49a) is satisfied for all $z > z_\sigma$, with*

$$z_\sigma \triangleq \sigma + \rho \cdot \left(1 + \frac{1}{D}\right), \quad (55)$$

and A and D defined in (52).

Proof See Appendix F. ■

Combining Lemma 16 and 17 we conclude that, for any given $\gamma > L_\Sigma$, under conditions (44) and (45), (49a)-(49b) are satisfied by the choices of z such that

$$z > \bar{z}, \quad \text{with} \quad \bar{z} \triangleq \max\{z_\rho, z_\sigma\}. \quad (56)$$

In the remaining of this section, we will simplify the expression of \bar{z} by making a specific choice of γ . We do not aim to find the optimal value of γ but a convenient feasible choice yielding an insightful expression of the rate. We set

$$\gamma = L_\Sigma + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1 - \rho)^4}, \quad (57)$$

where $C_3 > 0$ is some absolute constant (to be determined). Define

$$\sigma_0 \triangleq \frac{1 - \kappa^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - C_1 s\tau_g/L_\Sigma}, \quad \kappa \triangleq \frac{L_\Sigma}{\mu_\Sigma}. \quad (58)$$

Proposition 18 *Let γ be chosen according to (57). Under the following conditions:*

$$1 - 2C_1 s\tau_g/L_\Sigma > 0, \quad (59)$$

$$\mu_\Sigma > 4C_1 s(\tau_\mu + \tau_g) + \frac{\ell_\Sigma}{\mu_\Sigma} \frac{8C_2^2 c_m^2 \rho}{(1 - \rho)^2} s\tau_\ell + \rho \cdot \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell, \quad (60)$$

$$\text{and} \quad \rho \leq 1/2, \quad (61)$$

\bar{z} can be upperbounded as

$$\bar{z} \leq \max \left\{ \rho + \frac{\sqrt{2}C_2}{\sqrt{C_3}} \rho(1 - \rho) + \frac{2C_2^2}{C_3} \sqrt{\rho}(1 - \rho)^3, \right. \\ \left. \sigma_0 + \sqrt{\rho} \cdot \frac{18C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - 2C_1 s\tau_g/L_\Sigma} \right\}. \quad (62)$$

Proof See Appendix G. ■

To further simplify (62), we require the second argument in the max-expression in (62) to be no larger than

$$\frac{1 - (2\kappa)^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - 2C_1 s\tau_g/L_\Sigma}. \quad (63)$$

Corollary 19 *Let γ be chosen as*

$$\gamma = L_\Sigma + 4 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1 - \rho)^4}. \quad (64)$$

Under the following conditions:

$$\mu_\Sigma > 36C_1 s(\tau_\mu + \tau_g), \quad (65)$$

$$\rho \leq \left\{ 2 \left(72c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma^2} + \frac{\ell_\Sigma}{\mu_\Sigma^2} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2}{\ell_\Sigma} s\tau_\ell \right) \right\}^{-2}, \quad (66)$$

it holds

$$\bar{z} \leq \bar{z}_{up} \triangleq \max \left\{ \rho + C_2 \rho + \frac{C_2^2}{2} \sqrt{\rho}, \frac{1 - (2\kappa)^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - 2C_1 s\tau_g/L_\Sigma} \right\}. \quad (67)$$

Proof See Appendix H. ■

3.3.2 BOUNDING $B(z)$ [CF. (47)] AND Δ_{stat} [CF. (48)]

According to Eq. (94), under the setting of Proposition 18, any z satisfying $z \geq \bar{z}_{up}$ is lower bounded by

$$z \geq \rho + \sqrt{\rho}(1 - \rho)^3 \frac{C_2^2}{C_3}. \quad (68)$$

Substituting the expression of γ [cf. (64)] and ϵ [cf. (50)] into $B(z)$ [cf. (47)] and Δ_{stat} [cf. (48)], together with (68) we obtain the following bound.

Proposition 20 *Under the setting of Corollary 19 and assume $C_2 \geq 1$, then*

$$\frac{B(z)}{m} \leq \frac{C_4}{m} \left(\|\Theta^0 - \hat{\Theta}\|_F^2 + \|\Theta^*\|_F^2 + \|\Theta^{1/2} - \hat{\Theta}\|_F^2 + \frac{mc_m \sqrt{\rho}}{L_\Sigma} \left(\frac{c_m s c_g^2}{\mu_\Sigma} + c_g \nu \right) \right) \triangleq B, \quad (69)$$

$$\bar{\Delta}_{\text{stat}} \leq \frac{\rho}{2L_\Sigma} \left(\frac{\ell_\Sigma}{\mu_\Sigma} \frac{5C_2^2 c_m^2}{(1 - \rho)^2} \right) \tau_\ell \nu^2 + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2 \triangleq \Delta_{\text{stat}} \quad (70)$$

for some $C_4 > 0$, where c_g and ν are defined as

$$c_g = \max_{j \in [m]} \|\nabla \mathcal{L}_j(\theta^*)\|_\infty + \|\nabla \mathcal{L}(\theta^*)\|_\infty \quad \text{and} \quad \nu = 2\|\hat{\theta} - \theta^*\|_1 + 2\sqrt{s}\|\hat{\theta} - \theta^*\|. \quad (71)$$

Proof See Appendix I. ■

3.3.3 LINEAR CONVERGENCE UP TO $o(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2)$

We prove Theorem 6 and Corollary 7. Recall below the network connectivity condition (16) and rate expression λ (13) given in Theorem 6 for convenience:

$$\rho \leq \left\{ 2 \left(75c_m^2 C_2^2 \frac{\ell_\Sigma^2}{\mu_\Sigma^2} + \frac{\ell_\Sigma}{\mu_\Sigma^2} \cdot 6C_2^2 c_m^2 s \tau_\ell \right) \right\}^{-2} \quad \lambda = \frac{1 - (2\kappa)^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - 2C_1 s \tau_g/L_\Sigma}.$$

Based on Lemma 15 and the results in Sec. 3.3.1 & 3.3.2, it suffices to show (16) implies (66) and $\bar{z}_{\text{up}} \leq \lambda < 1$.

- (16) \implies (66): the claim holds since $C_2 > 1$, $\ell_\Sigma \geq \mu_\Sigma$, and $c_m \geq 1$.
- $\bar{z}_{\text{up}} \leq \lambda < 1$: it is not difficult to check that $\mu_\Sigma > 36C_1 s(\tau_\mu + \tau_g)$ implies $\lambda < 1 - (4\kappa)^{-1}$. Next we prove

$$\rho + C_2 \rho + \frac{C_2^2}{2} \sqrt{\rho} \leq \lambda,$$

and thus $\bar{z}_{\text{up}} = \lambda$. It is sufficient to check that the following chain of inequalities hold:

$$\begin{aligned} & \left(\rho + C_2 \rho + \frac{C_2^2}{2} \sqrt{\rho} \right) (1 - 2C_1 s \tau_g/L_\Sigma) \\ & \leq 3C_2^2 \sqrt{\rho} \stackrel{(16)}{\leq} \frac{1}{50} \cdot \frac{\mu_\Sigma^2}{c_m^2 \ell_\Sigma^2} \leq \frac{1}{2\kappa} \leq 1 - \frac{1}{2\kappa} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma, \end{aligned} \quad (72)$$

where we used $\rho < 1$, $c_m, C_2 \geq 1$, $\ell_\Sigma \geq L_\Sigma$ and $\kappa \geq 1$. This completes the proof of Theorem 6.

Lastly, we prove Corollary 7. Using (Agarwal et al., 2012, Lemma 5) yields $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq 2\sqrt{s}\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|$. Using $\rho \leq 1/4$ and conditions in (18) it is not difficult to verify that Δ_{stat} in (12) is $o(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2)$.

3.4 Linear convergence for sparse vector regression (Proof of Theorem 9)

We now customize Theorem 6 for the random design setting satisfying Assumption 8. We first state the following high probability bounds for the RSC/RSM properties.

Proposition 21 (Raskutti et al. (2010)) *Given m i.i.d. distributed random matrices $\mathbf{X}_i \in \mathbb{R}^{n \times d}$ drawn from the Σ -Gaussian ensemble. Let $\mathbf{X} = [\mathbf{X}_1^\top, \dots, \mathbf{X}_m^\top]^\top \in \mathbb{R}^{N \times d}$, with $N = n \cdot m$. Then, there exist universal constants $c_0, c_1 > 0$ such that*

$$\begin{aligned} & \frac{\|\mathbf{X}\mathbf{u}\|^2}{N} \geq \frac{1}{2} \|\Sigma^{1/2} \mathbf{u}\|^2 - c_1 \zeta \frac{\log d}{N} \|\mathbf{u}\|_1^2, \\ & \frac{\|\mathbf{X}\mathbf{u}\|^2}{N} \leq 2 \|\Sigma^{1/2} \mathbf{u}\|^2 + c_1 \zeta \frac{\log d}{N} \|\mathbf{u}\|_1^2, \end{aligned} \quad \forall \mathbf{u} \in \mathbb{R}^d; \quad (73)$$

with probability greater than $1 - \exp(-c_0 N)$, where $\zeta = \max_j \Sigma_{jj}$.

By slightly modifying the proof of (Raskutti et al., 2010, Theorem 1), provided in Appendix K for completeness, we can obtain the following bound for the local RSM condition.

Proposition 22 *Given a random matrix $\mathbf{X}_i \in \mathbb{R}^{n \times d}$ drawn from the Σ -Gaussian ensemble, there exists universal constants $c_0, c_1 > 0$ such that*

$$\frac{\|\mathbf{X}_i \mathbf{u}\|^2}{n} \leq 16m \|\Sigma^{1/2} \mathbf{u}\|^2 + c_1 \zeta \frac{m \log d}{n} \|\mathbf{u}\|_1^2, \quad \forall \mathbf{u} \in \mathbb{R}^d, i \in [m] \quad (74)$$

with probability greater than $1 - \exp(-c_0 N)$, where $\zeta = \max_j \Sigma_{jj}$.

Next we derive a high probability bound for $c_g = \max_{j \in [m]} \|\nabla \mathcal{L}_j(\boldsymbol{\theta}^*)\|_\infty + \|\nabla \mathcal{L}(\boldsymbol{\theta}^*)\|_\infty$.

Lemma 23 *Consider the linear model (1) under Assumption 8. Then*

$$\mathbb{P} \left(\max_{j \in [m]} \left\| \frac{\mathbf{X}_j^\top \mathbf{n}_j}{n} \right\|_\infty \geq \sqrt{\zeta} \sigma \cdot \max \left\{ \sqrt{\frac{2 \log md}{c_3} \frac{2 \log md}{n}}, \frac{2 \log md}{c_3} \frac{2 \log md}{n} \right\} \right) \leq 2 \exp(-\log d) \quad (75)$$

and

$$\mathbb{P} \left(\left\| \frac{\mathbf{X}^\top \mathbf{n}}{N} \right\|_\infty \geq \sqrt{\zeta} \sigma \cdot \max \left\{ \sqrt{\frac{2 \log d}{c_3} \frac{2 \log d}{N}}, \frac{2 \log d}{c_3} \frac{2 \log d}{N} \right\} \right) \leq 2 \exp(-\log d) \quad (76)$$

for some $c_3 > 0$.

Proof See Appendix J. ■

According to (1) we have $\|\nabla \mathcal{L}(\boldsymbol{\theta}^*)\|_\infty = \|\mathbf{X}^\top \mathbf{n}/N\|_\infty$ and $\|\nabla \mathcal{L}_j(\boldsymbol{\theta}^*)\|_\infty = \|\mathbf{X}_j^\top \mathbf{n}_j/n\|_\infty$. Using Lemma 23 we can readily obtain the following bound for c_g^2 .

Corollary 24 *Reintate the conditions of Lemma 23; there holds:*

$$\begin{aligned} c_g^2 &\leq 4\zeta \sigma^2 \left(\frac{2 \log md}{c_3} \frac{2 \log md}{n} + \left(\frac{2 \log md}{c_3} \frac{2 \log md}{n} \right)^2 + \frac{2 \log d}{c_3} \frac{2 \log d}{N} + \left(\frac{2 \log d}{c_3} \frac{2 \log d}{N} \right)^2 \right) \\ &\leq 4\zeta \sigma^2 \left(\frac{4 m \log md}{c_3} \frac{4 m \log md}{N} + \left(\frac{4 m \log md}{c_3} \frac{4 m \log md}{N} \right)^2 \right), \end{aligned} \quad (77)$$

with probability greater than $1 - 4 \exp(-c_4 \log d)$, for some $c_4 > 0$.

Leveraging Theorem 6 with the parameters $(\mu_\Sigma, L_\Sigma, \ell_\Sigma)$, $(\tau_\mu, \tau_g, \tau_\ell)$, and c_g guaranteed by Proposition 21, Proposition 22, and Corollary 24 we are ready to prove Theorem 9.

Specifically, let

$$\begin{aligned} \mu_\Sigma &= \frac{1}{2} \sigma_{\min}(\Sigma) \quad \text{and} \quad \tau_\mu = c_1 \zeta \frac{\log d}{N} \quad (\text{for the RSC}), \\ L_\Sigma &= 2 \sigma_{\max}(\Sigma) \quad \text{and} \quad \tau_g = c_1 \zeta \frac{\log d}{N} \quad (\text{for the RSM}), \end{aligned} \quad (78)$$

and

$$\ell_\Sigma = 16m \sigma_{\max}(\Sigma) \quad \text{and} \quad \tau_\ell = c_1 \zeta m^2 \frac{\log d}{N} \quad (\text{for the local RSM}); \quad (79)$$

Propositions 21 & 22 and Corollary 24 guarantee that, with the above choices for $(\mu_\Sigma, L_\Sigma, \ell_\Sigma)$ and $(\tau_\mu, \tau_g, \tau_\ell)$, the global RSC and RSM conditions (cf. Assumption 1), the local RSM property (cf. Assumption 2), and the bound (77) on c_g all hold with probability at least $1 - c_8 \exp(-c_9 \log d)$, where we invoked the union bound and used $\log d \leq s \log d \leq c_5 \frac{\mu_\Sigma}{\zeta} N \leq c_5 N$, due to $\frac{\zeta}{\mu_\Sigma} \frac{s \log d}{N} < c_5$.

The rest of the proof is to (i) show (20) implies (15)-(16) in Theorem 6 and (ii) simplify the expressions of the rate expression (13), the statistical error (12), and B in (69) using (78)-(79). The calculation is left to Appendix L.

4. Numerical Results

This section provides some numerical results that validate our theoretical findings.

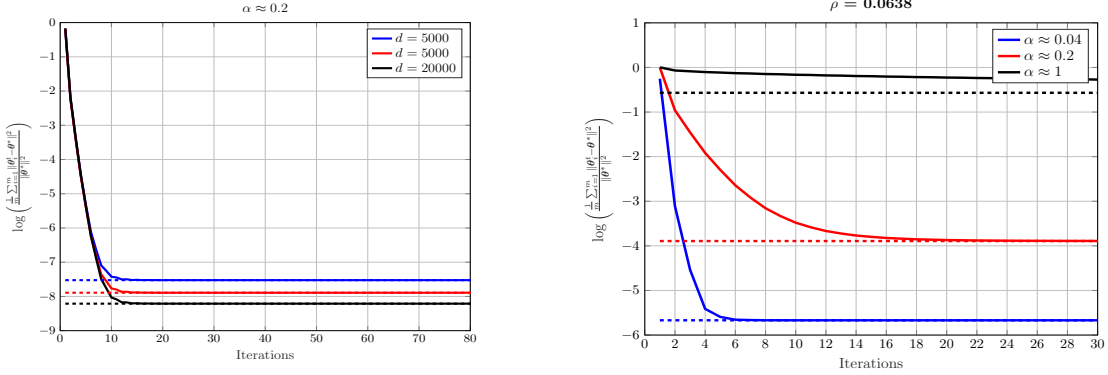
We consider the following problem setup. Given the distributed linear regression model (1), the unknown s -sparse vector θ^* is generated with the first s elements being i.i.d. $\mathcal{N}(0, 1)$; the remaining coordinates are set to zero. Each row of $\mathbf{X}_i \in \mathbb{R}^{n \times d}$ is generated according to $\mathcal{N}(\mathbf{0}, \mathbf{I})$; each element of \mathbf{n}_i follows $\mathcal{N}(0, 0.25)$; and r is set $r = \|\theta^*\|_1$. Unless otherwise specified, we the communication network is modeled as a Erdős-Rényi graph $G(m, p)$, where m is the number of nodes and each edge is included in the graph with probability p , independently from the others. The specific values of problem parameters (N, d, s) and network parameters (m, p) are specified in each simulation.

Aiming at validating Theorem 9, the reported experiments show the following: (i) the scalability of the algorithm with respect to the total sample size N and problem dimension d , for a fixed network; and (ii) the impact of network topology on the algorithm. Finally, we validate our findings on real data sets.

4.1 Scalability with respect to problem parameters for a fixed network

Theorem 9 states that, for a fixed, sufficiently connected network, the average optimization error of DGT among all agents decreases at a linear, up to a tolerance depending on the statistical precision $\|\hat{\theta} - \theta^*\|$. Moreover, the expression of the convergence rate and the tolerance term depend exclusively on the ratio $\alpha = \frac{s \log d}{N}$. To validate numerically these theoretical findings, we conduct the following two simulations over an Erdős-Rényi graph with $m = 50$ nodes and edge activation probability $p = 0.5$. The reported results are the average of 100 trials with the same graph and independently randomly generated data sets.

(i) Fixed α and increasing dimension d : We set $\alpha \approx 0.2$ and consider three sets of problem parameters, namely: $d = 5000, 10000, 20000$ and $s = \lceil \sqrt{d} \rceil$. Consequently, the number of total samples N and local sample size n is set to $N = 3050, 4700, 7100$ and $n = 61, 94, 142$, respectively. Observe that $N = m \cdot n$ with $m = 50$. Since the expression of the proximal parameter γ may be conservative in practice, we tune it manually as described next. We first generate one instance of the LASSO problem with $d = 5000$ and run DGT using the γ that yields the fastest convergence rate. This value is selected via grid search. The same γ is used for all problem instances. Fig. 3a plots the average normalized estimation error $(1/m) \sum_{i=1}^m \|\theta_i^t - \theta^*\|^2 \|\theta^*\|^{-2}$ versus the number of iterations, for $d = 5000, 10000, 20000$. The dashed-line curves represent the normalized statistical precision $(\|\hat{\theta} - \theta^*\|^2 \|\theta^*\|^{-2})$ for the three sets of problems. The figure shows that, in all cases, the estimation error decays



(a) Log-normalized error for $\alpha \approx 0.2$. In these problems $d \in \{5000, 10000, 20000\}$, $s = \lceil \sqrt{d} \rceil$, $m = n \in \{61, 94, 142\}$. (b) Log-normalized error for different values of α . In these problems $d = 20000$, $s = \lceil \log(d)/2 \rceil$, $m = 50$ and $\alpha \in \{0.04, 0.2, 1\}$.

Figure 3: Scalability of DGT for fixed networks. The network is composed of $m = 50$ agents with connectivity $\rho = 0.0638$. Dashed lines indicate estimation error if using $\hat{\theta}$.

at a linear rate before reaching the statistical precision. Further, the rate is invariant to the change of d as long as the ratio α remains fixed. This is exactly what Theorem 9 predicts.

(ii) Increasing α : This experiment explores the relation between the convergence rate of DGT and α . We set $d = 20000$, $s = \lceil (\log d)/2 \rceil$ and consider three local sample sizes, $n = 1, 5$, and 25 . Correspondingly, we have $\alpha \approx \{1, 0.2, 0.04\}$. For each value of n , we choose the proximal parameter γ via grid search. The one that yields fastest convergence rate is kept fixed for all independent trials. Fig. 3b shows that the algorithm’s performance degrades as α increases—both the convergence rate of DGT and the statistical precision get worse. Note that this behavior is in agreement with Theorem 9.

4.2 Impact of network scale and connectivity

This set of simulations investigates the scalability of DGT with respect to network parameters (m, ρ) , for fixed problem parameters (N, d, s) . In particular, Theorem 9 established the convergence of DGT under condition $\rho \lesssim m^{-8} \kappa^{-4}$, which requires better connectivity for a network involving more agents. To test the necessity of this condition, we run DGT on networks with increasing size $m = 50, 625, 1250, 2500$. Correspondingly we varied the edge connectivity probability $p = 0.87, 0.4, 0.23, 0.15$ in order to keep ρ roughly constant, equal to ≈ 0.18 . The problem parameters are set to be $N = 2500$, $d = 5000$, $s = \lceil \sqrt{d} \rceil$. The step size is chosen via grid search and set to be the one yielding the fastest convergence. All curves reported for this simulation in one trial are conducted on the same problem instance for different network configurations. The result is averaged over 100 trials and is reported in Fig 4a. Observe that keeping the graph’s connectivity constant while increasing the number of agents leads to a degradation of the achievable rate and estimation error.

Additionally, we empirically demonstrate that for Erdős-Rényi graphs with connectivity $\rho = \mathcal{O}(1/\sqrt{m})$ —achieved by fixing the network connectivity probability p —DGT converges under a growing network size. The result is reported in Fig. 4b. The simulation is conducted

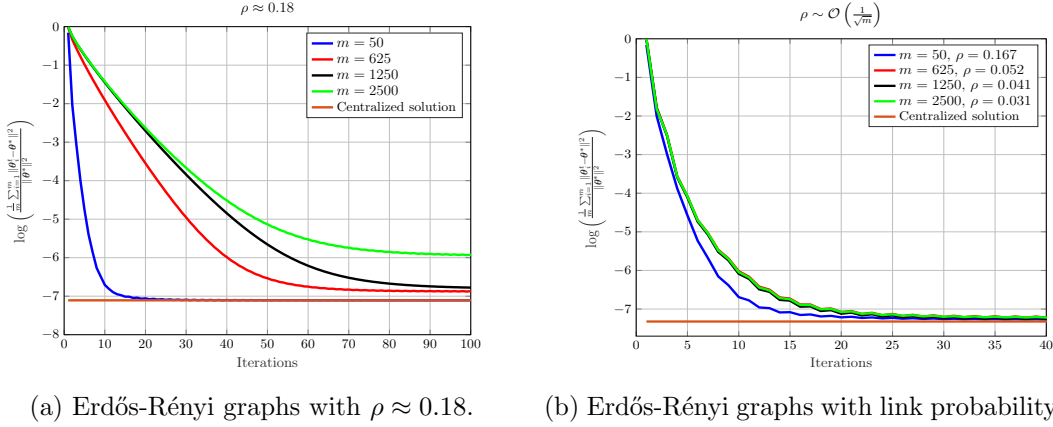


Figure 4: Plot of the log-normalized estimation error for $m \in \{50, 625, 1250, 2500\}$. In this problem $d = 5000$, $s = \lceil \sqrt{d} \rceil$, $n \in \{50, 4, 2, 1\}$ and $N = 2500$. All curves in each panel are obtained using the same data points. Black horizontal lines indicate the normalized estimation error if using $\hat{\theta}$.

	Erdős-Rényi ($p = 0.87$)	Line graph
$m = 50$	18	23775
$m = 625$	18	6115118
$m = 1250$	18	27094153
$m = 2500$	19	118911225

Figure 5: Number of communication rounds required for each network to satisfy $\rho^k \leq m^{-8}$.

under the same setting as Fig. 4a except for p set to be 0.87. The simulation suggests that the convergence condition $\rho \lesssim m^{-8}\kappa^{-4}$ might be more restrictive than what is needed in practice, which could be an artifact of our analysis.

To conclude the assessment of the algorithm performance as function of the network connectivity parameter ρ , in Table 5, we report the number of communication rounds required per iteration to satisfy the condition $\rho^K \leq m^{-8}$. Observe that the number remains approximately constant for Erdős-Rényi graphs while increasing dramatically for line-graphs. Similarly, in Fig. 6, we display the number of total communication rounds required to solve a problem instance with $d = 5000$, $s = \lceil \sqrt{d} \rceil$, $m \in \{50, 625, 1250, 2500\}$, and $n \in \{50, 4, 2, 1\}$. The results are consistent with those displayed in Table 5 and demonstrate that the quantity ρ is predictive of the number of communication rounds required for convergence.

4.3 Comparison with other distributed algorithms

We test the performance of DGT when compared to other decentralized algorithms with known guarantees in the high-dimensional regime. Specifically, we consider DGD-CTA (Ji et al., 2023a), DGD-ATC (Ji et al., 2023b), and DGD² proposed in our companion work (Maros and Scutari, 2022). DGD-ATC, and DGD² are known to achieve centralized statistical precision at a rate comparable that of the centralized PGD, provided the network is sufficiently connected. Conversely, DGD-CTA tends to slow down when aiming for central-

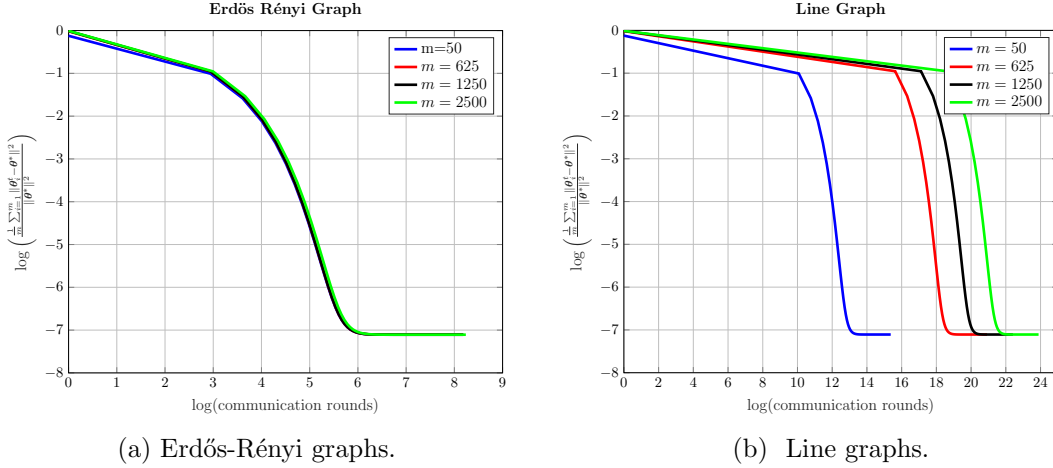


Figure 6: Plots of the log-normalized estimation error for Erdős-Rényi and line graphs, with $m \in \{50, 625, 1250, 2500\}$. In this problem, $d = 5000$, $s = \lceil \sqrt{d} \rceil$, $n \in \{50, 4, 2, 1\}$ and $N = 2500$.

ized statistical precision. Given that, except for DGD-CTA, all algorithms can potentially perform equivalently in terms of iterations—contingent on meeting specific network connectivity requirements—we focus on the number of communication rounds required to reach centralized statistical precision rather than just the total iteration count. For this analysis, we create problem instances with $d \in \{400, 2000, 4000, 20000\}$, $s \in \{5, 4, 7, 4\}$, and $N \in \{240, 240, 480, 360\}$. This yields $\alpha \approx 0.11$ for all instances. We solve each problem instance up to statistical precision using the centralized PGD and record the total number of performed iterations T_{cent} for each problem. In the decentralized setting, we generate a base mixing matrix W with $m = 120$ and $p = 0.02$, which results in a poorly connected base mixing matrix. For DGD-ATC, DGD² and DGT we select a mixing matrix $\bar{W} = W^K$ with the smallest possible K that allows the algorithms to achieve centralized statistical precision in at most $2T_{\text{cent}}$ iterations. For DGD-ATC, DGD² and DGT we measure the total communication rounds as iterations $\times K$. For DGD-CTA we set the network to be fully connected with the largest possible step-size that yields convergence to the order of centralized statistical precision and record the total number of network uses, which in this case also corresponds to the number of iterations. The results of this experiment are reported in Fig. 7. We observe that despite employing a fully connected network, DGD-CTA requires the most communication rounds of all schemes. This is because the network connectivity does not significantly aid the performance of DGD-CTA. Meanwhile, we observe that for DGD-ATC the number of communication rounds increases moderately with the dimension d while for DGD² and DGT the number of communication rounds remains invariant as the problem dimension increases. The statistical error achieved by DGD² and DGT while of similar order is beneficial towards DGT.

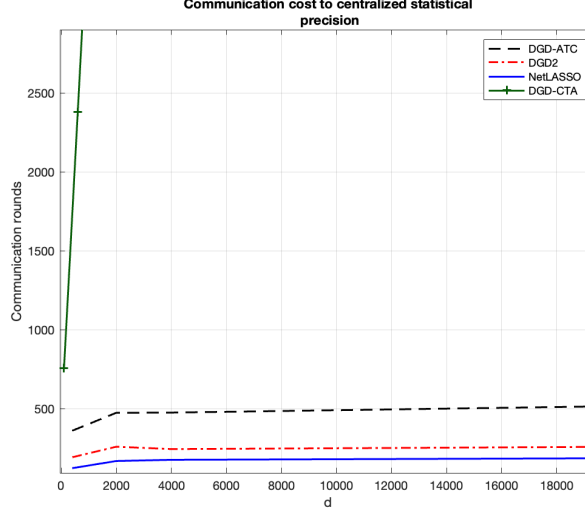


Figure 7: Number of communication rounds for $m = 120$, $d \in \{400, 2000, 4000, 20000\}$, $s \in \{5, 4, 7, 4\}$, and $N \in \{240, 240, 480, 360\}$. All curves are obtained using the same data points and consist of 20 averages each.

4.4 Simulations on real data

We test the DGT on the **Communities and Crime** data set (UC Irvine Machine Learning Repository) where we have removed data points with missing attributes (covariates), and removed attributes corresponding to community number/name or zip-code (attributes 1 to 4 in the data set) yielding a regression problem with $d = 123$ and total sample size of 123 which we have split into $N_{\text{train}} = 82$ and $N_{\text{test}} = 41$. For the decentralized implementation, we have a network of $m = 41$ agents corresponding to $n = 2$ samples per agent. The agents communicate via a network built upon an Erdős-Rényi where the probability of an edge being present is set to $p = 0.5$. Because the resulting graph has connectivity $\bar{\rho} = 0.5236$ we have agents perform three communication rounds per gradient computation yielding an effective connectivity $\rho = 0.1435 \approx \frac{1}{\sqrt{m}}$ which corresponds to the connectivity that is sufficient to yield centralized performance according to numerical simulations (c.f. Figure 4b). Observe that our theoretical results yield more conservative values for ρ ; consequently, employing the connectivity predicted by the theory can only yield improved results. Our results are reported in Figure 8 where the performance of the centralized projected gradient descent (PGD) is reported to the left, and the performance of the DGT is reported on the right panel. Denote by $(X_{\text{train}}, y_{\text{train}})$ the data pairs included in the training data set, and analogously by $(X_{\text{test}}, y_{\text{test}})$ the data pairs included in the testing data set. The training error is measured as $\frac{1}{2N_{\text{train}}} \|X_{\text{train}} \theta^t - y_{\text{train}}\|_2^2$ where θ^t are iterates of PGD generated by minimizing the function $\frac{1}{2N_{\text{train}}} \|X_{\text{train}} \theta - y_{\text{train}}\|^2$ over the set $\|\theta\|_1 \leq 0.85$ using step-size 0.09. The testing error is measured by $\frac{1}{2N_{\text{test}}} \|X_{\text{test}} \theta^t - y_{\text{test}}\|^2$. The right panel reports the performance of DGT with step-size set to be 0.05. The training and testing errors are measured by $\frac{1}{m} \sum_{i=1}^m \frac{1}{2N_{\text{train}}} \|X_{\text{train}} \theta_i^t - y_{\text{train}}\|^2$ and $\frac{1}{m} \sum_{i=1}^m \|X_{\text{test}} \theta_i^t - y_{\text{test}}\|^2$, respectively. We report in the right panel the difference between the centralized and decentralized train and




Figure 8: Plots of train and test error on the **Communities** and **Crime** data set for centralized projected gradient descent (PGD) (left) and DGT (right).

test errors respectively. We conclude that the performance of centralized and decentralized schemes is practically identical (thanks to multiple but finite consensus rounds), both in training and testing errors.

5. Conclusion

We studied sparse linear regression over mesh networks. We established statistical and computational guarantees in high-dimensions for DGT, a decentralization of the PGD based on gradient tracking. We proved that (near) optimal sample complexity $\mathcal{O}(s \log d/N)$ for the distributed estimator is achievable over networks, even when the local sample size is not sufficient for statistical consistency. Convergence at a linear rate was proved, which is of the same order as that of PGD. This improves on DGD-like methods (Ji et al., 2023a,b), which instead suffer from the speed-accuracy dilemma.

Our work, together with the study of DGD (Ji et al., 2023a,b) and DGD² (Maros and Scutari, 2022), reveal the (non)-convergence of decentralized algorithms is significantly affected by the gradient estimator used in the local optimization step. A future investigation will be the study of other distributed methods; of particular interest is understanding the role of other forms of gradient corrections towards statistical, computation, and communication tradeoffs.

Acknowledgments

Maros and Scutari would like to acknowledge support for this project from the Office of Naval Research (ONR Grant N. N000142112673 and ONR Grant N. N000142412751). Cheng acknowledges support from National Science Foundation (NSF-CNS 2247795), and the Office of Naval Research (ONR Grant N. N00014-22-1-2680).

Appendix A. Proof of Proposition 11

First note that the optimality of $\boldsymbol{\theta}_i^{t+\frac{1}{2}}$ implies:

$$\left(\boldsymbol{\theta} - \boldsymbol{\theta}_i^{t+\frac{1}{2}}\right)^\top (\mathbf{g}_i^t + \gamma \Delta \boldsymbol{\theta}_i^t) \geq 0, \quad \forall \|\boldsymbol{\theta}\|_1 \leq r, \quad (80)$$

where we recall $\Delta \boldsymbol{\theta}_i^t \triangleq \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t$.

We use now the global RSC/RSM property of \mathcal{L} (cf. Assumption 1) to compute the reduction of optimality gap $\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}(\hat{\boldsymbol{\theta}})$ by executing step (3b).

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_i^{t+\frac{1}{2}}) &= \mathcal{L}(\boldsymbol{\theta}_i^t) + \nabla \mathcal{L}(\boldsymbol{\theta}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t) + \frac{1}{2} (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t)^\top \left(\frac{\mathbf{X}^\top \mathbf{X}}{N} \right) (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t) \\ &\stackrel{(RSM)}{\leq} \mathcal{L}(\boldsymbol{\theta}_i^t) + \nabla \mathcal{L}(\boldsymbol{\theta}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t) + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + \frac{\tau_g}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|_1^2 \\ &= \mathcal{L}(\hat{\boldsymbol{\theta}}) + \nabla \mathcal{L}(\boldsymbol{\theta}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) - \frac{1}{2} (\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}})^\top \left(\frac{\mathbf{X}^\top \mathbf{X}}{N} \right) (\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}) \\ &\quad + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + \frac{\tau_g}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|_1^2 \\ &\stackrel{(RSC)}{\leq} \mathcal{L}(\hat{\boldsymbol{\theta}}) + \nabla \mathcal{L}(\boldsymbol{\theta}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) - \frac{\mu_\Sigma}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 \\ &\quad + \frac{\tau_\mu}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|_1^2 + \frac{\tau_g}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|_1^2. \end{aligned}$$

Letting $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$ in (80) and adding to the above inequality we have

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}_i^{t+\frac{1}{2}}) &\leq \mathcal{L}(\hat{\boldsymbol{\theta}}) - \frac{\gamma}{2} \left(\|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 - \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 \right) + (\nabla \mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) \\ &\quad - \frac{\mu_\Sigma}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + \frac{\tau_\mu}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|_1^2 + \frac{\tau_g}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|_1^2. \end{aligned}$$

Now applying the triangle inequality to the last term and Lemma 10 yields:

$$\begin{aligned} &\mathcal{L}(\boldsymbol{\theta}_i^{t+\frac{1}{2}}) \\ &\leq \mathcal{L}(\hat{\boldsymbol{\theta}}) - \frac{\gamma}{2} \left(\|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 - \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 \right) + (\nabla \mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) \\ &\quad - \frac{\mu_\Sigma}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + \frac{\tau_\mu}{2} (8s \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + 2\nu^2) \\ &\quad + \frac{\tau_g}{2} (16s \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 + 16s \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + 8\nu^2) \\ &\leq \mathcal{L}(\hat{\boldsymbol{\theta}}) - \frac{\gamma}{2} \left(\|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 - \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 \right) + (\nabla \mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}) \\ &\quad - \frac{\mu_\Sigma}{2} \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 + \frac{L_\Sigma}{2} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + 8s\tau_g \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 + (4s\tau_\mu + 8s\tau_g) \|\boldsymbol{\theta}_i^t - \hat{\boldsymbol{\theta}}\|^2 \\ &\quad + (\tau_\mu + 4\tau_g) \nu^2. \end{aligned}$$

Rearranging terms and using the fact that $\mathcal{L}(\boldsymbol{\theta}_i^{t+\frac{1}{2}}) \geq \mathcal{L}(\widehat{\boldsymbol{\theta}})$, yields

$$\begin{aligned} \left(1 - \frac{C_1 s \tau_g}{\gamma}\right) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 &\leq \left(1 - \frac{\mu \Sigma}{\gamma} + \frac{C_1 s (\tau_\mu + \tau_g)}{\gamma}\right) \|\boldsymbol{\theta}_i^t - \widehat{\boldsymbol{\theta}}\|^2 \\ &\quad - \left(1 - \frac{L \Sigma}{\gamma}\right) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \boldsymbol{\theta}_i^t\|^2 + \frac{2}{\gamma} \cdot (\nabla \mathcal{L}(\boldsymbol{\theta}_i^t) - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) + \frac{C_1 (\tau_\mu + \tau_g)}{\gamma} \nu^2. \end{aligned} \quad (81)$$

Summing over $i = 1, \dots, m$, completes the proof. \square

Appendix B. Proof of Proposition 12

The difference $\frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\boldsymbol{\theta}_i^t) - \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\boldsymbol{\theta}_j^t)$ can be expressed as

$$\frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\boldsymbol{\theta}_i^t) - \frac{1}{m} \sum_{j=1}^m \nabla \mathcal{L}_j(\boldsymbol{\theta}_j^t) = \frac{1}{N} \mathbf{X}^\top \mathbf{X} (\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}^t) + \frac{1}{m} \sum_{j=1}^m \frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j (\bar{\boldsymbol{\theta}}^t - \boldsymbol{\theta}_j^t).$$

Substituting into the expression of δ^t we can split it into three error terms:

$$\begin{aligned} \delta^t &= \underbrace{\sum_{i=1}^m (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\frac{1}{N} \mathbf{X}^\top \mathbf{X} \right) (\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}^t)}_{T_1} + \underbrace{\frac{1}{m} \sum_{i=1}^m \sum_{j=1}^m (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) (\bar{\boldsymbol{\theta}}^t - \boldsymbol{\theta}_j^t)}_{T_2} \\ &\quad + \underbrace{\sum_{i=1}^m (\bar{\mathbf{g}}^t - \mathbf{g}_i^t)^\top (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})}_{T_3}. \end{aligned} \quad (82)$$

We are now ready to bound δ^t . For this, we substitute the expressions of $\boldsymbol{\Theta}_\perp^t$ and \mathbf{G}_\perp^t [cf. (28a)-(28b)] into (82). We then use the RSM (cf. Assumption 1) to upper bound δ^t by quantities that depend exclusively on the sequences $\|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|$ and $\|\Delta \boldsymbol{\Theta}^t\|$.

1) Bounding T_1 : Iterating (28a) telescopically to $t = 0$, yields

$$\boldsymbol{\Theta}_\perp^t = (\mathbf{W} - \mathbf{J}) \boldsymbol{\Theta}_\perp^0 + \sum_{s=0}^{t-1} (\mathbf{W} - \mathbf{J})^{t-s} \Delta \boldsymbol{\Theta}^s, \quad t = 1, 2, \dots, \quad (83)$$

This shows that $\boldsymbol{\Theta}_\perp^t$ lies in the span of $\{\boldsymbol{\Theta}_\perp^0, \Delta \boldsymbol{\Theta}^0, \dots, \Delta \boldsymbol{\Theta}^{t-1}\}$. Therefore,

$$\boldsymbol{\theta}_i^t - \bar{\boldsymbol{\theta}}^t = \sum_{j=1}^m \ell_{ij}^{(t)} (\boldsymbol{\theta}_{j,\perp}^0) + \sum_{j=1}^m \sum_{s=0}^{t-1} \ell_{ij}^{(t-s)} \Delta \boldsymbol{\theta}_j^s, \quad t = 1, 2, \dots, \quad (84)$$

where $\ell_{ij}^{(t)}$ denotes the ij -th element of the matrix $(\mathbf{W} - \mathbf{J})^t$.

Substituting (84) in the expression of T_1 , we can bound it as

$$\begin{aligned}
 T_1 &= \sum_{i=1}^m (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\frac{1}{N} \mathbf{X}^\top \mathbf{X} \right) \left(\sum_{j=1}^m \ell_{ij}^{(t)} \boldsymbol{\theta}_{j,\perp}^0 + \sum_{j=1}^m \sum_{s=0}^{t-1} \ell_{ij}^{(t-s)} \Delta \boldsymbol{\theta}_j^s \right) \\
 &\leq \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \left\| \frac{1}{\sqrt{N}} \mathbf{X} (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \right\| \left\| \frac{1}{\sqrt{N}} \mathbf{X} \boldsymbol{\theta}_{j,\perp}^0 \right\| \\
 &\quad + \sum_{i=1}^m \sum_{j=1}^m \sum_{s=0}^{t-1} |\ell_{ij}^{(t-s)}| \left\| \frac{1}{\sqrt{N}} \mathbf{X} (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \right\| \left\| \frac{1}{\sqrt{N}} \mathbf{X} \Delta \boldsymbol{\theta}_j^s \right\| \\
 &\leq \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \left(\frac{\epsilon}{2} \left\| \frac{1}{\sqrt{N}} \mathbf{X} (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \right\|^2 + \frac{1}{2\epsilon} \left\| \frac{1}{\sqrt{N}} \mathbf{X} \boldsymbol{\theta}_{j,\perp}^0 \right\|^2 \right) \\
 &\quad + \sum_{i=1}^m \sum_{j=1}^m \sum_{s=0}^{t-1} |\ell_{ij}^{(t-s)}| \left(\frac{\epsilon}{2} \left\| \frac{1}{\sqrt{N}} \mathbf{X} (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \right\|^2 + \frac{1}{2\epsilon} \left\| \frac{1}{\sqrt{N}} \mathbf{X} \Delta \boldsymbol{\theta}_j^s \right\|^2 \right), \tag{85}
 \end{aligned}$$

where we have applied Cauchy-Schwartz and Young's inequality with $\epsilon > 0$. Invoking the RSM (Assumption 1) along with Assumption 4 and Lemma 10, we can bound the RHS of (85) as (cf. Appendix C.1): for all $t \geq 1$,

$$\begin{aligned}
 T_1 &\leq \frac{\rho \cdot c_m \epsilon}{1 - \rho} (L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|_F^2 + \frac{c_m L_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \boldsymbol{\theta}^s\|_F^2 \\
 &\quad + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\theta}^{s+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_F^2 + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\theta}^s - \widehat{\boldsymbol{\theta}}\|_F^2 \\
 &\quad + \underbrace{\frac{\rho \cdot m c_m}{1 - \rho} \cdot 4\tau_g \nu^2 (\epsilon + \epsilon^{-1}) + \frac{1}{2\epsilon} c_m \rho^t \left(L_\Sigma \|\boldsymbol{\theta}_\perp^0\|_F^2 + \tau_g \sum_{j=1}^m \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right)}_{=0 \text{ by initialization}}.
 \end{aligned}$$

2) Bounding T_2 : Recall that

$$T_2 = \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^m \underbrace{(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) (\bar{\boldsymbol{\theta}}^t - \boldsymbol{\theta}_j^t)}_{\triangleq T_2^{ij}}.$$

Focusing on the ij -th summand and using (84) we can write

$$T_2^{ij} = (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) \left(\sum_{k=1}^m (-\ell_{jk}^{(t)}) \boldsymbol{\theta}_{k,\perp}^0 + \sum_{k=1}^m \sum_{s=0}^{t-1} (-\ell_{jk}^{(t-s)}) \Delta \boldsymbol{\theta}_k^s \right). \tag{86}$$

Note that T_2^{ij} has the same form as T_1 and thus following similar steps as in (??) while using the local RSM (Assumption 2), we can bound T_2 as (cf. Appendix C.2): for all $t \geq 1$,

$$\begin{aligned}
T_2 &\leq \frac{\rho \cdot c_m \epsilon}{1 - \rho} (\ell_\Sigma + 8s\tau_\ell) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}} \right\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \boldsymbol{\Theta}^s\|_F^2 \\
&\quad + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^{s+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^s - \widehat{\boldsymbol{\Theta}}\|_F^2 \\
&\quad + \underbrace{\frac{\rho \cdot mc_m}{1 - \rho} 4\tau_\ell \nu^2 \cdot (\epsilon + \epsilon^{-1}) + \frac{1}{2\epsilon} c_m \rho^t \left(\ell_\Sigma \|\boldsymbol{\Theta}_\perp^0\|_F^2 + \tau_\ell \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right)}_{=0}.
\end{aligned} \tag{87}$$

3) Bounding T_3 : Using (28b) we can express $(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top (\bar{\mathbf{g}}^t - \mathbf{g}_i^t)$ in T_3 as

$$\begin{aligned}
&(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top (\bar{\mathbf{g}}^t - \mathbf{g}_i^t) \\
&= (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}})^\top \left(\sum_{j=1}^m (-\ell_{ij}^{(t)}) \mathbf{g}_{j,\perp}^0 + \sum_{s=0}^{t-1} \sum_{j=1}^m (-\ell_{ij}^{(t-s)}) \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) (\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s) \right) \\
&\leq \sum_{j=1}^m |\ell_{ij}^{(t)}| \underbrace{\left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right)^\top \mathbf{g}_{j,\perp}^0 \right|}_{T_{3,1}} + \sum_{s=0}^{t-1} \sum_{j=1}^m |\ell_{ij}^{(t-s)}| \underbrace{\left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right)^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) (\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s) \right|}_{T_{3,2}}.
\end{aligned}$$

Following again similar steps as in the bounds of T_1 and T_2 , we can obtain the following bound for T_3 (cf. Appendix C.3):

$$\begin{aligned}
T_3 &\leq \frac{3}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1 - \rho} (2\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^{s+1} - \boldsymbol{\Theta}^s\|_F^2 \\
&\quad + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^{s+1} - \widehat{\boldsymbol{\Theta}}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^s - \widehat{\boldsymbol{\Theta}}\|_F^2 \\
&\quad + mc_m \rho^t (A_3 \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) + \frac{\rho \cdot mc_m}{1 - \rho} 4\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}),
\end{aligned}$$

where

$$\begin{aligned}
A_3 &\triangleq \frac{1}{2} (\ell_\Sigma \|\boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}^*\|_1^2 + \frac{1}{m} \ell_\Sigma \|\boldsymbol{\Theta}^*\|_F^2 + \tau_\ell \|\boldsymbol{\theta}^*\|_1^2), \\
c_g &\triangleq \|\nabla \mathcal{L}_j(\boldsymbol{\theta}^*)\|_\infty + \|\nabla \mathcal{L}(\boldsymbol{\theta}^*)\|_\infty.
\end{aligned} \tag{88}$$

4) Combining the bounds of T_1 , T_2 , and T_3 : Adding up the bounds of T_1 , T_2 , and T_3 , and using the following two inequalities

$$\|\boldsymbol{\Theta}^{t+1} - \widehat{\boldsymbol{\Theta}}\|_F^2 = \|\mathbf{W}(\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}})\|_F^2 \leq \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 \tag{89}$$

and

$$\|\boldsymbol{\Theta}^{s+1} - \boldsymbol{\Theta}^s\|_F^2 \stackrel{(24a)}{=} \|\mathbf{W}(\boldsymbol{\Theta}^s + \Delta \boldsymbol{\Theta}^s) - \boldsymbol{\Theta}^s\|_F^2 \leq 4\|\boldsymbol{\Theta}_\perp^s\|_F^2 + 2\|\Delta \boldsymbol{\Theta}^s\|_F^2, \tag{90}$$

yield the desired bound, detailed calculations are provided in Appendix C.4. \square

Appendix C. Supplement to Appendix B

In this section, we upperbound the tracking errors T_1 , T_2 , and T_3 in (82).

Appendix C.1 Bounding T_1 in (85)

Invoking the RSC property (Assumption 1) and Lemma 10, we have

$$\begin{aligned} \left\| \frac{1}{\sqrt{N}} \mathbf{X}(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \right\|^2 &\leq L_\Sigma \left\| \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|^2 + \tau_g \left\| \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|_1^2 \\ &\leq (L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|^2 + 2\tau_g\nu^2. \end{aligned}$$

Similarly, the other two terms in (85) can be bounded respectively as

$$\left\| \frac{1}{\sqrt{N}} \mathbf{X}\boldsymbol{\theta}_{j,\perp}^0 \right\|^2 \leq L_\Sigma \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|^2 + \tau_g \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|_1^2$$

and

$$\begin{aligned} \left\| \frac{1}{\sqrt{N}} \mathbf{X}\Delta\boldsymbol{\theta}_j^s \right\|^2 &\leq L_\Sigma \|\Delta\boldsymbol{\theta}_j^s\|^2 + \tau_g \|\boldsymbol{\theta}_j^{s+\frac{1}{2}} - \boldsymbol{\theta}_j^s\|_1^2 \\ &\leq L_\Sigma \|\Delta\boldsymbol{\theta}_j^s\|^2 + 2\tau_g \|\boldsymbol{\theta}_j^{s+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_1^2 + 2\tau_g \|\boldsymbol{\theta}_j^s - \widehat{\boldsymbol{\theta}}\|_1^2 \\ &\leq L_\Sigma \|\Delta\boldsymbol{\theta}_j^s\|^2 + 16s\tau_g \|\boldsymbol{\theta}_j^{s+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 16s\tau_g \|\boldsymbol{\theta}_j^s - \widehat{\boldsymbol{\theta}}\|^2 + 8\tau_g\nu^2. \end{aligned}$$

Summing up the terms and using $\|(\mathbf{W} - \mathbf{J})^t\|_\infty \leq c_m\rho^t$ (due to Assumption 4), we arrive to the following bound for T_1 :

$$\begin{aligned} T_1 &\leq \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \frac{\epsilon}{2} \left((L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|^2 + 2\tau_g\nu^2 \right) \\ &\quad + \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \frac{1}{2\epsilon} \left(L_\Sigma \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|^2 + \tau_g \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right) \\ &\quad + \sum_{i=1}^m \sum_{j=1}^m \sum_{s=0}^{t-1} |\ell_{ij}^{(t-s)}| \frac{\epsilon}{2} \left((L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right\|^2 + 2\tau_g\nu^2 \right) \\ &\quad + \sum_{i=1}^m \sum_{j=1}^m \sum_{s=0}^{t-1} |\ell_{ij}^{(t-s)}| \frac{1}{2\epsilon} \left(L_\Sigma \|\Delta\boldsymbol{\theta}_j^s\|^2 + 16s\tau_g \|\boldsymbol{\theta}_j^{s+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 16s\tau_g \|\boldsymbol{\theta}_j^s - \widehat{\boldsymbol{\theta}}\|^2 + 8\tau_g\nu^2 \right) \\ &\leq \frac{\epsilon}{2} c_m \rho^t (L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}} \right\|_F^2 + \frac{\epsilon}{2} \cdot m c_m \rho^t \cdot 2\tau_g\nu^2 \\ &\quad + \frac{1}{2\epsilon} c_m \rho^t \left(L_\Sigma \|\boldsymbol{\Theta}_\perp^0\|_F^2 + \tau_g \sum_{j=1}^m \|\boldsymbol{\theta}_j^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right) \\ &\quad + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{\epsilon}{2} (L_\Sigma + 8s\tau_g) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}} \right\|_F^2 + \sum_{s=0}^{t-1} m c_m \rho^{t-s} \cdot \frac{\epsilon}{2} \cdot 2\tau_g\nu^2 \end{aligned}$$

$$\begin{aligned}
& + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{1}{2\epsilon} L_\Sigma \|\Delta \Theta^s\|_F^2 + \sum_{s=0}^{t-1} c_m \rho^{t-s} \cdot \frac{1}{2\epsilon} \cdot 16s\tau_g \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\
& + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{1}{2\epsilon} \cdot 16s\tau_g \|\Theta^s - \widehat{\Theta}\|_F^2 + \sum_{s=0}^{t-1} m c_m \rho^{t-s} \frac{1}{2\epsilon} 8\tau_g \nu^2.
\end{aligned}$$

Combining terms yields

$$\begin{aligned}
T_1 & \leq \frac{\rho \cdot c_m \epsilon}{1 - \rho} (L_\Sigma + 8s\tau_g) \left\| \Theta^{t+\frac{1}{2}} - \widehat{\Theta} \right\|_F^2 + \frac{c_m L_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 \\
& + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
& + \frac{\rho \cdot m c_m}{1 - \rho} \cdot 4\tau_g \nu^2 (\epsilon + \epsilon^{-1}) + \frac{1}{2\epsilon} c_m \rho^t \left(L_\Sigma \|\Theta_\perp^0\|_F^2 + \tau_g \sum_{j=1}^m \|\theta_j^0 - \bar{\theta}^0\|_1^2 \right).
\end{aligned}$$

Appendix C.2 Bounding T_2 in (82)

Using the local RSM (cf. Assumption 2) and Lemma 10, we can bound the summands in T_2^{ij} [cf. (86)] as

$$\begin{aligned}
& (\theta_i^{t+\frac{1}{2}} - \widehat{\theta})^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) \ell_{jk}^{(t)} \theta_{k,\perp}^0 \\
& \leq |\ell_{jk}^{(t)}| \left\| \frac{1}{\sqrt{n}} \mathbf{X}_j (\theta_i^{t+\frac{1}{2}} - \widehat{\theta}) \right\| \cdot \left\| \frac{1}{\sqrt{n}} \mathbf{X}_j (\theta_k^0 - \bar{\theta}^0) \right\| \\
& \leq |\ell_{jk}^{(t)}| \frac{\epsilon}{2} \left((\ell_\Sigma + 8s\tau_\ell) \|\theta_i^{t+\frac{1}{2}} - \widehat{\theta}\|^2 + 2\tau_\ell \nu^2 \right) + |\ell_{jk}^{(t)}| \frac{1}{2\epsilon} \left(\ell_\Sigma \|\theta_k^0 - \bar{\theta}^0\|^2 + \tau_\ell \|\theta_k^0 - \bar{\theta}^0\|_1^2 \right)
\end{aligned}$$

and

$$\begin{aligned}
& (\theta_i^{t+\frac{1}{2}} - \widehat{\theta})^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) \ell_{jk}^{(t-s)} \Delta \theta_k^s \\
& \leq |\ell_{jk}^{(t-s)}| \left\| \frac{1}{\sqrt{n}} \mathbf{X}_j (\theta_i^{t+\frac{1}{2}} - \widehat{\theta}) \right\| \left\| \frac{1}{\sqrt{n}} \mathbf{X}_j \Delta \theta_k^s \right\| \\
& \leq |\ell_{jk}^{(t-s)}| \frac{\epsilon}{2} \left((\ell_\Sigma + 8s\tau_\ell) \|\theta_i^{t+\frac{1}{2}} - \widehat{\theta}\|^2 + 2\tau_\ell \nu^2 \right) \\
& + |\ell_{jk}^{(t-s)}| \frac{1}{2\epsilon} \left(\ell_\Sigma \|\Delta \theta_k^s\|^2 + 16s\tau_\ell \|\theta_k^{s+\frac{1}{2}} - \widehat{\theta}\|^2 + 16s\tau_\ell \|\theta_k^s - \widehat{\theta}\|^2 + 8\tau_\ell \nu^2 \right),
\end{aligned}$$

respectively.

Summing over i, j, k and using $\|(\mathbf{W} - \mathbf{J})^t\|_\infty \leq c_m \rho^t$, we can bound T_2 as:

$$\begin{aligned}
& m \cdot T_2 \\
& \leq \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m |\ell_{jk}^{(t)}| \frac{\epsilon}{2} \left((\ell_\Sigma + 8s\tau_\ell) \|\theta_i^{t+\frac{1}{2}} - \widehat{\theta}\|^2 + 2\tau_\ell \nu^2 \right)
\end{aligned}$$

$$\begin{aligned}
 & + \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m |\ell_{jk}^{(t)}| \frac{1}{2\epsilon} \left(\ell_{\Sigma} \|\boldsymbol{\theta}_k^0 - \bar{\boldsymbol{\theta}}^0\|^2 + \tau_{\ell} \|\boldsymbol{\theta}_k^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right) \\
 & + \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m \sum_{s=0}^{t-1} |\ell_{jk}^{(t-s)}| \frac{\epsilon}{2} \left((\ell_{\Sigma} + 8s\tau_{\ell}) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 + 2\tau_{\ell}\nu^2 \right) \\
 & + \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m \sum_{s=0}^{t-1} |\ell_{jk}^{(t-s)}| \frac{1}{2\epsilon} \left(\ell_{\Sigma} \|\Delta\boldsymbol{\theta}_k^s\|^2 + 16s\tau_{\ell} \|\boldsymbol{\theta}_k^{s+\frac{1}{2}} - \hat{\boldsymbol{\theta}}\|^2 + 16s\tau_{\ell} \|\boldsymbol{\theta}_k^s - \hat{\boldsymbol{\theta}}\|^2 + 8\tau_{\ell}\nu^2 \right). \\
 \leq & mc_m \rho^t \cdot \frac{\epsilon}{2} (\ell_{\Sigma} + 8s\tau_{\ell}) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \hat{\boldsymbol{\Theta}} \right\|_F^2 + m^2 c_m \rho^t \cdot \tau_{\ell} \nu^2 \cdot \epsilon \\
 & + mc_m \rho^t \cdot \frac{1}{2\epsilon} \ell_{\Sigma} \|\boldsymbol{\Theta}_{\perp}^0\|_F^2 + mc_m \rho^t \cdot \frac{1}{2\epsilon} \tau_{\ell} \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \\
 & + \sum_{s=0}^{t-1} mc_m \rho^{t-s} \frac{\epsilon}{2} (\ell_{\Sigma} + 8s\tau_{\ell}) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \hat{\boldsymbol{\Theta}} \right\|_F^2 + \sum_{s=0}^{t-1} m^2 c_m \rho^{t-s} \tau_{\ell} \nu^2 \cdot \epsilon \\
 & + \sum_{s=0}^{t-1} mc_m \rho^{t-s} \frac{1}{2\epsilon} \ell_{\Sigma} \|\Delta\boldsymbol{\Theta}^s\|_F^2 + \sum_{s=0}^{t-1} mc_m \rho^{t-s} \frac{1}{\epsilon} 8s\tau_{\ell} \|\boldsymbol{\Theta}^{s+\frac{1}{2}} - \hat{\boldsymbol{\Theta}}\|_F^2 \\
 & + \sum_{s=0}^{t-1} mc_m \rho^{t-s} \frac{1}{\epsilon} 8s\tau_{\ell} \|\boldsymbol{\Theta}^s - \hat{\boldsymbol{\Theta}}\|_F^2 + \sum_{s=0}^{t-1} m^2 c_m \rho^{t-s} \frac{1}{\epsilon} 4\tau_{\ell} \nu^2 \\
 \leq & \frac{\rho \cdot mc_m \epsilon}{1 - \rho} (\ell_{\Sigma} + 8s\tau_{\ell}) \left\| \boldsymbol{\Theta}^{t+\frac{1}{2}} - \hat{\boldsymbol{\Theta}} \right\|_F^2 + \frac{mc_m \ell_{\Sigma}}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta\boldsymbol{\Theta}^s\|_F^2 \\
 & + mc_m \frac{1}{\epsilon} 8s\tau_{\ell} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^{s+\frac{1}{2}} - \hat{\boldsymbol{\Theta}}\|_F^2 + mc_m \frac{1}{\epsilon} 8s\tau_{\ell} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^s - \hat{\boldsymbol{\Theta}}\|_F^2 \\
 & + \frac{\rho \cdot m^2 c_m}{1 - \rho} 4\tau_{\ell} \nu^2 \cdot (\epsilon + \epsilon^{-1}) + \frac{1}{2\epsilon} mc_m \rho^t \left(\ell_{\Sigma} \|\boldsymbol{\Theta}_{\perp}^0\|_F^2 + \tau_{\ell} \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \bar{\boldsymbol{\theta}}^0\|_1^2 \right).
 \end{aligned}$$

Appendix C.3 Bounding T_3 in (82)

We bound first $T_{3,1}$ and $T_{3,2}$. Invoking the definition of $\mathbf{g}_{j,\perp}^0$ we have

$$\begin{aligned}
 T_{3,1} & = \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}} \right)^{\top} \left(\nabla \mathcal{L}_j(\boldsymbol{\theta}_j^0) - \frac{1}{m} \sum_{k=1}^m \nabla \mathcal{L}_k(\boldsymbol{\theta}_k^0) \right) \right| \\
 & = \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}} \right)^{\top} \left(\frac{1}{n} \mathbf{X}_j^{\top} (\mathbf{X}_j \boldsymbol{\theta}_j^0 - \mathbf{y}_j) - \frac{1}{mn} \sum_{k=1}^m \mathbf{X}_k^{\top} (\mathbf{X}_k \boldsymbol{\theta}_k^0 - \mathbf{y}_k) \right) \right| \\
 & = \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}} \right)^{\top} \left(\frac{1}{n} \mathbf{X}_j^{\top} (\mathbf{X}_j \boldsymbol{\theta}_j^0 - \mathbf{X}_j \boldsymbol{\theta}^* + \mathbf{X}_j \boldsymbol{\theta}^* - \mathbf{y}_j) - \frac{1}{mn} \sum_{k=1}^m \mathbf{X}_k^{\top} (\mathbf{X}_k \boldsymbol{\theta}_k^0 - \mathbf{X} \boldsymbol{\theta}^* + \mathbf{X} \boldsymbol{\theta}^* - \mathbf{y}_k) \right) \right| \\
 & \leq \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \hat{\boldsymbol{\theta}} \right)^{\top} \left(\frac{1}{n} \mathbf{X}_j^{\top} \mathbf{X}_j (\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*) - \frac{1}{mn} \sum_{k=1}^m \mathbf{X}_k^{\top} \mathbf{X}_k (\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*) \right) \right|
 \end{aligned}$$

$$+ \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right)^\top \nabla \mathcal{L}_j(\boldsymbol{\theta}^*) \right| + \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right)^\top \nabla \mathcal{L}(\boldsymbol{\theta}^*) \right|.$$

Therefore,

$$\begin{aligned}
T_{3,1} &\leq \frac{\epsilon}{2} \left(\ell_\Sigma \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + \tau_\ell \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_1^2 \right) + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|_1^2 \right) \\
&\quad + \frac{1}{m} \sum_{k=1}^m \left(\frac{\epsilon}{2} \left(\ell_\Sigma \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + \tau_\ell \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_1^2 \right) + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*\|_1^2 \right) \right) \\
&\quad + \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_1 \underbrace{(\|\nabla \mathcal{L}_j(\boldsymbol{\theta}^*)\|_\infty + \|\nabla \mathcal{L}(\boldsymbol{\theta}^*)\|_\infty)}_{c_g} \\
&\stackrel{(a)}{\leq} \frac{\epsilon}{2} \left((\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 2\tau_\ell \nu^2 \right) + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|_1^2 \right) \\
&\quad + \frac{1}{m} \sum_{k=1}^m \left(\frac{\epsilon}{2} \left((\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 2\tau_\ell \nu^2 \right) \right) \tag{91} \\
&\quad + \frac{1}{2\epsilon} \frac{1}{m} \left(\ell_\Sigma \|\boldsymbol{\Theta}^0 - \boldsymbol{\Theta}^*\|_F^2 + \tau_\ell \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*\|_1^2 \right) + c_g \left(2\sqrt{s} \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\| + \nu \right) \\
&\stackrel{(b)}{\leq} \epsilon \left((\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 2\tau_\ell \nu^2 \right) \\
&\quad + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|_1^2 + \frac{1}{m} \ell_\Sigma \|\boldsymbol{\Theta}^0 - \boldsymbol{\Theta}^*\|_F^2 + \tau_\ell \frac{1}{m} \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*\|_1^2 \right) \\
&\quad + \epsilon \ell_\Sigma \cdot \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu \\
&\leq \epsilon \left((2\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + 2\tau_\ell \nu^2 \right) \\
&\quad + \underbrace{\frac{1}{\epsilon} \frac{1}{2} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|^2 + \tau_\ell \|\boldsymbol{\theta}_j^0 - \boldsymbol{\theta}^*\|_1^2 + \frac{1}{m} \ell_\Sigma \|\boldsymbol{\Theta}^0 - \boldsymbol{\Theta}^*\|_F^2 + \tau_\ell \frac{1}{m} \sum_{k=1}^m \|\boldsymbol{\theta}_k^0 - \boldsymbol{\theta}^*\|_1^2 \right)}_{A_3} \\
&\quad + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu,
\end{aligned}$$

where (a) follows from Lemma 10 while in (b) used the bound $2 \cdot (c_g \sqrt{s}) \cdot \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\| \leq \epsilon \ell_\Sigma \cdot \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1}$.

Similarly, we can bound $T_{3,2}$ as

$$\begin{aligned}
 T_{3,2} &= \left| \left(\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}} \right)^\top \left(\frac{1}{n} \mathbf{X}_j^\top \mathbf{X}_j \right) (\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s) \right| \\
 &\leq \frac{\epsilon}{2} \left(\ell_\Sigma \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 + \tau_\ell \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|_1^2 \right) + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s\|^2 + \tau_\ell \|\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s\|_1^2 \right) \\
 &\leq \frac{\epsilon}{2} (\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 \\
 &\quad + \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s\|^2 + 16s\tau_\ell \|\boldsymbol{\theta}_j^{s+1} - \widehat{\boldsymbol{\theta}}\|^2 + 16s\tau_\ell \|\boldsymbol{\theta}_j^s - \widehat{\boldsymbol{\theta}}\|^2 \right) + \tau_\ell \nu^2 (\epsilon + 4\epsilon^{-1})
 \end{aligned} \tag{92}$$

Using (91) and (92) along with $\|(\mathbf{W} - \mathbf{J})^t\|_\infty \leq c_m \rho^t$, we can bound T_3 as follows:

$$T_3 = \sum_{i=1}^m (\mathbf{g}^t - \mathbf{g}_i^t)^T (\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}) \leq \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \cdot T_{3,1} + \sum_{s=0}^{t-1} \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t-s)}| \cdot T_{3,2}$$

and thus

$$\begin{aligned}
 T_3 &\leq \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| \left(\epsilon (2\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 \right) \\
 &\quad + \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t)}| (\epsilon \cdot 2\tau_\ell \nu^2 + \epsilon^{-1} A_3 + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) \\
 &\quad + \sum_{s=0}^{t-1} \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t-s)}| \left(\frac{\epsilon}{2} (\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\theta}_i^{t+\frac{1}{2}} - \widehat{\boldsymbol{\theta}}\|^2 \right) \\
 &\quad + \sum_{s=0}^{t-1} \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t-s)}| \frac{1}{2\epsilon} \left(\ell_\Sigma \|\boldsymbol{\theta}_j^{s+1} - \boldsymbol{\theta}_j^s\|^2 + 16s\tau_\ell \|\boldsymbol{\theta}_j^{s+1} - \widehat{\boldsymbol{\theta}}\|^2 + 16s\tau_\ell \|\boldsymbol{\theta}_j^s - \widehat{\boldsymbol{\theta}}\|^2 \right) \\
 &\quad + \sum_{s=0}^{t-1} \sum_{i=1}^m \sum_{j=1}^m |\ell_{ij}^{(t-s)}| \tau_\ell \nu^2 (\epsilon + 4\epsilon^{-1}) \\
 &\leq c_m \rho^t \epsilon (2\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 + m c_m \rho^t (\epsilon \cdot 2\tau_\ell \nu^2 + \epsilon^{-1} A_3 + s c_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) \\
 &\quad + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{\epsilon}{2} (\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 \\
 &\quad + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{1}{2\epsilon} \ell_\Sigma \|\boldsymbol{\Theta}^{s+1} - \boldsymbol{\Theta}^s\|_F^2 + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{1}{\epsilon} 8s\tau_\ell \|\boldsymbol{\Theta}^{s+1} - \widehat{\boldsymbol{\Theta}}\|^2 \\
 &\quad + \sum_{s=0}^{t-1} c_m \rho^{t-s} \frac{1}{\epsilon} 8s\tau_\ell \|\boldsymbol{\Theta}^s - \widehat{\boldsymbol{\Theta}}\|^2 + \sum_{s=0}^{t-1} m c_m \rho^{t-s} \tau_\ell \nu^2 (\epsilon + 4\epsilon^{-1}) \\
 &\leq \frac{3}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1 - \rho} (2\ell_\Sigma + 8s\tau_\ell) \|\boldsymbol{\Theta}^{t+\frac{1}{2}} - \widehat{\boldsymbol{\Theta}}\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\boldsymbol{\Theta}^{s+1} - \boldsymbol{\Theta}^s\|_F^2
 \end{aligned}$$

$$\begin{aligned}
& + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \widehat{\Theta}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
& + mc_m \rho^t (A_3 \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) + \frac{\rho \cdot mc_m}{1-\rho} 4\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) \\
\leq & \frac{3}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1-\rho} (2\ell_\Sigma + 8s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \Theta^s\|_F^2 \\
& + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \widehat{\Theta}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
& + mc_m \rho^t (A_3 \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + c_g \nu) + \frac{\rho \cdot mc_m}{1-\rho} 4\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}).
\end{aligned}$$

Appendix C.4 Bounding δ^t in (82)

Using the bounds on T_1 , T_2 , and T_3 , we can finally bound δ^t :

$$\begin{aligned}
\delta^t &= \sum_{i=1}^m (\nabla \mathcal{L}(\theta_i^t) - \mathbf{y}_i^t)^\top (\theta_i^{t+\frac{1}{2}} - \widehat{\theta}) \\
&\leq T_1 + T_2 + T_3 \\
&\leq \frac{\rho \cdot c_m \epsilon}{1-\rho} (L_\Sigma + 8s\tau_g) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{c_m L_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 \\
&\quad + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m \cdot 8s\tau_g \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
&\quad + \frac{\rho \cdot mc_m}{1-\rho} 4\tau_g \nu^2 (\epsilon + \epsilon^{-1}) \\
&\quad + \frac{\rho \cdot c_m \epsilon}{1-\rho} (\ell_\Sigma + 8s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 \\
&\quad + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
&\quad + \frac{\rho \cdot mc_m}{1-\rho} 4\tau_\ell \nu^2 \cdot (\epsilon + \epsilon^{-1}) \\
&\quad + \frac{3}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1-\rho} (2\ell_\Sigma + 8s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \Theta^s\|_F^2 \\
&\quad + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \widehat{\Theta}\|_F^2 + c_m 8s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
&\quad + mc_m \rho^t (A_3 \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + 1) + \frac{\rho \cdot mc_m}{1-\rho} 4\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) + \frac{\rho \cdot mc_m}{1-\rho} c_g^2 \nu^2.
\end{aligned}$$

$$\delta^t \stackrel{(a)}{\leq} \frac{7}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1-\rho} (2\ell_\Sigma + 8s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2$$

$$\begin{aligned}
 & + \frac{c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 + \frac{c_m \ell_\Sigma}{2\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+1} - \Theta^s\|_F^2 \\
 & + c_m 24\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m 24\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^s - \widehat{\Theta}\|_F^2 \\
 & + mc_m \rho^t (A_3 \cdot \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + 1) + \frac{\rho \cdot mc_m}{1-\rho} 12\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) + \frac{\rho \cdot mc_m}{1-\rho} c_g^2 \nu^2 \\
 & \stackrel{(b)}{\leq} \frac{7}{2} \cdot \frac{\rho \cdot c_m \epsilon}{1-\rho} (2\ell_\Sigma + 8s\tau_\ell) \|\Theta^{t+\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\
 & + \frac{2c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|_F^2 + \frac{2c_m \ell_\Sigma}{\epsilon} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta_\perp^s\|_F^2 \\
 & + c_m 24s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s+\frac{1}{2}} - \widehat{\Theta}\|_F^2 + c_m 24s\tau_\ell \epsilon^{-1} \sum_{s=0}^{t-1} \rho^{t-s} \|\Theta^{s-\frac{1}{2}} - \widehat{\Theta}\|_F^2 \\
 & + mc_m \rho^t (A_3 \cdot \epsilon^{-1} + sc_g^2 \ell_\Sigma^{-1} \epsilon^{-1} + 1) + \frac{\rho \cdot mc_m}{1-\rho} 12\tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) + \frac{\rho \cdot mc_m}{1-\rho} c_g^2 \nu^2,
 \end{aligned}$$

where in (a) we used $L_\Sigma \leq \ell_\Sigma$, $\tau_g \leq \tau_\ell$, and (89); and in (b) we used again (89) and (90). \square

Appendix D. Proof of Proposition 13

Using Assumptions 3 and 4 on (83) we have

$$\|\Theta_\perp^t\| \leq \rho^t \|\Theta_\perp^0\| + \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|.$$

Squaring both sides gives

$$\begin{aligned}
 \|\Theta_\perp^t\|^2 & \leq \left(\rho^t \|\Theta_\perp^0\| + \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\| \right)^2 \leq 2\rho^{2t} \|\Theta_\perp^0\|^2 + 2 \left(\sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\| \right)^2 \\
 & \leq 2\rho^{2t} \|\Theta_\perp^0\|^2 + \frac{2\rho}{1-\rho} \cdot \sum_{s=0}^{t-1} \rho^{t-s} \|\Delta \Theta^s\|^2,
 \end{aligned}$$

where the last inequality follows from Cauchy-Schwartz inequality. \square

Appendix E. Proof of Lemma 16

Recall $G_{\text{net}}(z) = \frac{\rho}{z-\rho}$. Since (49b) is quadratic in $G_{\text{net}}(z)$, we obtain the following feasible range of $G_{\text{net}}(z)$:

$$\begin{aligned} G_{\text{net}}(z) &\leq \frac{\frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} - \sqrt{\left(\frac{C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon}\right)^2 + \frac{4C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \frac{2\rho}{1-\rho} \left(1 - \frac{L_\Sigma}{\gamma}\right)}}{-\frac{2C_2}{\gamma} \frac{c_m \ell_\Sigma}{\epsilon} \frac{2\rho}{1-\rho}} \\ &= \frac{1-\rho}{4\rho} \left(\sqrt{1 + \frac{8\rho}{1-\rho} \left(1 - \frac{L_\Sigma}{\gamma}\right) \left(\frac{\gamma}{C_2} \frac{\epsilon}{c_m \ell_\Sigma}\right)} - 1 \right) \\ &\stackrel{(50)}{=} \frac{1-\rho}{4\rho} \left(\sqrt{1 + (\gamma - L_\Sigma) \left(\frac{\mu_\Sigma}{\ell_\Sigma^2} \frac{2\rho}{C_2^2 c_m^2}\right)} - 1 \right). \end{aligned}$$

Therefore, condition (49b) is fulfilled for all z in the form

$$z = \rho + \frac{\rho}{G_{\text{net}}(z)} \geq \rho + \frac{4\rho^2}{1-\rho} \left(\sqrt{1 + (\gamma - L_\Sigma) \left(\frac{\mu_\Sigma}{\ell_\Sigma^2} \frac{2\rho}{C_2^2 c_m^2}\right)} - 1 \right)^{-1}.$$

□

Appendix F. Proof of Lemma 17

We substitute $G_{\text{net}}(z) = \frac{\rho}{z-\rho}$ into (54) and solve the equation

$$z = \frac{A + \frac{\rho}{z-\rho}}{D - \frac{\rho}{z-\rho}}. \quad (93)$$

Let Δ denote the discriminant of the second order equation in z (54). Since

$$\Delta = (D\rho + \rho + A)^2 - 4D(A-1)\rho = (D\rho + \rho - A)^2 + 4(A+D)\rho \geq 0,$$

Eq. (93) has two real roots ($z_1 \leq z_2$)

$$z_{1,2} = \frac{(D+1)\rho + A \pm \sqrt{(D\rho + \rho + A)^2 - 4D(A-1)\rho}}{2D}.$$

By inspection, we have that $z_1 \geq 0$. Hence, (54) is satisfied for all $z \in (0, z_1] \cup [z_2, 1)$.

Next we show $z_1 \leq \rho$ and $z_2 \geq \rho$. As $z > \rho$ must be satisfied, $z \in [z_2, 1)$. Suppose on the contrary that $z_1 > \rho$, which implies that

$$\begin{aligned} &\rho + A - D\rho > \sqrt{(D\rho + \rho + A)^2 - 4D(A-1)\rho} \\ \Leftrightarrow &(\rho + A - D\rho)^2 > (\rho + A + D\rho)^2 - 4D(A-1)\rho \quad \text{and} \quad \rho + A - D\rho > 0. \end{aligned}$$

After simplification, the first inequality becomes $4D\rho^2 + 4D\rho < 0$, which is impossible. Therefore we have proved that $z_1 \leq \rho$.

As both z_1 and z_2 are nonnegative, we can simply upperbound z_2 as $z_2 \leq z_1 + z_2$. This completes the proof. □

Appendix G. Proof of Proposition 18

Substituting the expression of γ given by (57) into (51) we obtain:

$$\begin{aligned}
 z_\rho &= \rho + \sqrt{\rho}(1-\rho)^3 \frac{C_2^2}{C_3} \left(\sqrt{1 + \frac{\rho\sqrt{\rho}}{(1-\rho)^4} \left(\frac{2C_3}{C_2^2} \right)} + 1 \right) \\
 &\leq \rho + \sqrt{\rho}(1-\rho)^3 \frac{C_2^2}{C_3} \left(\frac{\sqrt{\rho\sqrt{\rho}}}{(1-\rho)^2} \cdot \frac{\sqrt{2C_3}}{C_2} + 2 \right) \\
 &\leq \rho + \frac{\sqrt{2}C_2}{\sqrt{C_3}} \rho(1-\rho) + \frac{2C_2^2}{C_3} \sqrt{\rho}(1-\rho)^3 \triangleq \bar{z}_\rho.
 \end{aligned} \tag{94}$$

The rest of the proof bounds z_σ based on Lemma 17. Substituting γ into the expression of A and D we get

$$\begin{aligned}
 \frac{A}{D} &= \frac{L_\Sigma - \mu_\Sigma + C_1 s(\tau_\mu + \tau_g) + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}}{L_\Sigma - C_1 s\tau_g - \frac{\rho}{2} \mu_\Sigma - \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}} \\
 &= \frac{1 - \kappa^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}}{1 - C_1 s\tau_g/L_\Sigma - \frac{\rho}{2} \kappa^{-1} - \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}}.
 \end{aligned} \tag{95}$$

Dividing the numerator and denominator by $1 - C_1 s\tau_g/L_\Sigma$ and by the definition of σ_0 :

$$\begin{aligned}
 \frac{A}{D} &= \frac{\sigma_0 + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} (1 - C_1 s\tau_g/L_\Sigma)^{-1}}{1 + \left(-\frac{\rho}{2} \kappa^{-1} - \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} \right) (1 - C_1 s\tau_g/L_\Sigma)^{-1}} \\
 &\stackrel{A \leq D}{\leq} \frac{\sigma_0 + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} (1 - C_1 s\tau_g/L_\Sigma)^{-1} + \left(\frac{\rho}{2} \kappa^{-1} + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right) (1 - C_1 s\tau_g/L_\Sigma)^{-1}}{1 + \left(C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} \right) (1 - C_1 s\tau_g/L_\Sigma)^{-1}} \\
 &\leq \sigma_0 + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} (1 - C_1 s\tau_g/L_\Sigma)^{-1} + \frac{\rho}{2} \left(\kappa^{-1} + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right) (1 - C_1 s\tau_g/L_\Sigma)^{-1} \\
 &= \sigma_0 + \sqrt{\rho} \cdot \frac{C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2}{(1-\rho)^4} + \frac{1}{2} \left(\kappa^{-1} + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right)}{1 - C_1 s\tau_g/L_\Sigma} \\
 &\stackrel{\rho \leq 1/2}{\leq} \sigma_0 + \sqrt{\rho} \cdot \frac{16C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{1}{2} \left(\kappa^{-1} + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right)}{1 - C_1 s\tau_g/L_\Sigma},
 \end{aligned} \tag{96}$$

where for $A \leq D$ to hold, we require

$$\left(1 - \frac{\rho}{2} \right) \mu_\Sigma \geq C_1 s\tau_g + C_1 s(\tau_\mu + \tau_g) + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell. \tag{97}$$

Similarly, we can bound $\rho + \rho/D$ as

$$\rho + \frac{\rho}{D} = \rho + \frac{\rho \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_2^2 c_m^2}{1-\rho} s\tau_\ell}{1 - C_1 s\tau_g/L_\Sigma + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} - \left\{ \frac{\rho}{2} \kappa^{-1} + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right\}}$$

$$\begin{aligned}
& \stackrel{\rho \leq D}{\leq} \rho + \frac{\rho \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_2^2 c_m^2}{1-\rho} s\tau_\ell + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - C_1 s\tau_g/L_\Sigma - \frac{\rho}{2} \kappa^{-1} + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}} \\
& \stackrel{\rho \kappa^{-1} \leq 1}{\leq} \rho + \frac{\rho \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_2^2 c_m^2}{1-\rho} s\tau_\ell + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{\frac{1}{2} - C_1 s\tau_g/L_\Sigma + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4}},
\end{aligned} \tag{98}$$

where $\rho \leq D$ requires

$$L_\Sigma + C_3 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} \geq \rho \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2^2 c_m^2}{1-\rho} s\tau_\ell + C_1 s\tau_g + \frac{\rho}{2} \mu_\Sigma + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell.$$

Using the fact that $\rho \leq 1$ and $L_\Sigma \geq \mu_\Sigma$, it suffices

$$\mu_\Sigma \geq \rho \frac{\ell_\Sigma}{\mu_\Sigma} \frac{4C_2^2 c_m^2}{1-\rho} s\tau_\ell + 2C_1 s\tau_g + \rho \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell. \tag{99}$$

Note that (60) implies both (97) and (99).

Finally, we further simplify the bound in (98) as follows. Multiplying both the numerator and denominator by $(1-\rho)^4$ we get

$$\rho + \frac{\rho}{D} \leq \rho + \frac{\rho \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 2C_2^2 c_m^2 s\tau_\ell (1-\rho)^3 + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell (1-\rho)^4}{(\frac{1}{2} - C_1 s\tau_g/L_\Sigma)(1-\rho)^4 + C_3 c_m^2 \sqrt{\rho} \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma}}.$$

Define function

$$f(\rho) = \left(\frac{1}{2} - C_1 s\tau_g/L_\Sigma \right) (1-\rho)^4 + C_3 c_m^2 \sqrt{\rho} \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma}.$$

Differentiating f with respect to ρ yields

$$f'(\rho) = -4 \left(\frac{1}{2} - C_1 s\tau_g/L_\Sigma \right) (1-\rho)^3 + C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \frac{1}{2\sqrt{\rho}} > 0,$$

for $C_3 \geq 4$, due to $c_m \geq 1$, $\ell_\Sigma \geq \mu_\Sigma$, and $\ell_\Sigma \geq L_\Sigma$. Therefore, $f(\rho)$ is monotonically increasing in $(0, 1)$ and it holds

$$\rho + \frac{\rho}{D} \leq \rho + \rho \cdot \frac{\frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - 2C_1 s\tau_g/L_\Sigma}. \tag{100}$$

Combining (96) and (100), z_σ can be readily bounded as

$$\begin{aligned}
z_\sigma & \leq \sigma_0 + \sqrt{\rho} \cdot \frac{16C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{1}{2} \left(\kappa^{-1} + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right)}{1 - C_1 s\tau_g/L_\Sigma} + \rho + \rho \cdot \frac{\frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - 2C_1 s\tau_g/L_\Sigma} \\
& \leq \sigma_0 + \sqrt{\rho} \cdot \frac{18C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - 2C_1 s\tau_g/L_\Sigma},
\end{aligned} \tag{101}$$

where in the last inequality we have used the fact that $C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} \geq 1$, $\rho \leq 1$ and $\kappa^{-1} \leq 1$. To make the bound not vacuous, we need to assume $1 - 2C_1 s\tau_g/L_\Sigma > 0$ [cf. (59)]. Combining (94) and (101) completes the proof. \square

Appendix H. Proof of Corollary 19

Recall the definition of σ_0 as in (58):

$$\sigma_0 \triangleq \frac{1 - \kappa^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - C_1 s\tau_g/L_\Sigma}, \quad \kappa \triangleq \frac{L_\Sigma}{\mu_\Sigma}.$$

By requiring

$$\sqrt{\rho} \leq \left\{ 2\kappa \left(18C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell \right) \right\}^{-1} \quad (102)$$

we have

$$\sigma_0 + \sqrt{\rho} \cdot \frac{18C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2\mu_\Sigma}{\ell_\Sigma L_\Sigma} s\tau_\ell}{1 - 2C_1 s\tau_g/L_\Sigma} \leq \frac{1 - (2\kappa)^{-1} + C_1 s(\tau_\mu + \tau_g)/L_\Sigma}{1 - 2C_1 s\tau_g/L_\Sigma}.$$

Choosing $C_3 = 4$ we obtain an upper bound of RHS of (62) given by (67), under network connectivity condition .

Thus far, we can conclude that (49a)-(49b) are satisfied by any z and γ such that $z > \bar{z}_{\text{up}}$, with γ and \bar{z}_{up} defined in (64) and (67), respectively, as long as the following conditions hold [note that conditions (44) and (59) are already implied by (64) and (60), respectively]:

- step size lower bound (45), with ϵ chosen according to (50);
- RSC/RSM tolerance parameters condition (60);
- and network connectivity condition (102).

Finally, we simplify these conditions. First, (45), with ϵ given by (50), is a consequence of (60) and (64). In fact, substituting the expression of γ and ϵ into (45), (45) reads

$$\begin{aligned} L_\Sigma + 4 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} &> C_1 s\tau_g + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} (\ell_\Sigma + s\tau_\ell) = C_1 s\tau_g + \frac{\rho}{2} \mu_\Sigma + \frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma} (s\tau_\ell) \\ \Leftrightarrow \quad 2L_\Sigma - \rho \mu_\Sigma + 8 \frac{\ell_\Sigma^2}{\mu_\Sigma} \frac{c_m^2 \sqrt{\rho}}{(1-\rho)^4} &> 2C_1 s\tau_g + \rho \frac{\mu_\Sigma}{\ell_\Sigma} (s\tau_\ell). \end{aligned}$$

Since $\rho \leq 1$ and $\mu_\Sigma \leq L_\Sigma$, clearly (60) implies (45). Second, we claim that (60) is implied by (102) and (65). Using the facts that $c_m \geq 1$ and $\ell_\Sigma/\mu_\Sigma \geq 1$, (102) implies the following: $\rho \leq 1/4$ and

$$2\sqrt{\rho} \left(\frac{\ell_\Sigma}{\mu_\Sigma^2} \cdot 4C_2^2 c_m^2 s\tau_\ell + \frac{2}{\ell_\Sigma} s\tau_\ell \right) \leq 1. \quad (103)$$

Consequently, we can upperbound the second and third summands on the RHS of (60) as

$$\begin{aligned} &\frac{\ell_\Sigma}{\mu_\Sigma} \frac{8C_2^2 c_m^2 \rho}{(1-\rho)^2} s\tau_\ell + \rho \cdot \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell \\ &\leq \frac{\ell_\Sigma}{\mu_\Sigma} \frac{8C_2^2 c_m^2 \rho}{9/16} s\tau_\ell + \rho \cdot \frac{\mu_\Sigma}{\ell_\Sigma} s\tau_\ell \\ &\leq \frac{16}{9} \cdot \rho \left(\frac{\ell_\Sigma}{\mu_\Sigma} 8C_2^2 c_m^2 s\tau_\ell + \frac{4\mu_\Sigma}{\ell_\Sigma} s\tau_\ell \right) \leq \frac{16}{9} \sqrt{\rho} \mu_\Sigma \leq \frac{8}{9} \mu_\Sigma. \end{aligned}$$

This together with (65) implies (60). \square

Appendix I. Proof of Proposition 20

We first bound $B(z)$. Substituting the γ [cf. (64)] and ϵ [cf. (50)] into (47):

$$\begin{aligned}
B(z) &\leq \frac{C_2 s \tau_\ell}{L_\Sigma} \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2 c_m^2}{1-\rho} \frac{\rho}{z-\rho} \|\Theta^0 - \hat{\Theta}\|^2 \\
&\quad + \frac{C_2}{L_\Sigma} m c_m \left(A_3 \cdot \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + s c_g^2 \frac{1}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + c_g \nu \right) \frac{\rho}{z-\rho} + \|\Theta^{1/2} - \hat{\Theta}\|^2 \\
&\stackrel{(68)}{\leq} \frac{s \tau_\ell}{L_\Sigma} \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_3 c_m^2 \sqrt{\rho}}{(1-\rho)^4} \|\Theta^0 - \hat{\Theta}\|^2 \\
&\quad + \frac{1}{L_\Sigma} m c_m \sqrt{\rho} \left(A_3 \cdot \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + s c_g^2 \frac{1}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + c_g \nu \right) \frac{C_3}{(1-\rho)^3 C_2} + \|\Theta^{1/2} - \hat{\Theta}\|^2.
\end{aligned}$$

Substituting the expression of A_3 [cf. Prop. 12] in the above bound yields

$$\begin{aligned}
B(z) &\leq \sqrt{\rho} \frac{s \tau_\ell}{L_\Sigma} \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_3 c_m^2}{(1-\rho)^4} \|\Theta^0 - \hat{\Theta}\|^2 \\
&\quad + m \sqrt{\rho} \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_3 c_m^2}{(1-\rho)^4} (\ell_\Sigma \|\theta^*\|^2 + \tau_\ell \|\theta^*\|_1^2) \\
&\quad + \frac{1}{L_\Sigma} m c_m \sqrt{\rho} \left(s c_g^2 \frac{1}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + c_g \nu \right) \frac{C_3}{(1-\rho)^3 C_2} + \|\Theta^{1/2} - \hat{\Theta}\|^2 \\
&\leq \sqrt{\rho} \frac{s \tau_\ell}{L_\Sigma} \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_3 c_m^2}{(1-\rho)^4} \|\Theta^0 - \hat{\Theta}\|^2 \\
&\quad + m \sqrt{\rho} \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_3 c_m^2}{(1-\rho)^4} (\ell_\Sigma \|\theta^*\|^2 + s \tau_\ell \|\theta^*\|^2) \quad (\theta^* \text{ is } s\text{-sparse}) \\
&\quad + \frac{1}{L_\Sigma} m c_m \sqrt{\rho} \left(s c_g^2 \frac{1}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} + c_g \nu \right) \frac{C_3}{(1-\rho)^3 C_2} + \|\Theta^{1/2} - \hat{\Theta}\|^2.
\end{aligned}$$

Since (66) is equivalent to

$$2\sqrt{\rho} \kappa \left(18C_3 c_m^2 \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} + \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot 4C_2^2 c_m^2 s \tau_\ell + \frac{2\mu_\Sigma}{\ell_\Sigma L_\Sigma} s \tau_\ell \right) \leq 1,$$

with $C_3 = 4$, together with the fact that $\kappa \geq 1$, we have

$$\sqrt{\rho} \frac{\ell_\Sigma^2}{\mu_\Sigma L_\Sigma} C_3 c_m^2 \leq \frac{1}{36} \quad \text{and} \quad \sqrt{\rho} \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \cdot C_2^2 c_m^2 s \tau_\ell \leq \frac{1}{8}.$$

In addition, from (66) we get $\rho \leq 1/4$, and thus we obtain (69).

We bound now $\bar{\Delta}_{\text{stat}}$. Using the expression of γ [cf. (64)] and ϵ [cf. (50)], we have

$$\begin{aligned}
\bar{\Delta}_{\text{stat}} &= \frac{C_2 \rho \cdot c_m}{\gamma} \tau_\ell \nu^2 (\epsilon + \epsilon^{-1}) + \frac{C_1(\tau_\mu + \tau_g)}{\gamma} \nu^2 \\
&\leq \frac{C_2 \rho \cdot c_m}{L_\Sigma} \tau_\ell \nu^2 \left(\frac{\mu_\Sigma}{\ell_\Sigma} \frac{1-\rho}{2C_2 c_m} + \frac{\ell_\Sigma}{\mu_\Sigma} \frac{2C_2 c_m}{1-\rho} \right) + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2 \\
&= \left(\frac{\rho}{2} \frac{\mu_\Sigma}{\ell_\Sigma L_\Sigma} + \rho \frac{\ell_\Sigma}{\mu_\Sigma L_\Sigma} \frac{2C_2^2 c_m^2}{(1-\rho)^2} \right) \tau_\ell \nu^2 + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2 \\
&\leq \frac{\rho}{2L_\Sigma} \left(1 + \frac{\ell_\Sigma}{\mu_\Sigma} \frac{4C_2^2 c_m^2}{(1-\rho)^2} \right) \tau_\ell \nu^2 + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2 \leq \Delta_{\text{stat}} \quad (\text{given } C_2 \geq 1).
\end{aligned}$$

□

Appendix J. Proof of Lemma 23

We first bound $\max_{j \in [m]} \|\mathbf{X}_j^\top \mathbf{n}_j / n\|_\infty$. Each column of \mathbf{X}_j is an n -dimensional Gaussian random vector with independent entries. The maximum variance of these entries is bounded by $\zeta = \max_{j \in [d]} \Sigma_{jj}$. Since \mathbf{X}_j is independent of \mathbf{n}_j and the elements of \mathbf{n}_j are i.i.d. σ^2 -sub-Gaussian, each element of $\mathbf{X}_j^\top \mathbf{n}_j$ is the sum of n independent sub-exponential random variables with sub-exponential norm no larger than $\sqrt{\zeta}\sigma$. Applying the Bernstein's inequality and the union bound we get

$$\mathbb{P} \left(\max_{j \in [m]} \left\| \frac{\mathbf{X}_j^\top \mathbf{n}_j}{n} \right\|_\infty \geq t \right) \leq 2 \exp \left(-c_3 \min \left\{ \frac{t^2}{\zeta \sigma^2}, \frac{t}{\sqrt{\zeta} \sigma} \right\} n + \log md \right),$$

for some $c_3 > 0$. When $\log md \geq (c_3/2) \cdot n$, we choose $t = (2/c_3) \sqrt{\zeta} \sigma \cdot \log md / n \geq \sqrt{\zeta} \sigma$, which leads to

$$\mathbb{P} \left(\max_{j \in [m]} \left\| \frac{\mathbf{X}_j^\top \mathbf{n}_j}{n} \right\|_\infty \geq \sqrt{\zeta} \sigma \cdot \frac{2 \log md}{c_3 n} \right) \leq 2 \exp(-\log md).$$

On the other hand, when $\log md < (c_3/2) \cdot n$, we choose $t = \sqrt{\zeta} \sigma \cdot \sqrt{(2/c_3) \log md / n} < \sqrt{\zeta} \sigma$, which leads to

$$\mathbb{P} \left(\max_{j \in [m]} \left\| \frac{\mathbf{X}_j^\top \mathbf{n}_j}{n} \right\|_\infty \geq \sqrt{\zeta} \sigma \cdot \sqrt{\frac{2 \log md}{c_3 n}} \right) \leq 2 \exp(-\log md).$$

Combining the two cases we obtain (75). Applying the same procedure to $\|\nabla \mathcal{L}(\boldsymbol{\theta}^*)\|_\infty$, proves (76). □

Appendix K. Proof of Proposition 22

The proof is a modification of (Raskutti et al., 2010). Next, we only highlight the key differences, for completeness. Define the set

$$V(r) = \{r \in \mathbb{R}^d \mid \|\Sigma^{\frac{1}{2}} v\|_2 = 1, \|v\|_1 \leq r\},$$

for fixed radius $r > 0$. Accordingly, define the random variable

$$M(r, \mathbf{X}_i) = \sup_{v \in V(r)} \frac{\|\mathbf{X}_i v\|}{\sqrt{n}}.$$

Following similar steps as in (Raskutti et al., 2010, Lemma 1), and using the Sudakov-Fernique inequality for Gaussian processes, we can bound the expectation of $M(r, \mathbf{X}_i)$ as

$$\mathbb{E} M(r, \mathbf{X}_i) \leq 1 + 3\rho(\Sigma) \sqrt{\frac{\log d}{n}} r \triangleq t(r), \quad (104)$$

where $\rho^2(\Sigma) = \max_j \Sigma_{jj}$. Then applying the concentration inequality in the same way as in (Raskutti et al., 2010, Lemma 2) we obtain

$$\mathbb{P}(M(r, \mathbf{X}_i) > t(r) + \sqrt{mt(r)}) \leq 2 \exp(-Nt(r)^2/2) \leq 2 \exp(-n((\sqrt{m}+1)t(r))^2/8). \quad (105)$$

Finally, define the event

$$\mathcal{T}_i = \{\exists v \in \mathbb{R}^d \text{ s.t. } \|\Sigma^{1/2}v\|_2 = 1 \text{ and } \|\mathbf{X}_i v\|_2/\sqrt{n} > 2(\sqrt{m}+1)t(\|v\|_1)\}.$$

Applying (Raskutti et al., 2010, Lemma 3) with

$$f(v, \mathbf{X}_i) = \|\mathbf{X}_i v\|_2/\sqrt{n}, \quad h(v) = \|v\|_1, \quad g(r) = (\sqrt{m}+1)t(r), \quad a_n = n/8 \quad \text{and} \quad \mu = \sqrt{m},$$

we have

$$\mathbb{P}(\mathcal{T}_i^c) \geq 1 - c'_2 \exp(-c'_0 N),$$

for some $c'_0 \geq 1/2$.

Let us use now the union bound:

$$\mathbb{P}(\cup_i \mathcal{T}_i) \leq \sum_{i=1}^m \mathbb{P}(\mathcal{T}_i) \leq c'_2 m \exp(-c'_0 N) = c'_2 \exp(-c'_0 N + \log m).$$

Therefore

$$\frac{\|\mathbf{X}_i v\|}{\sqrt{n}} \leq 2(\sqrt{m}+1)\|\Sigma^{1/2}v\| + 6(\sqrt{m}+1)\rho(\Sigma)\sqrt{\frac{\log d}{n}}\|v\|_1, \quad \forall v \in \mathbb{R}^d, \quad \forall i \in [m], \quad (106)$$

with probability larger than $1 - c'_2 \exp(-c'_0 N + \log m)$. Since $N/2 \geq m/2 > \log m$, it holds $1 - \exp(-c'_0 N + \log m) \geq 1 - \exp(-c''_0 N)$, for some $c''_0 > 0$. \square

Appendix L. Supplement to the proof of Theorem 9

(i) We show (20), i.e.,

$$s \log d/N < c_5 \cdot \frac{\mu_\Sigma}{\zeta}, \quad \text{and} \quad \rho < (c_6 m^8 \kappa^4)^{-1}$$

implies (15)-(16). Substituting the expression of τ_μ and τ_g it is not hard to verify (20) implies (15): $\mu_\Sigma > 36C_1 s(\tau_\mu + \tau_g)$, with c_5 sufficiently small. Substituting τ_ℓ into (16) we can see a sufficient condition for (16) is

$$\rho \leq C_6 \left(m^3 \kappa^2 + m^4 \kappa^2 \frac{\zeta}{L_\Sigma} \frac{s \log d}{N} \right)^{-2}, \quad (107)$$

for some $C_6 > 0$; (107) is implied by (20) for sufficiently large $c_6 > 0$.

(ii) We use (78)-(79) to further simplify the expressions of the rate expression (13), the statistical error (12), and B in (69).

The rate λ in (13) can be bounded as

$$\lambda \leq \left(1 - (2\kappa)^{-1} + 2C_1 c_1 \frac{\zeta}{L_\Sigma} \frac{s \log d}{N} \right) \left(1 - 2C_1 c_1 \frac{\zeta}{L_\Sigma} \frac{s \log d}{N} \right)^{-1}.$$

The statistical error can be bounded as

$$\begin{aligned}
 \Delta_{\text{stat}} &= \frac{\rho}{2L_\Sigma} \left(\frac{\ell_\Sigma}{\mu_\Sigma} \frac{5C_2^2 c_m^2}{(1-\rho)^2} \right) \tau_\ell \nu^2 + \frac{C_1(\tau_\mu + \tau_g)}{L_\Sigma} \nu^2 \\
 &= \rho \left(\frac{4}{\mu_\Sigma} \frac{5C_2^2 m^2}{(1-\rho)^2} \right) c_1 \zeta \frac{m^2 \log d}{N} \nu^2 + \frac{C_1}{L_\Sigma} \left(2c_1 \zeta \frac{\log d}{N} \right) \nu^2 \\
 &\leq C_5 \left(\frac{m^4 \rho \kappa}{(1-\rho)^2} \frac{\zeta}{L_\Sigma} \frac{\log d}{N} + \frac{\zeta}{L_\Sigma} \frac{\log d}{N} \right) \nu^2 \\
 &\leq c_7 \left(\frac{\zeta}{L_\Sigma} \frac{s \log d}{N} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2,
 \end{aligned} \tag{108}$$

for some $C_5, c_7 > 0$, where in the last inequality we used $c_6 \rho m^8 \kappa^4 < 1$, $\nu^2 = \mathcal{O}(s \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|^2)$ (due to $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq 2\sqrt{s} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|$ (Agarwal et al., 2012, Lemma 5)), and $\rho < (3C_2^2)^{-2}$ implied by (16).

Finally, for B in (69), we have

$$\begin{aligned}
 B &= \frac{C_4}{m} \left(\|\boldsymbol{\Theta}^0 - \hat{\boldsymbol{\Theta}}\|^2 + \|\boldsymbol{\Theta}^*\|^2 + \|\boldsymbol{\Theta}^{1/2} - \hat{\boldsymbol{\Theta}}\|^2 + \frac{mc_m \sqrt{\rho}}{L_\Sigma} \left(\frac{c_m s c_g^2}{\mu_\Sigma} + c_g \nu \right) \right) \\
 &\leq \frac{C_4}{m} \left(\|\boldsymbol{\Theta}^0 - \hat{\boldsymbol{\Theta}}\|^2 + \|\boldsymbol{\Theta}^*\|^2 + \|\boldsymbol{\Theta}^{1/2} - \hat{\boldsymbol{\Theta}}\|^2 \right) \\
 &\quad + C_4 \frac{4\zeta \sigma^2}{\mu_\Sigma L_\Sigma} m s \sqrt{\rho} \left(\frac{4}{c_3} \frac{m \log md}{N} + \left(\frac{4}{c_3} \frac{m \log md}{N} \right)^2 \right) \\
 &\quad + C_4 \frac{12\sqrt{\zeta} \sigma}{L_\Sigma} \sqrt{m \rho s} \left(\sqrt{\frac{4}{c_3} \frac{m \log md}{N}} + \frac{4}{c_3} \frac{m \log md}{N} \right) \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|,
 \end{aligned} \tag{109}$$

where we used again $\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq 2\sqrt{s} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|$. \square

Appendix M. Properties of The z -Transform

Proof of Lemma 14

(i): For the time-shifted sequence $\{a(t)\}_{t=2}^{K+1}$, it holds

$$\begin{aligned}
 \sum_{t=1}^K a(t+1) z^{-t} &= z \sum_{t=1}^K a(t+1) z^{-t-1} \\
 &= z \left(\sum_{t=1}^K a(t) z^{-t} + a(K+1) z^{-K-1} - a(1) z^{-1} \right) \\
 &= z A^K(z) - a(1) + a(K+1) z^{-K} \geq z A^K(z) - a(1).
 \end{aligned}$$

(ii): For the sequence $\{\sum_{s=0}^{t-1} \rho^{t-s} a(s)\}_{t=1}^K$, we have $\forall z \in (\rho, 1)$:

$$\sum_{t=1}^K \sum_{s=0}^{t-1} z^{s-t} \rho^{t-s} a(s) z^{-s} = \sum_{s=0}^{K-1} a(s) z^{-s} \sum_{t=s+1}^K z^{s-t} \rho^{t-s} \leq \frac{\rho}{z-\rho} \cdot (A^K(z) + a(0)).$$

(iii): Similarly, the sequence $\{\sum_{s=0}^{t-1} \rho^{t-s} a(s+1)\}_{t=1}^K$, we have

$$\sum_{t=1}^K \sum_{s=0}^{t-1} \rho^{t-s} a(s+1) z^{-t} = \sum_{s=0}^{K-1} a(s+1) z^{-s-1} \sum_{t=s+1}^K \rho^{t-s} z^{-t+s+1} \leq z \cdot \frac{\rho}{z-\rho} A^K(z).$$

□

Proof of Lemma 15

Since $a(t) \geq 0$, (41) implies that for all $K \geq 1$ and $z \in (\bar{z}, 1)$,

$$a(K) z^{-K} \leq B + c \cdot \sum_{t=1}^K z^{-t}.$$

Therefore, for all $z \in (\bar{z}, 1)$,

$$a(K) \leq B \cdot z^K + c \cdot \sum_{t=1}^K z^{K-t} \leq B \cdot z^K + \frac{c}{1-z}.$$

Letting $z \rightarrow \bar{z}^+$ concludes the proof.

□

References

- Alekh Agarwal, Sahand Negahban, Martin J Wainwright, et al. Fast global convergence of gradient methods for high-dimensional statistical recovery. *The Annals of Statistics*, 40(5):2452–2482, 2012.
- Heather Battey, Jianqing Fan, Han Liu, Junwei Lu, and Ziwei Zhu. Distributed testing and estimation under sparse high dimensional models. *Annals of Statistics*, 46(3):1352, 2018.
- Raphael Berthier, Francis Bach, and Pierre Gaillard. Accelerated gossip in networks of given dimension using jacobi polynomial iterations. *SIAM J. on Mathematics of Data Science*, 1:24–47, 2020.
- Stephen Boyd, Neal Parikh, and Eric Chu. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc, 2011.
- Bernadette Charron-Bost. Geometric bounds for convergence rates of averaging algorithms. *arXiv preprint arXiv:2007.04837*, 2020.
- Annie I. Chen and A. Ozdaglar. A fast distributed proximal-gradient method. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 601–608, October 2012.
- Jianshu Chen and Ali H. Sayed. Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, 60(8):4289–4305, August 2012.

- Xi Chen, Weidong Liu, and Yichen Zhang. Quantile regression under memory constraint. *The Annals of Statistics*, 47(6):3244–3273, 2019.
- Xueying Chen and Min-ge Xie. A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica*, pages 1655–1684, 2014.
- Paolo Di Lorenzo and Gesualdo Scutari. NEXT: In-network nonconvex optimization. *IEEE Transactions on Signal and Information Processing over Networks*, 2(2):120–136, June 2016.
- Jianqing Fan, Dong Wang, Kaizheng Wang, and Ziwei Zhu. Distributed estimation of principal eigenspaces. *Annals of statistics*, 47(6):3009, 2019.
- Jianqing Fan, Yongyi Guo, and Kaizheng Wang. Communication-efficient accurate statistical estimation. *J. of the American Statistical Association*, 0(0):1–11, 2021.
- Hadrien Hendrikx, Lin Xiao, Sebastien Bubeck, Francis Bach, and Laurent Massoulié. Statistically preconditioned accelerated gradient method for distributed optimization. In *International Conference on Machine Learning*, pages 4203–4227. PMLR, 2020.
- Yao Ji, Gesualdo Scutari, Ying Sun, and Harsha Honnappa. Distributed sparse regression via penalization. *J. of Machine Learning Research*, (24):1–62, 2023a.
- Yao Ji, Gesualdo Scutari, Ying Sun, and Harsha Honnappa. Distributed (ATC) gradient descent for high dimension sparse regression. *IEEE Transactions on Information Theory*, (69):5253–5276, 2023b.
- Michael I Jordan, Jason D Lee, and Yun Yang. Communication-efficient distributed statistical inference. *J. of the American Statistical Association*, 2018.
- Ariel Kleiner, Ameet Talwalkar, Purnamrita Sarkar, and Michael I Jordan. A scalable bootstrap for massive data. *J. of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4):795–816, 2014.
- Jason D Lee, Qiang Liu, Yuekai Sun, and Jonathan E Taylor. Communication-efficient sparse regression. *J. of Machine Learning Research*, 18(1):115–144, 2017.
- Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- Qihang Lin and Lin Xiao. An adaptive accelerated proximal gradient method and its homotopy continuation for sparse optimization. In *International Conference on Machine Learning*, pages 73–81. PMLR, 2014.
- Lester Mackey, Michael Jordan, and Ameet Talwalkar. Divide-and-conquer matrix factorization. In *Advances in Neural Information Processing Systems*, volume 24, pages 1134–1142. Curran Associates, Inc., 2011.

- Marie Maros and Gesualdo Scutari. Dgd²: A linearly convergent distributed algorithm for high-dimensional statistical recovery. In *Advances in Neural Information Processing Systems*, volume 35, pages 3475–3487. PMLR, 2022.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 1273–1282. PMLR, 20–22 Apr 2017.
- Angelia Nedic. Distributed gradient methods for convex machine learning problems in networks. *IEEE Signal Processing Magazine*, 10:92–101, 2020.
- Angelia Nedić and Alex Olshevsky. Distributed optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 60(3):601–615, 2014.
- Angelia Nedić and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, January 2009.
- Angelia Nedić, Asuman Ozdaglar, and Pablo A. Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, April 2010.
- Angelia Nedić, Alex Olshevsky, and Wei Shi. Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM J. on Optimization*, 27:2597–2633, July 2016.
- Angelia Nedić, Alex Olshevsky, and G. Rabbat, Michael. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106:953–976, 2018.
- Willie Neiswanger, Chong Wang, and Eric P. Xing. Asymptotically exact, embarrassingly parallel mcmc. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, UAI’14, page 623–632, 2014.
- Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Restricted eigenvalue properties for correlated gaussian designs. *J. of Machine Learning Research*, 11:2241–2259, 2010.
- Ali H. Sayed. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning*, 7:311–801, January 2014.
- Kevin Scaman, Francis Bach, Sébastien Bubeck, Yin Tat Lee, and Laurent Massoulié. Optimal algorithms for smooth and strongly convex distributed optimization in networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 3027–3036, 2017.
- Zuofeng Shang and Guang Cheng. Computational limits of a distributed algorithm for smoothing spline. *J. of Machine Learning Research*, 18(108):1–37, 2017.

- Chengchun Shi, Wenbin Lu, and Rui Song. A massive data framework for M-estimators with cubic-rate. *J. of the American Statistical Association*, 113(524):1698–1709, 2018.
- Wei Shi, Qing Ling, Gang Wu, and Wotao Yin. A proximal gradient algorithm for decentralized composite optimization. *IEEE Transactions on Signal Processing*, 63(22):6013–6023, 2015.
- Ying Sun, Marie Maros, Gesualdo Scutari, and Guang Chang. High-dimensional inference over networks: Linear convergence and statistical guarantees. *Arxiv preprint: arXiv:2201.08507*, pages 1–50, Jan. 2022a. doi: <https://doi.org/10.48550/arXiv.2201.08507>.
- Ying Sun, Gesualdo Scutari, and Amir Daneshmand. Distributed optimization based on gradient tracking revisited: Enhancing convergence rate via surrogation. *SIAM Journal on Optimization*, 32(2):354–385, 2022b.
- Ying Sun, Gesualdo Scutari, and Amir Daneshmand. Distributed optimization based on gradient tracking revisited: Enhancing convergence rate via surrogation. *SIAM J. on Optimization*, 32(2):354–385, 2022c.
- Stanislav Volgushev, Shih-Kang Chao, and Guang Cheng. Distributed inference for quantile regression processes. *The Annals of Statistics*, 47(3):1634–1662, 2019.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, NY, 1st ed edition, 2019. ISBN 9781108627771.
- Jialei Wang, Mladen Kolar, Nathan Srebro, and Tong Zhang. Efficient distributed learning with sparsity. In *International Conference on Machine Learning*, pages 3636–3645. PMLR, 2017.
- Zhaoran Wang, Han Liu, and Tong Zhang. Optimal computational and statistical rates of convergence for sparse nonconvex learning problems. *Annals of statistics*, 42(6):2164, 2014.
- Auzinger Wien. *Iterative solution of large linear systems*. Lecture Notes, TU Wien, 2011.
- Lin Xiao and Tong Zhang. A proximal-gradient homotopy method for the sparse least-squares problem. *SIAM J. on Optimization*, 23(2):1062–1091, 2013.
- Jinming Xu, Shanying Zhu, Yeng Chai Soh, and Lihua Xie. Convergence of asynchronous distributed gradient methods over stochastic networks. *IEEE Transactions on Automatic Control*, 63(2):434–448, 2018.
- Jinming Xu, Ye Tian, Ying Sun, and Gesualdo Scutari. Distributed algorithms for composite optimization: Unified framework and convergence analysis. *IEEE Transactions on Signal Processing*, 69:3555–3570, 2021.
- Kun Yuan, Bicheng Ying, Xiaochuan Zhao, and Ali H Sayed. Exact diffusion for distributed optimization and learning—part I: Algorithm development. *IEEE Transactions on Signal Processing*, 67(3):708–723, 2018.

- Yuchen Zhang, John Duchi, and Martin Wainwright. Divide and conquer kernel ridge regression. In *Conference on Learning Theory*, pages 592–617. PMLR, 2013a.
- Yuchen Zhang, John C Duchi, and Martin J Wainwright. Communication-efficient algorithms for statistical optimization. *J. of Machine Learning Research*, 14(1):3321–3363, 2013b.
- Tianqi Zhao, Guang Cheng, and Han Liu. A partially linear framework for massive heterogeneous data. *Annals of Statistics*, 44(4):1400, 2016.