

FAST DECENTRALIZED NONCONVEX FINITE-SUM
OPTIMIZATION WITH RECURSIVE VARIANCE REDUCTION*RAN XIN[†], USMAN A. KHAN[‡], AND SOUMMYA KAR[†]

Abstract. This paper considers decentralized minimization of $N := nm$ smooth nonconvex cost functions equally divided over a directed network of n nodes. Specifically, we describe a stochastic first-order gradient method, called **GT-SARAH**, that employs a **SARAH**-type variance reduction technique and gradient tracking (GT) to address the stochastic and decentralized nature of the problem. We show that **GT-SARAH**, with appropriate algorithmic parameters, finds an ϵ -accurate first-order stationary point with $\mathcal{O}(\max\{N^{1/2}, n(1-\lambda)^{-2}, n^{2/3}m^{1/3}(1-\lambda)^{-1}\}L\epsilon^{-2})$ gradient complexity, where $(1-\lambda) \in (0, 1]$ is the spectral gap of the network weight matrix and L is the smoothness parameter of the cost functions. This gradient complexity outperforms that of the existing decentralized stochastic gradient methods. In particular, in a big-data regime such that $n = \mathcal{O}(N^{1/2}(1-\lambda)^3)$, this gradient complexity furthers reduces to $\mathcal{O}(N^{1/2}L\epsilon^{-2})$, independent of the network topology, and matches that of the centralized near-optimal variance-reduced methods. Moreover, in this regime **GT-SARAH** achieves a *nonasymptotic linear speedup* in that the total number of gradient computations at each node is reduced by a factor of $1/n$ compared to the centralized near-optimal algorithms that perform all gradient computations at a single node. To the best of our knowledge, **GT-SARAH** is the first algorithm that achieves this property. In addition, we show that appropriate choices of local minibatch size balance the trade-offs between the gradient and communication complexity of **GT-SARAH**. Over infinite time horizon, we establish that all nodes in **GT-SARAH** asymptotically achieve consensus and converge to a first-order stationary point in the almost sure and mean-squared sense.

Key words. decentralized optimization, stochastic optimization, nonconvex optimization, variance reduction, gradient tracking

AMS subject classifications. 90C26, 90C15, 93A16

DOI. 10.1137/20M1361158

1. Introduction. We consider decentralized finite-sum minimization of $N := nm$ cost functions that takes the following form:

$$(1.1) \quad \min_{\mathbf{x} \in \mathbb{R}^p} F(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}), \quad f_i(\mathbf{x}) := \frac{1}{m} \sum_{j=1}^m f_{i,j}(\mathbf{x}),$$

where each $f_i : \mathbb{R}^p \rightarrow \mathbb{R}$, further decomposed as the average of m component costs $\{f_{i,j}\}_{j=1}^m$, is available only at the i th node in a network of n nodes. The network is abstracted as a directed graph $\mathcal{G} := \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V} := \{1, \dots, n\}$ is the set of node indices and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the collection of ordered pairs (i, r) , $i, r \in \mathcal{V}$, such that node r sends information to node i . We adopt the convention that $(i, i) \in \mathcal{E} \forall i \in \mathcal{V}$. Each node in the network is restricted to local computation and communication with its neighbors. Throughout the paper, we focus on the case where each $f_{i,j}$ is differentiable, not necessarily convex, and F is bounded below. This formulation often appears in decentralized empirical risk minimization, where each local cost f_i can be considered as an empirical risk computed over a finite number of m local data samples [48], and lies at the heart of many modern machine learning problems [4, 22, 52]. Examples

*Received by the editors August 20, 2020; accepted for publication (in revised form) September 1, 2021; published electronically January 5, 2022.

<https://doi.org/10.1137/20M1361158>

Funding: This work has been partially supported by the NSF under awards 1513936, 1903972, and 1935555.

[†]Department of Electrical and Computer Engineering (ECE), Carnegie Mellon University, Pittsburgh, PA 15213 USA (ranx@andrew.cmu.edu, soummyak@andrew.cmu.edu).

[‡]Department of ECE, Tufts University, Medford, MA 02155 USA (khan@ece.tufts.edu).

include nonconvex linear models and neural networks. When the local data size m is large, evaluating the exact gradient ∇f_i of each local cost at each iteration becomes computationally expensive, and methods that efficiently sample each local data batch are preferable. We are thus interested in designing fast stochastic gradient algorithms to find an ϵ -accurate first-order stationary point $\hat{\mathbf{x}} \in \mathbb{R}^p$ such that $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon^2$.

Towards problem (1.1), DSGD [7, 8, 33, 55], a decentralized version of stochastic gradient descent (SGD) [4, 12, 24], is often used to address the large-scale and decentralized nature of the data. DSGD is popular for several inference and learning tasks due to its simplicity of implementation and speedup in comparison to its centralized counterparts [20]. DSGD and its variants have been extensively studied for different computation and communication needs, e.g., momentum [43], directed graphs [3], escaping saddle-points [39, 41], zeroth-order schemes [44], swarming-based implementations [29], and constrained problems [53].

1.1. Challenges with DSGD. The performance of DSGD for the nonconvex problem (1.1), however, suffers from three major challenges: (i) the nondegenerate variance of the stochastic gradients at each node, (ii) the dissimilarity among the local functions across the nodes, and (iii) the transient time to reach the network topology independent region. To elaborate these issues, we recap DSGD for problem (1.1) and its convergence results as follows. Let $\mathbf{x}_i^k \in \mathbb{R}^p$ denote the iterate of DSGD at node i and iteration k . At each node i , DSGD performs [7, 33]

$$(1.2) \quad \mathbf{x}_i^{k+1} = \sum_{r=1}^n w_{ir} \mathbf{x}_r^k - \alpha \cdot \mathbf{g}_i^k, \quad k \geq 0,$$

where $\mathbf{W} = \{w_{ir}\} \in \mathbb{R}^{n \times n}$ is a weight matrix that respects the network topology, while $\mathbf{g}_i^k \in \mathbb{R}^p$ is a stochastic gradient such that $\mathbb{E}[\mathbf{g}_i^k | \mathbf{x}_i^k] = \nabla f_i(\mathbf{x}_i^k)$. Assuming the *bounded variance* of each local stochastic gradient \mathbf{g}_i^k , the *bounded dissimilarity* between the local and the global gradient [20], i.e., for some $\nu > 0$ and $\zeta > 0$,

$$(1.3) \quad \sup_{i \in \mathcal{V}, k \geq 0} \mathbb{E}[\|\mathbf{g}_i^k - \nabla f_i(\mathbf{x}_i^k)\|^2] \leq \nu^2 \text{ and } \sup_{\mathbf{x} \in \mathbb{R}^p} \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\mathbf{x}) - \nabla F(\mathbf{x})\|^2 \leq \zeta^2,$$

and L -smoothness of each f_i , it is shown in [20] that, for small enough $\alpha > 0$,

$$(1.4) \quad \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E}[\|\nabla F(\bar{\mathbf{x}}^k)\|^2] = \mathcal{O}\left(\frac{F(\bar{\mathbf{x}}^0) - F^*}{\alpha K} + \frac{\alpha L \nu^2}{n} + \frac{\alpha^2 L^2 \nu^2}{1 - \lambda} + \frac{\alpha^2 L^2 \zeta^2}{(1 - \lambda)^2}\right),$$

where $\bar{\mathbf{x}}^k := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^k$ and $(1 - \lambda) \in (0, 1]$ is the spectral gap of the weight matrix \mathbf{W} . It then follows that [20] for K large enough (see (iii) below) and with an appropriate step-size α , DSGD finds an ϵ -accurate first-order stationary point of F in $\mathcal{O}(\nu^2 L \epsilon^{-4})$ stochastic gradient computations across all nodes and therefore achieves *asymptotic* linear speedup compared to the centralized SGD [4, 12] that executes at a single node. Clearly, there are three issues with the convergence properties of DSGD:

- (i) Due to the nondegenerate stochastic gradient variance, the gradient complexity of DSGD does not match that of the centralized near-optimal variance-reduced methods when minimizing a finite sum of smooth nonconvex functions [11, 28, 45].
- (ii) The bounded dissimilarity assumption on the local and global gradients [3, 20, 41] or the coercivity of each local function [39] is essential for establishing the convergence of DSGD. In fact, a counterexample has been shown in [6] that *DSGD diverges for any constant step-size* when these types of assumptions are violated. Furthermore, the practical performance of DSGD degrades significantly when the local and the global

gradients are substantially different, i.e., when the data distributions across the nodes are largely heterogeneous [40, 50, 54].

(iii) DSGD achieves linear speedup only *asymptotically*, i.e., after a finite number of transient iterations that is a polynomial function of n, ν, ζ, L , and $(1 - \lambda)$ [20, 31, 43].

1.2. Main contributions. This paper proposes GT-SARAH, a novel decentralized stochastic variance-reduced gradient method that provably addresses the aforementioned challenges posed by DSGD. GT-SARAH is based on a *local SARAH-type gradient estimator* [11, 28], which removes the variance incurred by the local stochastic gradients, and *global gradient tracking (GT)* [10, 35, 51], which fuses the gradient estimators across the nodes such that the bounded dissimilarity or the coercivity-type assumptions are not required. Our main technical contributions are summarized in the following.

(i) We show that GT-SARAH, under appropriate algorithmic parameters, finds an ϵ -accurate first-order stationary point $\hat{\mathbf{x}}$ of F such that $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon^2$ in at most $\mathcal{H}_R := \mathcal{O}(\max\{N^{1/2}, n(1 - \lambda)^{-2}, n^{2/3}m^{1/3}(1 - \lambda)^{-1}\}L\epsilon^{-2})$ component gradient computations across all nodes. The gradient complexity \mathcal{H}_R significantly outperforms that of the existing decentralized stochastic gradient algorithms for problem (1.1); see Table 1 for a formal comparison.

(ii) In a big-data regime such that $n = \mathcal{O}(N^{1/2}(1 - \lambda)^3)$, the gradient complexity \mathcal{H}_R of GT-SARAH reduces to $\tilde{\mathcal{H}}_R := \mathcal{O}(N^{1/2}L\epsilon^{-2})$. We emphasize that $\tilde{\mathcal{H}}_R$ is independent of the network topology and matches that of the centralized near-optimal variance-reduced methods [11, 28, 45] under a slightly stronger smoothness assumption; see Remark 3.1 for details. Furthermore, since GT-SARAH computes n gradients in parallel at each iteration, its per-node gradient complexity in this regime is $\mathcal{O}(N^{1/2}n^{-1}\epsilon^{-2})$, demonstrating a *nonasymptotic linear speedup* compared with the aforementioned centralized near-optimal methods [11, 28, 45] that perform all gradient computations at a single node. To the best of our knowledge, GT-SARAH is the first decentralized method that achieves this property for Problem (1.1).

(iii) We show that choosing the local minibatch size of GT-SARAH judiciously balances the trade-offs between the gradient and communication complexity; see Corollary 3.9 and subsection 3.3.1 for details.

(iv) We establish that all nodes in GT-SARAH asymptotically achieve consensus and converge to a first-order stationary point of F over infinite time horizon in the almost sure and mean-squared sense.

1.3. Related work. Several algorithms have been proposed to improve certain aspects of DSGD. For example, a stochastic variant of EXTRA [36], Exact Diffusion [54], and NIDS [19], called D2 [40], removes the bounded dissimilarity assumption in DSGD based on a bias-correction principle. DSGT [50], introduced in [30] for smooth and strongly convex problems, achieves a similar theoretical performance as D2 via gradient tracking [10, 23, 32, 42] but with more general choices of weight matrices. Reference [17] establishes asymptotic properties of a decentralized stochastic primal-dual algorithm for smooth convex problems. Reference [16] develops decentralized primal-dual communication sliding algorithms that achieve communication efficiency for convex and possibly nonsmooth problems. These methods, however, are subject to the nondegenerate variance of the stochastic gradients. Inspired by the variance reduction techniques for centralized stochastic optimization [1, 5, 9, 11, 26, 27, 28, 34, 45, 47, 58], decentralized variance-reduced methods for smooth and strongly convex problems have been proposed recently, e.g., in [18, 21, 48, 49, 56]; in particular, the integration of gradient tracking and variance reduction described in this paper was introduced in [48, 49] to obtain linear convergence.

A recent paper [38] proposes D-GET for problem (1.1), which also considers local SARAH-type variance reduction and gradient tracking. In the following, we compare our work to [38] from a few major technical aspects.¹ *First*, the gradient complexity \mathcal{H}_R of GT-SARAH improves that of D-GET in terms of the dependence on n and m ; see Table 1. In particular, in a big-data regime, $n = \mathcal{O}(N^{1/2}(1 - \lambda)^3)$, \mathcal{H}_R matches the gradient complexity of the centralized near-optimal methods [11, 28, 45]; in contrast, the gradient complexity of D-GET is worse than that of the centralized near-optimal methods by a factor of $n^{1/2}$ even if the network is fully connected. *Second*, the complexity results of D-GET are attained with a specific local minibatch size $m^{1/2}$. Conversely, we establish general complexity bounds of GT-SARAH with arbitrary local minibatch size and characterize the computation-communication trade-offs induced by different choices of the minibatch size. *Third*, the Lyapunov function based convergence analysis of D-GET does not show explicit dependence of several important problem parameters, such as $(1 - \lambda)$ and L , while the analysis in this work reveals explicitly the dependence of all problem related parameters and sheds light on their implications. *Fourth*, we note that both GT-SARAH and D-GET achieve a worst case communication complexity of the form $\mathcal{O}((1 - \lambda)^{-a} L^b \epsilon^{-2})$, independent of m and n , for some $a, b \in \mathbb{R}^+$. Since the dependence of a and b in D-GET is not explicit, it is unclear which algorithm achieves a lower communication complexity. *Finally*, [38] presents a variant of D-GET that is applicable to a more general online setting such as expected risk minimization.

TABLE 1

A comparison of the gradient complexities of the state-of-the-art decentralized stochastic gradient methods to minimize a sum of $N = nm$ smooth nonconvex functions equally divided among n nodes. The gradient complexity is in terms of the total number of component gradient computations across all nodes to find a first-order stationary point $\hat{\mathbf{x}} \in \mathbb{R}^p$ such that $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon^2$. In the table, ν^2 denotes the bounded variance of the stochastic gradients described in (1.3), $(1 - \lambda) \in (0, 1]$ is the spectral gap of the network weight matrix, and L is the smoothness parameter of the cost functions. We note that the complexities of DSGD, D2, and DSGT in the table are established in the setting of stochastic first-order oracles, which is more general than the finite-sum formulation considered in this paper. Moreover, the complexities of DSGD, D2, and DSGT in the table are stated in the regime that ϵ is small enough for simplicity; see [20, 40, 50] for their precise expressions. Finally, we note that only the best possible gradient complexity of GT-SARAH, in the sense of Theorem 3.6, is presented in the table for conciseness; see Corollary 3.9 and subsection 3.3.1 for detailed discussion on balancing the trade-offs between the gradient and communication complexity of GT-SARAH.

Algorithm	Gradient complexity	Remarks
DSGD [20]	$\mathcal{O}\left(\frac{\nu^2 L}{\epsilon^4}\right)$	Bounded variance, bounded dissimilarity
D2 [40]	$\mathcal{O}\left(\frac{\nu^2 L}{\epsilon^4}\right)$	Bounded variance
DSGT [50]	$\mathcal{O}\left(\frac{\nu^2 L}{\epsilon^4}\right)$	Bounded variance
D-GET [38]	$\mathcal{O}\left(\frac{n^{1/2} N^{1/2} L^b}{(1 - \lambda)^a \epsilon^2}\right)$	$a, b \in \mathbb{R}^+$ are not explicitly shown in [38]
GT-SARAH (this work)	$\mathcal{O}\left(\max\left\{N^{1/2}, \frac{n}{(1 - \lambda)^2}, \frac{n^{2/3} m^{1/3}}{1 - \lambda}\right\} \frac{L}{\epsilon^2}\right)$	See Theorem 3.6 and Corollary 3.9

¹Note that [38] uses $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon$ as the performance metric, while we use $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon^2$ in this paper. We state the complexities of D-GET established in [38] under our metric for consistency.

1.4. Paper outline and notation. The proposed GT-SARAH algorithm is developed in section 2. We present the convergence results of GT-SARAH and discuss their implications in section 3. Section 4 presents the convergence analysis. Section 5 presents numerical experiments, while section 6 concludes the paper.

The set of positive integers and real numbers are denoted by \mathbb{Z}^+ and \mathbb{R}^+ , respectively. For any $a \in \mathbb{R}$, $\lfloor a \rfloor$ denotes the largest integer i such that $i \leq a$; similarly, $\lceil a \rceil$ denotes the smallest integer i such that $i \geq a$; We use lowercase bold letters to denote column vectors and uppercase bold letters to denote matrices. The matrix, \mathbf{I}_d , represents the $d \times d$ identity; $\mathbf{1}_d$ and $\mathbf{0}_d$ are the d -dimensional column vectors of all ones and zeros, respectively. The Kronecker product of two matrices \mathbf{A} and \mathbf{B} is denoted by $\mathbf{A} \otimes \mathbf{B}$. We use $\|\cdot\|$ to denote the Euclidean norm of a vector or the spectral norm of a matrix. For a matrix \mathbf{X} , we use $\rho(\mathbf{X})$ to denote its spectral radius, $\lambda_2(\mathbf{X})$ to denote its second largest singular value, and $\det(\mathbf{X})$ to denote its determinant. Matrix inequalities are interpreted in the entrywise sense. We use $\sigma(\cdot)$ to denote the σ -algebra generated by the random variables and/or sets in its argument. The empty set is denoted by \emptyset .

2. Algorithm development: GT-SARAH. We now systematically build the proposed algorithm GT-SARAH and provide the basic intuition. We recall that the performance (1.4) of DSGD, in addition to the first term which is similar to that of the centralized batch gradient descent, has three additional bias terms. The second and third bias terms in (1.4) depend on the variance ν^2 of local stochastic gradients. A variance-reduced gradient estimation procedure of SARAH-type [11, 28], employed locally at each node i in GT-SARAH, removes ν^2 . The last bias term in (1.4) is due to the dissimilarity ζ^2 between the local gradients $\{\nabla f_i\}_{i=1}^n$ and the global gradient ∇F . A dynamic fusion mechanism, called gradient tracking [10, 15, 23, 32, 51], removes ζ^2 by tracking the average of the local gradient estimators in GT-SARAH to learn the global gradient at each node. The resulting algorithm is illustrated in Figure 1.

2.1. Detailed implementation. The complete implementation of GT-SARAH is summarized in Algorithm 2.1, where we assume that all nodes start from the same point $\bar{\mathbf{x}}^{0,1} \in \mathbb{R}^p$. GT-SARAH can be interpreted as a double loop method with an outer loop, indexed by s , and an inner loop, indexed by t . At the beginning of each outer loop s , GT-SARAH computes the local batch gradient $\mathbf{v}_i^{0,s} := \nabla f_i(\mathbf{x}_i^{0,s})$ at each node i . These batch gradients are then used to compute the first iteration of the global gradient tracker $\mathbf{y}_i^{1,s}$ and the state update $\mathbf{x}_i^{1,s}$. The three quantities, $\mathbf{v}_i^{0,s}$, $\mathbf{y}_i^{1,s}$, $\mathbf{x}_i^{1,s}$, set up the subsequent inner loop iterations. At each inner loop iteration $t \geq 1$, each node i samples two minibatch stochastic gradients from its local data that are used

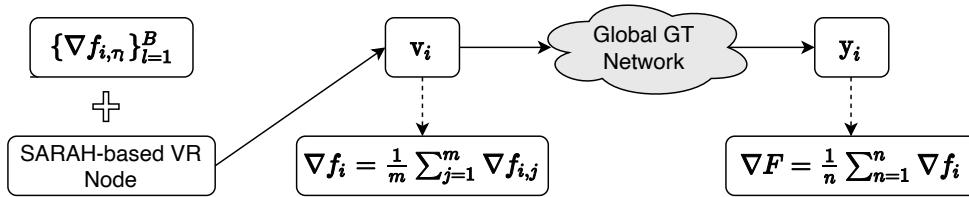


FIG. 1. Each node i samples a minibatch of stochastic gradients $\{\nabla f_{i,\tau_l}\}_{l=1}^B$ at each iteration from its local data batch and computes an estimator \mathbf{v}_i of its local batch gradient ∇f_i via a SARAH-type variance reduction (VR) procedure. These local gradient estimators \mathbf{v}_i 's are then fused over the network via a gradient tracking (GT) technique to obtain \mathbf{y}_i 's that approximate the global gradient ∇F .

to construct the gradient estimator $\mathbf{v}_i^{t,s}$. We note that the gradient estimator is of recursive nature, i.e., it depends on $\mathbf{v}_i^{t-1,s}$ and the minibatch stochastic gradients evaluated at the current and the past states $\mathbf{x}_i^{t,s}$ and $\mathbf{x}_i^{t-1,s}$. The next step is to update $\mathbf{y}_i^{t+1,s}$ based on the gradient tracking protocol. Finally, the state $\mathbf{x}_i^{t+1,s}$ at each node i is computed as a convex combination of the states of the neighboring nodes followed by a descent in the direction of the gradient tracker $\mathbf{y}_i^{t+1,s}$. The latest updates $\mathbf{x}_i^{q+1,s}$, $\mathbf{y}_i^{q+1,s}$, and $\mathbf{v}_i^{q,s}$ then set up the next inner-outer loop cycle of GT-SARAH.

Algorithm 2.1 GT-SARAH at each node i

Require: $\mathbf{x}_i^{0,1} = \bar{\mathbf{x}}^{0,1} \in \mathbb{R}^p$, $\alpha \in \mathbb{R}^+$, $q \in \mathbb{Z}^+$, $S \in \mathbb{Z}^+$, $B \in \mathbb{Z}^+$, $\{\underline{w}_{ir}\}_{r=1}^n$, $\mathbf{y}_i^{0,1} = \mathbf{0}_p$, $\mathbf{v}_i^{-1,1} = \mathbf{0}_p$.

1: **for** $s = 1, 2, \dots, S$ **do**

2: $\mathbf{v}_i^{0,s} = \nabla f_i(\mathbf{x}_i^{0,s}) = \frac{1}{m} \sum_{j=1}^m \nabla f_{i,j}(\mathbf{x}_i^{0,s})$; ▷ batch gradient computation

3: $\mathbf{y}_i^{1,s} = \sum_{r=1}^n \underline{w}_{ir} \mathbf{y}_i^{0,s} + \mathbf{v}_i^{0,s} - \mathbf{v}_i^{-1,s}$; ▷ gradient tracking

4: $\mathbf{x}_i^{1,s} = \sum_{r=1}^n \underline{w}_{ir} \mathbf{x}_r^{0,s} - \alpha \mathbf{y}_i^{1,s}$; ▷ state update

5: **for** $t = 1, 2, \dots, q$ **do**

6: for l in $\{1, \dots, B\}$, choose $\tau_{i,l}^{t,s}$ uniformly at random from $\{1, \dots, m\}$; ▷ sampling

7: $\mathbf{v}_i^{t,s} = \frac{1}{B} \sum_{l=1}^B \left(\nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t-1,s}) \right) + \mathbf{v}_i^{t-1,s}$; ▷ SARAH

8: $\mathbf{y}_i^{t+1,s} = \sum_{r=1}^n \underline{w}_{ir} \mathbf{y}_r^{t,s} + \mathbf{v}_i^{t,s} - \mathbf{v}_i^{t-1,s}$; ▷ gradient tracking

9: $\mathbf{x}_i^{t+1,s} = \sum_{r=1}^n \underline{w}_{ir} \mathbf{x}_r^{t,s} - \alpha \mathbf{y}_i^{t+1,s}$; ▷ state update

10: **end for**

11: Set $\mathbf{x}_i^{0,s+1} = \mathbf{x}_i^{q+1,s}$; $\mathbf{y}_i^{0,s+1} = \mathbf{y}_i^{q+1,s}$; $\mathbf{v}_i^{-1,s+1} = \mathbf{v}_i^{q,s}$. ▷ next cycle

12: **end for**

3. Main results. In this section, we present the main convergence results of GT-SARAH and discuss their implications.

3.1. Assumptions. We make the following assumptions to establish the convergence properties of GT-SARAH in this paper.

Assumption 3.1. Each local component cost $f_{i,j}$ is differentiable, and $\{f_{i,j}\}_{j=1}^m$ satisfies a mean-squared smoothness property, i.e., for some $L > 0$,

$$(3.1) \quad \frac{1}{m} \sum_{j=1}^m \|\nabla f_{i,j}(\mathbf{x}) - \nabla f_{i,j}(\mathbf{y})\|^2 \leq L^2 \|\mathbf{x} - \mathbf{y}\|^2 \quad \forall i \in \mathcal{V}, \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^p.$$

In addition, the global cost F is bounded below, i.e., $F^* := \inf_{\mathbf{x} \in \mathbb{R}^p} F(\mathbf{x}) > -\infty$.

It is clear that under Assumption 3.1, each f_i and F are L -smooth. We note that Assumption 3.1 is weaker than requiring each $f_{i,j}$ to be L -smooth.

Remark 3.1. The local mean-squared smoothness assumption (3.1), which is also used in the existing work [38], is slightly stronger than the smoothness assumption required by the existing lower bound $\Omega(N^{1/2}L\epsilon^{-2})$ [11, 57] and the centralized near-optimal methods [11, 28, 45] for finite-sum problems in the following sense. If we view problem (1.1) as a centralized optimization problem, that is, all $f_{i,j}$'s are available

at a single node, then the aforementioned lower bound and the convergence of the centralized near-optimal methods are established under the following assumption:

$$(3.2) \quad \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \|\nabla f_{i,j}(\mathbf{x}) - \nabla f_{i,j}(\mathbf{y})\|^2 \leq L^2 \|\mathbf{x} - \mathbf{y}\|^2 \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^p.$$

Clearly, (3.2) is implied by (3.1) but not vice versa. Due to this subtle difference, it is unclear whether the existing lower bound $\Omega(N^{1/2}L\epsilon^{-2})$ [11, 57] established under (3.2) remains valid under (3.1). Finally, we note that a lower bound result for decentralized deterministic first-order algorithms in the case of $m = 1$ can be found in [37].

Assumption 3.2. The family $\{\tau_{i,l}^{t,s} : t \in [1, q], s \geq 1, i \in \mathcal{V}, l \in [1, B]\}$ of random variables is independent.

Assumption 3.2 is standard in the stochastic optimization literature, e.g., [4, 11].

Assumption 3.3. The nonnegative weight matrix $\underline{\mathbf{W}} := \{\underline{w}_{ir}\} \in \mathbb{R}^{n \times n}$ associated with the network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has positive diagonals and is primitive. Moreover, $\underline{\mathbf{W}}$ is doubly stochastic, i.e., $\underline{\mathbf{W}}\mathbf{1}_n = \mathbf{1}_n$ and $\mathbf{1}_n^\top \underline{\mathbf{W}} = \mathbf{1}_n^\top$.

An important consequence of Assumption 3.3 is that [32]

$$(3.3) \quad \lambda := \|\underline{\mathbf{W}} - \frac{1}{n}\mathbf{1}_n\mathbf{1}_n^\top\| = \lambda_2(\underline{\mathbf{W}}) \in [0, 1],$$

where $\lambda_2(\underline{\mathbf{W}})$ denotes the second largest singular value of $\underline{\mathbf{W}}$.² We term $(1 - \lambda)$ as the spectral gap of $\underline{\mathbf{W}}$ that characterizes the connectivity of the network [22].

Remark 3.2. Weight matrices satisfying Assumption 3.3 may be designed for the family of strongly connected directed graphs that admit doubly stochastic weights: (i) towards the primitivity requirement in Assumption 3.3, we note that if a graph is strongly connected, then its associated weight matrix $\underline{\mathbf{W}}$ is irreducible [14, Theorem 6.2.14, 6.2.24] and $\underline{\mathbf{W}}$ is further primitive since it is nonnegative with positive diagonals [14, Lemma 8.5.4]; (ii) towards the doubly stochastic requirement in Assumption 3.3, we refer the readers to [13] for necessary and sufficient conditions under which a strongly connected directed graph admits doubly stochastic weights.

An important special case of this family is undirected connected graphs where doubly stochastic weights always exist and can be constructed in an efficient and decentralized manner, for instance, by the lazy Metropolis rule [22]. Hence, Assumption 3.3 is *more general* than the one required by EXTRA-based algorithms for decentralized optimization. For example, the weight matrix of D2 needs to be symmetric and meet certain spectral properties [40] and is therefore not applicable to directed graphs.

In the rest of the paper, we fix a rich enough probability space $(\Xi, \mathcal{F}, \mathbb{P})$ where all random variables generated by GT-SARAH are properly defined. We formally state the convergence results of GT-SARAH next, the proofs of which are deferred to subsection 4.2.

3.2. Asymptotic almost sure and mean-squared convergence. The following theorem shows the asymptotic convergence of GT-SARAH.

THEOREM 3.3. *Let Assumptions 3.1–3.3 hold. Suppose that the step-size α , mini-batch size B , and the inner loop length q of GT-SARAH follow*

$$0 < \alpha \leq \min \left\{ \frac{(1 - \lambda^2)^2}{4\sqrt{42}}, \left(\frac{nB}{6q} \right)^{1/2}, \left(\frac{4nB}{7nB + 24q} \right)^{1/4} \frac{1 - \lambda^2}{6} \right\} \frac{1}{2L},$$

²We note that the relation in (3.3) may be established by following the definition of the spectral norm with the help of the primitivity and doubly stochasticity of $\underline{\mathbf{W}}$ and $\underline{\mathbf{W}}^\top \underline{\mathbf{W}}$, the Perron–Frobenius theorem, and the spectral decomposition of $\underline{\mathbf{W}}^\top \underline{\mathbf{W}}$ [14, 32].

where $B \in [1, m]$. Then we have, $\forall t \in [0, q]$, $\forall i \in \mathcal{V}$,

$$\begin{aligned} \mathbb{P}\left(\lim_{s \rightarrow \infty} \|\nabla F(\mathbf{x}_i^{t,s})\| = 0\right) &= 1 \quad \text{and} \quad \lim_{s \rightarrow \infty} \mathbb{E}\left[\|\nabla F(\mathbf{x}_i^{t,s})\|^2\right] = 0, \\ \mathbb{P}\left(\lim_{s \rightarrow \infty} \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\| = 0\right) &= 1 \quad \text{and} \quad \lim_{s \rightarrow \infty} \mathbb{E}\left[\|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2\right] = 0, \end{aligned}$$

where $\bar{\mathbf{x}}^{t,s} := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i^{t,s}$.

In addition to the mean-squared convergence that is standard in the stochastic optimization literature, the almost sure convergence in Theorem 3.3 guarantees that all nodes in GT-SARAH asymptotically achieve consensus and converge to a first-order stationary point of F on almost every sample path.

3.3. Complexities of GT-SARAH for finding first-order stationary points.

We measure the outer loop complexity of GT-SARAH in the following sense.

DEFINITION 3.4. Consider the sequence of random state vectors $\{\mathbf{x}_i^{t,s}\}$ generated by GT-SARAH at each node i . We say that GT-SARAH finds an ϵ -accurate first-order stationary point of F in S outer loop iterations if

$$(3.4) \quad \frac{1}{S(q+1)} \sum_{s=1}^S \sum_{t=0}^q \frac{1}{n} \sum_{i=1}^n \mathbb{E}\left[\|\nabla F(\mathbf{x}_i^{t,s})\|^2 + L^2 \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2\right] \leq \epsilon^2.$$

This is a standard metric that is concerned with the minimum of the stationary gaps and consensus errors over iterations in the mean-squared sense at each node [11, 20, 28, 40, 45]. In particular, if (3.4) holds and the output $\hat{\mathbf{x}}$ of GT-SARAH is chosen uniformly at random from the set $\{\mathbf{x}_i^{t,s} : 0 \leq t \leq q, 1 \leq s \leq S, i \in \mathcal{V}\}$, then we have $\mathbb{E}[\|\nabla F(\hat{\mathbf{x}})\|^2] \leq \epsilon^2$. In the following, we first provide the outer loop iteration complexity of GT-SARAH.

THEOREM 3.5. Let Assumptions 3.1–3.3 hold. Suppose that the step-size α , mini-batch size B , and the inner loop length q of GT-SARAH follow

$$0 < \alpha \leq \min \left\{ \frac{(1-\lambda^2)^2}{4\sqrt{42}}, \left(\frac{nB}{6q} \right)^{1/2}, \left(\frac{4nB}{7nB+24q} \right)^{1/3} \frac{1-\lambda^2}{6} \right\} \frac{1}{2L},$$

where $B \in [1, m]$. Then the number of the outer loop iterations S required by GT-SARAH to find an ϵ -accurate stationary point of F is at most

$$\frac{1}{(q+1)\alpha L \epsilon^2} \left(4L(F(\bar{\mathbf{x}}^{0,1}) - F^*) + \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\bar{\mathbf{x}}^{0,1})\|^2 \right).$$

With Theorem 3.5 at hand, the gradient and communication complexities of GT-SARAH can be readily established.

THEOREM 3.6. Let Assumptions 3.1–3.3 hold. Suppose that the step-size α and the length q of the inner loop of GT-SARAH are chosen as³

$$(3.5) \quad q = \mathcal{O}\left(\frac{m}{B}\right) \quad \text{and} \quad \alpha = \mathcal{O}\left(\min \left\{ (1-\lambda)^2, \frac{n^{1/2}B}{m^{1/2}}, \frac{n^{1/3}B^{2/3}(1-\lambda)}{m^{1/3}} \right\} \frac{1}{L} \right),$$

³The \mathcal{O} notation only hides universal constants that are independent of problem parameters.

where $B \in [1, m]$. Then **GT-SARAH** finds an ϵ -accurate stationary point of F in

$$\mathcal{H}_B := \mathcal{O}\left(\max\left\{\frac{nB}{(1-\lambda)^2}, N^{1/2}, \frac{m^{1/3}n^{2/3}B^{1/3}}{1-\lambda}\right\}\frac{\Delta}{\epsilon^2}\right)$$

component gradient computations across all nodes and

$$\mathcal{K}_B := \mathcal{O}\left(\max\left\{\frac{1}{(1-\lambda)^2}, \frac{m^{1/2}}{n^{1/2}B}, \frac{m^{1/3}}{n^{1/3}B^{2/3}(1-\lambda)}\right\}\frac{\Delta}{\epsilon^2}\right)$$

rounds of communication, where $\Delta := L(F(\bar{\mathbf{x}}^{0,1}) - F^*) + \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\bar{\mathbf{x}}^{0,1})\|^2$.

Remark 3.7. Theorem 3.6 holds for an arbitrary minibatch size $B \in [1, m]$.

Remark 3.8. The gradient complexity at each node of **GT-SARAH** is \mathcal{H}_B/n .

In view of Theorem 3.6, as the minibatch size B increases, the gradient complexity \mathcal{H}_B (resp., the communication complexity \mathcal{K}_B) of **GT-SARAH** is nondecreasing (resp., nonincreasing). The following corollary may be obtained from Theorem 3.6 by standard algebraic manipulations and shows that choosing the minibatch size B appropriately leads to favorable computation and communication trade-offs.

COROLLARY 3.9. *Let Assumptions 3.1–3.3 hold. Suppose that the step-size α and the inner loop length q of **GT-SARAH** are chosen according to (3.5). We have the following complexity results.*

(i) *If $B \in [1, \lfloor R \rfloor]$, where $R := \max\{m^{1/2}n^{-1/2}(1-\lambda)^3, 1\}$, then **GT-SARAH** attains the best possible, in the sense of Theorem 3.6, gradient complexity*

$$(3.6) \quad \mathcal{H}_R := \mathcal{O}\left(\max\left\{\frac{n}{(1-\lambda)^2}, N^{1/2}, \frac{m^{1/3}n^{2/3}}{1-\lambda}\right\}\frac{\Delta}{\epsilon^2}\right);$$

moreover, when $B = \lfloor R \rfloor$, the corresponding communication complexity of **GT-SARAH** is

$$(3.7) \quad \mathcal{K}_R := \mathcal{O}\left(\max\left\{\frac{1}{(1-\lambda)^2}, \min\left\{\frac{m^{1/2}}{n^{1/2}}, \frac{1}{(1-\lambda)^3}\right\}, \min\left\{\frac{m^{1/3}}{n^{1/3}(1-\lambda)}, \frac{1}{(1-\lambda)^3}\right\}\right\}\frac{\Delta}{\epsilon^2}\right).$$

(ii) *If $B \in [\lceil C \rceil, m]$, where $C := \max\{m^{1/2}n^{-1/2}(1-\lambda)^{3/2}, 1\}$, then **GT-SARAH** attains the best possible, in the sense of Theorem 3.6, communication complexity*

$$(3.8) \quad \mathcal{K}_C := \mathcal{O}\left(\frac{1}{(1-\lambda)^2}\frac{\Delta}{\epsilon^2}\right);$$

moreover, when $B = \lceil C \rceil$, the corresponding gradient complexity of **GT-SARAH** is

$$(3.9) \quad \mathcal{H}_C := \mathcal{O}\left(\max\left\{\frac{n}{(1-\lambda)^2}, \frac{N^{1/2}}{(1-\lambda)^{1/2}}, \frac{m^{1/3}n^{2/3}}{1-\lambda}\right\}\frac{\Delta}{\epsilon^2}\right).$$

Comparing (3.6) and (3.7) with (3.9) and (3.8), we clearly have $\mathcal{H}_R \leq \mathcal{H}_C$ and $\mathcal{K}_R \geq \mathcal{K}_C$.

3.3.1. Two regimes of practical significance. We now discuss the implications of the complexity results in Corollary 3.9 and the corresponding computation-communication trade-offs in the following regimes of practical significance.

• **Big-data regime:** $n = \mathcal{O}(N^{1/2}(1-\lambda)^3)$. In this regime, typical of large-scale machine learning, i.e., the total number of data samples N is very large, it can be verified

that \mathcal{H}_R reduces to $\tilde{\mathcal{H}}_R := \mathcal{O}(N^{1/2}\Delta\epsilon^{-2})$ and \mathcal{K}_R reduces to $\tilde{\mathcal{K}}_R := \mathcal{O}((1-\lambda)^{-3}\Delta\epsilon^{-2})$. It is worth noting that $\tilde{\mathcal{H}}_R$ is independent of the network topology and matches the gradient complexity of the centralized near-optimal variance-reduced methods [11, 28, 45] for this problem class up to constant factors, under a slightly stronger smoothness assumption; see Remark 3.1. Moreover, $\tilde{\mathcal{H}}_R$ demonstrates a nonasymptotic linear speedup in that the number of component gradient computations required *at each node* to achieve an ϵ -accurate stationary point of F is reduced by a factor of $1/n$, compared to the aforementioned centralized near-optimal algorithms [11, 28, 45] that perform all gradient computations at a single node.

On the other hand, it is straightforward to verify that \mathcal{H}_C reduces to $\tilde{\mathcal{H}}_C := \mathcal{O}(N^{1/2}(1-\lambda)^{-1/2}\Delta\epsilon^{-2})$. In other words, in this big-data regime, choosing a large minibatch size $B = \lceil C \rceil$ improves the communication complexity from $\tilde{\mathcal{K}}_R$ to \mathcal{K}_C while deteriorating the gradient complexity from $\tilde{\mathcal{H}}_R$ to $\tilde{\mathcal{H}}_C$, demonstrating an interesting trade-off between computation and communication.

• **Large-scale network regime:** $n = \Omega(N^{1/2}(1-\lambda)^{3/2})$. In this regime, typical of ad hoc Internet of Things networks, i.e., the number of the nodes n and the network spectral gap inverse $(1-\lambda)^{-1}$ are large compared with the total number of samples N , it can be verified that $R = C = 1$ and consequently $\mathcal{H}_R = \mathcal{H}_C$ reduce to $\mathcal{O}(n(1-\lambda)^{-2}\Delta\epsilon^{-2})$, while $\mathcal{K}_R = \mathcal{K}_C$ reduce to $\mathcal{O}((1-\lambda)^{-2}\Delta\epsilon^{-2})$. In other words, in this large-scale network regime, the minibatch size $B = \mathcal{O}(1)$ is preferred since it attains the best possible gradient and communication complexity simultaneously, in the sense of Theorem 3.6.

Remark 3.10 (characterization of the big-data regime). We note that the number of nodes n may be interpreted as the intrinsic minibatch size of GT-SARAH. We recall that the centralized near-optimal variance-reduced algorithms [11, 28, 45] for this problem class retain their best possible gradient complexity if their minibatch size does not exceed $N^{1/2}$ [28]. Thus, the aforementioned big-data regime $n = \mathcal{O}(N^{1/2}(1-\lambda)^3)$ approaches the centralized one as the network connectivity improves and matches the centralized one when the network is fully connected, i.e., $\lambda = 0$.

4. Convergence analysis. In this section, we present the proof pipeline for Theorems 3.3, 3.5, and 3.6. The analysis framework is novel and general and may be applied to other decentralized algorithms built around variance reduction and gradient tracking. To proceed, we first write GT-SARAH in a matrix form. Recall that GT-SARAH is a double loop method, where the outer loop index is $s \in \{1, \dots, S\}$ and the inner loop index is $t \in \{0, \dots, q\}$. It is straightforward to verify that GT-SARAH can be equivalently written as, $\forall s \geq 1$ and $t \in [0, q]$,

$$(4.1a) \quad \mathbf{y}^{t+1,s} = \mathbf{W}\mathbf{y}^{t,s} + \mathbf{v}^{t,s} - \mathbf{v}^{t-1,s},$$

$$(4.1b) \quad \mathbf{x}^{t+1,s} = \mathbf{W}\mathbf{x}^{t,s} - \alpha\mathbf{y}^{t+1,s}$$

with $\mathbf{v}^{t,s}$, $\mathbf{x}^{t,s}$, and $\mathbf{y}^{t,s}$, in \mathbb{R}^{np} , that concatenate local gradient estimators $\{\mathbf{v}_i^{t,s}\}_{i=1}^n$, states $\{\mathbf{x}_i^{t,s}\}_{i=1}^n$, and gradient trackers $\{\mathbf{y}_i^{t,s}\}_{i=1}^n$, respectively, and $\mathbf{W} := \underline{\mathbf{W}} \otimes \mathbf{I}_p$. We recall that $\mathbf{x}^{0,s+1} = \mathbf{x}^{q+1,s}$, $\mathbf{y}^{0,s+1} = \mathbf{y}^{q+1,s}$, $\mathbf{v}^{-1,s+1} = \mathbf{v}^{q,s} \forall s \geq 1$, and $\mathbf{v}^{-1,1} = \mathbf{0}_{np}$ from Algorithm 2.1 under the vector notation. Under Assumption 3.3, we have [14]

$$\mathbf{J} := \lim_{k \rightarrow \infty} \mathbf{W}^k = \left(\frac{1}{n} \mathbf{1}_n \mathbf{1}_n^\top \right) \otimes \mathbf{I}_p,$$

i.e., the power limit of the network weight matrix \mathbf{W} is the exact averaging matrix \mathbf{J} . We also introduce the following notation for convenience:

$$\nabla \mathbf{f}(\mathbf{x}^{t,s}) := [\nabla f_1(\mathbf{x}_1^{t,s})^\top, \dots, \nabla f_n(\mathbf{x}_n^{t,s})^\top]^\top, \quad \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) := \frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p) \nabla \mathbf{f}(\mathbf{x}^{t,s}),$$

$$\bar{\mathbf{x}}^{t,s} := \frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p) \mathbf{x}^{t,s}, \quad \bar{\mathbf{y}}^{t,s} = \frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p) \mathbf{y}^{t,s}, \quad \bar{\mathbf{v}}^{t,s} := \frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p) \mathbf{v}^{t,s}.$$

In particular, we note that $\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2 := \sum_{i=1}^n \|\nabla f_i(\bar{\mathbf{x}}^{0,1})\|^2$. In the rest of the paper, we assume that Assumptions 3.1, 3.2, and 3.3 hold without explicitly stating them.

4.1. Auxiliary relationships. First, as a consequence of the gradient tracking update (4.1b), it is straightforward to show by induction the following result.

LEMMA 4.1. $\bar{\mathbf{y}}^{t+1,s} = \bar{\mathbf{v}}^{t,s} \ \forall s \geq 1$ and $t \in [0, q]$.

Proof. See Appendix A.1. \square

The above lemma states that the average of gradient trackers preserves the average of local gradient estimators. Under Assumption 3.3, we obtain that the weight matrix \mathbf{W} is a contraction operator [32].

LEMMA 4.2. $\|\mathbf{W}\mathbf{x} - \mathbf{J}\mathbf{x}\| \leq \lambda \|\mathbf{x} - \mathbf{J}\mathbf{x}\| \ \forall \mathbf{x} \in \mathbb{R}^{np}$ for $\lambda \in [0, 1)$ defined in (3.3).

Lemmas 4.1 and 4.2 are standard in decentralized optimization and gradient tracking [23, 32]. The L -smoothness of F leads to the following quadratic upper bound [25]:

$$(4.2) \quad F(\mathbf{y}) \leq F(\mathbf{x}) + \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{L}{2} \|\mathbf{y} - \mathbf{x}\|^2 \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^p.$$

Consequently, the following descent-type lemma on the iterates generated by **GT-SARAH** may be established by setting $\mathbf{y} = \bar{\mathbf{x}}^{t+1,s}$ and $\mathbf{x} = \bar{\mathbf{x}}^{t,s}$ in (4.2) and taking a telescoping sum across all iterations of **GT-SARAH** with the help of Lemma 4.1 and the L -smoothness of each f_i .

LEMMA 4.3. *If the step-size follows $0 < \alpha \leq \frac{1}{2L}$, then we have*

$$\begin{aligned} \mathbb{E}[F(\bar{\mathbf{x}}^{q+1,S})] &\leq F(\bar{\mathbf{x}}^{0,1}) - \frac{\alpha}{2} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2] - \frac{\alpha}{4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s}\|^2] \\ &\quad + \alpha \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s})\|^2] + \alpha L^2 \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}\left[\frac{\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2}{n}\right]. \end{aligned}$$

Proof. See Appendix B. \square

In light of Lemma 4.3, our analysis approach is to derive the range of the step-size α of **GT-SARAH** such that

$$\frac{1}{4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s}\|^2] - \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s})\|^2] - L^2 \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}\left[\frac{\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2}{n}\right]$$

is nonnegative and therefore establishes the convergence of **GT-SARAH** to a first-order stationary point following the standard arguments in *batch* gradient descent for nonconvex problems [4, 25]. To this aim, we need to derive upper bounds for two error terms in the above expression: (i) $\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s})\|^2$, the gradient estimation error; and (ii) $\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2$, the state consensus error. We quantify these two errors next and then return to Lemma 4.3. The following lemma is obtained with similar probabilistic arguments for **SARAH**-type [11, 28, 45] estimators, however, with subtle modifications due to the decentralized network effect.

LEMMA 4.4. *We have, $\forall s \geq 1$,*

$$\sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s})\|^2 \right] \leq \frac{3q\alpha^2 L^2}{nB} \sum_{t=0}^{q-1} \mathbb{E} [\|\bar{\mathbf{v}}^{t,s}\|^2] + \frac{6qL^2}{nB} \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2}{n} \right].$$

Proof. See Appendix C. \square

Note that Lemma 4.4 shows that the accumulated gradient estimation error over one inner loop may be bounded by the accumulated state consensus error and the norm of the gradient estimators. Lemma 4.4 thus may be used to simplify the right-hand side of the descent inequality in Lemma 4.3. Naturally, what is left is to seek an upper bound for the state consensus error in terms of $\mathbb{E}[\|\bar{\mathbf{v}}^{t,s}\|^2]$. This result is presented in the following lemma.

LEMMA 4.5. *If the step-size follows $0 < \alpha \leq \frac{(1-\lambda^2)^2}{8\sqrt{42}L}$, then*

$$\sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2}{n} \right] \leq \frac{64\alpha^2}{(1-\lambda^2)^3} \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n} + \frac{1536\alpha^4 L^2}{(1-\lambda^2)^4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} [\|\bar{\mathbf{v}}^{t,s}\|^2].$$

Proof. See Appendix D. \square

Establishing Lemma 4.5 requires a careful analysis; here, we provide a brief sketch. Recall the GT-SARAH algorithm in (4.1a)–(4.1b), and note that the state vector $\mathbf{x}^{t,s}$ is coupled with the gradient tracker $\mathbf{y}^{t,s}$. Thus, in order to quantify the state consensus error $\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2$, we need to establish its relationship with the gradient tracking error $\|\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2$. In fact, we show that these coupled errors jointly formulate a linear time-invariant (LTI) system dynamics whose system matrix is stable under a certain range of the step-size α . Solving this LTI yields Lemma 4.5.

Finally, it is straightforward to use Lemmas 4.4 and 4.5 to refine the descent inequality in Lemma 4.3 to obtain the following result.

LEMMA 4.6. *If $0 < \alpha \leq \bar{\alpha} := \min \left\{ \frac{(1-\lambda^2)^2}{4\sqrt{42}}, \left(\frac{nB}{6q} \right)^{1/2}, \left(\frac{4nB}{7nB+24q} \right)^{1/4} \frac{1-\lambda^2}{6} \right\} \frac{1}{2L}$, then*

$$\begin{aligned} L^2 \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2}{n} \right] &+ \frac{1}{n} \sum_{i=1}^n \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} [\|\nabla F(\mathbf{x}_i^{t,s})\|^2] \\ &\leq \frac{4(F(\bar{\mathbf{x}}^{0,1}) - F^*)}{\alpha} + \left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{256\alpha^2 L^2}{(1-\lambda^2)^3} \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n}. \end{aligned}$$

Proof. See Appendix E. \square

We note that the descent inequality in Lemma 4.6 that characterizes the convergence of GT-SARAH is independent of the variance of local gradient estimators and of the difference between the local and the global gradient. In fact, it has similarities to that of the centralized *batch* gradient descent [4, 25]; see also the discussion on DSGD in section 1. This is a consequence of the joint use of the local variance reduction and the global gradient tracking. This is essentially why we are able to match the gradient complexity of the centralized near-optimal methods for finite-sum problems and obtain the almost sure convergence guarantee of GT-SARAH to a stationary point.

4.2. Proofs of the main theorems. With the refined descent inequality in Lemma 4.6 at hand, Theorems 3.3, 3.5, and 3.6 are now straightforward to prove.

Proof of Theorem 3.3. We observe from Lemma 4.6 that if $0 < \alpha \leq \bar{\alpha}$, then $\sum_{s=1}^{\infty} \sum_{t=0}^q \mathbb{E} [\|\nabla F(\mathbf{x}_i^{t,s})\|^2 + L^2 \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2] < \infty \forall i \in \mathcal{V}$, which implies all nodes

achieve consensus and converge to a stationary point in the mean-squared sense. Further, by the monotone convergence theorem [46], we exchange the order of the expectation and the series to obtain $\mathbb{E} \left[\sum_{s=1}^{\infty} \sum_{t=0}^q (\|\nabla F(\mathbf{x}_i^{t,s})\|^2 + L^2 \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2) \right] < \infty \forall i \in \mathcal{V}$, which leads to $\mathbb{P} \left(\sum_{s=1}^{\infty} \sum_{t=0}^q (\|\nabla F(\mathbf{x}_i^{t,s})\|^2 + L^2 \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2) < \infty \right) = 1 \forall i \in \mathcal{V}$, i.e., the consensus and convergence to a stationary point in the almost sure sense. \square

Proof of Theorem 3.5. We recall the metric of the outer loop complexity in Definition 3.4, and we divide the descent inequality in Lemma 4.6 by $S(q+1)$ from both sides. It is then clear that to find an ϵ -accurate stationary point of F , it suffices to choose the total number of the outer loop iterations S such that

$$(4.3) \quad \frac{4(F(\bar{\mathbf{x}}^{0,1}) - F^*)}{S(q+1)\alpha} + \left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{256\alpha^2 L^2}{S(q+1)(1-\lambda^2)^3} \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n} \leq \epsilon^2.$$

The proof follows by the fact that if $0 < \alpha \leq \left(\frac{4nB}{7nB+24q} \right)^{1/3} \frac{1-\lambda^2}{12L}$ then $\left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{256\alpha^2 L^2}{(1-\lambda^2)^3} \leq \frac{1}{\alpha L}$ and by solving for the lower bound on S such that (4.3) holds. \square

Proof of Theorem 3.6. During each inner loop, GT-SARAH incurs $n(m + 2qB)$ component gradient computations across all nodes and q rounds of communication of the network. Hence, to find an ϵ -accurate stationary point of F , GT-SARAH requires, according to Theorem 3.5, at most

$$\mathcal{H} = \mathcal{O} \left(\frac{n(m + qB)}{q\alpha L \epsilon^2} \left(L(F(\bar{\mathbf{x}}^{0,1}) - F^*) + \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n} \right) \right)$$

component gradient computations across all nodes and

$$\mathcal{K} = \mathcal{O} \left(\frac{1}{\alpha L \epsilon^2} \left(L(F(\bar{\mathbf{x}}^{0,1}) - F^*) + \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n} \right) \right)$$

rounds of communication of the network. The proof follows by setting the step-size α as its upper bound in Theorem 3.5 and the length of the inner loop as $q = \mathcal{O}(\frac{m}{B})$. \square

5. Numerical experiments. In this section, we illustrate, by numerical experiments, our main theoretical claim that GT-SARAH finds a first-order stationary point of problem (1.1) with a significantly improved gradient complexity compared to the existing decentralized stochastic gradient methods.

5.1. Setup. We consider a *nonconvex* logistic regression model [2] for binary classification over a decentralized network of n nodes with m data samples at each node: $\min_{\mathbf{x} \in \mathbb{R}^p} F(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n \frac{1}{m} \sum_{j=1}^m (f_{i,j}(\mathbf{x}) + r(\mathbf{x}))$ such that the logistic loss $f_{i,j}(\mathbf{x})$ and the nonconvex regularization $r(\mathbf{x})$ are given by

$$(5.1) \quad f_{i,j}(\mathbf{x}) := \log \left[1 + \exp \left\{ -(\mathbf{x}^\top \boldsymbol{\theta}_{i,j}) \xi_{i,j} \right\} \right] \text{ and } r(\mathbf{x}) := R \sum_{d=1}^p [\mathbf{x}]_d^2 (1 + [\mathbf{x}]_d^2)^{-1},$$

where $[\mathbf{x}]_d$ denotes the d th coordinate of \mathbf{x} . In (5.1), note that $\boldsymbol{\theta}_{i,j} \in \mathbb{R}^p$ is the j th data sample at the i th node and $\xi_{i,j} \in \{-1, +1\}$ is the corresponding binary label. The details of the datasets under consideration are provided in Table 2. We normalize each data sample such that $\|\boldsymbol{\theta}_{i,j}\| = 1$ for all i, j and set the regularization parameter as $R = 10^{-3}$. The primitive doubly stochastic weight matrices associated with the networks are generated by the lazy Metropolis rule [22]. We characterize the performance of the algorithms in comparison in terms of the decrease of the network stationary gap versus epochs, where the stationary gap is defined as $\|\nabla F(\bar{\mathbf{x}})\| + \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \bar{\mathbf{x}}\|$, where \mathbf{x}_i is the estimate of the stationary point of F at node i and $\bar{\mathbf{x}} := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$, and each epoch represents m component gradient computations at each node.

TABLE 2
Datasets used in numerical experiments, available at <https://www.openml.org/>.

Dataset	Number of samples ($N = nm$)	Dimension (p)
covertype	100,000	54
MiniBooNE	100,000	51
KDD98	82,000	477
w8a	60,000	300
a9a	48,800	124
Fashion-MNIST (T-shirt versus dress)	10,000	784

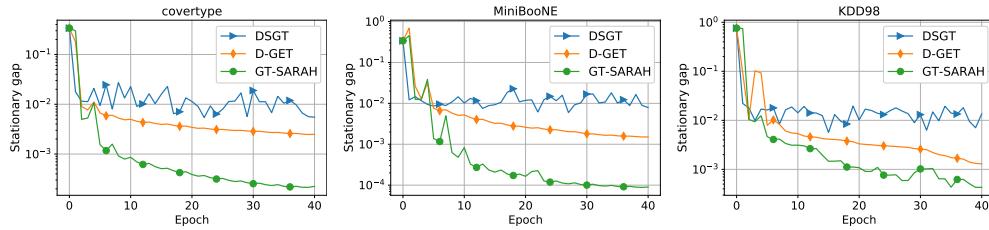


FIG. 2. Performance comparison of GT-SARAH, DSGT, and D-GET over a 10-node exponential graph on the covertype, MiniBooNE, and KDD98 datasets.

5.2. Performance comparisons. We compare the performance of GT-SARAH with DSGT [50] and D-GET [38]; we note that D2 [40] and DSGD [20] are not presented here for conciseness, since in general the former achieves a similar performance with DSGT and the latter underperforms DSGT and D2 [6, 40, 48]. Towards the parameter selection of each algorithm, we use the following setup: (i) for GT-SARAH, we choose its minibatch size as $B = 1$ and its inner loop length as $q = m$ in light of Corollary 3.9; (ii) for D-GET, we choose its minibatch size and inner loop length as $\lfloor m^{1/2} \rfloor$ under which its convergence is established; see Theorem 1 in [38]; and (iii) we manually optimize the step-sizes for GT-SARAH, D-GET, and DSGT across all experiments.

We first compare the performances of GT-SARAH, DSGT, and D-GET in the big-data regime, that is, the number of samples m at each node is relatively large. To this aim, we distribute the covertype, MiniBooNE, and KDD98 datasets over a 10-node exponential graph [22] whose associated second largest singular value $\lambda \approx 0.71$. The experimental results are presented in Figure 2, where GT-SARAH outperforms DSGT and D-GET. We also observe that D-GET outperforms DSGT in this case since the performance of the latter is deteriorated by the large variance of the stochastic gradients as the number of the samples m at each node is large.

We next consider the large-scale network regime, where the network spectral gap inverse $(1 - \lambda)^{-1}$ and the number of the nodes n are relatively large compared with the local sample size m . We distribute the w8a, a9a, and Fashion-MNIST datasets over the $n = 10 \times 10$ grid graph whose associated second largest eigenvalue $\lambda \approx 0.99$. The performance comparison of the algorithms is shown in Figure 3, where we observe that GT-SARAH still outperforms DSGT and D-GET. Besides, it is worth noting that D-GET underperforms DSGT in this case. We provide an explanation about this phenomenon in the following. In the regime where m is relatively small, the variance of the stochastic gradients is relatively small, and as a consequence DSGT performs well. On the other hand, the minibatch size $\lfloor m^{1/2} \rfloor$ of D-GET is too large in this regime to achieve a satisfactory performance; see the related discussion in subsection 3.3.1.

6. Conclusions. In this paper, we propose GT-SARAH to minimize a finite sum of N smooth nonconvex functions, equally distributed over a decentralized network of n nodes. With appropriate algorithmic parameters, GT-SARAH achieves significantly

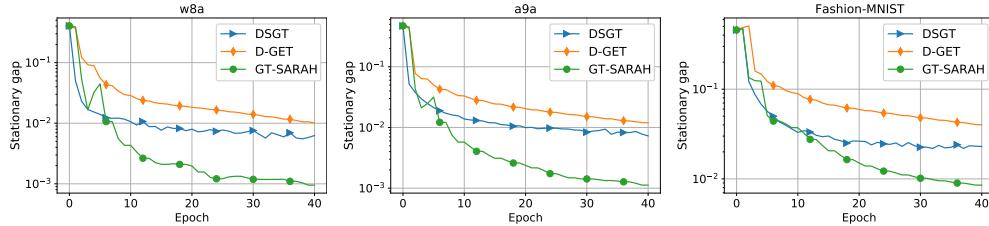


FIG. 3. Performance comparison of GT-SARAH, DSGT, and D-GET over the 10×10 grid graph on the w8a, a9a, and Fashion-MNIST datasets.

improved gradient complexity compared with the existing decentralized stochastic gradient methods. In particular, in a big-data regime $n = \mathcal{O}(N^{1/2}(1 - \lambda)^3)$, the gradient complexity of GT-SARAH reduces to $\mathcal{O}(N^{1/2}L\epsilon^{-2})$ which matches that of the centralized near-optimal variance-reduced methods such as SPIDER [11, 45] and SARAH [28] for problem (1.1), where L is the smoothness parameter and $(1 - \lambda)$ is the spectral gap of the network weight matrix. Moreover, GT-SARAH in this regime achieves nonasymptotic linear speedup compared with the centralized near-optimal approaches that perform all gradient computations at a single node. Compared with the minibatch implementations of SPIDER and SARAH over server-worker architectures [5], the decentralized GT-SARAH enjoys the same nonasymptotic linear speedup in terms of the gradient complexity but admits sparser and more flexible communication topology.

Appendix A. Preliminaries of the convergence analysis. In this section, we present the preliminaries for the proofs of the technical Lemmas 4.1, 4.3, 4.4, 4.5, and 4.6. We first define the natural filtration associated with the probability space, an increasing family of sub- σ -algebras of \mathcal{F} , as

$$\mathcal{F}^{t,s} := \sigma \left(\sigma(\tau_{i,l}^{t-1,s} : i \in \mathcal{V}, l \in [1, B]), \mathcal{F}^{t-1,s} \right), \quad t \in [2, q+1], s \geq 1,$$

where $\mathcal{F}^{1,s} := \mathcal{F}^{0,s} := \mathcal{F}^{q+1,s-1}$, $s \geq 2$, and $\mathcal{F}^{1,1} := \mathcal{F}^{0,1} := \{\phi, \Xi\}$. It can be verified by induction that $\mathbf{x}^{t,s}, \mathbf{y}^{t,s}$ are $\mathcal{F}^{t,s}$ -measurable and $\mathbf{v}^{t,s}$ is $\mathcal{F}^{t+1,s}$ -measurable $\forall s \geq 1$ and $t \in [0, q]$. We assume that the starting point $\bar{\mathbf{x}}^{0,1}$ of GT-SARAH is a constant vector. We next present some standard results in the context of decentralized optimization and gradient tracking methods. The following lemma provides an upper bound on the difference between the exact global gradient and the average of local batch gradients in terms of the state consensus error, as a result of the L -smoothness of each f_i .

LEMMA A.1. $\|\bar{\nabla}f(\mathbf{x}^{t,s}) - \nabla F(\bar{\mathbf{x}}^{t,s})\|^2 \leq \frac{L^2}{n} \|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2$, $\forall s \geq 1$ and $t \in [0, q]$.

Proof. Observe that, $\forall s \geq 1$ and $t \in [0, q]$,

$$\begin{aligned} \|\bar{\nabla}f(\mathbf{x}^{t,s}) - \nabla F(\bar{\mathbf{x}}^{t,s})\|^2 &= \frac{1}{n^2} \left\| \sum_{i=1}^n (\nabla f_i(\mathbf{x}_i^{t,s}) - \nabla f_i(\bar{\mathbf{x}}^{t,s})) \right\|^2 \\ &= \frac{1}{n} \sum_{i=1}^n \|\nabla f_i(\mathbf{x}_i^{t,s}) - \nabla f_i(\bar{\mathbf{x}}^{t,s})\|^2 \\ &\leq \frac{L^2}{n} \sum_{i=1}^n \|\mathbf{x}_i^{t,s} - \bar{\mathbf{x}}^{t,s}\|^2, \end{aligned}$$

where the last line is due to the L -smoothness of each f_i . The proof is complete. \square

The following are some standard inequalities on the state consensus error.

LEMMA A.2. *The following inequalities hold: $\forall s \geq 1$ and $t \in [0, q]$,*

$$(A.1) \quad \|\mathbf{x}^{t+1,s} - \mathbf{J}\mathbf{x}^{t+1,s}\|^2 \leq \frac{1+\lambda^2}{2} \|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2 + \frac{2\alpha^2}{1-\lambda^2} \|\mathbf{y}^{t+1,s} - \mathbf{J}\mathbf{y}^{t+1,s}\|^2,$$

$$(A.2) \quad \|\mathbf{x}^{t+1,s} - \mathbf{J}\mathbf{x}^{t+1,s}\|^2 \leq 2 \|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2 + 2\alpha^2 \|\mathbf{y}^{t+1,s} - \mathbf{J}\mathbf{y}^{t+1,s}\|^2.$$

Proof. Using (4.1b) and the fact that $\mathbf{JW} = \mathbf{J}$, we have, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$\begin{aligned} \|\mathbf{x}^{t+1,s} - \mathbf{Jx}^{t+1,s}\|^2 &= \|\mathbf{Wx}^{t,s} - \alpha\mathbf{y}^{t+1,s} - \mathbf{J}(\mathbf{Wx}^{t,s} - \alpha\mathbf{y}^{t+1,s})\|^2 \\ (A.3) \quad &= \|\mathbf{Wx}^{t,s} - \mathbf{Jx}^{t,s} - \alpha(\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s})\|^2 \end{aligned}$$

We apply Young's inequality, $\|\mathbf{a} + \mathbf{b}\|^2 \leq (1 + \eta)\|\mathbf{a}\|^2 + (1 + \eta^{-1})\|\mathbf{b}\|^2 \ \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^{np}$, $\forall \eta > 0$, and Lemma 4.2 to (A.3) to obtain, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$\|\mathbf{x}^{t+1,s} - \mathbf{Jx}^{t+1,s}\|^2 \leq (1 + \eta)\lambda^2\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2 + (1 + \eta^{-1})\alpha^2\|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2.$$

Setting η as $\frac{1-\lambda^2}{2\lambda^2}$ and 1, respectively, yields (A.1) and (A.2). \square

A.1. Proof of Lemma 4.1. Recall from Assumption 3.3 that $\mathbf{1}_n^\top \mathbf{W} = \mathbf{1}_n^\top$. We multiply (4.1a) by $\frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p)$ to obtain, $\forall s \geq 1$ and $t \in [0, q]$,

$$\begin{aligned} \bar{\mathbf{y}}^{t+1,s} &= \bar{\mathbf{y}}^{t,s} + \bar{\mathbf{v}}^{t,s} - \bar{\mathbf{v}}^{t-1,s} \\ &= \bar{\mathbf{y}}^{t-1,s} + \bar{\mathbf{v}}^{t,s} - \bar{\mathbf{v}}^{t-2,s} \\ &\dots \\ &= \bar{\mathbf{y}}^{0,s} + \bar{\mathbf{v}}^{t,s} - \bar{\mathbf{v}}^{-1,s} \\ &= \bar{\mathbf{y}}^{q+1,s-1} + \bar{\mathbf{v}}^{t,s} - \bar{\mathbf{v}}^{q,s-1} \\ &\dots \\ &= \bar{\mathbf{y}}^{0,1} + \bar{\mathbf{v}}^{t,s} - \bar{\mathbf{v}}^{-1,1} = \bar{\mathbf{v}}^{t,s}, \end{aligned}$$

where the above series of equalities follows directly from the updates of **GT-SARAH**.

Appendix B. Proof of Lemma 4.3. We multiply (4.1b) by $\frac{1}{n}(\mathbf{1}_n^\top \otimes \mathbf{I}_p)$ and then use Lemma 4.1 to obtain the recursion of the mean state $\bar{\mathbf{x}}^{t,s}$ as follows:

$$\bar{\mathbf{x}}^{t+1,s} = \bar{\mathbf{x}}^{t,s} - \alpha\bar{\mathbf{y}}^{t+1,s} = \bar{\mathbf{x}}^{t,s} - \alpha\bar{\mathbf{v}}^{t,s} \quad \forall s \geq 1 \text{ and } t \in [0, q].$$

Setting $\mathbf{y} = \bar{\mathbf{x}}^{t+1,s}$ and $\mathbf{x} = \bar{\mathbf{x}}^{t,s}$ in (4.2), we have, $\forall s \geq 1$ and $t \in [0, q]$,

$$(B.1) \quad F(\bar{\mathbf{x}}^{t+1,s}) \leq F(\bar{\mathbf{x}}^{t,s}) - \alpha\langle \nabla F(\bar{\mathbf{x}}^{t,s}), \bar{\mathbf{v}}^{t,s} \rangle + \frac{\alpha^2 L}{2}\|\bar{\mathbf{v}}^{t,s}\|^2.$$

Applying $\langle \mathbf{a}, \mathbf{b} \rangle = 0.5(\|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - \|\mathbf{a} - \mathbf{b}\|^2) \ \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^p$ to (B.1), we obtain an inequality that characterizes the descent of the network mean state over one inner loop iteration: $\forall s \geq 1$ and $t \in [0, q]$,

$$\begin{aligned} F(\bar{\mathbf{x}}^{t+1,s}) &\leq F(\bar{\mathbf{x}}^{t,s}) - \frac{\alpha}{2}\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2 - \frac{\alpha(1-\alpha L)}{2}\|\bar{\mathbf{v}}^{t,s}\|^2 + \frac{\alpha}{2}\|\bar{\mathbf{v}}^{t,s} - \nabla F(\bar{\mathbf{x}}^{t,s})\|^2, \\ &\leq F(\bar{\mathbf{x}}^{t,s}) - \frac{\alpha}{2}\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2 - \frac{\alpha}{4}\|\bar{\mathbf{v}}^{t,s}\|^2 + \alpha\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla f}(\mathbf{x}^{t,s})\|^2 \\ (B.2) \quad &+ \alpha\|\bar{\nabla f}(\mathbf{x}^{t,s}) - \nabla F(\bar{\mathbf{x}}^{t,s})\|^2 \end{aligned}$$

$$\begin{aligned} &\leq F(\bar{\mathbf{x}}^{t,s}) - \frac{\alpha}{2}\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2 - \frac{\alpha}{4}\|\bar{\mathbf{v}}^{t,s}\|^2 + \alpha\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla f}(\mathbf{x}^{t,s})\|^2 \\ (B.3) \quad &+ \frac{\alpha L^2}{n}\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2, \end{aligned}$$

where (B.2) is due to $0 < \alpha \leq \frac{1}{2L}$ and (B.3) is due to Lemma A.1. We then take the telescoping sum of (B.3) over t from 0 to q to obtain, $\forall s \geq 1$,

$$\begin{aligned} (B.4) \quad F(\bar{\mathbf{x}}^{0,s+1}) &\leq F(\bar{\mathbf{x}}^{0,s}) - \frac{\alpha}{2}\sum_{t=0}^q\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2 - \frac{\alpha}{4}\sum_{t=0}^q\|\bar{\mathbf{v}}^{t,s}\|^2 \\ &+ \alpha\sum_{t=0}^q\|\bar{\mathbf{v}}^{t,s} - \bar{\nabla f}(\mathbf{x}^{t,s})\|^2 + \frac{\alpha L^2}{n}\sum_{t=0}^q\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2. \end{aligned}$$

The proof then follows by taking the telescoping sum of (B.4) over s from 1 to S and taking the expectation of the resulting inequality.

Appendix C. Proof of Lemma 4.4. We first provide a useful result.

LEMMA C.1. *The following inequality holds: $\forall s \geq 1, \forall t \in [1, q], \forall i \in \mathcal{V}, \forall l \in [1, B]$,*

$$\mathbb{E} \left[\left\| \nabla f_{i, \tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i, \tau_{i,l}^{t-1,s}}(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \leq L^2 \|\mathbf{x}_i^{t,s} - \mathbf{x}_i^{t-1,s}\|^2.$$

Proof. In the following, we denote $\mathbb{1}\{A\}$ as the indicator function of an event $A \in \mathcal{F}$. Observe that, $\forall s \geq 1, \forall t \in [1, q], \forall i \in \mathcal{V}, \forall l \in [1, B]$,

$$\begin{aligned} & \mathbb{E} \left[\left\| \nabla f_{i, \tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i, \tau_{i,l}^{t-1,s}}(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\ &= \mathbb{E} \left[\left\| \sum_{j=1}^m \mathbb{1}\{\tau_{i,l}^{t,s} = j\} \left(\nabla f_{i,j}(\mathbf{x}_i^{t,s}) - \nabla f_{i,j}(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\ &= \sum_{j=1}^m \mathbb{E} \left[\mathbb{1}\{\tau_{i,l}^{t,s} = j\} \left\| \nabla f_{i,j}(\mathbf{x}_i^{t,s}) - \nabla f_{i,j}(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\ &= \sum_{j=1}^m \mathbb{E} \left[\mathbb{1}\{\tau_{i,l}^{t,s} = j\} \middle| \mathcal{F}^{t,s} \right] \left\| \nabla f_{i,j}(\mathbf{x}_i^{t,s}) - \nabla f_{i,j}(\mathbf{x}_i^{t-1,s}) \right\|^2 \\ &= \frac{1}{m} \sum_{j=1}^m \left\| \nabla f_{i,j}(\mathbf{x}_i^{t,s}) - \nabla f_{i,j}(\mathbf{x}_i^{t-1,s}) \right\|^2, \end{aligned}$$

where the last line uses that $\tau_{i,l}^{t,s}$ is independent of $\mathcal{F}^{t,s}$, i.e., $\mathbb{E}[\mathbb{1}\{\tau_{i,l}^{t,s} = j\} | \mathcal{F}^{t,s}] = \frac{1}{m}$. The proof follows by using Assumption 3.1. \square

Next, we derive an upper bound on the estimation error of the average of local SARAH gradient estimators across the nodes at each inner loop iteration.

LEMMA C.2. *The following inequality holds: $\forall s \geq 1$ and $t \in [1, q]$,*

$$\mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \right] \leq \frac{3\alpha^2 L^2}{nB} \sum_{u=0}^{t-1} \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{u,s} \right\|^2 \right] + \frac{6L^2}{n^2 B} \sum_{u=0}^t \mathbb{E} \left[\left\| \mathbf{x}^{u,s} - \mathbf{J} \mathbf{x}^{u,s} \right\|^2 \right].$$

Proof. For the ease of exposition, we denote, $\forall t \in [1, q], \forall s \geq 1, \forall i \in \mathcal{V}, \forall l \in [1, B]$,

$$(C.1) \quad \hat{\nabla}_{i,l}^{t,s} := \nabla f_{i, \tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i, \tau_{i,l}^{t-1,s}}(\mathbf{x}_i^{t-1,s}), \quad \hat{\nabla}_i^{t,s} := \frac{1}{B} \sum_{l=1}^B \hat{\nabla}_{i,l}^{t,s}.$$

Since $\mathbf{x}_i^{t,s}$ and $\mathbf{x}_i^{t-1,s}$ are $\mathcal{F}^{t,s}$ -measurable, we have

$$(C.2) \quad \mathbb{E} \left[\hat{\nabla}_{i,l}^{t,s} \middle| \mathcal{F}^{t,s} \right] = \mathbb{E} \left[\hat{\nabla}_i^{t,s} \middle| \mathcal{F}^{t,s} \right] = \nabla f_i(\mathbf{x}_i^{t,s}) - \nabla f_i(\mathbf{x}_i^{t-1,s}).$$

With the notations in (C.1), the local recursive update of the gradient estimator $\mathbf{v}_i^{t,s}$ described in Algorithm 2.1 may be written as

$$\mathbf{v}_i^{t,s} = \hat{\nabla}_i^{t,s} + \mathbf{v}_i^{t-1,s} \quad \forall t \in [1, q], \forall s \geq 1, \forall i \in \mathcal{V}.$$

In the light of (C.2), we have the following: $\forall s \geq 1$ and $t \in [1, q]$,

$$\begin{aligned}
 & \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \left(\hat{\nabla}_i^{t,s} + \mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \left(\hat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) + \mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \left(\hat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &\quad + \left\| \frac{1}{n} \sum_{i=1}^n \left(\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \\
 &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n \left(\hat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 (C.3) \quad &+ \left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2,
 \end{aligned}$$

where the third equality is due to (C.2) and the fact that $\sum_{i=1}^n (\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s}))$ is $\mathcal{F}^{t,s}$ -measurable. To proceed from (C.3), we note that since the collection of random variables $\{\tau_{i,l}^{t,s} : i \in \mathcal{V}, l \in [1, B]\}$ are independent of each other and of the filtration $\mathcal{F}^{t,s}$, by (C.2), we have, $\forall t \in [1, q]$ and $s \geq 1$,

$$(C.4) \quad \mathbb{E} \left[\left\langle \hat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}), \hat{\nabla}_r^{t,s} - \nabla f_r(\mathbf{x}_r^{t,s}) + \nabla f_r(\mathbf{x}_r^{t-1,s}) \right\rangle \middle| \mathcal{F}^{t,s} \right] = 0$$

whenever $i, r \in \mathcal{V}$ such that $i \neq r$. Similarly, we have, $\forall t \in [1, q]$ and $s \geq 1$, $\forall i \in \mathcal{V}$,

$$(C.5) \quad \mathbb{E} \left[\left\langle \hat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}), \hat{\nabla}_{i,h}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\rangle \middle| \mathcal{F}^{t,s} \right] = 0$$

whenever $l, h \in [1, m]$ such that $l \neq h$. With the help of (C.4) and (C.5), we may simplify (C.3) in the following: $\forall s \geq 1$ and $t \in [1, q]$,

$$\begin{aligned}
 & \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\left\| \hat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] + \left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2 \\
 &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[\left\| \frac{1}{B} \sum_{l=1}^B \left(\hat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &\quad + \left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2 \\
 (C.6) \quad &= \frac{1}{(nB)^2} \sum_{i=1}^n \sum_{l=1}^B \mathbb{E} \left[\left\| \hat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &\quad + \left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2,
 \end{aligned}$$

where the first line is due to (C.4) and the last line is due to (C.5). To proceed from (C.6), we observe that $\forall t \in [1, q], \forall s \geq 1, \forall i \in \mathcal{V}, \forall l \in [1, B]$,

$$\begin{aligned}
 \mathbb{E} \left[\left\| \widehat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] &= \mathbb{E} \left[\left\| \widehat{\nabla}_{i,l}^{t,s} - \mathbb{E} \left[\widehat{\nabla}_{i,l}^{t,s} \middle| \mathcal{F}^{t,s} \right] \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 &\leq \mathbb{E} \left[\left\| \widehat{\nabla}_{i,l}^{t,s} \right\|^2 \middle| \mathcal{F}^{t,s} \right] \\
 (C.7) \quad &\leq L^2 \left\| \mathbf{x}_i^{t,s} - \mathbf{x}_i^{t-1,s} \right\|^2,
 \end{aligned}$$

where the last line uses Lemma C.1. Applying (C.7) to (C.6) yields, $\forall s \geq 1, t \in [1, q]$,

$$(C.8) \quad \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \middle| \mathcal{F}^{t,s} \right] \leq \frac{L^2}{n^2 B} \left\| \mathbf{x}^{t,s} - \mathbf{x}^{t-1,s} \right\|^2 + \left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2.$$

We next bound the first term on the right-hand side of (C.8). Observe that $\forall s \geq 1$ and $t \in [1, q+1]$,

$$\begin{aligned}
 \left\| \mathbf{x}^{t,s} - \mathbf{x}^{t-1,s} \right\|^2 &= \left\| \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t,s} + \mathbf{J} \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t-1,s} + \mathbf{J} \mathbf{x}^{t-1,s} - \mathbf{x}^{t-1,s} \right\|^2 \\
 &\leq 3 \left\| \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t,s} \right\|^2 + 3n \left\| \bar{\mathbf{x}}^{t,s} - \bar{\mathbf{x}}^{t-1,s} \right\|^2 + 3 \left\| \mathbf{x}^{t-1,s} - \mathbf{J} \mathbf{x}^{t-1,s} \right\|^2 \\
 (C.9) \quad &= 3 \left\| \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t,s} \right\|^2 + 3n\alpha^2 \left\| \bar{\mathbf{v}}^{t-1,s} \right\|^2 + 3 \left\| \mathbf{x}^{t-1,s} - \mathbf{J} \mathbf{x}^{t-1,s} \right\|^2.
 \end{aligned}$$

Applying (C.9) to (C.8) and taking the expectation of the resulting inequality leads to, $\forall s \geq 1$ and $t \in [1, q]$,

$$\begin{aligned}
 \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \right] &\leq \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t-1,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t-1,s}) \right\|^2 \right] + \frac{3\alpha^2 L^2}{nB} \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t-1,s} \right\|^2 \right] \\
 (C.10) \quad &+ \frac{3L^2}{n^2 B} \mathbb{E} \left[\left\| \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t,s} \right\|^2 \right] + \frac{3L^2}{n^2 B} \mathbb{E} \left[\left\| \mathbf{x}^{t-1,s} - \mathbf{J} \mathbf{x}^{t-1,s} \right\|^2 \right].
 \end{aligned}$$

We recall the initialization of each inner loop that $\bar{\mathbf{v}}^{0,s} = \bar{\nabla} \mathbf{f}(\mathbf{x}^{0,s}) \forall s \geq 1$ and take the telescoping sum of (C.10) over t from 1 to z to obtain, $\forall s \geq 1$ and $\forall z \in [1, q]$,

$$\begin{aligned}
 \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{z,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{z,s}) \right\|^2 \right] &\leq \frac{3\alpha^2 L^2}{nB} \sum_{t=1}^z \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t-1,s} \right\|^2 \right] + \frac{3L^2}{n^2 B} \sum_{t=1}^z \mathbb{E} \left[\left\| \mathbf{x}^{t,s} - \mathbf{J} \mathbf{x}^{t,s} \right\|^2 \right] \\
 (C.11) \quad &+ \frac{3L^2}{n^2 B} \sum_{t=1}^z \mathbb{E} \left[\left\| \mathbf{x}^{t-1,s} - \mathbf{J} \mathbf{x}^{t-1,s} \right\|^2 \right].
 \end{aligned}$$

The proof follows by merging the last two terms on the right-hand side of (C.11). \square

Proof of Lemma 4.4. Summing up Lemma C.2 over t from 1 to q gives, $\forall s \geq 1$,

$$\begin{aligned}
 \sum_{t=1}^q \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{t,s} - \bar{\nabla} \mathbf{f}(\mathbf{x}^{t,s}) \right\|^2 \right] &\leq \frac{3\alpha^2 L^2}{nB} \sum_{t=1}^q \sum_{u=0}^{t-1} \mathbb{E} \left[\left\| \bar{\mathbf{v}}^{u,s} \right\|^2 \right] \\
 (C.12) \quad &+ \frac{6L^2}{n^2 B} \sum_{t=1}^q \sum_{u=0}^t \mathbb{E} \left[\left\| \mathbf{x}^{u,s} - \mathbf{J} \mathbf{x}^{u,s} \right\|^2 \right].
 \end{aligned}$$

The proof follows by relaxing the right-hand side of (C.12) on the summations and the initialization of each inner loop that $\bar{\mathbf{v}}^{0,s} = \bar{\nabla} \mathbf{f}(\mathbf{x}^{0,s}) \forall s \geq 1$. \square

Appendix D. Proof of Lemma 4.5.

D.1. Gradient tracking error. We first provide some useful bounds on the gradient estimator tracking errors. These bounds will later be coupled with (A.1) to formulate a dynamical system to characterize the error evolution of GT-SARAH.

The following lemma establishes an upper bound on the sum of the local gradient estimation errors across the nodes. Its proof is similar to that of Lemma C.2.

LEMMA D.1. *The following inequality holds $\forall s \geq 1$ and $t \in [1, q]$:*

$$\mathbb{E} \left[\|\mathbf{v}^{t,s} - \nabla \mathbf{f}(\mathbf{x}^{t,s})\|^2 \right] \leq \frac{3n\alpha^2 L^2}{B} \sum_{u=0}^{t-1} \mathbb{E} \left[\|\bar{\mathbf{v}}^{u,s}\|^2 \right] + \frac{6L^2}{B} \sum_{u=0}^t \mathbb{E} \left[\|\mathbf{x}^{u,s} - \mathbf{Jx}^{u,s}\|^2 \right].$$

Proof. See Appendix F. \square

We note that Lemma D.1 does not follow directly from the results of Lemma 4.4 because $\mathbf{v}_i^{t,s}$ is not a conditionally unbiased estimator of $\nabla f_i(\mathbf{x}_i^{t,s})$ with respect to $\mathcal{F}^{t,s}$. With Lemma D.1 at hand, we now quantify the gradient tracking errors.

LEMMA D.2. *We have the following three statements.*

- (i) *It holds that $\|\mathbf{y}^{1,1} - \mathbf{Jy}^{1,1}\|^2 \leq \|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2$.*
- (ii) *If $0 < \alpha \leq \frac{1-\lambda^2}{4\sqrt{3}L}$, the following inequality holds: $\forall s \geq 1$ and $t \in [1, q]$,*

$$\begin{aligned} \mathbb{E} \left[\frac{\|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2}{nL^2} \right] &\leq \frac{3 + \lambda^2}{4} \mathbb{E} \left[\frac{\|\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2}{nL^2} \right] \\ &\quad + \frac{18}{1 - \lambda^2} \mathbb{E} \left[\frac{\|\mathbf{x}^{t-1,s} - \mathbf{Jx}^{t-1,s}\|^2}{n} \right] + \frac{6\alpha^2}{1 - \lambda^2} \mathbb{E} \left[\|\bar{\mathbf{v}}^{t-1,s}\|^2 \right]. \end{aligned}$$

- (iii) *If $0 < \alpha \leq \frac{1-\lambda^2}{4\sqrt{6}L}$, the following inequality holds: $\forall s \geq 2$,*

$$\begin{aligned} \mathbb{E} \left[\frac{\|\mathbf{y}^{1,s} - \mathbf{Jy}^{1,s}\|^2}{nL^2} \right] &\leq \frac{3 + \lambda^2}{4} \mathbb{E} \left[\frac{\|\mathbf{y}^{q+1,s-1} - \mathbf{Jy}^{q+1,s-1}\|^2}{nL^2} \right] \\ &\quad + \frac{18}{1 - \lambda^2} \mathbb{E} \left[\frac{\|\mathbf{x}^{q,s-1} - \mathbf{Jx}^{q,s-1}\|^2}{n} \right] + \frac{12\alpha^2}{1 - \lambda^2} \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s-1}\|^2 \right] \\ &\quad + \frac{42}{1 - \lambda^2} \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s-1} - \mathbf{Jx}^{t,s-1}\|^2}{n} \right]. \end{aligned}$$

Proof. (i) Recall that $\mathbf{v}^{-1,1} = \mathbf{0}_{np}$, $\mathbf{y}^{0,1} = \mathbf{0}_{np}$, and $\mathbf{v}^{0,1} = \nabla \mathbf{f}(\mathbf{x}^{0,1})$. Using the gradient tracking update at iteration (1, 1) and $\|\mathbf{I}_{np} - \mathbf{J}\| = 1$, we have

$$\|\mathbf{y}^{1,1} - \mathbf{Jy}^{1,1}\|^2 = \|(\mathbf{I}_{np} - \mathbf{J})(\mathbf{W}\mathbf{y}^{0,1} + \mathbf{v}^{0,1} - \mathbf{v}^{-1,1})\|^2 \leq \|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2,$$

which proves the first statement in the lemma. In the following, we prove the second and the third statements. We have, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$\begin{aligned} \|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2 &= \|\mathbf{W}\mathbf{y}^{t,s} + \mathbf{v}^{t,s} - \mathbf{v}^{t-1,s} - \mathbf{J}(\mathbf{W}\mathbf{y}^{t,s} + \mathbf{v}^{t,s} - \mathbf{v}^{t-1,s})\|^2 \\ (D.1) \quad &= \|\mathbf{W}\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s} + (\mathbf{I}_{np} - \mathbf{J})(\mathbf{v}^{t,s} - \mathbf{v}^{t-1,s})\|^2. \end{aligned}$$

We apply the inequality that $\|\mathbf{a} + \mathbf{b}\|^2 \leq (1 + \eta)\|\mathbf{a}\|^2 + (1 + \frac{1}{\eta})\|\mathbf{b}\|^2 \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^{np}$ with $\eta = \frac{1-\lambda^2}{2\lambda^2}$ and that $\|\mathbf{I}_{np} - \mathbf{J}\| = 1$ to (D.1) to obtain, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$\begin{aligned} \|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2 &\leq \frac{1+\lambda^2}{2\lambda^2} \|\mathbf{W}\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2 + \frac{1+\lambda^2}{1-\lambda^2} \|\mathbf{v}^{t,s} - \mathbf{v}^{t-1,s}\|^2 \\ (D.2) \quad &\leq \frac{1+\lambda^2}{2} \|\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2 + \frac{2}{1-\lambda^2} \|\mathbf{v}^{t,s} - \mathbf{v}^{t-1,s}\|^2, \end{aligned}$$

where the last line is due to Lemma 4.2. Next, we derive upper bounds for the last term in (D.2) under different ranges of t and s .

(ii) $\forall t \in [1, q]$ and $\forall s \geq 1$. By the update of each local $\mathbf{v}_i^{t,s}$, we have that

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{v}^{t,s} - \mathbf{v}^{t-1,s}\|^2 \mid \mathcal{F}^{t,s} \right] &= \sum_{i=1}^n \mathbb{E} \left[\left\| \frac{1}{B} \sum_{l=1}^B \left(\nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i,\tau_{i,l}^{t-1,s}}(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \mid \mathcal{F}^{t,s} \right] \\ &\leq \frac{1}{B} \sum_{i=1}^n \sum_{l=1}^B \mathbb{E} \left[\left\| \nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i,\tau_{i,l}^{t-1,s}}(\mathbf{x}_i^{t-1,s}) \right\|^2 \mid \mathcal{F}^{t,s} \right] \\ (D.3) \quad &\leq L^2 \|\mathbf{x}^{t,s} - \mathbf{x}^{t-1,s}\|^2, \end{aligned}$$

where the last line is due to Lemma C.1. To proceed, we further use (C.9) and (A.2) to refine (D.3) as follows: $\forall s \geq 1$ and $\forall t \in [1, q]$,

$$\begin{aligned} &\mathbb{E} \left[\|\mathbf{v}^{t,s} - \mathbf{v}^{t-1,s}\|^2 \mid \mathcal{F}^{t,s} \right] \\ &\leq 3L^2 \|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2 + 3n\alpha^2 L^2 \|\bar{\mathbf{v}}^{t-1,s}\|^2 + 3L^2 \|\mathbf{x}^{t-1,s} - \mathbf{Jx}^{t-1,s}\|^2 \\ (D.4) \quad &\leq 3n\alpha^2 L^2 \|\bar{\mathbf{v}}^{t-1,s}\|^2 + 9L^2 \|\mathbf{x}^{t-1,s} - \mathbf{Jx}^{t-1,s}\|^2 + 6\alpha^2 L^2 \|\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2. \end{aligned}$$

We take the expectation of (D.4) and use it in (D.2) to obtain, $\forall s \geq 1$ and $\forall t \in [1, q]$,

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2 \right] &\leq \left(\frac{1+\lambda^2}{2} + \frac{12\alpha^2 L^2}{1-\lambda^2} \right) \mathbb{E} \left[\|\mathbf{y}^{t,s} - \mathbf{Jy}^{t,s}\|^2 \right] \\ &\quad + \frac{18L^2}{1-\lambda^2} \mathbb{E} \left[\|\mathbf{x}^{t-1,s} - \mathbf{Jx}^{t-1,s}\|^2 \right] + \frac{6n\alpha^2 L^2}{1-\lambda^2} \mathbb{E} \left[\|\bar{\mathbf{v}}^{t-1,s}\|^2 \right]. \end{aligned}$$

The second statement in the lemma follows by the fact that $\frac{1+\lambda^2}{2} + \frac{12\alpha^2 L^2}{1-\lambda^2} \leq \frac{3+\lambda^2}{4}$ if $0 < \alpha \leq \frac{1-\lambda^2}{4\sqrt{3}L}$.

(iii) $t = 0$ and $\forall s \geq 2$. By the update of GT-SARAH, we observe that, $\forall s \geq 2$,

$$\begin{aligned} \|\mathbf{v}^{0,s} - \mathbf{v}^{-1,s}\|^2 &= \|\nabla \mathbf{f}(\mathbf{x}^{q+1,s-1}) - \mathbf{v}^{q,s-1}\|^2 \\ &= \|\nabla \mathbf{f}(\mathbf{x}^{q+1,s-1}) - \nabla \mathbf{f}(\mathbf{x}^{q,s-1}) + \nabla \mathbf{f}(\mathbf{x}^{q,s-1}) - \mathbf{v}^{q,s-1}\|^2 \\ &\leq 2L^2 \|\mathbf{x}^{q+1,s-1} - \mathbf{x}^{q,s-1}\|^2 + 2 \|\nabla \mathbf{f}(\mathbf{x}^{q,s-1}) - \mathbf{v}^{q,s-1}\|^2, \\ &\leq 6L^2 \|\mathbf{x}^{q+1,s-1} - \mathbf{Jx}^{q+1,s-1}\|^2 + 6n\alpha^2 L^2 \|\bar{\mathbf{v}}^{q,s-1}\|^2 \\ &\quad + 6L^2 \|\mathbf{x}^{q,s-1} - \mathbf{Jx}^{q,s-1}\|^2 + 2 \|\nabla \mathbf{f}(\mathbf{x}^{q,s-1}) - \mathbf{v}^{q,s-1}\|^2 \\ &\leq 18L^2 \|\mathbf{x}^{q,s-1} - \mathbf{Jx}^{q,s-1}\|^2 + 6n\alpha^2 L^2 \|\bar{\mathbf{v}}^{q,s-1}\|^2 \\ (D.5) \quad &\quad + 12\alpha^2 L^2 \|\mathbf{y}^{q+1,s-1} - \mathbf{Jy}^{q+1,s-1}\|^2 + 2 \|\nabla \mathbf{f}(\mathbf{x}^{q,s-1}) - \mathbf{v}^{q,s-1}\|^2, \end{aligned}$$

where the first inequality uses the L -smoothness of each f_i , the second inequality uses (C.9), and the last inequality uses (A.2). Taking the expectation of (D.5) and then using Lemma D.1 gives, $\forall s \geq 2$,

$$\begin{aligned} \mathbb{E} \left[\|\mathbf{v}^{0,s} - \mathbf{v}^{-1,s}\|^2 \right] &\leq 18L^2 \mathbb{E} \left[\|\mathbf{x}^{q,s-1} - \mathbf{Jx}^{q,s-1}\|^2 \right] + 6n\alpha^2 L^2 \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s-1}\|^2 \right] \\ &\quad + 12\alpha^2 L^2 \mathbb{E} \left[\|\mathbf{y}^{q+1,s-1} - \mathbf{Jy}^{q+1,s-1}\|^2 \right] \\ (D.6) \quad &\quad + \frac{12L^2}{B} \sum_{t=0}^q \mathbb{E} \left[\|\mathbf{x}^{t,s-1} - \mathbf{Jx}^{t,s-1}\|^2 \right]. \end{aligned}$$

We recall from (D.2) that $\forall s \geq 2$,

$$(D.7) \quad \|\mathbf{y}^{1,s} - \mathbf{Jy}^{1,s}\|^2 \leq \frac{1+\lambda^2}{2} \|\mathbf{y}^{q+1,s-1} - \mathbf{Jy}^{q+1,s-1}\|^2 + \frac{2}{1-\lambda^2} \|\mathbf{v}^{0,s} - \mathbf{v}^{-1,s}\|^2.$$

We finally apply (D.6) to (D.7) to obtain, $\forall s \geq 2$,

$$\begin{aligned} \mathbb{E}[\|\mathbf{y}^{1,s} - \mathbf{Jy}^{1,s}\|^2] &\leq \left(\frac{1+\lambda^2}{2} + \frac{24\alpha^2 L^2}{1-\lambda^2}\right) \mathbb{E}[\|\mathbf{y}^{q+1,s-1} - \mathbf{Jy}^{q+1,s-1}\|^2] \\ &\quad + \frac{36L^2}{1-\lambda^2} \mathbb{E}[\|\mathbf{x}^{q,s-1} - \mathbf{Jx}^{q,s-1}\|^2] + \frac{12n\alpha^2 L^2}{1-\lambda^2} \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s-1}\|^2] \\ &\quad + \frac{24L^2}{B(1-\lambda^2)} \sum_{t=0}^q \mathbb{E}[\|\mathbf{x}^{t,s-1} - \mathbf{Jx}^{t,s-1}\|^2]. \end{aligned}$$

We note that $\frac{1+\lambda^2}{2} + \frac{24\alpha^2 L^2}{1-\lambda^2} \leq \frac{3+\lambda^2}{4}$ if $0 < \alpha \leq \frac{1-\lambda^2}{4\sqrt{6}L}$ and then the third statement in the lemma follows. \square

D.2. GT-SARAH as an LTI system. With the help of (A.1) and Lemma D.2, we now abstract GT-SARAH with an LTI system to quantify jointly the state consensus and the gradient tracking error.

LEMMA D.3. *If the step-size α follows $0 < \alpha \leq \frac{1-\lambda^2}{4\sqrt{6}L}$, then we have*

$$(D.8) \quad \mathbf{u}^{t,s} \leq \mathbf{Gu}^{t-1,s} + \mathbf{b}^{t-1,s} \quad \forall s \in [1, S] \text{ and } t \in [1, q],$$

$$(D.9) \quad \mathbf{u}^{0,s} \leq \mathbf{Gu}^{q,s-1} + \mathbf{b}^{q,s-1} + \sum_{t=0}^q (\mathbf{b}^{t,s-1} + \mathbf{Hu}^{t,s-1}) \quad \forall s \in [2, S],$$

where, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$\begin{aligned} \mathbf{u}^{t,s} &:= \begin{bmatrix} \frac{1}{n} \mathbb{E}[\|\mathbf{x}^{t,s} - \mathbf{Jx}^{t,s}\|^2] \\ \frac{1}{nL^2} \mathbb{E}[\|\mathbf{y}^{t+1,s} - \mathbf{Jy}^{t+1,s}\|^2] \end{bmatrix}, \quad \mathbf{b} := \begin{bmatrix} 0 \\ 12\alpha^2 \\ 1-\lambda^2 \end{bmatrix}, \quad \mathbf{b}^{t,s} := \mathbf{b} \mathbb{E}[\|\bar{\mathbf{v}}^{t,s}\|^2], \\ \mathbf{G} &:= \begin{bmatrix} \frac{1+\lambda^2}{2} & \frac{2\alpha^2 L^2}{1-\lambda^2} \\ \frac{18}{1-\lambda^2} & \frac{3+\lambda^2}{4} \end{bmatrix}, \quad \mathbf{H} := \begin{bmatrix} 0 & 0 \\ 42 & 0 \end{bmatrix}. \end{aligned}$$

Proof. The proof of Lemma D.3 follows by writing the inequalities in (A.1) and Lemma D.2 jointly in a matrix form. \square

We next derive the range of the step-size α such that $\rho(\mathbf{G}) < 1$, i.e., the LTI system does not diverge, with the help of the following lemma.

LEMMA D.4 (see [14]). *Let $\mathbf{X} \in \mathbb{R}^{d \times d}$ be (entrywise) nonnegative and $\mathbf{x} \in \mathbb{R}^d$ be (entrywise) positive. If $\mathbf{Xx} < \mathbf{x}$ (entrywise), then $\rho(\mathbf{X}) < 1$.*

LEMMA D.5. *If the step-size α follows $0 < \alpha < \frac{(1-\lambda^2)^2}{8\sqrt{5}L}$, then $\rho(\mathbf{G}) < 1$, and therefore $\sum_{k=0}^{\infty} \mathbf{G}^k$ is convergent such that $\sum_{k=0}^{\infty} \mathbf{G}^k = (\mathbf{I}_2 - \mathbf{G})^{-1}$.*

Proof. In the light of Lemma D.4, we solve the range of α and a positive vector $\varepsilon = [\varepsilon_1, \varepsilon_2]^\top$ such that $\mathbf{G}\varepsilon < \varepsilon$, which is equivalent to the following two inequalities:

$$(D.10) \quad \begin{cases} \frac{1+\lambda^2}{2} \varepsilon_1 + \frac{2\alpha^2 L^2}{1-\lambda^2} \varepsilon_2 < \varepsilon_1, \\ \frac{18}{1-\lambda^2} \varepsilon_1 + \frac{3+\lambda^2}{4} \varepsilon_2 < \varepsilon_2, \end{cases} \iff \begin{cases} \alpha^2 < \frac{(1-\lambda^2)^2}{4L^2} \frac{\varepsilon_1}{\varepsilon_2}, \\ \frac{\varepsilon_1}{\varepsilon_2} < \frac{(1-\lambda^2)^2}{72}. \end{cases}$$

According to the second inequality of (D.10), we set $\varepsilon_1/\varepsilon_2 = (1-\lambda^2)^2/80$, and the proof of Lemma D.5 follows by using it in the first inequality of (D.10) to solve for the range of α . \square

Based on Lemma D.5, the LTI system is stable under an appropriate step-size α , and therefore we can solve the LTI system to obtain the following lemma, the proof of which is deferred to Appendix G for ease of exposition.

LEMMA D.6. *If $0 < \alpha < \frac{(1-\lambda^2)^2}{8\sqrt{5}L}$, then the following inequality holds:*

$$(\mathbf{I}_2 - (\mathbf{I}_2 - \mathbf{G})^{-1}\mathbf{H}) \sum_{s=1}^S \sum_{t=0}^q \mathbf{u}^{t,s} \leq (\mathbf{I}_2 - \mathbf{G})^{-1}\mathbf{u}^{0,1} + 2(\mathbf{I}_2 - \mathbf{G})^{-1} \sum_{s=1}^S \sum_{t=0}^q \mathbf{b}^{t,s}.$$

Proof. See Appendix G. \square

In the following lemma, we compute $(\mathbf{I}_2 - \mathbf{G})^{-1}$ and $(\mathbf{I}_2 - \mathbf{G})^{-1}\mathbf{b}$.

LEMMA D.7. *If $0 < \alpha \leq \frac{(1-\lambda^2)^2}{24L}$, then the following entrywise inequality holds:*

$$(\mathbf{I}_2 - \mathbf{G})^{-1} \leq \begin{bmatrix} \frac{4}{1-\lambda^2} & \frac{32\alpha^2 L^2}{(1-\lambda^2)^3} \\ \frac{288}{(1-\lambda^2)^3} & \frac{8}{1-\lambda^2} \end{bmatrix}, \quad (\mathbf{I}_2 - \mathbf{G})^{-1}\mathbf{b} \leq \begin{bmatrix} \frac{384\alpha^4 L^2}{(1-\lambda^2)^4} \\ \frac{96\alpha^2}{(1-\lambda^2)^2} \end{bmatrix}.$$

Proof. We first derive a lower bound for $\det(\mathbf{I}_2 - \mathbf{G})$. Note that if $0 < \alpha \leq \frac{(1-\lambda^2)^2}{24L}$, then $\det(\mathbf{I}_2 - \mathbf{G}) = \frac{(1-\lambda^2)^2}{8} - \frac{36\alpha^2 L^2}{(1-\lambda^2)^2} \geq \frac{(1-\lambda^2)^2}{16}$ and therefore

$$(\mathbf{I}_2 - \mathbf{G})^{-1} \leq \frac{16}{(1-\lambda^2)^2} \begin{bmatrix} \frac{1-\lambda^2}{4} & \frac{2\alpha^2 L^2}{1-\lambda^2} \\ \frac{18}{1-\lambda^2} & \frac{1-\lambda^2}{2} \end{bmatrix} = \begin{bmatrix} \frac{4}{1-\lambda^2} & \frac{32\alpha^2 L^2}{(1-\lambda^2)^3} \\ \frac{288}{(1-\lambda^2)^3} & \frac{8}{1-\lambda^2} \end{bmatrix},$$

and the proof follows by the definition of \mathbf{b} in Lemma D.3. \square

D.3. Proof of Lemma 4.5. Using Lemma D.7, we have the following: if $0 < \alpha \leq \frac{(1-\lambda^2)^2}{8\sqrt{42}L}$,

$$(D.11) \quad \mathbf{I}_2 - (\mathbf{I}_2 - \mathbf{G})^{-1}\mathbf{H} \geq \begin{bmatrix} 1 - \frac{1344\alpha^2 L^2}{(1-\lambda^2)^4} & 0 \\ -\frac{336}{(1-\lambda^2)^2} & 1 \end{bmatrix} \geq \begin{bmatrix} \frac{1}{2} & 0 \\ -\frac{336}{(1-\lambda^2)^2} & 1 \end{bmatrix}.$$

Finally, we apply (D.11) and Lemma D.7 to Lemma D.6 to obtain

$$(D.12) \quad \begin{aligned} \frac{1}{2n} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2] &\leq \frac{32\alpha^2}{n(1-\lambda^2)^3} \mathbb{E}[\|\mathbf{y}^{1,1} - \mathbf{J}\mathbf{y}^{1,1}\|^2] \\ &\quad + \frac{768\alpha^4 L^2}{(1-\lambda^2)^4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E}[\|\bar{\mathbf{v}}^{t,s}\|^2]. \end{aligned}$$

The proof follows by applying the first statement in Lemma D.2 to (D.12).

Appendix E. Proof of Lemma 4.6. We have, $\forall s \geq 1$ and $\forall t \in [0, q]$,

$$(E.1) \quad \begin{aligned} \frac{1}{2n} \sum_{i=1}^n \mathbb{E}[\|\nabla F(\mathbf{x}_i^{t,s})\|^2] &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|\nabla F(\mathbf{x}_i^{t,s}) - \nabla F(\bar{\mathbf{x}}^{t,s})\|^2] + \mathbb{E}[\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2] \\ &\leq \frac{L^2}{n} \mathbb{E}[\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2] + \mathbb{E}[\|\nabla F(\bar{\mathbf{x}}^{t,s})\|^2], \end{aligned}$$

where the second line is due to the L -smoothness of F . Since F is bounded below by F^* , we may apply (E.1) to Lemma 4.3 to obtain the following: if $0 < \alpha \leq \frac{1}{2L}$,

$$\begin{aligned}
F^* &\leq F(\bar{\mathbf{x}}^{0,1}) - \frac{\alpha}{4n} \sum_{i=1}^n \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\nabla F(\mathbf{x}_i^{t,s})\|^2 \right] - \frac{\alpha}{4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s}\|^2 \right] \\
(E.2) \quad &+ \alpha \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s} - \nabla \mathbf{f}(\mathbf{x}^{t,s})\|^2 \right] + \frac{3\alpha L^2}{2} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2}{n} \right].
\end{aligned}$$

We then apply Lemma 4.4 to (E.2) to obtain the following: if $0 < \alpha \leq \frac{1}{2L}$,

$$\begin{aligned}
F^* &\leq F(\bar{\mathbf{x}}^{0,1}) - \frac{\alpha}{4n} \sum_{i=1}^n \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\nabla F(\mathbf{x}_i^{t,s})\|^2 \right] - \frac{\alpha}{8} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s}\|^2 \right] \\
&+ \alpha L^2 \left(\frac{3}{2} + \frac{6q}{nB} \right) \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2}{n} \right] \\
(E.3) \quad &- \frac{\alpha}{8} \left(1 - \frac{24q\alpha^2 L^2}{nB} \right) \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s}\|^2 \right].
\end{aligned}$$

If $0 < \alpha \leq \frac{\sqrt{nB}}{2\sqrt{6q}L}$, then $1 - \frac{24q\alpha^2 L^2}{nB} \geq 0$, and thus the last term in (E.3) may be dropped. We finally apply Lemma 4.5 to (E.3) to obtain the following: if $0 < \alpha \leq \min \left\{ \frac{(1-\lambda^2)^2}{4\sqrt{42}}, \sqrt{\frac{nB}{6q}} \right\} \frac{1}{2L}$,

$$\begin{aligned}
F^* &\leq F(\bar{\mathbf{x}}^{0,1}) - \frac{\alpha}{4n} \sum_{i=1}^n \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\nabla F(\mathbf{x}_i^{t,s})\|^2 \right] + \left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{64\alpha^3 L^2}{(1-\lambda^2)^3} \frac{\|\nabla \mathbf{f}(\mathbf{x}^{0,1})\|^2}{n} \\
&- \frac{\alpha L^2}{4} \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\frac{\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2}{n} \right] \\
&- \frac{\alpha}{8} \left(1 - \left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{12288\alpha^4 L^4}{(1-\lambda^2)^4} \right) \sum_{s=1}^S \sum_{t=0}^q \mathbb{E} \left[\|\bar{\mathbf{v}}^{t,s}\|^2 \right].
\end{aligned}$$

We observe that if $0 < \alpha \leq \left(\frac{4nB}{7nB+24q} \right)^{1/4} \frac{1-\lambda^2}{12L}$, then $1 - \left(\frac{7}{4} + \frac{6q}{nB} \right) \frac{12288\alpha^4 L^4}{(1-\lambda^2)^4} \geq 0$, and thus the last term in the above inequality may be dropped; the proof follows.

Appendix F. Proof of Lemma D.1. In the following, we use the notation in (C.1). Using the update of each local recursive gradient estimator $\mathbf{v}_i^{t,s}$, we have that, $\forall i \in \mathcal{V}, \forall s \geq 1$ and $t \in [1, q]$,

$$\begin{aligned}
&\mathbb{E} \left[\|\mathbf{v}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s})\|^2 \mid \mathcal{F}^{t,s} \right] \\
&= \mathbb{E} \left[\left\| \widehat{\nabla}_i^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) + \mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\|^2 \mid \mathcal{F}^{t,s} \right] \\
&= \mathbb{E} \left[\left\| \frac{1}{B} \sum_{l=1}^B \left(\widehat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right) \right\|^2 \mid \mathcal{F}^{t,s} \right] + \|\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s})\|^2, \\
&= \frac{1}{B^2} \sum_{l=1}^B \mathbb{E} \left[\left\| \widehat{\nabla}_{i,l}^{t,s} - \nabla f_i(\mathbf{x}_i^{t,s}) + \nabla f_i(\mathbf{x}_i^{t-1,s}) \right\|^2 \mid \mathcal{F}^{t,s} \right] + \|\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s})\|^2, \\
&\leq \frac{1}{B^2} \sum_{l=1}^B \mathbb{E} \left[\left\| \nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t,s}) - \nabla f_{i,\tau_{i,l}^{t,s}}(\mathbf{x}_i^{t-1,s}) \right\|^2 \mid \mathcal{F}^{t,s} \right] + \|\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s})\|^2, \\
&\leq \frac{L^2}{B} \|\mathbf{x}_i^{t,s} - \mathbf{x}_i^{t-1,s}\|^2 + \|\mathbf{v}_i^{t-1,s} - \nabla f_i(\mathbf{x}_i^{t-1,s})\|^2.
\end{aligned}$$

The above derivations follow a similar line of arguments as in the proof of Lemma C.2, and hence we omit the details here. Summing up the last inequality above over i from 1 to n and taking the expectation, we have, $\forall s \geq 1$ and $t \in [1, q]$,

$$(F.1) \quad \mathbb{E} \left[\|\mathbf{v}^{t,s} - \nabla \mathbf{f}(\mathbf{x}^{t,s})\|^2 \right] \leq \mathbb{E} \left[\frac{L^2}{B} \|\mathbf{x}^{t,s} - \mathbf{x}^{t-1,s}\|^2 + \|\mathbf{v}^{t-1,s} - \nabla \mathbf{f}(\mathbf{x}^{t-1,s})\|^2 \right].$$

Recall from (C.9) that $\forall s \geq 1$ and $t \in [1, q]$,

$$(F.2) \quad \|\mathbf{x}^{t,s} - \mathbf{x}^{t-1,s}\|^2 \leq 3\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2 + 3n\alpha^2 \|\bar{\mathbf{v}}^{t-1,s}\|^2 + 3\|\mathbf{x}^{t-1,s} - \mathbf{J}\mathbf{x}^{t-1,s}\|^2.$$

Applying (F.2) to (F.1) obtains, $\forall s \geq 1$ and $t \in [1, q]$,

$$(F.3) \quad \begin{aligned} \mathbb{E}[\|\mathbf{v}^{t,s} - \nabla \mathbf{f}(\mathbf{x}^{t,s})\|^2] &\leq \mathbb{E}[\|\mathbf{v}^{t-1,s} - \nabla \mathbf{f}(\mathbf{x}^{t-1,s})\|^2] + \frac{3n\alpha^2 L^2}{B} \mathbb{E}[\|\bar{\mathbf{v}}^{t-1,s}\|^2] \\ &\quad + \frac{3L^2}{B} \mathbb{E}[\|\mathbf{x}^{t,s} - \mathbf{J}\mathbf{x}^{t,s}\|^2] + \frac{3L^2}{B} \mathbb{E}[\|\mathbf{x}^{t-1,s} - \mathbf{J}\mathbf{x}^{t-1,s}\|^2]. \end{aligned}$$

Recall that $\mathbf{v}^{0,s} = \nabla \mathbf{f}(\mathbf{x}^{0,s}) \forall s \geq 1$, and we take the telescoping sum of (F.3) over t to obtain, $\forall s \geq 1$ and $t \in [1, q]$,

$$\begin{aligned} \mathbb{E}[\|\mathbf{v}^{t,s} - \nabla \mathbf{f}(\mathbf{x}^{t,s})\|^2] &\leq \frac{3n\alpha^2 L^2}{B} \sum_{u=1}^t \mathbb{E}[\|\bar{\mathbf{v}}^{u-1,s}\|^2] + \frac{3L^2}{B} \sum_{u=1}^t \mathbb{E}[\|\mathbf{x}^{u,s} - \mathbf{J}\mathbf{x}^{u,s}\|^2] \\ &\quad + \frac{3L^2}{B} \sum_{u=1}^t \mathbb{E}[\|\mathbf{x}^{u-1,s} - \mathbf{J}\mathbf{x}^{u-1,s}\|^2]. \end{aligned}$$

The proof follows by merging the last two terms on the right-hand side of the inequality above.

Appendix G. Proof of Lemma D.6.

G.1. Step 1: A loopless dynamical system. For the ease of calculations, we first write the LTI system in Lemma D.3 in an equivalent *loopless* form. To do this, we unroll the original *double loop* sequences $\{\mathbf{u}^{t,s}\}$ and $\{\mathbf{b}^{t,s}\}$, where $t \in [0, q]$ and $s \in [1, S]$, respectively, as *loopless* sequences $\{\mathbf{u}^k\}$ and $\{\mathbf{b}^k\}$, where $k \in [0, S(q+1)-1]$, as follows:

$$(G.1) \quad \mathbf{u}^k := \mathbf{u}^{t,s}, \quad \mathbf{b}^k := \mathbf{b}^{t,s}, \quad \text{where } k = t + (s-1)(q+1)$$

for $t \in [0, q]$ and $s \in [1, S]$. Reversely, given \mathbf{u}^k and \mathbf{b}^k for $k \in [0, S(q+1)-1]$, we can find their positions in the original double loop sequence, $\mathbf{u}^{t,s}$ and $\mathbf{b}^{t,s}$, by

$$(G.2) \quad t = \text{mod}(k, q+1) \text{ and } s = \lfloor k/(q+1) \rfloor + 1 \quad \text{for } k \in [0, S(q+1)-1].$$

This one-on-one correspondence is visualized in Table 3.

With (G.1) and (G.2) at hand, it can be verified that the following single loop system is equivalent to the double loop system in (D.8) and (D.9). For $k \in [1, S(q+1)-1]$,

$$(G.3) \quad \mathbf{u}^k \leq \mathbf{G}\mathbf{u}^{k-1} + \mathbf{b}^{k-1} \quad \text{if } \text{mod}(k, q+1) \neq 0,$$

$$(G.4) \quad \mathbf{u}^{z(q+1)} \leq \mathbf{G}\mathbf{u}^{z(q+1)-1} + \mathbf{b}^{z(q+1)-1} + \sum_{r=(z-1)(q+1)}^{z(q+1)-1} \mathbf{h}^r \quad \forall z \in [1, S-1],$$

where $\mathbf{h}^k := \mathbf{b}^k + \mathbf{H}\mathbf{u}^k \forall k \in [0, S(q+1)-1]$. The system in (G.3) and (G.4) can be further written equivalently as the following: $\forall k \in [1, S(q+1)-1]$,

$$(G.5) \quad \mathbf{u}^k \leq \mathbf{G}\mathbf{u}^{k-1} + \mathbf{d}^k,$$

where $\mathbf{d}^k := \mathbf{b}^{k-1} + \mathbb{1}\{\text{mod}(k, q+1) = 0\} \sum_{r=k-(q+1)}^{k-1} \mathbf{h}^r$, $\mathbb{1}\{\cdot\}$ is the indicator function of an event, and $\sum_{r=k-(q+1)}^{k-1} \mathbf{h}^r := 0$ for $k \in [1, q]$.

TABLE 3

The one-on-one mapping between the single loop sequences $\{\mathbf{u}^k\}, \{\mathbf{b}^k\}$ for $k \in [0, S(q+1)-1]$ and the double loop sequences $\{\mathbf{u}^{t,s}\}, \{\mathbf{b}^{t,s}\}$ for $s \in [1, S]$ and $t \in [0, q]$.

k	(t, s)
$0, \dots, q$	$(0, 1), \dots, (q, 1)$
$q+1, \dots, 2q+1$	$(0, 2), \dots, (q, 2)$
\dots	\dots
$(S-1)(q+1), \dots, S(q+1)-1$	$(0, S), \dots, (q, S)$

G.2. Step 2: Analyzing the recursion. We recursively apply (G.5) over k to obtain, $\forall k \in [1, S(q+1)-1]$,

$$(G.6) \quad \mathbf{u}^k \leq \mathbf{G}^k \mathbf{u}^0 + \sum_{r=1}^k \mathbf{G}^{k-r} \mathbf{d}^r.$$

Summing up (G.6) over k from 0 to $S(q+1)-1$ gives the following: if $0 < \alpha \leq \frac{(1-\lambda^2)^2}{8\sqrt{5}L}$,

$$\begin{aligned} \sum_{k=0}^{S(q+1)-1} \mathbf{u}^k &\leq \sum_{k=0}^{S(q+1)-1} \mathbf{G}^k \mathbf{u}^0 + \sum_{k=1}^{S(q+1)-1} \sum_{r=1}^k \mathbf{G}^{k-r} \mathbf{d}^r \\ &\leq (\sum_{k=0}^{\infty} \mathbf{G}^k) \mathbf{u}^0 + \sum_{k=1}^{S(q+1)-1} (\sum_{r=1}^{\infty} \mathbf{G}^r) \mathbf{d}^k \\ (G.7) \quad &= (\mathbf{I}_2 - \mathbf{G})^{-1} \mathbf{u}^0 + (\mathbf{I}_2 - \mathbf{G})^{-1} \sum_{k=1}^{S(q+1)-1} \mathbf{d}^k. \end{aligned}$$

To proceed, we recall the definition of \mathbf{d}^k and \mathbf{h}^k in Appendix G.1 and observe that

$$\begin{aligned} \sum_{k=1}^{S(q+1)-1} \mathbf{d}^k &= \sum_{k=0}^{S(q+1)-2} \mathbf{b}^k + \sum_{k=1}^{S(q+1)-1} \left(\mathbb{1}\{\text{mod}(k, q+1) = 0\} \sum_{r=k-(q+1)}^{k-1} \mathbf{h}^r \right) \\ &= \sum_{k=0}^{S(q+1)-2} \mathbf{b}^k + \sum_{z=1}^{S-1} \left(\sum_{r=(z-1)(q+1)}^{z(q+1)-1} \mathbf{h}^r \right) \\ &= \sum_{k=0}^{S(q+1)-2} \mathbf{b}^k + \sum_{k=0}^{(S-1)(q+1)-1} \mathbf{h}^k \\ (G.8) \quad &\leq 2 \sum_{k=0}^{S(q+1)-1} \mathbf{b}^k + \sum_{k=0}^{(S-1)(q+1)-1} \mathbf{H} \mathbf{u}^k, \end{aligned}$$

where the first line and the last line are due to the definition of \mathbf{d}^k and \mathbf{h}^k , respectively. Finally, we use (G.8) in (G.7) to obtain the following: if $0 < \alpha \leq \frac{(1-\lambda^2)^2}{8\sqrt{5}L}$, then

$$\begin{aligned} \sum_{k=0}^{S(q+1)-1} \mathbf{u}^k &\leq (\mathbf{I}_2 - \mathbf{G})^{-1} \mathbf{u}^0 + 2(\mathbf{I}_2 - \mathbf{G})^{-1} \sum_{k=0}^{S(q+1)-1} \mathbf{b}^k \\ &\quad + (\mathbf{I}_2 - \mathbf{G})^{-1} \mathbf{H} \sum_{k=0}^{S(q+1)-1} \mathbf{u}^k, \end{aligned}$$

which is the same as

$$(\mathbf{I}_2 - (\mathbf{I}_2 - \mathbf{G})^{-1} \mathbf{H}) \sum_{k=0}^{S(q+1)-1} \mathbf{u}^k \leq (\mathbf{I}_2 - \mathbf{G})^{-1} \mathbf{u}^0 + 2(\mathbf{I}_2 - \mathbf{G})^{-1} \sum_{k=0}^{S(q+1)-1} \mathbf{b}^k.$$

We conclude the proof of Lemma D.6 by rewriting the above inequality in the original double loop form.

REFERENCES

- [1] Z. ALLEN-ZHU AND E. HAZAN, *Variance reduction for faster nonconvex optimization*, in Proceedings of the International Conference on Machine Learning, 2016, pp. 699–707.
- [2] A. ANTONIADIS, I. GIBBELS, AND M. NIKOLOVA, *Penalized likelihood regression for generalized linear models with non-quadratic penalties*, Ann. Inst. Statist. Math., 63 (2011), pp. 585–615.
- [3] M. ASSRAN, N. LOIZOU, N. BALLAS, AND M. RABBAT, *Stochastic gradient push for distributed deep learning*, in Proceedings of the 36th International Conference on Machine Learning, 2019, pp. 97: 344–353.
- [4] L. BOTTOU, F. E. CURTIS, AND J. NOCEDAL, *Optimization methods for large-scale machine learning*, SIAM Rev., 60 (2018), pp. 223–311.
- [5] S. CEN, H. ZHANG, Y. CHI, W. CHEN, AND T.-Y. LIU, *Convergence of distributed stochastic variance reduced methods without sampling extra data*, IEEE Trans. Signal Process., 68 (2020), pp. 3976–3989.
- [6] T.-H. CHANG, M. HONG, H.-T. WAI, X. ZHANG, AND S. LU, *Distributed learning in the nonconvex world: From batch data to streaming and beyond*, IEEE Signal Process. Mag., 37 (2020), pp. 26–38.
- [7] J. CHEN AND A. H. SAYED, *Diffusion adaptation strategies for distributed optimization and learning over networks*, IEEE Trans. Signal Process., 60 (2012), pp. 4289–4305.

- [8] J. CHEN AND A. H. SAYED, *On the learning behavior of adaptive networks—part I: Transient analysis*, IEEE Trans. Inform. Theory, 61 (2015), pp. 3487–3517.
- [9] A. DEFazio, F. BACH, AND S. LACOSTE-JULIEN, *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*, in Proceedings of the 28th Conference on Neural Information Processing Systems, 2014, pp. 1646–1654.
- [10] P. DI LORENZO AND G. SCUTARI, *NEXT: In-network nonconvex optimization*, IEEE Trans. Signal Inform. Process. Netw., 2 (2016), pp. 120–136.
- [11] C. FANG, C. J. LI, Z. LIN, AND T. ZHANG, *SPIDER: Near-optimal non-convex optimization via stochastic path-integrated differential estimator*, in Proceedings of the 32nd Conference on Neural Information Processing Systems, 2018, pp. 689–699.
- [12] S. GHADIMI AND G. LAN, *Stochastic first-and zeroth-order methods for nonconvex stochastic programming*, SIAM J. Optim., 23 (2013), pp. 2341–2368.
- [13] B. GHARESFARID AND J. CORTÉS, *Distributed strategies for generating weight-balanced and doubly stochastic digraphs*, Eur. J. Control, 18 (2012), pp. 539–557.
- [14] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 2012.
- [15] D. JAKOVETIĆ, *A unification and generalization of exact distributed first-order methods*, IEEE Trans. Signal Inform. Process. Netw., 5 (2018), pp. 31–46.
- [16] G. LAN, S. LEE, AND Y. ZHOU, *Communication-efficient algorithms for decentralized and stochastic optimization*, Math. Program., 180 (2020), pp. 237–284.
- [17] J. LEI, H. CHEN, AND H. FANG, *Asymptotic properties of primal-dual algorithm for distributed stochastic optimization over random networks with imperfect communications*, SIAM J. Control Optim., 56 (2018), pp. 2159–2188.
- [18] B. LI, S. CEN, Y. CHEN, AND Y. CHI, *Communication-efficient distributed optimization in networks with gradient tracking and variance reduction*, J. Mach. Learn. Res., 21 (2020), pp. 1–51.
- [19] Z. LI, W. SHI, AND M. YAN, *A decentralized proximal-gradient method with network independent step-sizes and separated convergence rates*, IEEE Trans. Signal Process., 67 (2019), pp. 4494–4506.
- [20] X. LIAN, C. ZHANG, H. ZHANG, C.-J. HSIEH, W. ZHANG, AND J. LIU, *Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent*, in Proceedings of the 31st Conference on Neural Information Processing Systems, 2017, pp. 5330–5340.
- [21] A. MOKHTARI AND A. RIBEIRO, *DSA: Decentralized double stochastic averaging gradient algorithm*, J. Mach. Learn. Res., 17 (2016), pp. 2165–2199.
- [22] A. NEDIĆ, A. OLSHEVSKY, AND M. G. RABBAT, *Network topology and communication-computation tradeoffs in decentralized optimization*, Proc. IEEE, 106 (2018), pp. 953–976.
- [23] A. NEDICH, A. OLSHEVSKY, AND W. SHI, *Achieving geometric convergence for distributed optimization over time-varying graphs*, SIAM J. Optim., 27 (2017), pp. 2597–2633.
- [24] A. NEMIROVSKI, A. JUDITSKY, G. LAN, AND A. SHAPIRO, *Robust stochastic approximation approach to stochastic programming*, SIAM J. Optim., 19 (2009), pp. 1574–1609.
- [25] Y. NESTEROV, *Lectures on Convex Optimization*, Springer, Cham, 2018.
- [26] L. M. NGUYEN, J. LIU, K. SCHEINBERG, AND M. TAKAC, *SARAH: A novel method for machine learning problems using stochastic recursive gradient*, in Proceedings of the 34th International Conference on Machine Learning., 2017, pp. 2613–2621.
- [27] L. M. NGUYEN, K. SCHEINBERG, AND M. TAKAC, *Inexact SARAH algorithm for stochastic optimization*, Optim. Methods Softw., 36 (2020), pp. 237–258.
- [28] N. H. PHAM, L. M. NGUYEN, D. T. PHAN, AND Q. TRAN-DINH, *ProxSARAH: An efficient algorithmic framework for stochastic composite nonconvex optimization*, J. Mach. Learn. Res., 21 (2020), pp. 1–48.
- [29] S. PU AND A. GARCIA, *Swarming for faster convergence in stochastic optimization*, SIAM J. Control Optim., 56 (2018), pp. 2997–3020.
- [30] S. PU AND A. NEDICH, *Distributed stochastic gradient tracking methods*, Math. Program., 187 (2021), pp. 409–457.
- [31] S. PU, A. OLSHEVSKY, AND I. C. PASHALIDIS, *A Sharp Estimate on the Transient Time of Distributed Stochastic Gradient Descent*, arXiv preprint, 2019, <https://arxiv.org/abs/1906.02702>.
- [32] G. QU AND N. LI, *Harnessing smoothness to accelerate distributed optimization*, IEEE Trans. Control. Netw. Syst., 5 (2017), pp. 1245–1260.
- [33] S. S. RAM, A. NEDIĆ, AND V. V. VEERAVALLI, *Distributed stochastic subgradient projection algorithms for convex optimization*, J. Optim. Theory Appl., 147 (2010), pp. 516–545.

- [34] S. J. REDDI, A. HEFNY, S. SRA, B. PÓCZOS, AND A. SMOLA, *Stochastic variance reduction for nonconvex optimization*, in Proceedings of the International Conference on Machine Learning, 2016, pp. 314–323.
- [35] G. SCUTARI AND Y. SUN, *Distributed nonconvex constrained optimization over time-varying digraphs*, Math. Program., 176 (2019), pp. 497–544.
- [36] W. SHI, Q. LING, G. WU, AND W. YIN, *EXTRA: An exact first-order algorithm for decentralized consensus optimization*, SIAM J. Optim., 25 (2015), pp. 944–966.
- [37] H. SUN AND M. HONG, *Distributed non-convex first-order optimization and information processing: Lower complexity bounds and rate optimal algorithms*, IEEE Trans. Signal process., 67 (2019), pp. 5912–5928.
- [38] H. SUN, S. LU, AND M. HONG, *Improving the sample and communication complexity for decentralized non-convex optimization: Joint gradient estimation and tracking*, in Proceedings of the International Conference on Machine Learning, 2020, pp. 9217–9228.
- [39] B. SWENSON, R. MURRAY, S. KAR, AND H. V. POOR, *Distributed Stochastic Gradient Descent and Convergence to Local Minima*, arXiv preprint, 2020, <https://arxiv.org/abs/2003.02818>.
- [40] H. TANG, X. LIAN, M. YAN, C. ZHANG, AND J. LIU, *D²: Decentralized training over decentralized data*, in Proceedings of the International Conference on Machine Learning, 2018, pp. 4848–4856.
- [41] S. VLASKI AND A. H. SAYED, *Distributed Learning in Non-Convex Environments—Part II: Polynomial Escape from Saddle-Points*, arXiv preprint, 2019, <https://arxiv.org/abs/1907.01849>.
- [42] H. WAI, J. LAFOND, A. SCAGLIONE, AND E. MOULINES, *Decentralized Frank–Wolfe algorithm for convex and nonconvex problems*, IEEE Trans. Autom. Control, 62 (2017), pp. 5522–5537.
- [43] J. WANG, V. TANTIA, N. BALLAS, AND M. RABBAT, *SlowMo: Improving Communication-Efficient Distributed SGD with Slow Momentum*, arXiv preprint, 2019 <https://arxiv.org/abs/1910.00643>.
- [44] Y. WANG, W. ZHAO, Y. HONG, AND M. ZAMANI, *Distributed subgradient-free stochastic optimization algorithm for nonsmooth convex functions over time-varying networks*, SIAM J. Control Optim., 57 (2019), pp. 2821–2842.
- [45] Z. WANG, K. JI, Y. ZHOU, Y. LIANG, AND V. TAROKH, *Spiderboost and momentum: Faster variance reduction algorithms*, in Proceedings of the 33rd Conference on Neural Information Processing Systems, 2019, pp. 2403–2413.
- [46] D. WILLIAMS, *Probability with Martingales*, Cambridge University Press, Cambridge, UK, 1991.
- [47] L. XIAO AND T. ZHANG, *A proximal stochastic gradient method with progressive variance reduction*, SIAM J. Optim., 24 (2014), pp. 2057–2075.
- [48] R. XIN, S. KAR, AND U. A. KHAN, *Decentralized stochastic optimization and machine learning: A unified variance-reduction framework for robust performance and fast convergence*, IEEE Signal Process. Mag., 37 (2020), pp. 102–113.
- [49] R. XIN, U. A. KHAN, AND S. KAR, *Variance-reduced decentralized stochastic optimization with accelerated convergence*, IEEE Trans. Signal Process., 68 (2020), pp. 6255–6271.
- [50] R. XIN, U. A. KHAN, AND S. KAR, *An improved convergence analysis for decentralized online stochastic non-convex optimization*, IEEE Trans. Signal Process., 69 (2021), pp. 1842–1858.
- [51] J. XU, S. ZHU, Y. C. SOH, AND L. XIE, *Augmented distributed gradient methods for multi-agent optimization under uncoordinated constant stepsizes*, in Proceedings of the 54th IEEE Conference on Decision and Control (CDC), 2015, pp. 2055–2060.
- [52] T. YANG, X. YI, J. WU, Y. YUAN, D. WU, Z. MENG, Y. HONG, H. WANG, Z. LIN, AND K. H. JOHANSSON, *A survey of distributed optimization*, Ann. Rev. Control, 47 (2019), pp. 278–305.
- [53] D. YUAN, D. W. HO, AND Y. HONG, *On convergence rate of distributed stochastic gradient algorithm for convex optimization with inequality constraints*, SIAM J. Control Optim., 54 (2016), pp. 2872–2892.
- [54] K. YUAN, S. A. ALGHUNAIM, B. YING, AND A. H. SAYED, *On the influence of bias-correction on distributed stochastic optimization*, IEEE Trans. Signal Process., 68 (2020), pp. 4352–4367.
- [55] K. YUAN, Q. LING, AND W. YIN, *On the convergence of decentralized gradient descent*, SIAM J. Optim., 26 (2016), pp. 1835–1854.
- [56] K. YUAN, B. YING, J. LIU, AND A. H. SAYED, *Variance-reduced stochastic learning by networked agents under random reshuffling*, IEEE Trans. Signal Process., 67 (2018), pp. 351–366.
- [57] D. ZHOU AND Q. GU, *Lower bounds for smooth nonconvex finite-sum optimization*, in Proceedings of the International Conference on Machine Learning, PMLR, 2019, pp. 7574–7583.
- [58] D. ZHOU, P. XU, AND Q. GU, *Stochastic nested variance reduction for nonconvex optimization*, J. Mach. Learn. Res., 21 (2020), pp. 1–63.