arXiv:2002.07378v3 [math.OC] 14 Jan 2022

Distributed Adaptive Newton Methods with Global Superlinear Convergence ⋆

Jiaqi Zhang [a], Keyou You [a,*], Tamer Başar [b]

[a] *Department of Automation, and BNRist, Tsinghua University, Beijing 100084, China.*

[b] *Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA.*

Abstract

This paper considers the distributed optimization problem where each node of a peer-to-peer network minimizes a finite sum of objective functions by communicating with its neighboring nodes. In sharp contrast to the existing literature where the fastest distributed algorithms converge either with a global linear or a local superlinear rate, we propose a distributed adaptive Newton (DAN) algorithm with a *global quadratic* convergence rate. Our key idea lies in the design of a finite-time set-consensus method with Polyak's adaptive stepsize. Moreover, we introduce a low-rank matrix approximation (LA) technique to compress the innovation of Hessian matrix so that each node only needs to transmit message of dimension $\mathcal{O}(p)$ (where p is the dimension of decision vectors) per iteration, which is essentially the same as that of first-order methods. Nevertheless, the resulting DAN-LA converges to an optimal solution with a *global superlinear* rate. Numerical experiments on logistic regression problems are conducted to validate their advantages over existing methods.

Key words: distributed optimization; Newton method; low-rank approximation; superlinear convergence.

1 Introduction

Distributed optimization entails solving the following problem over a peer-to-peer network system

$$\begin{aligned} \underset{\mathbf{x}_1,\ldots,\mathbf{x}_n}{\text{minimize}} \quad & F(\mathbf{x}_1,\ldots,\mathbf{x}_n) \triangleq \sum_{i=1}^{n} f_i(\mathbf{x}_i) \\ \text{subject to} \quad & \mathbf{x}_1 = \cdots = \mathbf{x}_n \in \mathbb{R}^p \end{aligned} \tag{1}$$

where each node i privately holds a local objective function f_i and updates its decision vector $\mathbf{x}_i$ via communicating with its neighboring nodes. Our goal is to design efficient distributed algorithms to find an optimal solution of (1). Many efforts have been devoted along this line, see e.g., Nedić et al. (2017), Xin & Khan (2019), Scaman et al. (2017), Qu & Li (2019). It is known that the fastest rate for first-order methods is linear (Nesterov 2018), and second-order methods are unavoidable for the superlinear convergence.

Unfortunately, the existing distributed methods cannot achieve *global superlinear* convergence rates. In contrast, our paper proposes two Newton-based distributed algorithms with global quadratic and superlinear convergence rates, respectively. A direct comparison with the existing literature can be found in Table 1. For second-order methods in Mokhtari et al. (2017), Mansoori & Wei (2020, 2017), Tutunov et al. (2019) and Zargham et al. (2014), an *inexact* penalization reformulation as an unconstrained optimization problem is adopted to solve (1). Their superlinear rate is restricted to a limited region, which does not include an optimal solution, and then reduced to a linear rate. Even though the inexact issue has been resolved in Eisen et al. (2017, 2019) and Varagnolo et al. (2016), their convergence rates are still linear.

Under a master-slave network configuration, some distributed quasi-Newton methods have been proposed in Shamir et al. (2014), Wang et al. (2018), Zhang & Lin (2015), and Soori et al. (2020). Though the master node of this setting uses all the information from the slave

⋆ This work was supported by the National Natural Science Foundation of China under grant no. 62033006, a grant from the Guoqiang Institute, Tsinghua University, and in part by the ARL under cooperative agreement W911NF-17-2-0196. The material in this paper was partially presented at the 59th IEEE Conference on Decision and Control.

* Corresponding author

Email addresses: zjq16@mails.tsinghua.edu.cn (Jiaqi Zhang), youky@tsinghua.edu.cn (Keyou You), basar1@illinois.edu (Tamer Başar).

nodes, the algorithms in Shamir et al. (2014), Zhang & Lin (2015) and Wang et al. (2018) still converge linearly, and only Soori et al. (2020) achieves a *local* superlinear convergence rate if the starting point is sufficiently close to an optimal solution.

To achieve a global superlinear convergence rate, two bottlenecks have to be resolved. The first is the use of the linear consensus algorithm whose convergence rate is at most linear. Thus, the resulting second-order methods are still constrained to the linear convergence rate. Similar to this work, finite-time consensus methods have been adopted in Qu et al. (2019) to achieve local quadratic convergence by using the pure Newton direction.

The second bottleneck lies in the design of a backtracking line search method to distributedly tune stepsizes of the second-order methods. Even though the pure Newton direction with a unit stepsize is key to the superlinear convergence, it may lead to divergence if the algorithm is far from an optimal solution (Boyd & Vandenberghe 2004). In fact, the line search in the early stage is vital to the global convergence of second-order methods and usually requires nodes to evaluate the objective function of F multiple times in a single iteration, which is not possible in our distributed setting as each node i can only access its individual objective function f_i. To circumvent this problem, local line search methods have been designed in Zargham et al. (2012) and Jadbabaie et al. (2009) for network flow problems, and fixed stepsizes have been used in Mokhtari et al. (2017), Mansoori & Wei (2020, 2017), Tutunov et al. (2019), Mokhtari et al. (2016), Eisen et al. (2019), Varagnolo et al. (2016). However, they are too conservative to achieve superlinear convergence.

In this work, we resolve the above issues and design distributed algorithms with global superlinear convergence rates. Building on the distributed flooding (DF) algorithm (Liu et al. 2007, Dias & Bruno 2013, Li et al. 2017), we propose a distributed selective flooding (DSF) algorithm to achieve exact consensus in a finite number of local communication rounds and improves the DF with provably lower communication complexity. Precisely, nodes in the DSF achieve consensus with at most $n-1$ rounds of local communication rather than $n+d_{\mathcal{G}}-1$ in the DF, where $d_{\mathcal{G}}$ is the diameter of the network with n nodes. By running the DSF at each iteration, the consensus achievement is no longer a bottleneck for superlinear convergence.

Then, we leverage the adaptive stepsize in Polyak & Tremba (2019), which is especially amenable to the distributed setting, and propose the Distributed Adaptive Newton algorithm (DAN) to maintain the quadratic convergence rate from any starting point. In DAN, each node i computes the gradient and Hessian of f_i, and then transmits to its neighbors to perform the DSF, after which the Newton direction and the adaptive stepsize are computed for updates. Note that Polyak & Tremba (2019) does not address any distributed issue.

To explicitly reduce the communication cost per iteration, we further propose a Low-rank Approximation (LA) technique to design a communication-efficient DAN-LA. The idea builds on the observation that the rank-1 matrix approximation of a symmetric matrix $W \in \mathbb{R}^{p\times p}$ with respect to (w.r.t.) the spectral norm can be expressed as $s \cdot \mathbf{w}\mathbf{w}^{\mathsf{T}}$, where $s \in \{-1, 1\}$ and $\mathbf{w} \in \mathbb{R}^p$. As first-order methods (Shi et al. 2015, Nedić et al. 2017, Scaman et al. 2017, Qu & Li 2019, Zhang & You 2019*b*), using such an approximation of W only requires transmitting a vector of dimension $\mathcal{O}(p)$. Then, we apply this idea to the "innovation" of Hessian, which is defined as the difference between the true Hessian and its estimate. Noticeably different from the quasi-Newton methods (e.g., BFGS (Nesterov 2018, Nocedal & Wright 2006)), here we can control the approximation error to co-design stepsizes to ensure global superlinear convergence. Note that the truncation method in Mokhtari et al. (2017) cannot achieve a global superlinear convergence rate. Finally, we test DAN and DAN-LA on a logistic regression problem over the Covertype dataset. The results confirm the advantages of DAN and DAN-LA and validate our theoretical results.

The rest of this paper is organized as follows. Section 2 formulates the problem. Section 3 introduces the finite-time DSF with theoretical guarantees. Section 4 first reviews the Newton method and Polyak's adaptive stepsize, and then proposes DAN. Section 5 presents DAN-LA with low-rank matrix approximation methods. We test the two algorithms in Section 6 on a logistic regression problem, and conclude the paper in Section 7. A conference version of this paper appears in Zhang et al. (2020), where none of the proofs have been included and the DSF method in Section 3.2 does not appear.

Notation: $\mathbf{x}^{\mathsf{T}}$ denotes the transpose of $\mathbf{x}$. $\|\cdot\|$ denotes the l_2-norm for vectors and the spectral norm for matrices. $|\mathcal{X}|$ denotes the cardinality of set $\mathcal{X}$. ∇f and $\nabla^2 f$ denote the gradient and Hessian of f, respectively. $\mathcal{O}(\cdot)$ denotes the big-O notation. $A \succeq B$ means $A - B$ is non-negative definite. The ceiling function $\lceil x \rceil$ returns the least integer greater than or equal to x. A sequence $\{e_k\}$ converges to 0 with a Q-linear rate if $\lim_{k\to\infty} |e_{k+1}|/|e_k| = \gamma$ for some $\gamma \in (0, 1)$, a Q-superlinear rate if $\lim_{k\to\infty} |e_{k+1}|/|e_k| = 0$, and a Q-quadratic rate if $\lim_{k\to\infty} |e_{k+1}|/|e_k^2| = \gamma$ for some $\gamma > 0$. A sequence $\{\varepsilon_k\}$ converges R-linearly (superlinearly, quadratically) if there exists a Q-linearly (superlinearly, quadratically) convergent sequence $\{e_k\}$ such that $|\varepsilon_k| \leq |e_k|, \forall k$. In this paper, the convergence rates are always in the sense of R-rate.

Table 1
A brief summary of distributed second-order methods.

Literature	Convergence rates	Main features
Mansoori & Wei (2017, 2020) Tutunov et al. (2019), Zargham et al. (2014) Mokhtari et al. (2017)	Linear [a], *inexact* [b]	Approximate the Hessian of a penalty function of (1) with truncated Taylor series
Mokhtari et al. (2016), Eisen et al. (2017, 2019)	Linear	Primal-dual methods
Varagnolo et al. (2016), Zhang et al. (2021)	Linear	Asymptotic average consensus
Shamir et al. (2014), Wang et al. (2018)	Linear	Master-slave networks
Soori et al. (2020)	Superlinear, *local* [c]	Master-slave networks
Qu et al. (2019)	Quadratic, *local* [c]	Finite-time average consensus
DAN of this work	Quadratic	Adaptive stepsize and finite-time consensus
DAN-LA of this work	Superlinear	Low-rank approximation to compress Hessians

[a] Superlinear rate may happen in a region that does not include the optimal solution, and then the rate reduces to linear.
[b] Cannot converge to an exact optimal solution
[c] The starting point should be sufficiently close to the optimal solution

2 Problem Formulation

The communication network among nodes is modeled by a directed network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{1, \cdots, n\}$ denotes the set of nodes and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges with $(i, j) \in \mathcal{E}$ if and only if node i can directly send messages to j. Let $\mathcal{N}_{\text{in}}^i = \{j | (j, i) \in \mathcal{E}\}$ denote the set of in-neighbors and $\mathcal{N}_{\text{out}}^i = \{j | (i, j) \in \mathcal{E}\}$ is the set of out-neighbors of node i. If $\mathcal{G}$ is undirected, i.e., $(i, j) \in \mathcal{E}$ implies $(j, i) \in \mathcal{E}$, we simply denote by $\mathcal{N}_i$ the neighbors of node i. A path from node i to node j is a sequence of consecutively directed edges from i to j. We say $\mathcal{G}$ is *strongly connected* if there exists a directed path between any pair of nodes. The distance between two nodes in $\mathcal{G}$ is the number of edges in a shortest path connecting them, and the diameter $d_{\mathcal{G}}$ of $\mathcal{G}$ is the largest distance between any pair of nodes. A *tree* is an undirected network that is strongly connected without any cycle.

Essentially, (1) is equivalent to the following optimization problem

$$\underset{\mathbf{x} \in \mathbb{R}^p}{\text{minimize}} \;\; f(\mathbf{x}) \triangleq \sum_{i=1}^{n} f_i(\mathbf{x}). \tag{2}$$

We invoke the following assumptions in the paper.

Assumption 1 (a) *f is twice continuously differentiable and μ-strongly convex, i.e., there exists a positive $\mu > 0$ such that $\nabla^2 f(\mathbf{x}) \succeq \mu I, \forall \mathbf{x}$.*
(b) *f has L-Lipschitz continuous Hessian, i.e.,*

$$\|\nabla^2 f(\mathbf{x}) - \nabla^2 f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|, \;\; \forall \mathbf{x}, \mathbf{y}.$$

Assumption 2 (a) *The communication network is strongly connected.*
(b) *Each node has a unique identifier.*

Assumption 1 is standard in the second-order methods (Boyd & Vandenberghe 2004), which implies that f has a unique global minimum point $\mathbf{x}^\star$, i.e., $f(\mathbf{x}^\star) = \min_{\mathbf{x}} f(\mathbf{x})$. Assumption 2 is common to handle directed networks (Xie et al. 2018, Xi et al. 2018) and is easily satisfied. For example, nodes are generally equipped with network interface cards (NIC) for communication, and each card is assigned to a unique MAC address. The MAC address then naturally serves as an identifier.

3 Distributed Finite-time Set-consensus

This section first introduces the concept of distributed finite-time set-consensus that plays a crucial role in both DAN and DAN-LA. Following that, we propose a Distributed Selective Flooding (DSF) algorithm that improves the communication efficiency of the existing DF.

3.1 Distributed Finite-time Set-consensus

Set-consensus aims to ensure that all nodes reach consensus on a set of values among nodes, which is different from value-consensus to agree on some value. Finite-time consensus achieves the consensus goal only after a *finite* number of rounds of communication with neighbors. Until now, most distributed algorithms (Shi et al. 2015, Nedić et al. 2017, Scaman et al. 2017, Qu & Li 2019, Zhang & You 2019*b*) to solve (1) have been built on asymptotic average-consensus methods with linear iterations, which can only attain linear convergence and become a bottleneck for the design of superlinearly convergent algorithms. To resolve this issue, we design the DSF to achieve set-consensus, which improves the DF in Liu et al. (2007), Dias & Bruno (2013), Li et al. (2017).

3.2 The Distributed Selective Flooding Algorithm

The DSF is given in Algorithm 1, where $\mathcal{I}_i(t)$ contains the information that node i has received after t communication rounds, and is updated iteratively. The implementation for undirected and directed networks are slightly different. We first focus on the former case. Let $\mathcal{S}_i$ be the information to be shared by node i, which can be a scalar, a vector, a matrix, or a set. Each node i initializes a set $\mathcal{I}_i(0) = \{\mathcal{S}_i\}$. At the k-th round, for each neighbor j, node i selects an element $e \in \mathcal{I}_i(k-1)$ such that (a) node i has not sent e to node j, and (b) node i has not received e from node j before. Then, node i sends e to node j (Line 3). Meanwhile, it receives an element e_j from node j and copies to $\mathcal{I}_i(k-1)$, i.e., $\mathcal{I}_i(k) = \mathcal{I}_i(k-1) \cup \{e_j | j \in \mathcal{N}_i\}$ (Line 6).

We show in Theorem 1 that each node i of DSF obtains a set containing the information of all nodes after $n-1$ rounds of communication, i.e., $\mathcal{I}_i(n-1) = \{\mathcal{S}_u | u \in \mathcal{V}\}$. That is, nodes reach set-consensus in finite time. In DSF, each message sent from node i to node j is 'new' to node j, and hence no messages shall be repeatedly transmitted over any link (i, j). To this end, unique identifiers are needed to distinguish messages (c.f. Assumption 2(b)), which requires additional $\lceil \log_2(n) \rceil$ bits of memory for each message than iterative consensus methods (e.g. Olfati-Saber & Murray (2004)). On a positive note, however, it is less than 16 bits even for $n \leq 6 \times 10^5$.

For directed networks, the in-neighbors of a node can be different from its out-neighbors. It is infeasible to check whether an element e has been received from an out-neighbor j. Then, the DSF is modified and a node selects e that is not sent to its out-neighbor j, i.e., Line 3 is replaced by Line 5 in Algorithm 1. However, this slows down the consensus seeking and duplicate transmissions may happen. In fact, the DSF in this case reduces to the DF, see Remark 3 for details.

Theorem 1 *Suppose Assumption 2 holds. If $\mathcal{G}$ is a tree, then each node in Algorithm 1 obtains a set containing all nodes' messages after $n-1$ rounds of local communication with its neighbors, i.e., $\mathcal{I}_i(n-1) = \{\mathcal{S}_u | u \in \mathcal{V}\}, \forall i \in \mathcal{V}$. If $\mathcal{G}$ is directed, then the maximum number of communication rounds is $n + d_{\mathcal{G}} - 1$.*

The proof is relegated to Appendix B. Although Theorem 1 only considers undirected trees, the result also holds for other types of undirected networks where we can first construct a spanning tree in a distributed manner and then apply the DSF on the selected tree.

Remark 1 *A spanning tree can be constructed in a distributed way via a simple breadth-first search, see e.g. Lynch (1996, Section 4.2), where $\mathcal{O}(d_{\mathcal{G}})$ communication rounds are needed and the total number of communication bits of all nodes is at most $\mathcal{O}(|\mathcal{E}| \log(n))$ for undirected graphs. Since we focus on static networks, the tree only needs to be constructed once. The DSF is also resilient against dropping some edges if the resulting time-varying graphs always contain a fixed spanning tree. As Theorem 1 dictates, the essential requirement of the graph lies in the existence of a spanning tree. Whether the network conditions can be relaxed to those supporting iterative consensus (Olfati-Saber & Murray 2004) is not addressed in this work.*

Algorithm 1 The DSF — from the viewpoint of node i

Input: A message $\mathcal{S}_i$, which can be a scalar, a vector, a matrix, or a set, and let $\mathcal{I}_i(0) = \{\mathcal{S}_i\}$.

1: **for** $k = 1, 2, \cdots, n-1$ **do**
2: **if** $\mathcal{G}$ is undirected **then**
3: For each node $j \in \mathcal{N}_i$, node i selects an element $e \in \mathcal{I}_i(k-1)$ such that (a) node i has not sent e to node j, and (b) node i has not received e from node j before. Then, node i sends e to node j.
4: **else if** $\mathcal{G}$ is directed **then**
5: For each node $j \in \mathcal{N}^i_{\text{out}}$, node i selects an element $e \in \mathcal{I}_i(k-1)$ that has not been sent to node j from node i, and sends e to node j.
6: Node i receives an element e_j from each neighbor j and copies to $\mathcal{I}_i(k-1)$, i.e.,

$$\mathcal{I}_i(k) = \mathcal{I}_i(k-1) \cup \{e_j | j \in \mathcal{N}_i\}.$$

Output: Each node obtains a set $\mathcal{I}_i(n-1) = \{\mathcal{S}_u | u \in \mathcal{V}\}, \forall i \in \mathcal{V}$ which contains messages of all nodes.

Remark 2 (Optimality for trees) *Since there exists at least one node with only one neighbor in a tree, such a node has to receive $n-1$ transmissions from its neighbor to achieve set-consensus. Thus, the set-consensus needs at least $n-1$ number of transmissions, showing the optimality of Theorem 1.*

Remark 3 (Improvement over the DF) *The DF is designed for directed networks and is equivalent to our DSF for directed networks. However, the DSF improves the DF on undirected networks by eliminating the duplicate transmissions of the same message over an edge. Since the diameter of a line network is $d_{\mathcal{G}} = n-1$, the DSF only requires a half number of communications than that of the DF. Despite the widespread use of the DF, its minimum number of communication rounds has only recently been established (Oliva et al. 2017, Theorem 1).*

4 The Distributed Adaptive Newton Method

We first review the centralized Newton method with a backtracking line search, and reveals its difficulty in the distributed setting. Then, we introduce an adaptive Newton method (Polyak & Tremba 2019), which does not require any line search while maintaining a quadratic convergence rate. Finally, we integrate the

adaptive Newton method with the DSF to propose the Distributed Adaptive Newton method (DAN) and provide its convergence result.

4.1 The Newton Method

The Newton method has the following update rule

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k(\nabla^2 f(\mathbf{x}_k))^{-1}\nabla f(\mathbf{x}_k). \tag{3}$$

If the stepsize $\alpha_k = 1$, then (3) is the *pure* Newton method, which converges quadratically only when the starting point $\mathbf{x}_0$ is sufficiently close to $\mathbf{x}^\star$, i.e., $\|\mathbf{x}_0 - \mathbf{x}^\star\| \leq \mu^2/L$ (Nesterov 2018). If $\mathbf{x}_0$ is far from $\mathbf{x}^\star$, the pure Newton method may diverge.

To ensure global convergence, α_k in (3) is usually determined by a backtracking line search method. For example, the popular Armijo rule chooses $\alpha_k = \beta^l$ where $\beta \in (0,1)$ and l is the smallest nonnegative integer satisfying

$$f(\mathbf{x}_k + \beta^l \mathbf{d}_k) - f(\mathbf{x}_k) \leq \gamma\beta^l \nabla f(\mathbf{x}_k)^\mathsf{T}\mathbf{d}_k \tag{4}$$

where $\gamma \in (0,1)$ and $\mathbf{d}_k = (\nabla^2 f(\mathbf{x}_k))^{-1}\nabla f(\mathbf{x}_k)$ is the Newton direction. Even though this often promises a global superlinear convergence rate, it needs at least $\mathcal{O}((l+1)(n-1))$ communication rounds in the distributed setting per iteration, leading to a huge communication overhead. If $\mathbf{x}_k$ is far away from an optimal point $\mathbf{x}^\star$, then l in (4) is often large, and thus the line search is not suitable for distributed implementation.

4.2 The Polyak's Adaptive Newton Method

In Polyak & Tremba (2019), a damped Newton method with adaptive stepsizes has been proposed to solve nonlinear equations, and has a basic form as follows

$$\begin{aligned} \alpha_k &= \min\left\{1, \frac{\mu^2}{L\|\nabla f(\mathbf{x}_k)\|}\right\} \\ \mathbf{x}_{k+1} &= \mathbf{x}_k - \alpha_k(\nabla^2 f(\mathbf{x}_k))^{-1}\nabla f(\mathbf{x}_k) \end{aligned} \tag{5}$$

where μ and L are defined in Assumption 1.

The idea in (5) is very natural. When $\mathbf{x}_k$ is far from the optimal point $\mathbf{x}^\star$, the algorithm is in the damped Newton phase with the stepsize inversely proportional to the size of $\nabla f(\mathbf{x}_k)$. When $\mathbf{x}_k$ is close to $\mathbf{x}^\star$, then (5) switches to the pure Newton method. Even though it does not involve any line search, it still achieves a global quadratic convergence rate. Particularly, the number of iterations to achieve $\|\mathbf{x}_k - \mathbf{x}^\star\| \leq \epsilon$ is $O\big(\log_2\log_2(1/\epsilon)\big)$, which matches the theoretical rate of the Newton method with line search in Boyd & Vandenberghe (2004).

Algorithm 2 The Distributed Adaptive Newton method (DAN) — from the view of node i

Input: Starting point $\mathbf{x}_i^0 = \mathbf{x}_0, \forall i$.

1: **for** $k = 0, 1, 2, \cdots$ **do**

2: Compute $\mathbf{g}_k^{(i)} = \nabla f_i(\mathbf{x}_k^{(i)})$ and $H_k^{(i)} = \nabla^2 f_i(\mathbf{x}_k^{(i)})$.

3: Obtain $\mathcal{S} = \{(\mathbf{g}_k^{(u)}, H_k^{(u)}), u \in \mathcal{V}\}$ by performing the DSF (Algorithm 1) via $n-1$ communication rounds with neighbors.

4: Compute the global gradient and Hessian $\bar{\mathbf{g}}_k = \sum_{u=1}^n \mathbf{g}_k^{(u)}$ and $\bar{H}_k = \sum_{u=1}^n H_k^{(u)}$.

5: Let $\alpha_k^{(i)} = \min\left\{1, \frac{\mu^2}{L\|\bar{\mathbf{g}}_k\|}\right\}$ and update

$$\mathbf{x}_{k+1}^{(i)} = \mathbf{x}_k^{(i)} - \alpha_k^{(i)}(\bar{H}_k)^{-1}\bar{\mathbf{g}}_k.$$

4.3 The Distributed Adaptive Newton Method

We use the DSF to extend the above idea to propose DAN in Algorithm 2 to solve the distributed optimization problem. At each iteration, each node i computes the local gradient and Hessian of f_i (Line 2), and then runs the DSF to obtain the gradient and Hessian of f (Lines 3 and 4). Finally, each node uses the aggregated gradient and Hessian to perform a Newton step with the adaptive Polyak's stepsize (Line 5). For brevity, all nodes are assumed to start from the same point $\mathbf{x}_0$, which can be relaxed by adding a finite-time value-consensus step.

Theorem 2 *Suppose that Assumptions 1 and 2 hold. Then, all sequences $\{\mathbf{x}_k^{(i)}\}, i \in \mathcal{V}$ in Algorithm 2 converge to an optimal point $\mathbf{x}^\star$, and $\|\nabla f(\mathbf{x}_k^{(i)})\|$ is monotonically decreasing. Moreover, let*

$$\begin{aligned} k_0 &= \max\left\{0, \left\lceil \frac{2L}{\mu^2}\|\nabla f(\mathbf{x}_0)\| \right\rceil - 2\right\}, \\ \gamma &= \frac{L}{2\mu^2}\|\nabla f(\mathbf{x}_0)\| - \frac{k_0}{4} \in [0, \frac{1}{2}). \end{aligned}$$

Then, it holds that

$$\|\nabla f(\mathbf{x}_k^{(i)})\| \leq \begin{cases} \|\nabla f(\mathbf{x}^0)\| - \frac{\mu^2}{2L}k, & k \leq k_0 \\ \frac{2\mu^2}{L}\gamma^{2^{(k-k_0)}}, & k > k_0 \end{cases}$$

and

$$\|\mathbf{x}_k^{(i)} - \mathbf{x}^\star\| \leq \begin{cases} \frac{\mu}{L}(k_0 - k + \frac{2\gamma}{1-\gamma}), & k \leq k_0, \\ \frac{2\mu\gamma^{2^{(k-k_0)}}}{L(1-\gamma^{2^{(k-k_0)}})}, & k > k_0. \end{cases}$$

Proof. One can use mathematical induction to easily show that $\mathbf{x}_k^{(1)} = \cdots = \mathbf{x}_k^{(n)}$ for all k, i.e., all nodes' states

are identical at any time. Let $\bar{\mathbf{x}}_k \triangleq \mathbf{x}_k^{(1)} = \cdots = \mathbf{x}_k^{(n)}$. Then, each node actually performs the following update

$$\alpha_k = \min\left\{1, \frac{\mu^2}{L\|\nabla f(\bar{\mathbf{x}}_k)\|}\right\}$$
$$\bar{\mathbf{x}}_{k+1} = \bar{\mathbf{x}}_k - \alpha_k(\nabla^2 f(\bar{\mathbf{x}}_k))^{-1}\nabla f(\bar{\mathbf{x}}_k)$$

which is exactly Polyak's adaptive Newton method (5). Thus, the result follows from Polyak & Tremba (2019, Theorem 4.1). □

By Theorem 2, $\|\nabla f(\mathbf{x}_k^{(i)})\|$ and $\|\mathbf{x}_k^{(i)} - \mathbf{x}^\star\|$ in each node of DAN decrease by at least a constant positive value at each iteration when $k \leq k_0$, after which the decreasing rate becomes quadratic. Specifically, the number of iterations for $\|\mathbf{x}_k^{(i)} - \mathbf{x}^\star\| \leq \epsilon$ in each node is given by

$$k_0 + \log_2 \log_{(1/\gamma)} \left(4\mu/L\epsilon\right) = O\big(\log_2 \log_2(1/\epsilon)\big).$$

It is worth noting that the above quadratic rate is *global* and *exact*, which is in sharp contrast with the existing distributed second-order methods (Qu et al. 2019, Mokhtari et al. 2017, Mansoori & Wei 2020, 2017, Tutunov et al. 2019, Zargham et al. 2014).

Remark 4 (Communication complexity) *Although each node in DAN involves $n-1$ rounds of communicating messages of size $\mathcal{O}(p^2)$ per iteration, the overall communication complexity could be much lower than that of first-order algorithms. In particular, the communication complexity in each node to achieve $\|\mathbf{x}_k^{(i)} - \mathbf{x}^\star\| \leq \epsilon$ is $O\big(np^2 \log_2 \log_2(1/\epsilon)\big)$ for DAN and is $\mathcal{O}(d_{\mathcal{G}} p \log_2 1/\epsilon)$ for the optimal first-order algorithm (Scaman et al. 2017). The latter is higher than the former if a high precision is desired, i.e., $\epsilon \leq \mathcal{O}(2^{-\frac{np}{d_{\mathcal{G}}}})$. A similar analysis can also be made for the computational complexity.*

Despite the global quadratic convergence, DAN may progress relatively slow in the early stage due to the use of small stepsizes in the damped Newton phase ($k \leq k_0$) and the DSF in Line 3 of Algorithm 2. In practice, we can start at a first-order method for a few iterations and then switch to DAN for quadratic convergence afterwards, which is similar to the practical use of the Newton method. In Section 5, we propose a novel communication-efficient DAN with low rank approximation.

5 The Communication-efficient DAN-LA

In this section, we propose a novel communication-efficient version of DAN in Algorithm 3 that reduces the transmitted Hessian to a vector of size $\mathcal{O}(p)$. This reduction is significant in applications with high dimensional decision vectors. Since it maintains a global superlinear convergence rate, it is much faster than the existing distributed Newton-type algorithms (Mokhtari et al. 2017, Mansoori & Wei 2020, 2017, Tutunov et al. 2019, Mokhtari et al. 2016, Eisen et al. 2019, Varagnolo et al. 2016).

To this end, we introduce two novel ideas in Algorithm 3: (a) a symmetric rank-1 matrix in $\mathbb{R}^{p\times p}$ can be represented by the outer product of a vector in $\mathbb{R}^p$ and itself with only a possible sign change; and (b) instead of directly compressing the Hessian, we approximate its *innovation* by a symmetric rank-1 matrix. If the decision vectors do not vary much between two consecutive iterations, the innovation is expected to be "small", which suggests that a rank-1 matrix approximation might not lose much information.

Specifically, suppose that each node has the same estimate of the Hessian of f in (2) at the $(k-1)$-th iteration, say $\hat{H}_{k-1}^{(i)}$, and an estimate $H_{k-1}^{(i)}$ of the Hessian of f_i such that $\hat{H}_{k-1}^{(i)} = \sum_{i=1}^n H_{k-1}^{(i)}$, which can be easily satisfied when $k=1$, e.g., let $\hat{H}_0^{(i)} = H_0^{(i)} = 0, \forall i \in \mathcal{V}$. At the next iteration, node i computes a local Hessian $\nabla^2 f_i(\mathbf{x}_k^{(i)})$ and approximates its innovation via a rank-1 matrix, i.e.,

$$\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)} \approx s_k^{(i)} \cdot \mathbf{h}_k^{(i)} (\mathbf{h}_k^{(i)})^\mathsf{T}$$

where $s_k^{(i)}$ and $\mathbf{h}_k^{(i)}$ are solved by the Eckart-Young-Mirsky Theorem (Markovsky 2012, Theorem 2.23). Particularly, let λ_i be the i-th largest eigenvalue of $\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)}$ in magnitude and $\mathbf{w}_i$ be the associated normalized right eigenvector, i.e., $(\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)})\mathbf{w}_i = \lambda_i \mathbf{w}_i$ and $\|\mathbf{w}_i\| = 1$. Then, $\mathbf{h}_k^{(i)} = \sqrt{|\lambda_1|}\mathbf{w}_1$, $s_k^{(i)} = \text{sign}(\lambda_1)$, $r_k^{(i)} = |\lambda_2|$ (Line 2). Clearly, this can be computed by the eigenvalue decomposition with complexity $\mathcal{O}(p^3)$ or a truncated singular value decomposition (SVD) with a complexity of $\mathcal{O}(p^2)$ in some cases (Allen-Zhu & Li 2016). Overall, the complexity is comparable to that of computing a Newton direction.

Moreover, node i updates an estimate of the local Hessian as $H_k^{(i)} = H_{k-1}^{(i)} + s_k^{(i)} \cdot \mathbf{h}_k^{(i)} (\mathbf{h}_k^{(i)})^\mathsf{T}$(Line 3). Via only communicating $s_k^{(i)}$ and $\mathbf{h}_k^{(i)}$ in the DSF (Line 4), it is able to compute an estimate of the Hessian of f by

$$\hat{H}_k^{(i)} = \hat{H}_{k-1}^{(i)} + \sum_{i=u}^n s_k^{(u)} \cdot \mathbf{h}_k^{(u)} (\mathbf{h}_k^{(u)})^\mathsf{T} \text{(Line 5)}.$$

The Newton direction is then computed based on the estimates $\hat{H}_k^{(i)}$ and $\hat{\mathbf{g}}_k^{(i)}$.

Another challenge is how to handle the error in approximating the Hessian of f. Intuitively, if the error is large, a pure Newton may lead to divergence, and we adaptively turn to a higher rank approximation for a better estimate. Otherwise, we can safely run a Newton update. To achieve it, each node i computes an upper bound on the estimation error $\hat{r}_k^{(i)}$ (Line 5). If the error is smaller than the explicit threshold $\underline{r}$ in (7) (Line 6), it updates via an adaptive Newton step (8) with our novel stepsize (Line 7). Note that the invertibility of $\hat{H}_k^{(i)}$ in (8) always holds, since $\|\hat{H}_k^{(i)} - \nabla^2 f(\mathbf{x}_k^{(i)})\| \leq \underline{r} \leq \mu/3$ and $\nabla^2 f(\mathbf{x}_k^{(i)}) \succeq \mu I$ (c.f. Assumption 1(a)).

More specifically, if $\alpha_k^{(i)} = 0$ at the k-th iteration, then $\mathbf{x}_{k+1}^{(i)} = \mathbf{x}_k^{(i)}$ and we obtain from Lines 2-3 in Algorithm 3 that in the $(k+1)$-th iteration, $\nabla^2 f_i(\mathbf{x}_k^{(i)}) - \left(H_{k-1}^{(i)} + s_k^{(i)}\mathbf{h}_k^{(i)}(\mathbf{h}_k^{(i)})^\top\right) \approx s_{k+1}^{(i)}\mathbf{h}_{k+1}^{(i)}(\mathbf{h}_{k+1}^{(i)})^\top$. Since $\mathbf{h}_k^{(i)}$ and $\mathbf{h}_{k+1}^{(i)}$ are two linearly independent vectors, it implies that $\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)}$ is in fact approximated by a rank-2 matrix $s_k^{(i)}\mathbf{h}_k^{(i)}(\mathbf{h}_k^{(i)})^\top + s_{k+1}^{(i)}\mathbf{h}_{k+1}^{(i)}(\mathbf{h}_{k+1}^{(i)})^\top$, and the approximation error is strictly reduced, i.e., $r_{k+1}^{(i)} < r_k^{(i)}$. Jointly with Lines 4-5, there must exist a smallest nonnegative integer $m \leq p-1$ such that $\hat{r}_{k+m}^{(i)} \leq \underline{r}$ which implies that $\alpha_{k+m}^{(i)} > 0$, and eventually $\mathbf{x}_k^{(i)}$ is updated to a new vector. One can easily see that the above is equivalent to the direct use of rank-$(m+1)$ approximation in (6).

We prove the convergence of DAN-LA under the standard assumption in quasi-Newton methods, see e.g. Nocedal & Wright (2006).

Assumption 3 *The Hessian of f in (2) is upper bounded by M, i.e., $\nabla^2 f(\mathbf{x}) \preceq MI, \forall \mathbf{x} \in \mathbb{R}^p$.*

Theorem 3 *Suppose Assumptions 1-3 hold, and let $\{\mathbf{x}_k^{(i)}\}$ be generated by DAN-LA. For any $c > 0$ in Line 6 of Algorithm 3 and any $i \in \mathcal{V}$, $\|\mathbf{x}_k^{(i)} - \mathbf{x}^\star\|$ converges superlinearly to 0, where $\mathbf{x}^\star$ is the optimal point of (2).*

Similar to quasi-Newton methods (Nocedal & Wright 2006), it is very difficult to explicitly quantify its convergence rate. Some remarks are provided after the proof.

Proof of Theorem 3. We first prove the convergence of $\{\mathbf{x}_k^{(i)}\}$ and $\{r_k^{(i)}\}$, and then invoke Dennis-Moré Theorem (Dennis & Moré 1974) to prove global superlinear convergence.

Step 1: A key inequality. It can readily be obtained by mathematical induction that $\mathbf{x}_k^{(1)} = \cdots = \mathbf{x}_k^{(n)} \triangleq \mathbf{x}_k$, $\hat{\mathbf{g}}_k^{(1)} = \cdots = \hat{\mathbf{g}}_k^{(n)} = \mathbf{g}_k = \nabla f(\mathbf{x}_k)$, $\alpha_k^{(1)} = \cdots = \alpha_k^{(n)} \triangleq \alpha_k$ and $\hat{H}_k^{(1)} = \cdots = \hat{H}_k^{(n)} \triangleq \hat{H}_k$. That is, all agents actually perform identical updates with aid of the set-consensus algorithm, and (8) becomes

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k(\hat{H}_k)^{-1}\mathbf{g}_k. \tag{9}$$

Moreover, it implies that $\hat{H}_k = \sum_{i=1}^n H_k^{(i)}$ (c.f. Lines 3 and 5). By the Eckart-Young-Mirsky Theorem (Markovsky 2012, Theorem 2.23), we have $\|\nabla^2 f_i(\mathbf{x}_k) - H_k^{(i)}\| = r_k^{(i)}$ and

$$\begin{aligned}\|\nabla^2 f(\mathbf{x}_k) - \hat{H}_k\| &= \Big\|\sum_{i=1}^n \nabla^2 f_i(\mathbf{x}_k) - H_k^{(i)}\Big\| \\ &\leq \sum_{i=1}^n \|\nabla^2 f_i(\mathbf{x}_k) - H_k^{(i)}\| = \hat{r}_k^{(i)}.\end{aligned} \tag{10}$$

Recall that $\mathbf{g}_k = \nabla f(\mathbf{x}_k)$, and let

$$\mathbf{z}_k = \mathbf{x}_{k+1} - \mathbf{x}_k = -\alpha_k(\hat{H}_k)^{-1}\mathbf{g}_k. \tag{11}$$

Then, we have

$$\begin{aligned}\mathbf{g}_{k+1} &= \mathbf{g}_k + \int_0^1 \nabla^2 f(\mathbf{x}_k + t\mathbf{z}_k)\mathbf{z}_k dt \\ &= \mathbf{g}_k + \int_0^1 \big(\nabla^2 f(\mathbf{x}_k + t\mathbf{z}_k) - \hat{H}_k\big)\mathbf{z}_k dt + \hat{H}_k\mathbf{z}_k \\ &= (1-\alpha_k)\mathbf{g}_k + \int_0^1 \Big(\nabla^2 f(\mathbf{x}_k) - \hat{H}_k\Big)\mathbf{z}_k dt \\ &\quad + \int_0^1 \Big(\nabla^2 f(\mathbf{x}_k + t\mathbf{z}_k) - \nabla^2 f(\mathbf{x}_k)\Big)\mathbf{z}_k dt\end{aligned} \tag{12}$$

We now show that $\alpha_k \in (0, 1]$ by proving that ϕ in (7) is positive. Let $\tilde{\phi}(\underline{r}) \triangleq \frac{L}{2(\mu-\underline{r})^2}\phi = \frac{\mu}{M+\mu} - \frac{\underline{r}}{\mu-\underline{r}}$. Then, $\tilde{\phi}(\underline{r})$ is strictly decreasing on $(0, \mu)$. Since $\underline{r} = \frac{1}{3}\sqrt{M_c^2 + 3\mu^2} - \frac{M_c}{3} = \mu^2/(\sqrt{M_c^2 + 3\mu^2} + M_c) < \mu^2/(M + 2\mu) < \mu$, then $\tilde{\phi}(\underline{r}) > \tilde{\phi}(\frac{\mu^2}{M+2\mu}) = 0$, i.e., $\phi > 0$.

It follows from (12) that

$$\begin{aligned}\|\mathbf{g}_{k+1}\| &\leq (1-\alpha_k)\|\mathbf{g}_k\| + \|\nabla^2 f(\mathbf{x}_k) - \hat{H}_k\|\|\mathbf{z}_k\| \\ &\quad + \int_0^1 \Big\|\nabla^2 f(\mathbf{x}_k + t\mathbf{z}_k) - \nabla^2 f(\mathbf{x}_k)\Big\|\|\mathbf{z}_k\| dt \\ &\leq (1-\alpha_k)\|\mathbf{g}_k\| + \hat{r}_k^{(i)}\|\mathbf{z}_k\| + \frac{L}{2}\|\mathbf{z}_k\|^2\end{aligned} \tag{13}$$

where we used (10) and Assumption 1 in the last inequality.

Step 2: prove that $\lim_{k\to\infty}\{\alpha_k\} = 1$ and $\{\mathbf{x}_k\}$ converges. Let $\mathcal{K} = \{k_l : l \geq 0\}$ be an increasing sequence such that $k \in \mathcal{K}$ if and only if $\alpha_k > 0$ at iteration k, which follows from (7) that $\hat{r}_k^{(i)} \leq \underline{r}$.

Algorithm 3 The DAN-LA — from the view of node i

Input: Starting point $\mathbf{x}_0^{(i)} = \mathbf{x}_0$, $\mathbf{g}_0^{(i)} = \nabla f_i(\mathbf{x}_0)$, $\hat{H}_{-1}^{(i)} = H_{-1}^{(i)} = 0$, $c > 0$, and μ, L, M in Assumptions 1 and 3.

1: **for** $k = 0, 1, 2, \cdots$ **do**

2: Let λ_j be the j-th largest eigenvalue in magnitude of $\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)}$, $\mathbf{w}_j$ be the associated unit eigenvector and

$$\mathbf{h}_k^{(i)} = \sqrt{|\lambda_1|}\mathbf{w}_1,\ s_k^{(i)} = \text{sign}(\lambda_1),\ r_k^{(i)} = |\lambda_2|. \tag{6}$$

3: Compute $H_k^{(i)} = H_{k-1}^{(i)} + s_k^{(i)}\mathbf{h}_k^{(i)}(\mathbf{h}_k^{(i)})^\mathsf{T}$ and $\mathbf{g}_k^{(i)} = \nabla f_i(\mathbf{x}_k^{(i)})$. *% $H_k^{(i)}$ is the rank-1 approximation of $\nabla^2 f_i(\mathbf{x}_k^{(i)})$*

4: Run Algorithm 1 to obtain $\mathcal{S} = \text{DSF}(r_k^{(i)}, s_k^{(i)}, \mathbf{g}_k^{(i)}, \mathbf{h}_k^{(i)})$, where $\mathcal{S} = \{(r_k^{(u)}, s_k^{(u)}, \mathbf{g}_k^{(u)}, \mathbf{h}_k^{(u)}) | u \in \mathcal{V}\}$. *% Finite-time set-consensus*

5: Use $\mathcal{S}$ to compute $\hat{\mathbf{g}}_k^{(i)} = \sum_{u=1}^n \mathbf{g}_k^{(u)}$, $\hat{H}_k^{(i)} = \hat{H}_{k-1}^{(i)} + \sum_{u=1}^n s_k^{(u)}\mathbf{h}_k^{(u)}(\mathbf{h}_k^{(u)})^\mathsf{T}$, and $\hat{r}_k^{(i)} = \sum_{u=1}^n r_k^{(u)}$. *% $\hat{H}_k^{(i)}$ is an approximation of $\nabla^2 f(\mathbf{x}_k^{(i)})$ and $\hat{r}_i^{(i)}$ bounds the approximation error*

6: Let $M_c = M + c$. Set the stepsize *% Update only if the approximation error is small*

$$\alpha_k^{(i)} = \begin{cases} \min\left\{1, \frac{\phi}{\|\hat{\mathbf{g}}_k^{(i)}\|}\right\}, \text{ where } \phi \triangleq \frac{2\mu(\mu-\underline{r})^2}{L(M+\mu)} - \frac{2\underline{r}(\mu-\underline{r})}{L} > 0, & \text{if } \hat{r}_k^{(i)} \le \underline{r} \triangleq \frac{1}{3}\sqrt{M_c^2 + 3\mu^2} - \frac{M_c}{3} \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

7: If $\alpha_k^{(i)} = 0$, set $\mathbf{x}_{k+1}^{(i)} = \mathbf{x}_k^{(i)}$. Otherwise, update *% Adaptive Newton step*

$$\mathbf{x}_{k+1}^{(i)} = \mathbf{x}_k^{(i)} - \alpha_k^{(i)}(\hat{H}_k^{(i)})^{-1}\hat{\mathbf{g}}_k^{(i)}. \tag{8}$$

We show that $k_{l+1} - k_l \le p$ and $\mathbf{x}_{k_l+1} = \mathbf{x}_{k_l+2} = \cdots = \mathbf{x}_{k_{l+1}}$, where p is the dimension of $\mathbf{x}_k$. For any $k_l < k < k_{l+1}$, it follows from (9) and $\alpha_k = 0$ that $\mathbf{x}_{k+1} = \mathbf{x}_k$, and the second part follows by induction. Suppose that there exists an l such that $k_{l+1} - k_l > p$. Let $\lambda_k^{(j)}, j = 1, \cdots, p$ be the j-th largest eigenvalue of $\nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)}$ in magnitude and $\mathbf{w}_k^{(j)}$ be the corresponding unit eigenvector. Define $\lambda_k^{(j)} = 0$ for any $j > p$. Since $\nabla^2 f_i(\mathbf{x}_k^{(i)}) = \nabla^2 f_i(\mathbf{x}_{k+1})$ for any $k_l < k < k_{l+1}$, it follows from (6) that

$$\begin{aligned} &\nabla^2 f_i(\mathbf{x}_{k+1}) - H_k^{(i)} \\ &= \nabla^2 f_i(\mathbf{x}_k) - H_{k-1}^{(i)} - \lambda_k^{(1)}\mathbf{w}_k^{(1)}(\mathbf{w}_k^{(1)})^\mathsf{T} \\ &= \sum_{j=2}^p \lambda_k^{(j)}\mathbf{w}_k^{(j)}(\mathbf{w}_k^{(j)})^\mathsf{T}, \end{aligned}$$

and hence $\lambda_{k+1}^{(j)} = \lambda_k^{(j+1)}$ for all $j > 0$, which further implies that $|\lambda_{k+1}^{(j)}| \le |\lambda_k^{(j)}|$. Then, for any $k \in [k_l + p, k_{l+1}) \neq \emptyset$, we have $r_k^{(i)} = |\lambda_k^{(2)}| \le |\lambda_{k_l+p}^{(2)}| = |\lambda_{k_l+1}^{(p+1)}| = 0, \forall i$. Thus, $\hat{r}_k^{(i)} = \sum_{i=1}^n r_k^{(i)} = 0$ and we must have $\alpha_k > 0$, and hence $k \in \mathcal{K}$, which leads to a contradiction.

Next, we study the sequence $\{\|\mathbf{g}_{k_l}\| : k_l \in \mathcal{K}\}$. Consider the first stage that $\alpha_{k_l} = \phi/\|\hat{\mathbf{g}}_{k_l}^{(i)}\| \le 1$. It follows from (10) and the definition of $\mathcal{K}$ that

$$\|\nabla^2 f(\mathbf{x}_{k_l}) - \hat{H}_{k_l}\| \le \hat{r}_{k_l}^{(i)} \le \underline{r} = \frac{\underline{r}\alpha_{k_l}\|\hat{\mathbf{g}}_{k_l}^{(i)}\|}{\phi} \tag{14}$$

which jointly with Assumption 3 implies that $\|\hat{H}_{k_l}\| \le \|\nabla^2 f(\mathbf{x}_{k_l})\| + \underline{r} \le M + \underline{r}$. Thus, it follows from (11) that $\|\mathbf{z}_{k_l}\| = \alpha_{k_l}\|(\hat{H}_{k_l})^{-1}\mathbf{g}_{k_l}\| = \alpha_{k_l}\sqrt{(\mathbf{g}_{k_l})^\mathsf{T}(\hat{H}_{k_l})^{-2}\mathbf{g}_{k_l}} \ge \alpha_{k_l}\|\mathbf{g}_{k_l}\|/(M + \underline{r})$ which, together with (14), implies that

$$\hat{r}_{k_l}^{(i)} \le \frac{\underline{r}\alpha_{k_l}\|\hat{\mathbf{g}}_{k_l}^{(i)}\|}{\phi} \le \frac{\underline{r}(M+\underline{r})\|\mathbf{z}_{k_l}\|}{\phi}. \tag{15}$$

Combining (15) and (13) yields that

$$\begin{aligned} &\|\mathbf{g}_{k_{l+1}}\| = \|\mathbf{g}_{k_l+1}\| \\ &\le (1 - \alpha_{k_l})\|\mathbf{g}_{k_l}\| + \left(\frac{L}{2} + \frac{\underline{r}(M+\underline{r})}{\phi}\right)\|\mathbf{z}_{k_l}\|^2 \\ &\le \|\mathbf{g}_{k_l}\| - \alpha_{k_l}\|\mathbf{g}_{k_l}\| + \frac{L\phi + 2\underline{r}(M+\underline{r})}{2\phi(\mu-\underline{r})^2}(\alpha_{k_l}\|\mathbf{g}_{k_l}\|)^2 \\ &= \|\mathbf{g}_{k_l}\| - \phi + \frac{L\phi + 2\underline{r}(M+\underline{r})}{2(\mu-\underline{r})^2}\phi \end{aligned} \tag{16}$$

where the first equality follows from $\mathbf{x}_{k_l+1} = \mathbf{x}_{k_l+2} = \cdots = \mathbf{x}_{k_{l+1}}$, the second inequality follows from (11) and (14), and the last equality follows from the definition of the first stage.

Let $\theta = \mu/(M + \mu)$. The definition of ϕ implies that

$$\frac{L\phi}{2(\mu-\underline{r})^2} = \theta - \frac{\underline{r}}{\mu-\underline{r}}. \tag{17}$$

Let $q(r) \triangleq -\frac{r}{\mu-r} + \frac{r(M+r)}{(\mu-r)^2} = \frac{2r^2+(M-\mu)r}{(\mu-r)^2}$ which is strictly increasing on $[0, \mu)$. Notice that $q(0) = 0$ and $q(r_0) = 1/2$ with $r_0 = 1/3 \cdot (\sqrt{M^2 + 3\mu^2} - M)$. Since

$0 < \underline{r} < r_0$, we have $q(\underline{r}) < q(r_0) = 1/2$. Since $M \geq \mu$, it jointly with (17) implies that

$$\frac{L\phi + 2\underline{r}(M+\underline{r})}{2(\mu-\underline{r})^2} < \frac{\mu}{M+\mu} + \frac{1}{2} \leq 1.$$

By (16), we obtain that

$$\|\mathbf{g}_{k_{l+1}}\| \leq \|\mathbf{g}_{k_l}\| - \Big(1 - \frac{L\phi + 2\underline{r}(M+\underline{r})}{2(\mu-\underline{r})^2}\Big)\phi. \tag{18}$$

Thus, the sequence $\{\|\mathbf{g}_{k_l}\|\}$ is monotonically decreasing by at least a constant at each step in the first stage. (18) implies that there exists a k_l such that the algorithm enters the second stage where $\alpha_{k_l} = 1$, i.e., $\|\mathbf{g}_{k_l}\| \leq \phi$. Similar to (16), it follows from (13) that

$$\begin{aligned}\|\mathbf{g}_{k_{l+1}}\| = \|\mathbf{g}_{k_l+1}\| &\leq \hat{r}_k^{(i)}\|\mathbf{z}_{k_l}\| + \frac{L}{2}\|\mathbf{z}_{k_l}\|^2 \\ &\leq \frac{\underline{r}}{\mu-\underline{r}}\|\mathbf{g}_{k_l}\| + \frac{L}{2(\mu-\underline{r})^2}\|\mathbf{g}_{k_l}\|^2.\end{aligned} \tag{19}$$

Let $e_{k_l} = \|\mathbf{g}_{k_l}\|/\phi \leq 1$. Then, it implies that

$$\begin{aligned}e_{k_{l+1}} &\leq \frac{\underline{r}}{\mu-\underline{r}}e_{k_l} + \frac{L\phi}{2(\mu-\underline{r})^2}(e_{k_l})^2 \\ &\leq \Big(\frac{\underline{r}}{\mu-\underline{r}} + \frac{L\phi}{2(\mu-\underline{r})^2}\Big)e_{k_l}.\end{aligned} \tag{20}$$

Recall from (17) that

$$\frac{\underline{r}}{\mu-\underline{r}} + \frac{L\phi}{2(\mu-\underline{r})^2} = \theta \leq \frac{1}{2}, \tag{21}$$

which, combined with (20), shows that $\{e_{k_l}\}$ and $\{\|\mathbf{g}_{k_l}\|\}$ are monotonically decreasing and converge to 0 at least linearly at the second stage. This also concludes that $\|\mathbf{g}_k\|$ will remain in the second stage for all $k \geq k_l$, and hence $\lim_{k\to\infty}\alpha^k = 1$.

Combining all above, it implies that $\{\|\mathbf{g}_{k_l}\| : k_l \in \mathcal{K}\}$ converges to 0. Since $k_{l+1} - k_l \leq p$ and $\mathbf{x}_{k_l+1} = \cdots = \mathbf{x}_{k_{l+1}}$, then $\|\mathbf{g}_k\|$ converges to 0 as well. Since $\|\mathbf{x}_k - \mathbf{x}^\star\| \leq \frac{1}{\mu}\|\mathbf{g}_k\|$, we have thus proved that $\mathbf{x}_k$ converges to $\mathbf{x}^\star$.

Step 3: prove the convergence of $\{r_k^{(i)}\}$ to 0. We now show that the approximation error $\{r_k^{(i)}\}$ converges to 0 for all i. Denote by $\sigma_j(A)$ the j-th largest singular value of $A \in \mathbb{R}^{p\times p}$. Let $E_k = \nabla^2 f_i(\mathbf{x}_k^{(i)}) - H_{k-1}^{(i)}$ and $\varepsilon_k = \|E_k\|_*$, where $\|A\|_* = \sum_{j=1}^p \sigma_j(A) \geq 0$ is the nuclear norm of A. Since the algorithm will enter the second stage after a finite number of iterations, it is sufficient to focus only on this stage. Denote by $k_l \in \mathcal{K}$ the first time step when the algorithm enters it. Then,

$$\begin{aligned}&\varepsilon_{k_{l+1}} \\ &\leq \varepsilon_{k_l+1} = \|\nabla^2 f_i(\mathbf{x}_{k_l+1}^{(i)}) - H_{k_l}^{(i)}\|_* \\ &= \|\nabla^2 f_i(\mathbf{x}_{k_l+1}^{(i)}) - \nabla^2 f_i(\mathbf{x}_{k_l}^{(i)}) + \nabla^2 f_i(\mathbf{x}_{k_l}^{(i)}) - H_{k_l}^{(i)}\|_* \\ &\leq \|\nabla^2 f_i(\mathbf{x}_{k_l}^{(i)}) - H_{k_l}^{(i)}\|_* + \|\nabla^2 f_i(\mathbf{x}_{k_l+1}^{(i)}) - \nabla^2 f_i(\mathbf{x}_{k_l}^{(i)})\|_* \\ &\leq \sum_{j=2}^p \sigma_j(E_{k_l}) + p\|\nabla^2 f_i(\mathbf{x}_{k_l+1}^{(i)}) - \nabla^2 f_i(\mathbf{x}_{k_l}^{(i)})\| \\ &\leq \varepsilon_{k_l} - \sigma_1(E_{k_l}) + \frac{pL}{\mu-\underline{r}}\|\mathbf{g}_{k_l}\| \\ &\leq (1-\frac{1}{p})\varepsilon_{k_l} + \frac{pL}{\mu-\underline{r}}\|\mathbf{g}_{k_l}\|\end{aligned} \tag{22}$$

where we used $\|A\|_* \leq p\|A\|$, Assumption 1, the rank-1 approximation property, and $0 \leq \varepsilon_k \leq p\sigma_1(E_k)$. Note that (19)-(21) show that $\|\mathbf{g}_{k_l}\|, k_l \in \mathcal{K}$ converges to 0 at least linearly. Then, it follows from (22) that $\lim_{k\to\infty}\varepsilon_k = 0$. The convergence of $\{r_k^{(i)}\}$ to 0 follows immediately by noticing $\varepsilon_k \geq r_k^{(i)}$.

Step 4: prove the superlinear convergence. Now we study the sequence $\hat{H}_k$. It follows from (10) and Assumption 1 that

$$\begin{aligned}\|\hat{H}_k - \nabla^2 f(\mathbf{x}^\star)\| &\leq \|\hat{H}_k - \nabla^2 f(\mathbf{x}_k)\| + L\|\mathbf{x}_k - \mathbf{x}^\star\| \\ &\leq \sum_{i=1}^n r_k^{(i)} + L\|\mathbf{x}_k - \mathbf{x}^\star\|.\end{aligned}$$

Since we have already shown that both $\|\mathbf{x}_k - \mathbf{x}^\star\|$ and $\sum_{i=1}^n r_k^{(i)}$ converge to 0, then $\hat{H}_k$ converges to $\nabla^2 f(\mathbf{x}^\star)$. In view of (9), it follows from Lemma 1 that the superlinear convergences of $\|\mathbf{x}_k - \mathbf{x}^\star\|$ and $\|\nabla f(\mathbf{x}_k)\|$. □

Remark 5 (Relation with BFGS) *Even though DAN-LA is also a quasi-Newton method, it does not imply that we can combine DSF and other quasi-Newton methods (e.g. BFGS (Nesterov 2018, Nocedal & Wright 2006)) to achieve the global superlinear convergence. As it usually requires using a line-search scheme, one of the main challenges lies in the design of stepsizes in the distributed setting. A striking difference from BFGS is that the Hessian approximation error in DAN-LA can be explicitly controlled, which is essential to our stepsize co-design (Line 6).*

In fact, BFGS has been originally adopted to avoid computing a Hessian, whereas the Hessian approximation in DAN-LA aims to reduce the communication cost and nodes still need to compute the Hessian.

Remark 6 (The effect of parameter c) *The parameter c in Line 6 aims to balance the communication*

cost and the computational cost. Specifically, from (7), a larger c means a better approximation of the Hessian of f and also suggests a larger stepsize. However, it may result in a larger number of un-updated iterations. How the parameter c affects the global convergence rate is related to the network bandwidth, parameters in Assumption 1, etc. In fact, DAN-LA reduces to DAN if c tends to infinity, and it is also feasible to set $c = 0$ if M is strictly *larger than μ. We will empirically show its effect in Section 6.*

Remark 7 (Comparison with DAN) *Compared to DAN, DAN-LA reduces the transmitted messages' size from $\mathcal{O}(p^2 + p)$ to $\mathcal{O}(2p + 1)$ per iteration, which is essentially identical to existing first-order methods, e.g., Nedić et al. (2017) requires nodes to transmit messages with size $\mathcal{O}(2p)$. Nevertheless, DAN-LA may require more iterations than DAN to achieve the same level of accuracy. It is difficult to conclude which one is always better from a theoretical point of view. In our experiments, we find that the number of total transmitted messages in bits of DAN-LA is much smaller than that of DAN to achieve the same level of accuracy, while the computation increases.*

Remark 8 (Importance of set-consensus) *Although the finite-time set-consensus step in DAN can be replaced by a finite-time* average *consensus (e.g. Charalambous & Hadjicostis 2018, Wang & Ong 2018, Yuan et al. 2013) to reduce memory size, it is indispensable in DAN-LA. In Line 5 of Algorithm 3, nodes need to compute the summation $\frac{1}{n}\sum_{i=1}^{n} s_k^{(i)}\mathbf{h}_k^{(i)}(\mathbf{h}_k^{(i)})^\mathsf{T}$, which requires transmitting messages in size $\mathcal{O}(p^2)$ to their neighbors if iterative consensus methods are adopted. In contrast, the transmitted messages are of size $\mathcal{O}(p)$ in the DSF.*

Remark 9 *Similar to DAN, it is suggested to initially perform several iterations of first-order methods to achieve a good starting point for DAN-LA. In view of Line 5, the complexity of obtaining the inverse in (8) can be reduced from $\mathcal{O}(p^3)$ to $\mathcal{O}(np^2)$ if $n < p$ by invoking the Sherman-Morrison-Woodbury formula (Horn & Johnson 2012, Section 0.7.4).*

6 Numerical Examples

In this section, we test DAN and DAN-LA by training a binary logistic regression classifier for the Covertype datatset from the UCI machine learning repository (Dheeru & Karra Taniskidou 2017), where the samples in Classes 3 and 7 are used. The optimization problem involved has the following form:

$$\min. l(\omega) \triangleq -\sum_{i=1}^{m} y_i \ln \sigma(z_i) + (1-y_i)\ln(1-\sigma(z_i)) + \frac{\rho}{2}\|\omega\|^2$$

where $\omega \in \mathbb{R}^{55}$ and $m = 56264$ is the number of samples; $z_i = \omega^\mathsf{T}\mathbf{x}_i$ where $\mathbf{x}_i \in \mathbb{R}^{55}$ is the feature of the i-th sample with each entry normalized to $[-1, 1]$, and $y_i \in \{0, 1\}$ is the corresponding label. The regularization parameter is chosen as $\rho = 0.01m$. The gradient and Hessian are respectively $\nabla l(\omega) = \sum_{i=1}^{m} \mathbf{x}_i(\sigma(z_i) - y_i) + r\omega$ and $\nabla^2 l(\omega) = \sum_{i=1}^{m} \mathbf{x}_i\mathbf{x}_i^\mathsf{T}\sigma(z_i)(1 - \sigma(z_i)) + rI$.

For distributed training, we randomly partition the dataset over $n = 10$ or $n = 100$ nodes with each one privately holding a local subset. We compare our algorithms with the four second-order methods: FNRC (Varagnolo et al. 2016), ESOM-3 (Mokhtari et al. 2016), Newton tracking (Zhang et al. 2021), D-BFGS (Eisen et al. 2017), and a first-order method: DIGing (Nedić et al. 2017). An undirected communication network is constructed by adopting the Erdős-Rényi model (Erdős & Rényi 1960), i.e., each pair of nodes is connected with probability $2\ln n/n$. For comparison, the edge weights are generated by the Metropolis method (Nedić et al. 2017, Shi et al. 2015).

We also implement the centralized gradient descent method as a baseline, where the training is conducted on a single node. In all algorithms, we set $\mu = 0.02m$, $L = m$ and $M = 0.04m$ in Assumptions 1 and 3 to guide the selection of stepsizes. For example, the stepsize in the centralized gradient method is set to the optimal one $2/(\mu+L)$ (Nesterov 2018) and the stepsize used in DIGing is a/L with a determined by a grid search. If the theoretical suggestions are not clear, we manually tune the stepsizes to obtain the numerically best one. The convergence behaviors of these algorithms are depicted in Figs. 1 and 2.

Fig. 1 plots the convergence rates for the cases of $n = 10$ and $n = 100$, respectively. The rate is measured by the decreasing speed of the global gradient's norm versus the number of evaluations of the local gradient or Hessian. It shows that both DAN and DAN-LA achieve superlinear convergence and outperform DIGing, FNRC, D-BFGS, Newton tracking, and ESOM-3. Though they progress slow in the early stages, they rapidly outpace others in the pure Newton phase as expected. In practice, one can adopt some first-order methods initially and then switch to DAN or DAN-LA (c.f. Remark 9). Fig. 1 validates Remark 6 where a larger c leads to a better approximation quality of the global Hessian and a larger stepsize.

Fig. 2 shows the convergence rates versus the averaged transmitted bits of a node (a number is stored and transmitted in 64-bit floating-point format). The DAN and DAN-LA behave similarly as in Fig. 1. Moreover, the transmission of Hessians results in large communication overhead of DAN before it enters the pure Newton phase. Thus, DAN seems more suitable for networks with high bandwidth in modest size, and DAN-LA with an appropriate c is much more communication efficient.

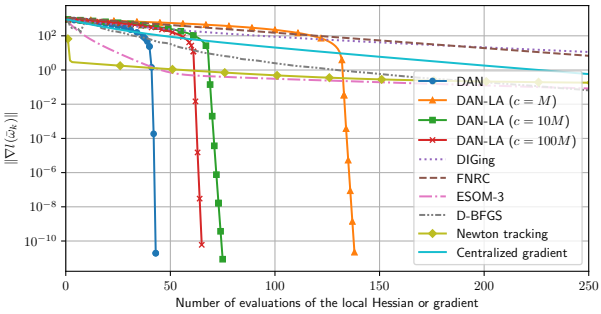

(a) $n = 10$

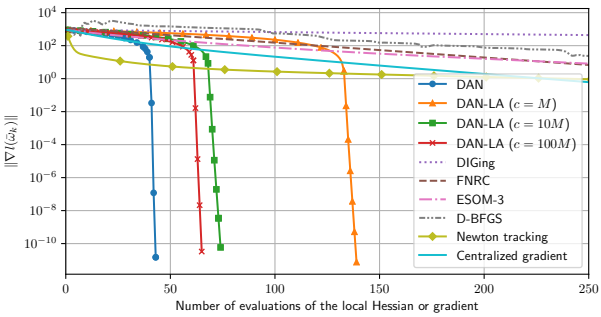

(b) $n = 100$

Fig. 1. The convergence rates w.r.t. the number of evaluations of local Hessian or gradient.

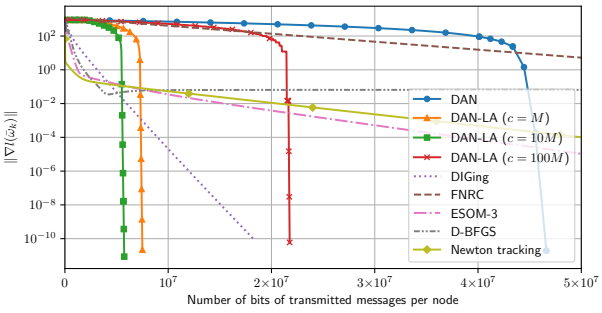

(a) $n = 10$

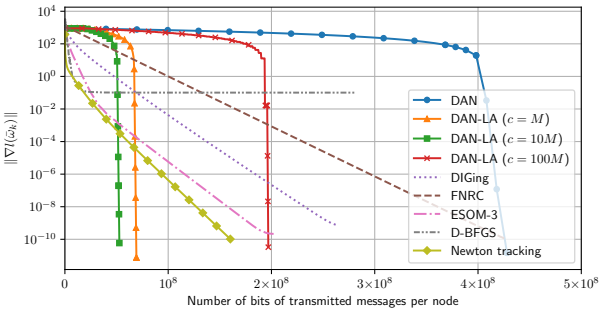

(b) $n = 100$

Fig. 2. The convergence rate w.r.t. the transmitted bits.

7 Conclusion

This paper has proposed two distributed second-order optimization algorithms with global superlinear convergence. The striking features lie in the use of (a) a finite-time set-consensus method, (b) an adaptive version of Newton method for global convergence, and (c) the low-rank matrix approximation methods to compress the Hessian for efficient communication. Future works can focus on asynchronous versions of the proposed algorithms as in (Zhang & You 2019*b*,*a*), the integration of the proposed algorithms with quasi-Newton methods, and developing more communication-efficient ones.

A Technical Lemma

Lemma 1 (Dennis & Moré 1974) *Suppose f is twice continuously differentiable and $\nabla^2 f(\mathbf{x}^\star)$ is nonsingular for some $\mathbf{x}^\star$. Let $\{B_k\}$ be a sequence of nonsingular matrices, and for some $\mathbf{x}_0$ let the sequence $\{\mathbf{x}_k\}$ converge to $\mathbf{x}^\star$, where $\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k (B_k)^{-1} \nabla f(\mathbf{x}_k)$. Then, $\{\mathbf{x}_k\}$ converges superlinearly to $\mathbf{x}^\star$ and $\nabla f(\mathbf{x}^\star) = 0$ if and only if $\lim_{k\to\infty} \alpha_k = 1$ and*

$$\lim_{k\to\infty} \frac{\|[B_k - \nabla^2 f(\mathbf{x}^\star)](\mathbf{x}_{k+1} - \mathbf{x}_k)\|}{\|\mathbf{x}_{k+1} - \mathbf{x}_k\|} = 0.$$

B Proof of Theorem 1

The proof for directed networks is given in Oliva et al. (2017, Theorem 1), and here we just deal with undirected trees. Consider a tree with the root an arbitrary node. Let $\mathcal{U}_v(k) = \{u | \mathcal{S}_u \in \mathcal{I}_v(k)\}$ be the set of nodes whose messages arrive at node v in k iterations, $\mathcal{R}_v$ be the set of *descendants* of v (not including v), and $\widetilde{\mathcal{U}}_v(k) = \mathcal{U}_v(k) \cap \mathcal{R}_v$ be the set of descendants of v whose messages are in $\mathcal{U}_v(k)$. It follows from the definition that $\widetilde{\mathcal{U}}_v(k) \subseteq \mathcal{R}_v$ and $\widetilde{\mathcal{U}}_v(k) \subseteq \mathcal{U}_v(k)$, and $\mathcal{R}_v = \emptyset$ for a leaf node v.

Claim 1: If v is not a leaf and u is a child of v, then $|\mathcal{U}_v(k) \cap \mathcal{R}_u| \leq k - 1$.

Proof: Since $\mathcal{G}$ is a tree, node v can get the messages of nodes in $\mathcal{R}_u$ only from node u. Note that node u sends the message of itself to node v at the first iteration. The result then follows from that only $k-1$ iterations can be used to transmit the messages of nodes in $\mathcal{R}_u$, and each iteration only transmits one message. □

Claim 2: $|\widetilde{\mathcal{U}}_v(k)| \geq \min\{|\mathcal{R}_v|, k\}$ for all $v \in \mathcal{V}$ and $k \in \mathbb{N}$.

Proof: We prove it by mathematical induction. It is clear that $|\widetilde{\mathcal{U}}_v(0)| = 0$ for all $v \in \mathcal{V}$. Suppose that $|\widetilde{\mathcal{U}}_v(t)| \geq \min\{|\mathcal{R}_v|, t\}$ for all $v \in \mathcal{V}$, and we next show that $|\widetilde{\mathcal{U}}_v(t+1)| \geq \min\{|\mathcal{R}_v|, t+1\}$. For some node v, if $|\widetilde{\mathcal{U}}_v(t)| > t$

or $|\widetilde{\mathcal{U}}_v(t)| = |\mathcal{R}_v|$, then the result follows immediately. Now consider $|\widetilde{\mathcal{U}}_v(t)| = t < |\mathcal{R}_v|$, which means that v cannot be a leaf. It follows that there exists a child u of node v and a nonempty subset $\mathcal{R}' \subset \mathcal{R}_u$ such that all elements in $\mathcal{R}'$ are not contained in $\mathcal{U}_v(k)$, i.e., $\mathcal{U}_v(k) \cap \mathcal{R}' = \emptyset$ (otherwise $|\widetilde{\mathcal{U}}_v(t)| = |\mathcal{R}_v|$). Then, the hypothesis implies that $|\mathcal{R}_u| \geq |\widetilde{\mathcal{U}}_u(t)| \geq \min\{|\mathcal{R}_u|, t\}$. If $|\widetilde{\mathcal{U}}_u(t)| = |\mathcal{R}_u|$, then $\mathcal{R}' \subset \mathcal{U}_u(t)$. Note that the selective operation ensures that no messages are transmitted more than once over a link (u, v), which implies that an element in $\mathcal{R}'$ will be sent to node v at iteration $t+1$, and hence $|\mathcal{U}_v(t+1)| \geq t+1$. If $|\widetilde{\mathcal{U}}_u(t)| \geq t$, which means $|\mathcal{R}_u| \geq t$, then it follows from Claim 1 that there must exist a node in $\widetilde{\mathcal{U}}_u(k)$ whose message is not contained in $\mathcal{U}_v(k)$, and hence this element will be sent to node v from node u at iteration $t+1$ due to the selective operation. Therefore, $|\mathcal{U}_v(t+1)| \geq t+1$. □

For any node i, let $k = n-1$ and designate i as the root. Then, it follows from Claim 2 that $|\widetilde{\mathcal{U}}_i(n-1)| \geq n-1$ and $\mathcal{R}_i \subseteq \mathcal{U}_i(n-1)$, which proves the result.

References

Allen-Zhu, Z. & Li, Y. (2016), LazySVD: even faster SVD decomposition yet without agonizing pain, *in* 'Advances in Neural Information Processing Systems', pp. 974–982.

Boyd, S. & Vandenberghe, L. (2004), *Convex Optimization*, Cambridge University Press.

Charalambous, T. & Hadjicostis, C. N. (2018), Laplacian-based matrix design for finite-time aveazge consensus in digraphs, *in* 'IEEE Conference on Decision and Control', IEEE, pp. 3654–3659.

Dennis, J. E. & Moré, J. J. (1974), 'A characterization of superlinear convergence and its application to quasi-newton methods', *Mathematics of Computation* **28**(126), 549–560.

Dheeru, D. & Karra Taniskidou, E. (2017), 'UCI machine learning repository'.
URL: *http://archive.ics.uci.edu/ml*

Dias, S. S. & Bruno, M. G. (2013), 'Cooperative target tracking using decentralized particle filtering and RSS sensors', *IEEE Transactions on Signal Processing* **61**(14), 3632–3646.

Eisen, M., Mokhtari, A. & Ribeiro, A. (2017), 'Decentralized quasi-newton methods', *IEEE Transactions on Signal Processing* **65**(10), 2613–2628.

Eisen, M., Mokhtari, A. & Ribeiro, A. (2019), 'A primal-dual quasi-Newton method for exact consensus optimization', *IEEE Transactions on Signal Processing* **67**(23), 5983–5997.

Erdős, P. & Rényi, A. (1960), 'On the evolution of random graphs', *Publ. Math. Inst. Hung. Acad. Sci* **5**(1), 17–60.

Horn, R. A. & Johnson, C. R. (2012), *Matrix analysis*, Cambridge University Press.

Jadbabaie, A., Ozdaglar, A. & Zargham, M. (2009), A distributed Newton method for network optimization, *in* 'IEEE Conference on Decision and Control held jointly with Chinese Control Conference', IEEE, pp. 2736–2741.

Li, T., Corchado, J. M. & Prieto, J. (2017), 'Convergence of distributed flooding and its application for distributed Bayesian filtering', *IEEE Transactions on Signal and Information Processing over Networks* **3**(3), 580–591.

Liu, H., Jia, X., Wan, P.-J., Liu, X. & Yao, F. F. (2007), 'A distributed and efficient flooding scheme using 1-hop information in mobile ad hoc networks', *IEEE Transactions on Parallel and Distributed Systems* **18**(5), 658–671.

Lynch, N. A. (1996), *Distributed algorithms*, Elsevier.

Mansoori, F. & Wei, E. (2017), Superlinearly convergent asynchronous distributed network Newton method, *in* 'Conference on Decision and Control', IEEE, pp. 2874–2879.

Mansoori, F. & Wei, E. (2020), 'A fast distributed asynchronous Newton-based optimization algorithm', *IEEE Transactions on Automatic Control* **65**(7), 2769–2784.

Markovsky, I. (2012), *Low Rank Approximation : Algorithms, Implementation, Applications*, Springer, London New York.

Mokhtari, A., Ling, Q. & Ribeiro, A. (2017), 'Network Newton distributed optimization methods', *IEEE Transactions on Signal Processing* **65**(1), 146–161.

Mokhtari, A., Shi, W., Ling, Q. & Ribeiro, A. (2016), 'A decentralized second-order method with exact linear convergence rate for consensus optimization', *IEEE Transactions on Signal and Information Processing over Networks* **2**(4), 507–522.

Nedić, A., Olshevsky, A. & Shi, W. (2017), 'Achieving geometric convergence for distributed optimization over time-varying graphs', *SIAM Journal on Optimization* **27**(4), 2597–2633.

Nesterov, Y. (2018), *Lectures on convex optimization*, Vol. 137, Springer.

Nocedal, J. & Wright, S. (2006), *Numerical optimization*, Springer Science & Business Media.

Olfati-Saber, R. & Murray, R. M. (2004), 'Consensus problems in networks of agents with switching topology and time-delays', *IEEE Transactions on Automatic Control* **49**(9), 1520–1533.

Oliva, G., Setola, R. & Hadjicostis, C. N. (2017), 'Distributed finite-time average-consensus with limited computational and storage capability', *IEEE Transactions on Control of Network Systems* **4**(2), 380–391.

Polyak, B. & Tremba, A. (2019), 'New versions of Newton method: step-size choice, convergence domain and under-determined equations', *Optimization Methods and Software* pp. 1–32.

Qu, G. & Li, N. (2019), 'Accelerated distributed nesterov gradient descent', *IEEE Transactions on Automatic Control* **65**(6), 2566–2581.

Qu, Z., Wu, X. & Lu, J. (2019), Finite-time-consensus-

based methods for distributed optimization, *in* 'Chinese Control Conference', IEEE, pp. 5764–5769.

Scaman, K., Bach, F., Bubeck, S., Lee, Y. T. & Massoulié, L. (2017), Optimal algorithms for smooth and strongly convex distributed optimization in networks, *in* 'International Conference on Machine Learning', pp. 3027–3036.

Shamir, O., Srebro, N. & Zhang, T. (2014), Communication-efficient distributed optimization using an approximate Newton-type method, *in* 'International Conference on Machine Learning', pp. 1000–1008.

Shi, W., Ling, Q., Wu, G. & Yin, W. (2015), 'EXTRA: An exact first-order algorithm for decentralized consensus optimization', *SIAM Journal on Optimization* **25**(2), 944–966.

Soori, S., Mishchenko, K., Mokhtari, A., Dehnavi, M. M. & Gurbuzbalaban, M. (2020), DAve-QN: A distributed averaged quasi-Newton method with local superlinear convergence rate, *in* 'International Conference on Artificial Intelligence and Statistics', PMLR, pp. 1965–1976.

Tutunov, R., Bou-Ammar, H. & Jadbabaie, A. (2019), 'Distributed Newton method for large-scale consensus optimization', *IEEE Transactions on Automatic Control* **64**(10), 3983–3994.

Varagnolo, D., Zanella, F., Cenedese, A., Pillonetto, G. & Schenato, L. (2016), 'Newton-Raphson consensus for distributed convex optimization', *IEEE Transactions on Automatic Control* **61**(4), 994–1009.

Wang, S., Roosta-Khorasani, F., Xu, P. & Mahoney, M. W. (2018), GIANT: Globally improved approximate Newton method for distributed optimization, *in* 'Advances in Neural Information Processing Systems', pp. 2332–2342.

Wang, Z. & Ong, C. J. (2018), 'Speeding up finite-time consensus via minimal polynomial of a weighted graph: A numerical approach', *Automatica* **93**, 415–421.

Xi, C., Mai, V., Xin, R., Abed, E. & Khan, U. (2018), 'Linear convergence in optimization over directed graphs with row-stochastic matrices', *IEEE Transactions on Automatic Control* **63**(10), 3558–3565.

Xie, P., You, K., Tempo, R., Song, S. & Wu, C. (2018), 'Distributed convex optimization with inequality constraints over time-varying unbalanced digraphs', *IEEE Transactions on Automatic Control* **63**(12), 4331–4337.

Xin, R. & Khan, U. A. (2019), 'Distributed heavy-ball: A generalization and acceleration of first-order methods with gradient tracking', *IEEE Transactions on Automatic Control* **65**(6), 2627–2633.

Yuan, Y., Stan, G.-B., Shi, L., Barahona, M. & Goncalves, J. (2013), 'Decentralised minimum-time consensus', *Automatica* **49**(5), 1227–1235.

Zargham, M., Ribeiro, A. & Jadbabaie, A. (2012), A distributed line search for network optimization, *in* 'American Control Conference', IEEE, pp. 472–477.

Zargham, M., Ribeiro, A., Ozdaglar, A. & Jadbabaie, A. (2014), 'Accelerated dual descent for network flow optimization', *IEEE Transactions on Automatic Control* **59**(4), 905–920.

Zhang, J., Ling, Q. & So, A. M.-C. (2021), 'A Newton tracking algorithm with exact linear convergence for decentralized consensus optimization', *IEEE Transactions on Signal and Information Processing over Networks* **7**, 346–358.

Zhang, J. & You, K. (2019*a*), 'AsySPA: An exact asynchronous algorithm for convex optimization over digraphs', *IEEE Transactions on Automatic Control* **65**(6), 2494–2509.

Zhang, J. & You, K. (2019*b*), 'Fully asynchronous distributed optimization with linear convergence in directed networks', *arXiv preprint arXiv:1901.08215* .

Zhang, J., You, K. & Başar, T. (2020), Achieving globally superlinear convergence for distributed optimization with adaptive newton method, *in* 'IEEE Conference on Decision and Control', IEEE, pp. 2329–2334.

Zhang, Y. & Lin, X. (2015), Disco: Distributed optimization for self-concordant empirical loss, *in* 'International conference on machine learning', PMLR, pp. 362–370.