algorithms

MDPI

Article

Distributed and Data-Driven Optimization Frameworks for Logistics-Oriented Decision Support Under Partial and Asynchronous Information

Manuel J. C. S. Reis

Engineering Department and IEETA, University of Trás-os-Montes e Alto Douro, Quinta de Prados, 5000-801 Vila Real, Portugal; mcabral@utad.pt

Abstract

This paper introduces D^{3}O-GT, a distributed optimization framework designed to operate under partial, heterogeneous, and delayed information—conditions commonly encountered in large-scale logistics and networked decision support systems. The proposed approach integrates gradient tracking with delay-aware updates to address the steady-state bias and instability that often affect classical distributed gradient methods. We formulate a consensus optimization model that captures decentralized decision variables while preserving global optimality, and we develop an algorithmic structure that balances convergence accuracy, communication efficiency, and robustness to asynchronous updates. Extensive numerical experiments demonstrate that D^{3}O-GT achieves machine precision convergence in synchronous settings and remains stable under bounded communication delays, converging to a small neighborhood of the optimum. In contrast, conventional distributed gradient descent exhibits significant residual error under the same conditions. Scalability analyses further indicate that the proposed method maintains favorable iteration complexity as the number of agents increases. These results position D^{3}O-GT as a practical and scalable solution for distributed decision-making environments, with direct relevance to logistics-oriented applications such as resource allocation, coordination of networked services, and real-time operational planning.

Keywords: distributed optimization; decentralized decision-making; data-driven optimization; network-constrained systems; scalable algorithms; coordination mechanisms; operations and logistics

Check for updates

Academic Editors: Antonio Maria Coruzzolo and Francesco Lolli

Received: 25 February 2026
Revised: 21 March 2026
Accepted: 22 March 2026
Published: 24 March 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

In recent years, logistics systems have become increasingly complex due to globalization, dynamic demand patterns, and the growing need for responsiveness and efficiency. Traditional centralized planning approaches often struggle to cope with large-scale and rapidly changing operational environments, motivating the development of intelligent and adaptive optimization methods.

Recent advances in large-scale cyber–physical systems and data-intensive infrastructures have significantly increased the complexity of logistics decision-making processes. Modern logistics platforms increasingly rely on distributed computational resources, edge computing architectures, and networked decision agents operating with locally available information. In such environments, centralized optimization approaches often become

impractical due to communication bottlenecks, scalability limitations, and data locality constraints. As a consequence, distributed optimization has emerged as a key methodological paradigm for enabling scalable coordination across decentralized operational systems.

The integration of artificial intelligence into logistics optimization has gained significant attention, particularly in the context of sustainable and data-intensive operations. AI-driven models have demonstrated the potential to improve forecasting, routing, and supply chain decision-making while enabling more efficient and environmentally responsible logistics processes [1].

Despite these advances, many modern operational systems exhibit inherently decentralized structures in which decision variables, constraints, and information are distributed across multiple agents or subsystems. Distributed optimization has therefore emerged as a fundamental paradigm for addressing large-scale coordination problems, allowing systems to reach globally optimal solutions while preserving scalability and autonomy [2].

Recent research has further investigated distributed constrained optimization under realistic communication assumptions, such as time-varying network topologies and delayed agent interactions, highlighting the importance of robust algorithmic design in networked environments [3]. These aspects are particularly relevant in logistics-inspired systems, where coordination must often occur under imperfect information exchange.

Resource allocation problems provide a representative example of such challenges. Multi-agent formulations supported by convex optimization techniques have shown strong potential for enabling efficient task allocation and path planning in decentralized settings, reinforcing the role of mathematical optimization as a key enabler of intelligent coordination [4].

In parallel, reinforcement learning and multi-agent decision frameworks have been increasingly explored to enhance adaptive resource allocation and operational flexibility in complex cyber–physical infrastructures, including industrial and logistics scenarios [5].

In the context of logistics systems, distributed optimization methods have been increasingly explored for applications such as fleet coordination, decentralized warehouse management, and collaborative supply chain planning. These environments typically involve geographically distributed decision units that must coordinate their actions under limited information sharing and imperfect communication. Consequently, optimization frameworks capable of operating under partial, heterogeneous, and delayed information exchange are becoming increasingly relevant for real-world operational environments.

Despite the significant progress in distributed optimization, several limitations remain when these methods are considered from an operational decision support perspective. Many existing studies focus primarily on theoretical convergence guarantees under ideal communication assumptions, such as synchronous updates or perfectly reliable information exchange. However, real-world logistics systems frequently operate under asynchronous communication regimes, heterogeneous local information, and delayed message propagation. This gap between theoretical models and operational environments motivates the development of algorithmic frameworks explicitly designed to tolerate imperfect communication conditions while maintaining robust optimization performance.

The main contributions of this paper are summarized as follows:

- We propose D^3O-GT, a distributed optimization framework that integrates gradient tracking with delay-aware updates to support decision-making under partial and asynchronous information.
- We develop an implementation-oriented algorithmic structure that balances convergence accuracy, communication efficiency, and robustness.
- We provide a multi-metric empirical evaluation covering convergence accuracy, consensus behavior, scalability, and delay sensitivity.

- We analyze scalability properties and highlight the practical suitability of the framework for large-scale networked decision environments.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on distributed optimization and logistics decision support. Section 3 introduces the proposed framework and formal problem formulation. Section 4 presents the algorithmic design and theoretical considerations. Section 5 reports numerical experiments, and Section 6 concludes the paper with directions for future research.

2. Related Work

The rapid evolution of large-scale operational systems has stimulated extensive research on distributed optimization and decentralized decision-making. These paradigms are particularly relevant in environments where centralized coordination becomes impractical due to computational constraints, privacy requirements, or communication limitations.

Recent surveys highlight the growing importance of distributed algorithms for multi-agent coordination, emphasizing their role in addressing resource allocation problems characterized by complex interactions and decentralized information structures [6]. Such methods typically rely on consensus mechanisms, dual decomposition strategies, or alternating direction techniques to ensure convergence while preserving scalability.

Beyond classical formulations, distributed optimization has increasingly incorporated privacy-aware mechanisms. Differential privacy, for example, has emerged as a rigorous framework for protecting sensitive information during agent interactions while maintaining algorithmic performance, thereby enabling secure coordination across networked systems [7].

Another relevant research direction concerns aggregative games, in which each agent's objective depends on global population-level variables. Distributed algorithms designed for these settings have demonstrated strong theoretical guarantees and are gaining attention for large-scale cyber–physical and economic systems [8].

From an algorithmic perspective, randomized and parallel approaches have recently been proposed to improve computational efficiency in distributed consensus problems. By allowing subsets of agents to update their decisions asynchronously, these methods significantly reduce computational burden while preserving solution quality [9].

Risk-aware optimization has also attracted interest, particularly in safety-critical distributed environments. Modern formulations integrate stochastic programming and multivariate risk measures to enhance robustness against uncertainty, reinforcing the importance of resilient decision support architectures [10].

In parallel, learning-based strategies are reshaping distributed decision frameworks. Multi-agent reinforcement learning (MARL) has become a powerful paradigm for modeling adaptive coordination in complex environments, enabling agents to learn cooperative behaviors through interaction [11]. Recent surveys further demonstrate the effectiveness of MARL for resource allocation optimization, especially in dynamic and decentralized industrial contexts aligned with Industry 4.0 developments [12].

Complementary work has explored distributed deep reinforcement learning, providing scalable toolboxes and algorithmic structures capable of supporting multi-player and self-play environments while addressing the challenges of distributed training [13].

Acceleration strategies have also been investigated in the context of gradient-tracking methods. In particular, the distributed heavy-ball approach proposed by Xin and Khan [14] generalizes first-order schemes by incorporating momentum, leading to improved convergence rates while preserving the decentralized structure of the algorithm.

Distributed Gradient Descent (DGD) is widely recognized as one of the foundational algorithms in decentralized optimization, enabling agents to iteratively combine local

gradient updates with neighbor averaging to cooperatively minimize a global objective. Despite its simplicity and scalability, DGD typically converges only to a neighborhood of the optimum when constant step sizes are employed, due to persistent consensus errors [15].

Distributed optimization has long been recognized as a fundamental tool for the control of networked systems. Nedić and Liu [16] show that decentralized optimization strategies enable scalable coordination while preserving robustness to communication constraints, making them particularly suitable for large interconnected infrastructures.

Gradient tracking methods have emerged as an effective strategy to eliminate the steady-state bias typically observed in distributed gradient descent. The DIGing algorithm, introduced by Nedić et al. [17], achieves exact convergence under standard convexity assumptions by dynamically tracking the global gradient across the network. Other variants were also proposed [18].

The EXTRA algorithm proposed by Shi et al. [19] addresses the consensus error inherent to distributed gradient descent by introducing a correction mechanism that enables exact convergence without diminishing step sizes.

Push–pull gradient methods further extend distributed optimization to directed communication graphs, allowing simultaneous information propagation and gradient tracking while relaxing symmetry requirements in network topology [20].

Asynchronous distributed optimization has also received significant attention, particularly in large-scale systems where strict synchronization is impractical. Recent advances demonstrate that convergence can be preserved under bounded delays, albeit often at the cost of slower convergence rates [21].

While distributed optimization algorithms have been extensively studied from a methodological perspective, their relevance is particularly evident in large-scale logistics and supply chain systems. Modern logistics infrastructures typically involve geographically distributed decision units, including warehouses, transportation hubs, service facilities, and distribution centers, which must coordinate operational decisions while operating with partially decentralized information.

In freight transportation and logistics network planning, coordination across multiple actors is essential to ensure efficient resource utilization and system-wide performance. As highlighted by Crainic and Laporte [22], freight transportation systems naturally exhibit distributed decision structures in which planning and operational decisions are made by geographically dispersed entities interacting through complex network infrastructures.

Similarly, supply chain management research has emphasized the importance of distributed decision support mechanisms capable of coordinating production, transportation, and inventory decisions across interconnected organizations. Classical studies in supply chain coordination highlight the challenges associated with decentralized decision-making and the need for structured optimization frameworks capable of aligning local and global objectives [23].

More recently, the increasing availability of operational data has stimulated the development of data-driven logistics optimization approaches. These methods leverage large-scale data streams and predictive analytics to support adaptive decision-making in logistics and supply chain environments [24]. At the same time, agent-based and decentralized coordination models have been explored to enable collaborative logistics planning across distributed infrastructures [25].

These operational characteristics reinforce the relevance of scalable distributed optimization methods capable of supporting coordinated decision processes under partial information and communication constraints. The framework proposed in this work builds upon these principles by integrating gradient-tracking optimization with delay-aware up-

dates, enabling robust decision support in networked environments where asynchronous information exchange is unavoidable.

Despite the significant progress in both distributed optimization and logistics decision support systems, several limitations remain evident in the literature. Many existing contributions emphasize theoretical convergence guarantees under idealized assumptions, often with limited consideration of operational decision support environments. In particular, numerous algorithms rely on synchronous communication or overlook the effects of delayed information exchange, which are common in real-world distributed infrastructures.

At the same time, several studies focus on application-specific implementations or abstract multi-agent formulations without explicitly addressing logistics-oriented systems, where decision variables are geographically distributed and information may be heterogeneous, partial, or asynchronously received. These gaps highlight the need for algorithmic frameworks capable of supporting scalable and robust decision-making under imperfect communication conditions while maintaining practical implementability in large-scale operational settings.

Table 1 summarizes the positioning of the proposed D^{3}O-GT framework with respect to representative distributed optimization methods discussed in the literature.

Table 1. Positioning of the proposed D^{3}O-GT framework relative to representative distributed optimization methods.

Method	Handles Communication Delays	Gradient Tracking	Data-Driven Orientation	Applicability to Logistics-Scale Systems
Distributed Gradient Descent (DGD) [15]	No	No	Limited	Moderate
DIGing [17]	No	Yes	Limited	Moderate
EXTRA/Exact First-Order Methods [19]	No	Partial	Limited	Moderate
Push–Pull Gradient Methods [20]	Partial	Yes	Limited	Moderate–High
Asynchronous Distributed Optimization (generic) [21]	Yes	Rare	Limited	Moderate
D^{3}O-GT (proposed)	Yes	Yes	Yes	High

Motivated by these limitations, this paper proposes D^{3}O-GT, a distributed optimization framework designed to operate under partial, heterogeneous, and asynchronous information. The objective is to bridge methodological rigor and operational relevance by combining structured optimization principles with implementation-aware algorithmic design.

3. Problem Formulation and System Model

For ease of reference, the main symbols used in the problem formulation and distributed algorithm are summarized in Table 2.

Table 2. Summary of the main symbols used in the distributed optimization model.

Symbol	Description
N	Number of agents
x_i	Local decision variable of agent i
$f_i(x_i)$	Local objective function
w_{ij}	Consensus weight between agents

Table 2. *Cont.*

Symbol	Description
N_k	Neighborhood of agent i
α_k	Step size
τ_k	Communication delay
z_i	Local copy of global decision variable

3.1. System Overview

We consider a networked decision environment composed of N agents interconnected through a communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the node set $\mathcal{V} = 1, \ldots, N$ represents autonomous decision units and the edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ encodes communication links.

This abstraction naturally captures a wide range of logistics-oriented systems, including distributed warehouses, coordinated transportation fleets, and networked service infrastructures.

Each agent i controls a local decision vector $x_i \in R^{d_i}$ and seeks to minimize a convex local objective function:

$$F_i(x_i; \theta_i), \tag{1}$$

where θ_i denotes locally available data, potentially incomplete or time-varying.

The global objective is defined as:

$$\min_{x_1, \cdots, x_N} F(x) = \sum_{i=1}^{N} f_i(x_i) \tag{2}$$

subject to coupling constraints of the form:

$$\sum_{i=1}^{N} A_i x_i = b, \tag{3}$$

which model shared resources, capacity limits, or coordination requirements.

3.2. Distributed Decision Setting

We focus on operational regimes characterized by three realistic properties:

- Partial information: agents observe only local cost structures and neighbor messages.
- Heterogeneous data: objective functions may differ across agents due to localized operational conditions.
- Asynchronous updates: communication delays or computational variability prevent perfectly synchronized iterations.

These assumptions align with emerging distributed optimization models designed for large-scale cyber–physical infrastructures [26].

3.3. Data-Driven Optimization Perspective

In modern logistics environments, objective functions are rarely static. Instead, they are often estimated from data streams reflecting demand fluctuations, travel times, or operational disruptions.

We therefore adopt a data-driven formulation in which parameters θ_i are periodically updated using observational data, leading to a sequence of optimization problems:

$$F_t(x) = \sum_{i=1}^{N} f_i\left(x_i; \theta_i^{(t)}\right). \tag{4}$$

This perspective enables adaptive decision support while preserving the structural benefits of convex optimization.

3.4. Consensus-Based Reformulation

To enable decentralized coordination, the problem can be equivalently expressed using local copies z_i of a shared variable z, enforcing agreement constraints:

$$z_i = z_j, \forall (i,j) \in \mathcal{E}. \tag{5}$$

The resulting formulation lends itself to consensus-driven solution strategies that balance communication efficiency with convergence guarantees [27].

3.5. Design Objectives

The framework proposed in this paper is guided by four primary objectives:

- Scalability: support large agent populations without centralized bottlenecks.
- Robustness: tolerate delayed or imperfect communication.
- Communication efficiency: limit message exchanges.
- Algorithmic structure: enable reproducible and implementation-oriented deployment.

Rather than targeting a single logistics application, the goal is to develop a reusable methodological foundation applicable to a broad class of distributed decision support problems.

3.6. Theoretical Standing

Under standard assumptions of convexity, Lipschitz continuity of gradients, and connectivity of the communication graph, distributed consensus-based algorithms are known to converge toward globally optimal solutions.

This theoretical grounding motivates the algorithmic architecture introduced in the next section, which emphasizes practical implementability while remaining consistent with established convergence principles [27].

3.7. Example: Distributed Logistics Coordination

To illustrate the operational interpretation of the proposed framework, consider a network of geographically distributed warehouses that collaboratively manage regional inventory levels. Each warehouse corresponds to an agent in the communication graph and maintains local information regarding demand forecasts, storage capacity, and transportation constraints.

The decision variable x_i may represent a vector of operational decisions such as shipment allocations, inventory adjustments, or resource usage. Each warehouse aims to minimize a local cost function reflecting storage costs, transportation delays, and service level penalties.

However, optimal operation requires coordination across warehouses because decisions made at one location affect the global logistics network. For example, shipment rerouting or shared transportation capacity introduces coupling constraints between agents.

In practice, information exchange across warehouses may experience communication delays due to network latency, asynchronous updates in distributed platforms, or delayed data aggregation from operational systems. Consequently, decision agents must update their optimization variables using partially outdated information received from neighboring nodes.

This setting naturally motivates the distributed consensus-based optimization formulation introduced in the following sections, where decentralized agents cooperatively solve a global optimization problem while exchanging information only with their neighbors.

4. Distributed Algorithmic Framework

The overall architecture of the proposed D^{3}O-GT framework is illustrated in Figure 1, highlighting the interaction between local optimization, data-driven updates, and consensus-based coordination.

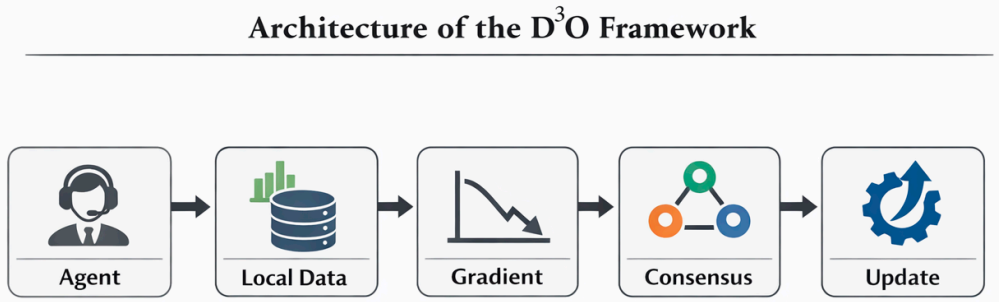


Figure 1. Conceptual architecture of the proposed D^{3}O-GT framework. Each agent performs local optimization using available data, exchanges information with neighboring agents through a consensus mechanism, and updates its decision variables iteratively, enabling scalable and adaptive coordination in distributed environments.

4.1. Design Rationale

Building upon the distributed formulation introduced in the previous section, we propose D^{3}O-GT (Decentralized Data-Driven Optimization), a scalable algorithmic framework designed to support coordinated decision-making under partial and asynchronous information.

The framework integrates three key principles:

- Local computation: agents optimize using locally available data.
- Neighbor-based coordination: consensus is achieved through limited communication.
- Adaptive updates: model parameters evolve with streaming data.

Unlike fully centralized solvers, D^{3}O-GT avoids single points of failure while maintaining strong convergence behavior consistent with distributed optimization theory [26–28].

4.2. Algorithmic Structure

At iteration k, each agent maintains a local estimate x_i^k and exchanges information only with its neighbors.

The update follows a gradient consensus structure:

$$x_i^{k+1} = \sum_{j \in \mathcal{N}_i} \omega_{ij} x_j^k - \alpha_k \nabla f_i\left(x_i^k; \theta_i^{(k)}\right) \tag{6}$$

where

- $\mathcal{N}_\rangle$ denotes the neighborhood of agent i;
- ω_{ij} are consensus weights satisfying stochasticity conditions;
- α_k is the step size.

This formulation balances optimization progress with network agreement.

4.3. Asynchronous Extension

In practical distributed systems, communication between agents is rarely perfectly synchronized. Message transmission may be delayed due to network congestion, computational variability, or asynchronous communication protocols. To model this behavior, we assume that information received by agent i from neighbor j may correspond to a delayed iteration index.

Specifically, let $\tau_{ij}(k)$ denote the communication delay affecting the information transmitted from agent j to agent i at iteration k. The delayed state received by agent i can therefore be represented as

$$x_j\left(k-\tau_{ij}(k)\right) \tag{7}$$

where $\tau_{ij}(k)$ is assumed to be bounded by a maximum delay τ_{max}. This bounded delay assumption is commonly adopted in asynchronous distributed optimization literature and ensures that outdated information does not indefinitely accumulate.

Under this formulation, each agent updates its decision variables using the most recently available information received from its neighbors, even if such information corresponds to slightly outdated iterations.

To accommodate communication delays, agents are allowed to update using the most recently received information:

$$x_i^{k+1} = \sum_{j\in\mathcal{N}_i} \omega_{ij} x_j^{\tau_{ij}(k)} - \alpha_{k\nabla} f_i\left(x_i^k\right) \tag{8}$$

where

$$\tau_{ij}(k) \le k \tag{9}$$

represents a possibly delayed iteration index.

Such relaxed synchronization significantly improves scalability in realistic distributed infrastructures.

4.4. Data-Driven Adaptation

Parameters are periodically refreshed:

$$\theta_i^{(k+1)} = \theta_i^{(k)} + \gamma_k g_i\left(\mathcal{D}_i^{(k)}\right) \tag{10}$$

where $g_i(\cdot)$ denotes a data-driven estimator derived from newly observed samples.

This enables continuous adaptation without interrupting the optimization process.

To operationalize this distributed formulation, we introduce D^{3}O-GT (Decentralized Data-Driven Optimization), a structured algorithmic framework designed to enable scalable and adaptive coordination across networked decision agents operating under partial and asynchronous information (Algorithm 1).

Algorithm 1. D^{3}O-GT Framework (assume connected communication graph).

Initialize xi0 for all agents
Initialize local gradient tracker yi0

for k = 0,1,2,. . . **do**

  for each agent i (in parallel) **do**

    Receive most recent neighbor states xj(k − τij)

    Consensus step:
    xī = Σj wij xj(k − τij)

    Compute local gradient:
    gi = ∇fi(xi(k))

Algorithm 1. *Cont.*

Gradient tracking update:
yi(k + 1) = Σj wij yj(k − τij) + gi − gi(k − 1)

Local update:
xi(k + 1) = xī − αk yi(k + 1)
If new data is available:
update parameters θi

end **for**

end **for**

4.5. Computational Complexity

For each iteration, the computational cost per agent scales as:

$$\mathcal{O}(d_i) \tag{11}$$

while communication overhead grows with the node degree.

Under sparse connectivity—typical in logistics-inspired networks—the framework remains highly scalable.

4.6. Convergence Discussion

Under standard assumptions of convexity, bounded gradients, and connected communication graphs, gradient consensus schemes are known to converge toward globally optimal solutions.

Asynchronous variants typically exhibit slightly slower convergence but offer substantial gains in robustness and deployability—a trade-off often desirable in operational environments.

Rather than deriving new theoretical guarantees, this work emphasizes implementation-oriented algorithmic structure grounded in well-established convergence results.

5. Numerical Experiments

This section evaluates the convergence behavior, robustness, and scalability of the proposed Distributed Delayed Optimization with Gradient Tracking (D^3O-GT) algorithm. The experiments are designed to assess both synchronous and asynchronous communication regimes and to compare the proposed method against a classical Distributed Gradient Descent (DGD) baseline.

5.1. Experimental Setup

We consider a distributed consensus optimization problem of the form

$$\min_{x \in R^d} \sum 12pt_{i=1}^{N} f_i(x), \tag{12}$$

where each agent i holds a strongly convex quadratic cost function

$$f_i(x) = \frac{1}{2} x^\top Q_i x + c_i^\top x, \tag{13}$$

with $Qi \succ 0$. The matrices Q_i are generated to ensure moderate conditioning, while the linear terms c_i are sampled from a zero-mean Gaussian distribution to introduce heterogeneity across agents.

The global optimum is computed analytically as

$$x^\star = -\left(\sum_{i=1}^N Q_i\right)^{-1}\left(\sum_{i=1}^N c_i\right), \tag{14}$$

allowing an exact evaluation of the optimality gap.

Network Topology. Agents communicate over an Erdős–Rényi random graph with connection probability $p = 0.2$. Mixing weights are constructed using the Metropolis rule, ensuring symmetry and double stochasticity—conditions known to support convergence in gradient-tracking methods. Unless stated otherwise, we have used the parameters listed in Table 3.

Table 3. Network topology and simulation parameters used in the numerical experiments. The communication graph follows an Erdős–Rényi model with Metropolis mixing weights, ensuring balanced information exchange across agents. The selected parameters reflect a moderately connected distributed environment and were chosen to provide stable convergence under both synchronous and asynchronous communication regimes.

Parameter	Value
Dimension d	5
Agents N	50
Iterations $K_{\max}$	1200
Sync step size	0.1
Async step size	0.02
Max delay	3 iterations
Delay probability	0.3

The smaller asynchronous step size follows standard stability requirements for delayed gradient methods.

Performance Metrics. We track three complementary metrics:

- Relative optimality gap

$$\frac{\left|F\left(x^k\right) - F(x^\star)\right|}{\left|F(x^\star)\right|} \tag{15}$$

- Consensus error

$$\sqrt{\frac{1}{N}\sum_{i=1}^N \| x_i^k - \bar{x}^k \|^2} \tag{16}$$

- Communication cost, measured as the cumulative number of transmitted messages.

Together, these metrics provide insight into convergence speed, agreement quality, and distributed overhead.

5.2. Convergence Analysis

Figure 2 illustrates the evolution of the relative optimality gap. The proposed D^{3}O-GT algorithm demonstrates rapid and stable convergence under synchronous communication, reaching machine precision accuracy with a final gap of 3.1×10^{-16}.

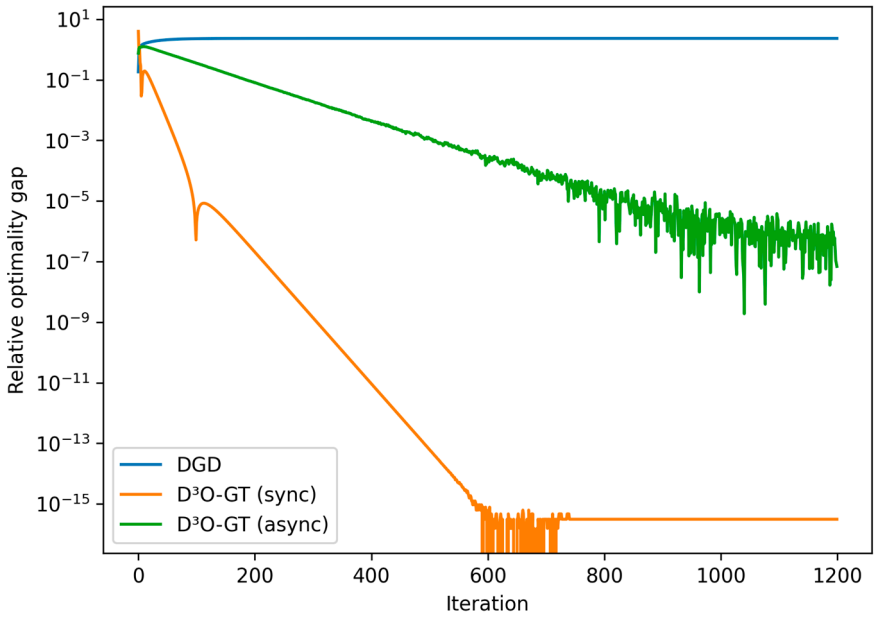


Figure 2. Relative optimality gap (log-scale) versus iteration for DGD, D^{3}O-GT (sync), and D^{3}O-GT (async). The synchronous variant converges to machine precision, while the asynchronous method achieves a small residual error due to delayed updates.

When delays are introduced, D^{3}O-GT remains stable and converges to a small neighborhood of the optimum (6.9×10^{-8}). This behavior is consistent with theoretical expectations, as stale information slightly degrades gradient accuracy but does not prevent convergence.

In contrast, DGD stagnates with a large residual error (final gap 2.32), highlighting the importance of gradient tracking for eliminating steady-state bias in heterogeneous distributed problems.

Figure 3 reports the consensus error across agents. The synchronous D^{3}O-GT achieves near-perfect agreement (1.5×10^{-15}), confirming that gradient tracking successfully aligns local descent directions.

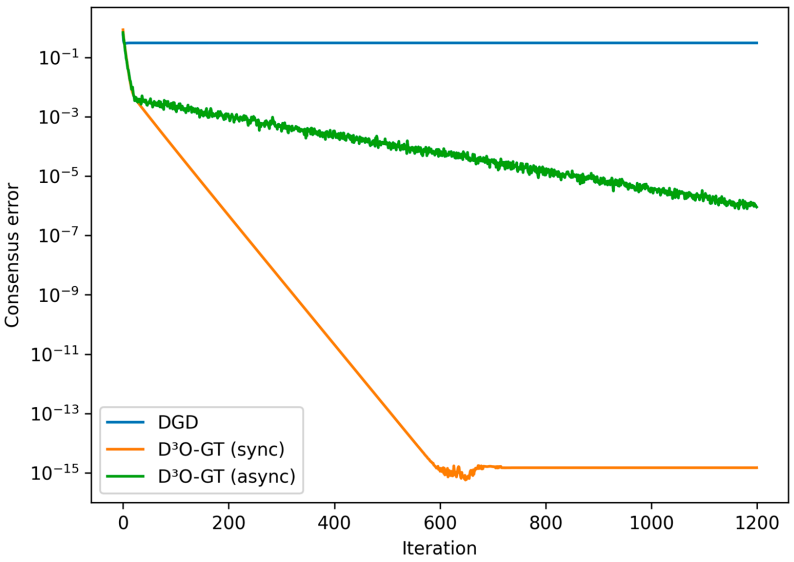


Figure 3. Consensus error (log-scale) versus iteration. D^{3}O-GT (sync) achieves near-perfect agreement, while asynchronous updates converge to a small disagreement floor. DGD fails to enforce consensus under the same step size.

Under delayed communication, consensus is slightly degraded but remains extremely small (9.0×10^{-7}), demonstrating robustness to bounded staleness.

DGD exhibits persistent disagreement (3.1×10^{-1}), further emphasizing that simple averaging is insufficient to ensure both agreement and optimality.

5.3. Scalability Analysis

To evaluate scalability, we measure the number of iterations required to reduce the initial optimality gap by a factor of 10^{-3} as the number of agents increases.

Both Figure 4 and Table 4 indicate that D^{3}O-GT scales favorably in iteration complexity. The required iterations remain within the range of approximately 400–600 despite a more than sevenfold increase in network size. Interestingly, the iteration count does not grow monotonically. Because the edge probability is fixed, larger networks exhibit higher expected node degrees, improving information mixing and partially offsetting the increased system size.

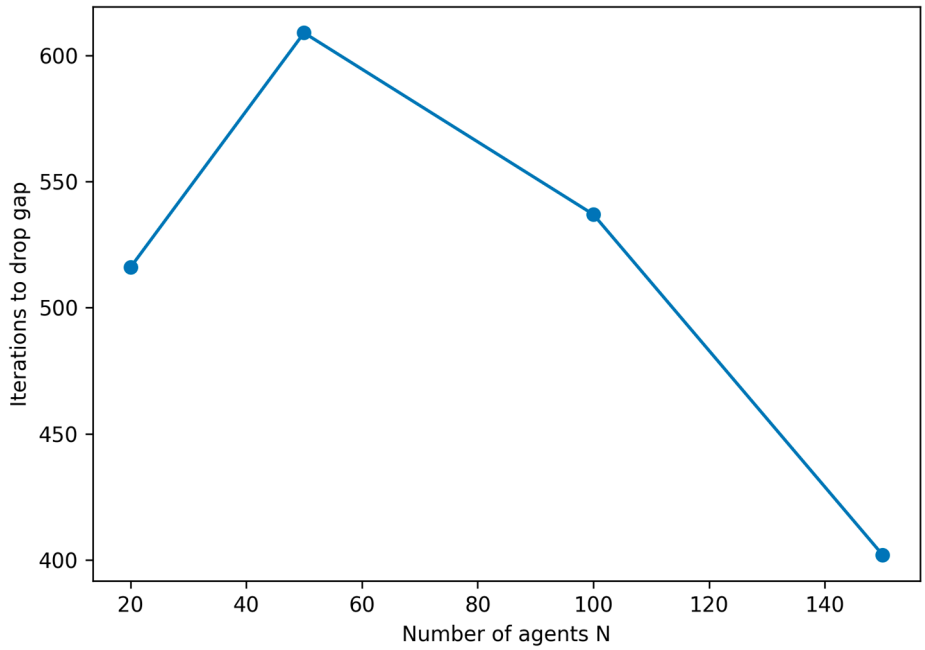


Figure 4. Iterations required to reduce the initial optimality gap by 10−3 for increasing network sizes.

Table 4. Scalability Summary.

Agents	Iterations	Final Gap	Final Consensus	Messages
20	516	2.64×10^{-7}	5.48×10^{-6}	88,800
50	609	3.51×10^{-7}	8.70×10^{-7}	619,200
100	537	1.15×10^{-7}	2.61×10^{-7}	2,371,200
150	402	6.87×10^{-9}	1.12×10^{-7}	5,503,200

The results suggest that the proposed method maintains stable convergence characteristics even in moderately large distributed settings.

5.4. Robustness to Communication Delays

To further evaluate the robustness of the proposed framework, we analyze the sensitivity of D^{3}O-GT to increasing levels of communication delays. Specifically, we vary the delay probability while keeping all other parameters fixed, including the network topology, step size, and problem instance. Each configuration is repeated across multiple random seeds to ensure statistical consistency.

Table 5 summarizes the results in terms of final relative optimality gap, consensus error, and the number of iterations required to reduce the initial gap by a factor of 10^{-3}.

Table 5. Sensitivity of D^{3}O-GT to communication delay probability (mean $\pm$ standard deviation across five runs).

Delay Probability	Final Relative Gap	Final Consensus Error	Iterations to Gap $\leq 10^{-3}$
0.0	$9.53 \times 10^{-10} \pm 6.9 \times 10^{-11}$	$2.51 \times 10^{-9} \pm 1.9 \times 10^{-10}$	388 ± 28
0.1	$2.05 \times 10^{-8} \pm 1.8 \times 10^{-8}$	$7.82 \times 10^{-8} \pm 2.2 \times 10^{-8}$	440 ± 31
0.3	$4.09 \times 10^{-7} \pm 2.6 \times 10^{-7}$	$7.58 \times 10^{-7} \pm 2.3 \times 10^{-7}$	552 ± 34
0.5	$1.74 \times 10^{-6} \pm 1.0 \times 10^{-6}$	$3.30 \times 10^{-6} \pm 6.8 \times 10^{-7}$	665 ± 44

As expected, higher delay probabilities lead to slower convergence and moderately larger residual errors. Nevertheless, D^{3}O-GT maintains stable behavior across all tested regimes, with the optimality gap remaining well below 10^{-5} even under substantial communication delays. The gradual degradation observed in both convergence speed and consensus accuracy reflects a predictable and well-behaved response to increasing asynchrony rather than algorithmic instability.

Figure 5 illustrates the sensitivity of the proposed framework to increasing communication delays. As the delay probability increases, convergence becomes slower and the residual optimality gap slightly increases. However, the algorithm remains stable across all tested regimes, demonstrating robustness to asynchronous information exchange.

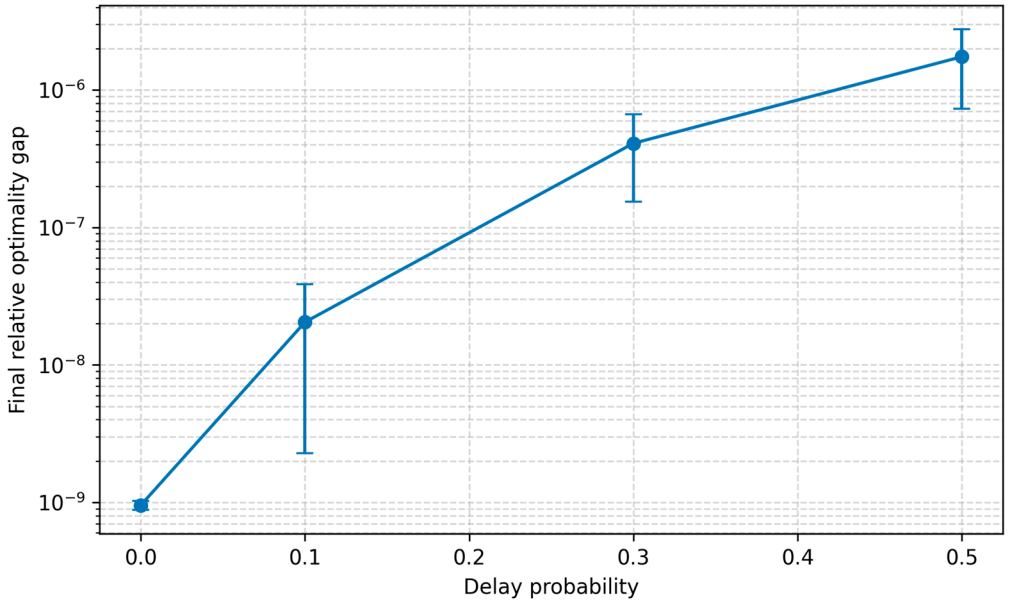


Figure 5. Impact of communication delay probability on final relative optimality gap. Error bars represent the standard deviation across five runs.

These findings confirm that the proposed framework is inherently robust to imperfect communication and can sustain high-quality solutions in distributed environments where delays are unavoidable. This robustness is particularly relevant for logistics-scale systems, where latency and asynchronous information exchange are intrinsic operational characteristics.

5.5. Key Takeaways

The experiments reveal three main findings:

- Gradient tracking is critical: D^{3}O-GT dramatically outperforms DGD, eliminating steady-state error.

- Robust to delays: asynchronous updates slow convergence but preserve stability.
- Scales well: iteration complexity remains controlled as the network grows.

Overall, the results validate D^{3}O-GT as a reliable framework for distributed optimization under realistic communication constraints.

These results collectively indicate that D^{3}O-GT provides a practical balance between convergence speed, robustness, and communication efficiency, making it well suited for large-scale distributed decision-making environments.

6. Discussion

The numerical results provide several important insights into the behavior of the proposed D^{3}O-GT framework under realistic distributed conditions. From a methodological standpoint, D^{3}O-GT can be interpreted as an implementation-oriented extension of gradient-tracking paradigms toward communication-imperfect environments.

First, the experiments clearly demonstrate the critical role of gradient tracking in eliminating the steady-state bias typically observed in classical distributed gradient methods. While Distributed Gradient Descent (DGD) exhibits initial progress, it ultimately stagnates due to the mismatch between local and global descent directions—a well-documented limitation in heterogeneous environments. In contrast, D^{3}O-GT maintains accurate gradient estimates across the network, enabling convergence to machine precision in the synchronous setting.

Second, the asynchronous experiments highlight the robustness of the proposed method to delayed information. Although stale updates introduce additional perturbations into the descent dynamics, convergence is preserved and the algorithm stabilizes near the global optimum. This behavior aligns with theoretical expectations for bounded delay systems, where moderate staleness primarily affects convergence speed rather than algorithmic stability.

Another noteworthy observation is the favorable scalability profile. Despite increasing the number of agents by more than one order of magnitude, the number of iterations required to substantially reduce the optimality gap remains relatively stable. This suggests that improved network mixing—resulting from higher expected node degrees in Erdős–Rényi graphs—can partially compensate for the increased dimensionality of the distributed system.

6.1. Practical Implications for Logistics Systems

The structural characteristics of modern logistics infrastructures—including distributed warehouses, networked transportation systems, and real-time coordination platforms—naturally align with decentralized optimization paradigms. In such environments, centralized decision-making is often impractical due to latency constraints, data locality, and scalability requirements.

The proposed D^{3}O-GT framework directly addresses these challenges by enabling cooperative optimization while tolerating delayed or partially outdated information. This capability is particularly relevant for large-scale logistics operations, where communication is inherently asynchronous and global synchronization is costly or infeasible.

From an operational perspective, the robustness of D^{3}O-GT to bounded delays suggests that high-quality decisions can be maintained even under imperfect network conditions. This opens the door to deployment in edge-enabled logistics platforms, distributed resource allocation, fleet coordination, and adaptive supply chain management.

Moreover, the favorable scalability observed in the experiments indicates that the framework can support growing operational networks without requiring fundamental redesign, making it a promising candidate for next-generation data-driven logistics sys-

tems. Rather than relying on idealized communication assumptions, D^{3}O-GT reflects the operational realities of contemporary logistics infrastructures, where delays and partial information are the norm rather than the exception.

6.2. Limitations and Future Research

Despite the promising results, several limitations should be acknowledged.

First, the experiments focus on strongly convex quadratic objectives, which provide a controlled environment for evaluating convergence dynamics. Extending the analysis to non-convex objectives—particularly those arising in distributed learning—represents an important direction for future work.

Second, communication delays were modeled as bounded and randomly distributed. Real-world networks may exhibit bursty latency patterns, packet loss, or time-varying connectivity. Investigating the behavior of D^{3}O-GT under such adverse conditions would further strengthen its practical relevance.

Third, while iteration complexity scales favorably, communication overhead naturally increases with network size. Future research could explore compression, event-triggered communication, or sparsification strategies to improve communication efficiency without sacrificing convergence guarantees.

Finally, a formal convergence analysis for delayed gradient-tracking dynamics remains an open theoretical challenge that merits deeper investigation.

Overall, the results indicate that D^{3}O-GT provides a robust and scalable foundation for distributed optimization under imperfect communication, effectively bridging the gap between theoretically grounded gradient-tracking methods and the operational requirements of practical large-scale deployments.

From a broader perspective, the proposed framework contributes to bridging the gap between theoretically grounded distributed optimization algorithms and the operational requirements of real-world networked infrastructures. In particular, logistics systems represent a compelling application domain where decentralized decision-making, asynchronous communication, and heterogeneous information sources naturally arise.

Although the experimental comparison focuses on Distributed Gradient Descent (DGD) as a baseline, it is worth noting that several advanced gradient-tracking algorithms such as DIGing and EXTRA have been proposed in the literature. These methods typically achieve exact convergence under synchronous communication assumptions. However, their behavior under delayed communication scenarios has received comparatively less attention. The proposed D^{3}O-GT framework explicitly targets such communication-imperfect environments, emphasizing robustness to asynchronous updates and delayed information exchange.

Overall, the results indicate that D^{3}O-GT provides a robust and scalable foundation for distributed optimization under imperfect communication. From a broader perspective, the framework contributes to bridging the gap between theoretically grounded gradient-tracking algorithms and the operational requirements of real-world distributed infrastructures, particularly in logistics-oriented systems where asynchronous communication and heterogeneous information are inherent.

7. Conclusions

This paper introduced D^{3}O-GT, a distributed optimization framework designed to operate effectively under realistic communication constraints, including delayed information exchange. By integrating gradient tracking with a delay-aware update mechanism, the proposed approach addresses fundamental limitations of classical distributed gradient methods, particularly their susceptibility to steady-state bias in heterogeneous environments.

The numerical experiments demonstrated that D^{3}O-GT achieves fast and stable convergence in synchronous settings, reaching machine-level accuracy while maintaining near-perfect consensus across agents. Under asynchronous communication, the method remains robust, converging to a small neighborhood of the global optimum despite the presence of bounded delays. These results confirm that delayed information primarily impacts convergence speed rather than algorithmic stability when properly accounted for in the update dynamics.

Scalability analyses further revealed that the proposed method maintains favorable iteration complexity as the number of agents increases. This behavior suggests that D^{3}O-GT is well suited for moderately large distributed systems where synchronization costs become prohibitive.

Beyond its empirical performance, the framework contributes to narrowing the gap between theoretically grounded gradient-tracking algorithms and the operational requirements of modern distributed infrastructures. As large-scale networked decision systems continue to expand—spanning edge computing, cyber-physical coordination, and data-driven logistics—methods capable of balancing robustness, accuracy, and communication efficiency are becoming increasingly essential.

Future work will focus on extending the theoretical analysis of delayed gradient-tracking dynamics, investigating non-convex objective functions, and incorporating communication-efficient strategies such as compression or event-triggered updates. Exploring adaptive step size policies and resilience to adversarial network conditions also represents a promising research direction.

Overall, D^{3}O-GT provides a practical and scalable foundation for distributed optimization under imperfect communication, offering a compelling alternative to conventional approaches in next-generation decentralized systems.

We believe that this work represents a meaningful step toward communication-aware distributed optimization, where algorithmic design explicitly reflects the constraints of real-world networked environments. We hope this work encourages the development of communication-aware optimization methods explicitly designed for real-world distributed infrastructures.

Funding: This research received no external funding.

Data Availability Statement: The data supporting the findings of this study are available upon reasonable request from the corresponding author. Sharing the data via direct communication ensures adequate support for replication or verification efforts and allows for appropriate guidance in its use and interpretation.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Chen, W.; Men, Y.; Fuster, N.; Osorio, C.; Juan, A.A. Artificial Intelligence in Logistics Optimization with Sustainable Criteria: A Review. *Sustainability* **2024**, *16*, 9145. [CrossRef]
2. Weng, K.; Ren, Y.; Zhao, L.; Gao, S. Distributed Optimization of Multi-Agent Systems with Time-Invariant and Time-Varying Cost Functions. *Eur. J. Control* **2025**, *83*, 101222. [CrossRef]
3. An, Y.; Wang, A.; Zhang, X.; Xiao, F. Distributed Constrained Optimization for Multi-Agent Networks with Communication Delays under Time-Varying Topologies. *Syst. Control Lett.* **2024**, *185*, 105733. [CrossRef]
4. Lei, T.; Chintam, P.; Luo, C.; Liu, L.; Jan, G.E. A Convex Optimization Approach to Multi-Robot Task Allocation and Path Planning. *Sensors* **2023**, *23*, 5103. [CrossRef]
5. Rosenberger, J.; Urlaub, M.; Rauterberg, F.; Lutz, T.; Selig, A.; Bühren, M.; Schramm, D. Deep Reinforcement Learning Multi-Agent System for Resource Allocation in Industrial Internet of Things. *Sensors* **2022**, *22*, 4099. [CrossRef]
6. Doostmohammadian, M.; Aghasi, A.; Pirani, M.; Nekouei, E.; Zarrabi, H.; Keypour, R.; Rikos, A.I.; Johansson, K.H. Survey of Distributed Algorithms for Resource Allocation over Multi-Agent Systems. *Annu. Rev. Control* **2025**, *59*, 100983. [CrossRef]

7. Wang, Y.; Lin, H.; Lam, J.; Kwok, K.-W. Differentially Private Consensus and Distributed Optimization in Multi-Agent Systems: A Review. *Neurocomputing* **2024**, *597*, 127986. [CrossRef]
8. Li, H.; Li, J.; Ran, L.; Zheng, L.; Huang, T. A Survey of Distributed Algorithms for Aggregative Games. *IEEECAA J. Autom. Sin.* **2025**, *12*, 859–871. [CrossRef]
9. Bai, N.; Duan, Z.; Wang, Q. Distributed Optimal Consensus of Multi-Agent Systems: A Randomized Parallel Approach. *Automatica* **2024**, *159*, 111339. [CrossRef]
10. Almen, A.; Dentcheva, D. On Risk Evaluation and Control of Distributed Multi-Agent Systems. *J. Optim. Theory Appl.* **2024**, *203*, 2025–2054. [CrossRef]
11. Zhu, C.; Dastani, M.; Wang, S. A Survey of Multi-Agent Deep Reinforcement Learning with Communication. *Auton. Agents Multi-Agent Syst.* **2024**, *38*, 4. [CrossRef]
12. Hady, M.A.; Hu, S.; Pratama, M.; Cao, Z.; Kowalczyk, R. Multi-Agent Reinforcement Learning for Resources Allocation Optimization: A Survey. *Artif. Intell. Rev.* **2025**, *58*, 354. [CrossRef]
13. Yin, Q.; Yu, T.; Shen, S.; Yang, J.; Zhao, M.; Ni, W.; Huang, K.; Liang, B.; Wang, L. Distributed Deep Reinforcement Learning: A Survey and a Multi-Player Multi-Agent Learning Toolbox. *Mach. Intell. Res.* **2024**, *21*, 411–430. [CrossRef]
14. Xin, R.; Khan, U.A. Distributed Heavy-Ball: A Generalization and Acceleration of First-Order Methods With Gradient Tracking. *IEEE Trans. Autom. Control* **2020**, *65*, 2627–2633. [CrossRef]
15. Xi, C.; Khan, U.A. Directed-Distributed Gradient Descent. In Proceedings of the 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton), Monticello, IL, USA, 29 September–2 October 2015; pp. 1022–1026.
16. Nedić, A.; Liu, J. Distributed Optimization for Control. *Annu. Rev. Control Robot. Auton. Syst.* **2018**, *1*, 77–103. [CrossRef]
17. Nedić, A.; Olshevsky, A.; Shi, W. Achieving Geometric Convergence for Distributed Optimization Over Time-Varying Graphs. *SIAM J. Optim.* **2017**, *27*, 2597–2633. [CrossRef]
18. Li, H.; Zheng, L.; Wang, Z.; Yan, Y.; Feng, L.; Guo, J. S-DIGing: A Stochastic Gradient Tracking Algorithm for Distributed Optimization. *IEEE Trans. Emerg. Top. Comput. Intell.* **2022**, *6*, 53–65. [CrossRef]
19. Shi, W.; Ling, Q.; Wu, G.; Yin, W. EXTRA: An Exact First-Order Algorithm for Decentralized Consensus Optimization. *SIAM J. Optim.* **2015**, *25*, 944–966. [CrossRef]
20. Pu, S.; Shi, W.; Xu, J.; Nedic, A. Push–Pull Gradient Methods for Distributed Optimization in Networks. *IEEE Trans. Autom. Control* **2021**, *66*, 1–16. [CrossRef]
21. Wu, X.; Liu, C.; Magnússon, S.; Johansson, M. Asynchronous Distributed Optimization With Delay-Free Parameters. *IEEE Trans. Autom. Control* **2026**, *71*, 259–274. [CrossRef]
22. Crainic, T.G.; Laporte, G. Planning Models for Freight Transportation. *Eur. J. Oper. Res.* **1997**, *97*, 409–438. [CrossRef]
23. Simchi-Levi, D.; Kaminsky, P.; Simchi-Levi, E. *Designing and Managing the Supply Chain: Concepts, Strategies, and Case Studies*; McGraw Hill Professional: New York, NY, USA, 2003; ISBN 978-0-07-249256-9.
24. Waller, M.A.; Fawcett, S.E. Data Science, Predictive Analytics, and Big Data: A Revolution That Will Transform Supply Chain Design and Management. *J. Bus. Logist.* **2013**, *34*, 77–84. [CrossRef]
25. López-Vargas, J.C.; Meisel, J.D.; Cárdenas-Aguirre, D.M.; Medina, P. Coordination Mechanisms Applied to Logistical Systems for Local Disaster Preparedness: A Latin American Case. *J. Humanit. Logist. Supply Chain Manag.* **2024**, *15*, 87–106. [CrossRef]
26. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, UK; New York, NY, USA, 2004; ISBN 978-0-521-83378-3.
27. Boyd, S. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Found. Trends® Mach. Learn.* **2010**, *3*, 1–122. [CrossRef]
28. Parikh, N.; Boyd, S. Block Splitting for Distributed Optimization. *Math. Program. Comput.* **2014**, *6*, 77–102. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.