arXiv:2610.25017v1 [cs.LG] 18 Oct 2026

Decentralized Online Exp-Concave Optimization

Hao Qiu[1*] Mengxiao Zhang[2*] Yuanyu Wan[3,4†]
qiuhaosai@gmail.com mengxiao-zhang@uiowa.edu wanyy@zju.edu.cn

[1]National University of Singapore, Singapore

[2]Department of Business Analytics, Tippie College of Business, University of Iowa, USA

[3]School of Software Technology, Zhejiang University, Ningbo, China

[4]State Key Laboratory of Blockchain and Data Security, Zhejiang University, Hangzhou, China

Abstract

Decentralized Online Convex Optimization (D-OCO) is a fundamental framework in which a set of local agents aims to minimize a sequence of global loss functions using only local computations and communications. While recent work has established near-optimal regret guarantees for convex and strongly convex losses, the case of exp-concave losses remains an open challenge: specifically, whether logarithmic regret is achievable. In this paper, we answer this question affirmatively by proposing Decentralized Online Newton Step (D-ONS), which adapts the idea of blocking and accelerated gossip to the Online Newton Step algorithm in a non-trivial manner. Concretely, our algorithm achieves a regret bound of $\mathcal{O}\left(Nd(1-\sigma_2)^{-1/2}\log^2(T)\right)$ where N is the number of agents, d is the dimension, T is the horizon, and $1-\sigma_2$ represents the spectral gap of the network. To our knowledge, this is the first logarithmic regret bound for decentralized online exp-concave optimization. We further establish a nearly matching lower bound of $\Omega\left(Nd(1-\sigma_2)^{-1/2}\log(T/N)\right)$ showing the near optimality of our upper bound. Finally, we extend our algorithm to decentralized online linear regression in which the feasible domain is unbounded and features and labels are locally observed. By gossiping covariance and label-feature statistics, we also achieve logarithmic regret.

1 Introduction

Decentralized online convex optimization (D-OCO) [Yan et al., 2013, Hosseini et al., 2013, Wan et al., 2024] provides a powerful framework for sequential decision making in distributed systems, where multiple agents *collaboratively* optimize a global objective while only observing *local* information and communicating through a network. In this setting, agents repeatedly choose decisions, incur global losses but only observe local feedback, and exchange information with neighboring agents in order to compete with the best fixed decision for the global losses in hindsight. This paradigm naturally captures a variety of applications including multi-robot systems [Shorinwa et al., 2024], sensor networks [Li et al., 2002], online packet routing [Awerbuch and Kleinberg, 2004], where centralized coordination is either infeasible due to communication constraints or undesirable due to privacy and robustness considerations.

For general convex losses, early studies by Yan et al. [2013], Hosseini et al. [2013] established decentralized online algorithms with regret bounds $\mathcal{O}(N^{5/4}(1-\sigma_2)^{-1/2}\sqrt{T})$ and $\mathcal{O}(N^{3/2}(1-\sigma_2)^{-1}\log(T))$ for convex and strongly convex losses, respectively. Here, N denotes the number of agents, T is the horizon, and $1-\sigma_2$ represents the spectral gap of the communication matrix. More recently, Wan et al. [2024, 2025] tightened these guarantees using accelerated gossip and block-based scheduling, achieving $\mathcal{O}(N(1-\sigma_2)^{-1/4}\sqrt{T\log(N)})$ regret for convex losses and $\mathcal{O}(N(1-\sigma_2)^{-1/2}\log(N)\log(T))$ for strongly convex losses, alongside nearly matching lower bounds.

*Equal contribution.
†Corresponding author.

Despite these advances, the landscape for decentralized OCO with exp-concave losses remains largely unexplored. This class of functions captures a wide range of applications, including online portfolio selection, logistic regression, and online linear regression on bounded domains. Similar to the strongly convex case, exp-concavity allows for significantly faster convergence rates in centralized settings than general convexity. Specifically, the seminal Online Newton Step (ONS) algorithm [Hazan et al., 2007] achieves $\mathcal{O}(d\log(T))$ regret for d-dimensional exp-concave losses. However, the possibility of achieving such fast rates with the decentralized communication remains an open challenge. This leads to the following central question of this work:

Can decentralized agents exploit exp-concavity to achieve logarithmic regret while communicating only through local gossip?

Our contributions. In this paper, we answer the above question affirmatively. More concretely:

- We propose Decentralized Online Newton Step (D-ONS), a decentralized variant of Online Newton Step for exp-concave losses under a gossip communication protocol. Specifically, we obtain a regret bound of $\mathcal{O}(Nd(1-\sigma_2)^{-1/2}\log^2(T))$. To the best of our knowledge, this is the first algorithm to achieve logarithmic regret for exp-concave losses under D-OCO, answering an open question proposed in Wan [2025]. Technically, our algorithm combines three ingredients: a blocking scheme that allows communication to be amortized over time, an accelerated gossip protocol to obtain a sharp dependency on the network spectral gap, and a delayed ONS update adapted to the exp-concave structure.
- We further complement our upper bound with a nearly matching lower bound, demonstrating that the polynomial dependence on the number of agents, the dimension, and the network spectral gap is essentially unavoidable. In particular, we adapt the hard instance of online exp-concave optimization from Ordentlich and Cover [1998] to the decentralized setting. This establishes that no algorithm can achieve a regret better than $\Omega\left(Nd(1-\sigma_2)^{-1/2}\log(T/N)\right)$, showing that our upper bound is near-optimal up to logarithmic factors.
- Finally, we extend our approach to decentralized online linear regression, where the decision domain for each agent is *unbounded* and thus not captured by standard exp-concave analysis. Specifically, we propose a decentralized online ridge regression algorithm with clipped predictions that achieves a regret bound of $\mathcal{O}(Nd(1-\sigma_2)^{-1/2}\log^2(T))$, demonstrating that logarithmic regret remains attainable even in the absence of a compact feasible set in online linear regression.

1.1 Related Works

Decentralized online convex optimization. Decentralized online convex optimization (D-OCO) studies sequential decision making over a network, where multiple agents collaboratively minimize a global objective without relying on a central coordinator. At each round, agents observe only local loss information and communicate with neighboring agents through a graph, making consensus and gossip-based methods central algorithmic tools. Foundational work in decentralized optimization was developed primarily in offline settings, where a standard gossip step Xiao and Boyd [2004] was used to enable decentralized consensus and optimization [Nedic and Ozdaglar, 2009, Duchi et al., 2011]. The online counterpart was first systematically studied by Yan et al. [2013], who proposed a Decentralized Online Gradient Descent and established sublinear regret guarantees. In particular, their algorithm achieves a regret bound of $\mathcal{O}(N^{5/4}(1-\sigma_2)^{-1/2}\sqrt{T})$. More recently, Wan et al. [2024, 2025] leverage an accelerated gossip strategy [Liu and Morse, 2011, Ye et al., 2023] together with a blocking update mechanism to develop new D-OCO algorithms that achieve regret bounds of $\mathcal{O}(N(1-\sigma_2)^{-1/4}\sqrt{T\log(N)})$ for convex functions and $\mathcal{O}(N(1-\sigma_2)^{-1/2}\log(N)\log(T))$ for strongly convex functions. They further establish matching lower bounds of $\Omega(N(1-\sigma_2)^{-1/4}\sqrt{T})$ and $\Omega(N(1-\sigma_2)^{-1/2}\log(T/N))$ for the convex and strongly convex settings, respectively, demonstrating the near-optimality of their algorithms. See Yuan et al. [2024] for a comprehensive overview.

Online learning with curved losses. Curvature assumptions on the loss function are well known to yield faster regret rates in online convex optimization. For strongly convex losses with G-bounded gradients, Online Gradient Descent achieves $\mathcal{O}(G^2\log(T))$ regret, while for d-dimensional exp-concave losses, the Online Newton Step of Hazan et al. [2007] attains $\mathcal{O}(d\log(T))$ regret. Closely related second-order methods for online linear regression include the Vovk–Azoury–Warmuth forecaster [Vovk, 2001]

and online ridge regression [Hoerl and Kennard, 1970]. Beyond these foundations, research has expanded into several advanced perspectives, including scale-free guarantees [Mayo et al., 2022], non-stationarity [Zhang et al., 2025], adaptive guarantees [Bartlett et al., 2007, Luo et al., 2022], and robustness to delayed feedback [Qiu et al., 2025]. In contrast to these established paradigms, our work investigates exp-concave online learning in a decentralized setting, where multiple agents must collaboratively minimize the global objective over a communication network. For a more extensive literature review on online learning with curved losses, we refer the reader to several comprehensive surveys [Hazan, 2016, Orabona, 2026].

2 Preliminaries

Throughout this paper, we denote the set $\{1, 2, ..., m\}$ for some positive integer m by $[m]$ and let $\mathbf{1}$ be an all-one vector in an appropriate dimension. For a matrix $M \in \mathbb{R}^{m_1 \times m_2}$, denote its (i, j)-th entry by M_{ij}.

In this work, we study decentralized online convex optimization (D-OCO). The framework consists of a group of N agents interacting with a dynamic environment and communicating over a network modeled by a connected, undirected graph $\mathcal{G} \triangleq (V, E)$. The vertex set $V = [N]$ represents the agents, while the edge set E specifies the available communication links. For each agent $i \in V$, we denote its neighborhood by $\mathcal{N}_i \triangleq \{j \in V : (i, j) \in E\} \cup \{i\}$. Throughout the paper, we use V and $[N]$ interchangeably. Each agent $i \in [N]$ encounters a sequence of local loss functions $\{f_{t,i}\}_{t=1}^{T}$, where each $f_{t,i} : \mathcal{X} \subset \mathbb{R}^d \to \mathbb{R}$ is chosen arbitrarily by an adversary and is unknown to the agents. Below, we list several standard assumptions [Yan et al., 2013, Hosseini et al., 2013, Wan et al., 2024] and will explicitly specify which are required for each of our results.

Assumption 1. *For each round $t \in [T]$, the local loss function $f_{t,i}$ of each agent $i \in [N]$ is G-Lipschitz continuous over $\mathcal{X}$ with respect to the ℓ_2 norm, i.e., $|f_{t,i}(\mathbf{x}) - f_{t,i}(\mathbf{y})| \leq G\|\mathbf{x} - \mathbf{y}\|_2$ for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$.*

Assumption 2. *The feasible set $\mathcal{X}$ is compact and contains the origin ($\mathbf{0} \in \mathcal{X}$). Furthermore, its diameter is bounded by D, such that $\|\mathbf{x} - \mathbf{y}\|_2 \leq D$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$.*

In this paper, we mainly focus on *exp-concave losses* Hazan et al. [2007], a function class that is broader than the class of strongly convex functions and also enjoys faster convergence rate in classic online optimization.

Assumption 3. *For each round $t \in [T]$, the loss function $f_{t,i} : \mathcal{X} \mapsto \mathbb{R}$ of each agent $i \in [N]$ is α-exp-concave over $\mathcal{X}$. That is, the mapping $\mathbf{x} \mapsto \exp(-\alpha f_{t,i}(\mathbf{x}))$ is concave for all $\mathbf{x} \in \mathcal{X}$.*

The learning protocol of D-OCO involves an interaction between the agents and the environments over T rounds. Without loss of generality, we assume that $T \geq N$. At each round $t \in [T]$, each agent $i \in [N]$ chooses a decision $\mathbf{x}_i(t)$ from the feasible set $\mathcal{X}$. Simultaneously, the adversary selects the local loss functions $f_{t,i}$ for all $i \in [N]$. Each agent i then observes only the gradient of her *local loss* $\nabla f_{t,i}(\mathbf{x}_i(t))$. Subsequently, agents are permitted to exchange information with their neighbors in the graph $\mathcal{G}$ to update their internal states for the next round. The objective of each agent i is then to minimize the regret in terms of *the global function*, defined as:

$$\mathrm{Reg}_{T,i} \triangleq \max_{\mathbf{x} \in \mathcal{X}} \mathrm{Reg}_{T,i}(\mathbf{x}) \triangleq \max_{\mathbf{x} \in \mathcal{X}} \left(\sum_{t=1}^{T} \sum_{j=1}^{N} f_{t,j}(\mathbf{x}_i(t)) - \sum_{t=1}^{T} \sum_{j=1}^{N} f_{t,j}(\mathbf{x}) \right) . \tag{1}$$

It remains to specify how agents communicate over the network $\mathcal{G}$. Following Hosseini et al. [2013], Yan et al. [2013], Wan et al. [2024], we model information exchange through a gossip protocol induced by a communication matrix $P \in \mathbb{R}^{N \times N}$. Intuitively, P encodes how each agent aggregates messages from its neighbors: the entry P_{ij} represents the weight that agent i assigns to the information received from agent j. Since communication is restricted by the network structure, P is supported on $\mathcal{G}$, and its symmetry and double stochasticity ensure that repeated local averaging preserves the global mean. We formalize these requirements below.

Definition 1. *The communication matrix $P \in \mathbb{R}^{N \times N}$ is supported on the graph $\mathcal{G} = ([N], E)$, symmetric, and doubly stochastic: (i) $P_{ij} > 0$ only if $(i, j) \in E$ or $i = j$; (ii)$\sum_{j=1}^{N} P_{ij} = 1$ for all $i \in [N]$; (iii) $\sum_{i=1}^{N} P_{ij} = 1$ for all $j \in [N]$. Moreover, P is positive semidefinite, and its second largest singular value*

$\sigma_2(P)$ satisfies $\sigma_2(P) < 1$. When the dependence on P is clear, we write σ_2. Here, for a matrix P, we denote by $\sigma_1(P) \geq \sigma_2(P) \geq \cdots$ its singular values in nonincreasing order.

Given a collection of local vectors $\{\mathbf{u}_i\}_{i=1}^N$ (for example, local gradients or iterates), the goal of gossip is to approximately recover the network average $\overline{\mathbf{u}} = \frac{1}{N}\sum_{i=1}^N \mathbf{u}_i$ for every agent using only neighbor-to-neighbor communication. To this end, agents run L rounds of gossip. In each round k, agent i updates its local state according to

$$\mathbf{u}_i^{k+1} = (1+\theta)\sum_{j\in\mathcal{N}_i} P_{ij}\mathbf{u}_j^k - \theta\mathbf{u}_i^{k-1}, \tag{2}$$

for $k \geq 0$, where $\mathbf{u}_i^0 = \mathbf{u}_i^{-1} = \mathbf{u}_i$. When $\theta = 0$, this reduces to the standard gossip iteration [Xiao and Boyd, 2004]; when $\theta > 0$, it becomes the accelerated gossip scheme of Liu and Morse [2011], which improves the convergence rate to $\overline{\mathbf{u}}$.

The following lemma extended from Ye et al. [2023] shows that accelerated gossip allows each local vector to approximate the global average vector with an error of order $\mathcal{O}((TN)^{-2})$ after $L = \Theta((1-\sigma_2)^{-1/2}\log(NT))$ by picking θ appropriately. We include the proof in Section A for completeness.

Lemma 1. *Let P be the communication matrix defined in Definition 1. Let*

$$\theta \triangleq \frac{1}{1+\sqrt{1-\sigma_2^2(P)}}, \qquad L \triangleq \left\lceil \frac{\sqrt{2}\log(T^2\sqrt{14}N^{5/2})}{(\sqrt{2}-1)\sqrt{1-\sigma_2(P)}} \right\rceil. \tag{3}$$

Consider arbitrary vectors $\mathbf{u}_1, \ldots, \mathbf{u}_N \in \mathbb{R}^d$ satisfying $\|\mathbf{u}_i\|_2 \leq R, i \in [N]$. Let

$$U^0 = U^{-1} = \begin{bmatrix} \mathbf{u}_1^\top \\ \vdots \\ \mathbf{u}_N^\top \end{bmatrix}, \qquad \overline{\mathbf{u}} = \frac{1}{N}\sum_{i=1}^N \mathbf{u}_i^0, \qquad \overline{U} = \begin{bmatrix} \overline{\mathbf{u}}^\top \\ \vdots \\ \overline{\mathbf{u}}^\top \end{bmatrix}.$$

Suppose accelerated gossip is run for L iterations with update $U^{k+1} = (1+\theta)PU^k - \theta U^{k-1}$ for $k \geq 0$. Then, for every agent $i \in [N]$, $\|\mathbf{u}_i^L - \overline{\mathbf{u}}\|_2 \leq \frac{2R}{T^2N^2}$.

Other notations. Define $\Delta_d \triangleq \{x \in \mathbb{R}^d : \sum_{j=1}^d x_j = 1, x_j \geq 0, \forall j \in [d]\}$ to be the $(d-1)$-dimensional simplex. I_d denotes the $d \times d$ identity matrix. Let $A_\mathcal{G}$ be the adjacency matrix of $\mathcal{G}$. Let $D_\mathcal{G} = \mathrm{diag}(d_1, \ldots, d_N)$ be the diagonal degree matrix, and let $d_{\max} = \max_{i\in[N]} d_i$ denote the maximum degree of $\mathcal{G}$. The graph Laplacian of $\mathcal{G}$ is defined as $L_\mathcal{G} \triangleq D_\mathcal{G} - A_\mathcal{G}$. For a positive semidefinite matrix $M \in \mathbb{R}^{d\times d}$ and $x \in \mathbb{R}^d$, we denote $\|x\|_M = \sqrt{x^\top Mx}$ to be the Mahalanobis norm induced by M and, if M is positive definite, let $\|x\|_{M^{-1}} = \sqrt{x^\top M^{-1}x}$ be the dual norm.

3 Decentralized Online Exp-concave Optimization

In this section, we introduce our algorithm for decentralized online exp-concave optimization. We first provide an algorithm that achieves a logarithmic regret upper bound (Section 3.1), followed by a nearly-matching lower bound (Section 3.3) to show the optimality of our algorithm.

3.1 Upper Bound

In this subsection, we present our algorithm for decentralized online exp-concave optimization. Our algorithm builds upon the block-based communication schedule and accelerated gossip mechanism developed in recent decentralized OCO reductions for the convex and strongly convex settings [Wan et al., 2024, Wan, 2025]. However, the exp-concave case cannot be handled by a direct black-box reduction, as unlike the strongly convex setting where the curvature is characterized by a fixed global parameter, the curvature information in exp-concave optimization is decision-dependent and may incur larger consensus errors during communication. To address this challenge, we propose Decentralized Online Newton Step (D-ONS), which gossips both block gradients and decisions before performing a delayed ONS update.

Algorithm 1: Decentralized Online Newton Step for each agent $i \in [N]$

1 **Input:** gossip parameter θ, epoch length L, regularization parameter β, η, communication matrix P

2 **Initialization:** set $\mathbf{x}_i(1) = \hat{\mathbf{x}}_i(1) = \mathbf{0}$.

3 **for** $z = 1, \ldots, T/L$ **do**

4 Compute

$$\hat{\mathbf{x}}_i(z+1) = \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{s=1}^{z-2} \left(\langle \mathbf{g}_i(s), \mathbf{x} \rangle + \frac{\beta}{2L} \langle \mathbf{g}_i(s), \mathbf{x}_i(s) - \mathbf{x} \rangle^2 \right) + \frac{\eta}{2} \|\mathbf{x}\|_2^2. \tag{4}$$

5 Set $k_x = 0$ and $k_g = 0$ and $\mathbf{x}_i^0(z+1) = \mathbf{x}_i^{-1}(z+1) = \hat{\mathbf{x}}_i(z+1)$

6 **for** $t = (z-1)L + 1, \ldots, zL$ **do**

7 Play $\mathbf{x}_i(z)$ and query $\nabla f_{t,i}(\mathbf{x}_i(z))$

8 **if** $z \geq 2$ **then** update $\mathbf{g}_i^{k_g+1}(z-1) = (1+\theta)\sum_{j\in\mathcal{N}_i} P_{ij}\mathbf{g}_j^{k_g}(z-1) - \theta\mathbf{g}_i^{k_g-1}(z-1)$;

9 Update $\mathbf{x}_i^{k_x+1}(z+1) = (1+\theta)\sum_{j\in\mathcal{N}_i} P_{ij}\mathbf{x}_j^{k_x}(z+1) - \theta\mathbf{x}_i^{k_x-1}(z+1)$

10 Set $k_x \leftarrow k_x + 1$ and $k_g \leftarrow k_g + 1$

11 Update $\mathbf{x}_i(z+1) = \arg\min_{\mathbf{x}\in\mathcal{X}} \|\mathbf{x}_i^L(z+1) - \mathbf{x}\|_2$

12 **if** $z \geq 2$ **then** set $\mathbf{g}_i(z-1) = \mathbf{g}_i^L(z-1)$;

13 Set $\mathbf{g}_i^0(z) = \mathbf{g}_i^{-1}(z) = \sum_{\tau\in\mathcal{T}_z} \nabla f_{\tau,i}(\mathbf{x}_i(z))$

More specifically, the algorithm partitions the time horizon T into T/L blocks indexed by z, where $\mathcal{T}_z \triangleq \{(z-1)L+1, \ldots, zL\}$ denotes the set of time steps in block z.[1] Before introducing the algorithm, for notational convenience, we define $\overline{\mathbf{g}}(z)$ and $\overline{\mathbf{x}}(z)$ as follows:

$$\overline{\mathbf{g}}(z) \triangleq \frac{1}{N} \sum_{t\in\mathcal{T}_z} \sum_{j=1}^{N} \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \qquad \overline{\mathbf{x}}(z) \triangleq \frac{1}{N} \sum_{j=1}^{N} \hat{\mathbf{x}}_j(z). \tag{5}$$

Specifically, $\overline{\mathbf{g}}(z)$ denotes the network average of the cumulative gradients collected during block z, while $\overline{\mathbf{x}}(z)$ denotes the average of the preparatory decisions. Algorithm 1 proceeds over these blocks as follows.

- **Step 1: Local Preparatory Update.** Before making a decision for the current block z, each agent i computes a local preparatory decision $\hat{\mathbf{x}}_i(z+1)$. This is obtained via a delayed ONS update (Line 4) using the gossiped aggregated gradient $\mathbf{g}_i(s)$ and decision information $\mathbf{x}_i(s)$ collected from blocks $s = 1$ to $z-2$.
- **Step 2: Decision Execution and Parallel Gossip.** Each agent i plays the fixed decision $\mathbf{x}_i(z) \in \mathcal{X}$ for all rounds $t \in \mathcal{T}_z$ and observes the local gradients $\nabla f_{t,i}(\mathbf{x}_i(z))$ (Line 7). Simultaneously, two accelerated gossip processes are executed in parallel:
  - *Gradient Gossip:* Agents run accelerated gossip on the gradients collected during the previous block $z-1$ (Line 8). According to Lemma 1, by the end of block z, the local state $\mathbf{g}_i^L(z-1)$ will approximate the network average of the cumulative gradients $\overline{\mathbf{g}}(z-1)$ within a poly$(1/T)$ error. This will serve as the feedback for future ONS updates.
  - *Decision Gossip:* Agents also run accelerated gossip on their local preparatory decisions $\hat{\mathbf{x}}_i(z+1)$. While one might intuitively use $\hat{\mathbf{x}}_i(z+1)$ directly as the decision, these local ONS outputs may *diverge* significantly across agents, complicating the regret analysis. To ensure consensus on the decision over all the agents, we also gossip these vectors so that $\mathbf{x}_i^L(z+1)$ converges toward the network average $\overline{\mathbf{x}}(z+1) = \frac{1}{N}\sum_j \hat{\mathbf{x}}_j(z+1)$ up to a poly$(1/T)$ consensus error.
- **Step 3: Projection and Finalization.** At the end of block z, the output of the decision-gossip process $\mathbf{x}_i^L(z+1)$ is projected back onto the feasible set $\mathcal{X}$. This projection is necessary because the linear combinations involved in accelerated gossip may produce iterates that lie outside the convex constraints.[2] The resulting projected vector is then used as the played decision $\mathbf{x}_i(z+1)$ for the next

[1] For notational convenience, we assume that T/L is an integer.

[2] We point out that the algorithm in Wan [2025] in fact also requires this projection to make sure that $\mathbf{x}_i(z)$ lies in the feasible set.

block. This ensures that the decisions made by each agent at each epoch are close to each other. For $z \geq 2$, the refined gradient estimate $\mathbf{g}_i(z-1)$ is fed into the local ONS subroutine to guide the next iteration.

Consequently, our algorithm can be viewed as operating with a constant *2-block delay*. Specifically, after agents collect their local cumulative gradients in block z, they perform a gossip process during block $z+1$ to approximate the network average of the cumulative gradients. Following the ONS update, a second accelerated gossip process is executed in block $z+2$ to synchronize the decisions for use in block $z+3$. Although this delayed information flow is adapted from the blocking framework of Wan [2025], the analysis must account for the gradient-dependent curvature of exp-concave losses and the distinct preparatory and executed decisions.

Formally, we state our regret guarantee for decentralized online exp-concave optimization as follows:

Theorem 1. *Suppose that the sequence of loss functions $\{f_{t,i}\}_{t\in[T],i\in[N]}$ satisfy Assumption 1, 2, and 3. Let P be the communication matrix defined in Definition 1. If every agent $i \in [N]$ runs Algorithm 1 with $\beta = \frac{1}{2}\min\{\frac{1}{4GD}, \alpha\}$ and with parameters θ and L chosen in Eq. (3), then, for any agent $i \in [N]$, the regret is bounded by:*

$$\mathrm{Reg}_{T,i}(\mathbf{x}) = \mathcal{O}\left(\eta N\|\mathbf{x}\|_2^2 + \frac{Nd(1+\beta GD)L}{\beta}\log\left(1+\frac{\beta G^2 T}{\eta d}\right)\right). \tag{6}$$

By substituting the choice of L from Eq. (3) into Eq. (6), we obtain a total regret of $\mathcal{O}\left(Nd(1-\sigma_2(P))^{-1/2}\log^2(T)\right)$. To our knowledge, this is the first logarithmic regret guarantee for decentralized online exp-concave optimization. We complement this result with a lower bound of $\Omega\left(Nd(1-\sigma_2(P))^{-1/2}\log(T/N)\right)$ in Section 3.3, demonstrating the near-optimality of our upper bound with respect to T, the number of agents N, the dimension d, and the network spectral gap $1-\sigma_2(P)$.

3.2 Analysis of Algorithm 1

In this subsection, we provide a proof sketch of Theorem 1. The complete proof is deferred to Appendix B. The analysis consists of two main steps. First, we provide a decentralized-to-block-level reduction, showing that the regret can be controlled by a sum of local ONS-type surrogate terms up to lower-order gossip errors. For any comparator $\mathbf{x} \in \mathcal{X}$, utilizing the Lipschitz property, the property of exp-concavity, and the consensus guarantee in Lemma 1, we obtain that

$$\begin{aligned} f_{t,j}\big(\mathbf{x}_i(z)\big) - f_{t,j}(\mathbf{x}) \leq \big\langle \nabla f_{t,j}\big(\mathbf{x}_j(z)\big), \overline{\mathbf{x}}(z) - \mathbf{x}\big\rangle - \frac{\beta}{2}\big\langle \nabla f_{t,j}\big(\mathbf{x}_j(z)\big), \mathbf{x}_j(z) - \mathbf{x}\big\rangle^2 \\ + \mathcal{O}\big(G\|\overline{\mathbf{x}}(z) - \mathbf{x}_i(z)\| + \|\overline{\mathbf{x}}(z) - \mathbf{x}_j(z)\|\big). \end{aligned} \tag{7}$$

Specifically, the last term of Eq. (7) can be bounded by $\mathcal{O}(GD/(TN)^2)$ since $\mathbf{x}_i(z)$ and $\mathbf{x}_j(z)$ are both projected approximations of the network-average decision $\overline{\mathbf{x}}(z)$, thus Lemma 1 implies that their deviations from $\overline{\mathbf{x}}(z)$ are of order $\mathcal{O}(D/(T^2N^2))$. Therefore, the last term in Eq. (7) contributes only a lower-order term after summing over all rounds and agents. For the first two terms, after taking a summation over all $j \in [N]$ and $t \in \mathcal{T}_z$ and applying Jensen's inequality, with some calculation, we know that they are bounded by

$$N\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle - \frac{\beta}{2NL}\Big(\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}(\mathbf{x}_j(z)), \mathbf{x}_j(z) - \mathbf{x}\rangle\Big)^2. \tag{8}$$

By definition of $\overline{\mathbf{x}}(z)$, the first term of Eq. (8) can be written as $\sum_{j=1}^{N}\langle \overline{\mathbf{g}}(z), \widehat{\mathbf{x}}_j(z) - \mathbf{x}\rangle$, which can be further bounded by $\sum_{j=1}^{N}\langle \mathbf{g}_j(z), \widehat{\mathbf{x}}_j(z) - \mathbf{x}\rangle$ plus an $\mathcal{O}(1/T)$ approximation error due to Lemma 1. Via some calculation, the second term of Eq. (8) can be bounded by $-\frac{\beta}{2L}\sum_{j=1}^{N}\langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x}\rangle^2 + \mathcal{O}(\sum_{j=1}^{N}(\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\|_2 + \|\mathbf{g}_j(z) - \overline{\mathbf{g}}(z)\|_2))$, where the second term is bounded by $\mathcal{O}(1/T)$ due to Lemma 1 again. Combining all of the above, we reduce bounding $\mathrm{Reg}_{T,i}(\mathbf{x})$ to bounding

$$\sum_{z=1}^{T/L}\sum_{j=1}^{N}\left(\langle \mathbf{g}_j(z), \widehat{\mathbf{x}}_j(z) - \mathbf{x}\rangle - \frac{\beta}{2L}\langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x}\rangle^2\right) \tag{9}$$

plus $\mathcal{O}(1)$ term induced by the gossip process. The second step is to bound Eq. (9). A natural approach would be to interpret each summand as a surrogate loss and then apply a standard FTRL or ONS regret bound to the preparatory decision $\widehat{\mathbf{x}}_j(z)$. However, unlike the strongly convex setting considered in Wan [2025], such a direct argument does not apply here. The difficulty is that the linear term is evaluated at the preparatory decision $\widehat{\mathbf{x}}_j(z)$, whereas the negative quadratic term depends on the executed decision $\mathbf{x}_j(z)$. This mismatch prevents the summand in Eq. (9) from being treated as the value of a standard ONS surrogate loss at a single decision point.

To overcome this issue, we analyze the delayed FTRL dynamics for $\widehat{\mathbf{x}}_j(z)$ more carefully. The key observation is that although $\widehat{\mathbf{x}}_j(z)$ is computed using only feedback up to block $z-3$, it remains close, in the adaptive second-order geometry induced by the accumulated gradients, to the "current" regularized leader that has access to feedback up to block z. This closeness allows us to control the price of the two-block delay by a logarithmic term using ellipsoid potential arguments. The following lemma formalizes this delayed ONS bound. Its proof uses the same local-norm stability and elliptical-potential tools as Qiu et al. [2025, Theorem 4.1]; in particular, their Lemma A.2 gives the intermediate stability comparison. Here, $\mathbf{g}_j(z)$ aggregates gradients evaluated at different agents' executed decisions, while the linear and quadratic terms involve $\widehat{\mathbf{x}}_j(z)$ and $\mathbf{x}_j(z)$, respectively. We use the lemma to examine this two-decision structure, rather than highlight it as a separate novel contribution. The self-contained derivation is retained in Appendix B.

Lemma 2. *Under the same conditions as Algorithm 1. For any $j \in [N]$ and every comparator $\mathbf{x} \in \mathcal{X}$,*

$$\sum_{z=1}^{T/L} \left(\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z) - \mathbf{x} \rangle - \tfrac{\beta}{2L} \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x} \rangle^2 \right) \leq \mathcal{O}\left(\eta \|\mathbf{x}\|_2^2 + \frac{d(1+\beta GD)L \log\left(1 + \frac{\beta G^2 T}{\eta d}\right)}{\beta} \right).$$

The full proof is deferred to Appendix B. Applying Lemma 2 to each $j \in [N]$ and combining it with the regret decomposition above yields the claimed bound in Theorem 1. This completes the proof sketch.

3.3 Lower Bound

In this section, we provide the lower bound of decentralized online exp-concave optimization. Specifically, we show that no algorithm can achieve a regret better than $\Omega(Nd(1-\sigma_2)^{-1/2} \log(T/N))$ in decentralized exp-concave online optimization.

Theorem 2. *Let $N = 4K + 2$ for some positive integer K, and let $\mathcal{X} = \Delta_d$.*[3] *There exists a graph $\mathcal{G} = ([N], E)$ with $P = I - \frac{1}{2(d_{\max}+1)} L_{\mathcal{G}}$ such that there exists a sequence of local convex α-exp-concave losses $\{f_{t,i} : \mathcal{X} \to \mathbb{R}\}_{i \in [N], t \in [T]}$ such that $\mathbb{E}[\mathrm{Reg}_{T,1}] \geq \Omega\left(\frac{Nd}{\alpha \sqrt{(1-\sigma_2(P))}} \log(T/N) \right)$.*

Compared to this lower bound, our upper bounds are optimal in their polynomial dependence on $1-\sigma_2(P)$, d, T, and N. While a logarithmic gap in T and N remains, we leave bridging this gap as a future direction.

We provide a proof sketch here and defer the full proof to Appendix B.1. Our proof is adapted from the construction in Wan et al. [2024]. We consider a carefully designed instance on a 1-connected cycle graph. The horizon is divided into blocks of length $\Theta(N)$. In each block, we place one hard α-exp-concave portfolio loss from Ordentlich and Cover [1998] on a contiguous middle segment of $\Theta(N)$ agents, while the remaining agents have identically zero local losses. Hence, over each block, the global loss is a $\Theta(N^2)$-scaled copy of one hard exp-concave loss.

Focusing on an agent u in the zero-loss segment, we observe that all agents carrying the non-zero loss are at graph distance $\Theta(N)$ from u. Therefore, the information about the current block loss cannot reach u before the block ends. After averaging the decisions of u within each block, its optimization problem reduces to a centralized online exp-concave optimization instance on a sequence of $\Theta(T/N)$ losses. By convexity, the distributed regret of agent u is at least $\Theta(N^2)$ times the regret in this centralized hard instance. Applying the portfolio lower bound of Ordentlich and Cover [1998] gives $\Omega(\alpha^{-1} N^2 d \log(T/N))$. Finally, for the cycle graph, the spectral gap satisfies $1/\sqrt{1-\sigma_2(P)} = \Theta(N)$, which yields the claimed lower bound.

[3] To align with Assumption 2 used in our upper bound, one can shift the simplex Δ_d by $\frac{1}{d}\mathbf{1}$ to include the origin. The α-exp-concavity is invariant under such constant shifts.

4 Decentralized Online Linear Regression

In this section, we investigate decentralized online linear regression (D-OLR) over an unbounded domain. The learning protocol proceeds as follows. At each round $t \in [T]$, before making its prediction, agent i observes only its local feature vector $\mathbf{h}_{t,i}$. Each agent i decides a model $\mathbf{x}_i(t) \in \mathbb{R}^d$ to generate its prediction. Crucially, after making this decision, agent i only observes its local true label $y_{t,i} \in [-Y, Y]$ corresponding to the feature vector $\mathbf{h}_{t,i}$, while the features and labels of other agents are not directly available to agent i. The agents communicate the sufficient statistics needed for their model updates through the network.

Because the true labels are bounded such that $y_{t,i} \in [-Y, Y]$ for all $t \in [T]$ and $i \in [N]$, we define the prediction for a given feature $\mathbf{h}$ and model $\mathbf{x}$ as the clipped inner product $\Pi_{[-Y,Y]}(\langle \mathbf{h}, \mathbf{x} \rangle)$, where $\Pi_{[-Y,Y]}(a) \triangleq \max\{-Y, \min\{a, Y\}\}$. Furthermore, let $\widehat{y}^i_{t,j} \triangleq \Pi_{[-Y,Y]}\left(\langle \mathbf{h}_{t,j}, \mathbf{x}_i(t) \rangle\right)$ denote the prediction of agent i's model evaluated on agent j's feature at round t. For $j \neq i$, this is a virtual prediction used only to define regret; its evaluation does not require agent i to observe $\mathbf{h}_{t,j}$ online. Clipping does not increase the squared loss because the true labels lie in $[-Y, Y]$. It also makes that loss $2Y$-Lipschitz in the unbounded raw prediction, which is essential for controlling the gossip error without imposing a bounded model domain.

We then measure the performance of each agent i using the regret evaluated on the squared loss of their predictions across all global features. Specifically, the regret of agent i with respect to a fixed comparator $\mathbf{x}$ is defined as follows:

$$\mathrm{Reg}^{\mathrm{OLR}}_{T,i}(\mathbf{x}) \triangleq \sum_{t=1}^{T}\sum_{j=1}^{N} \frac{1}{2}\left(\Pi_{[-Y,Y]}\left(\langle \mathbf{h}_{t,j}, \mathbf{x}_i(t) \rangle\right) - y_{t,j}\right)^2 - \sum_{t=1}^{T}\sum_{j=1}^{N} \frac{1}{2}\left(\langle \mathbf{h}_{t,j}, \mathbf{x} \rangle - y_{t,j}\right)^2 .$$

We also define the optimal baseline as $\mathrm{Reg}^{\mathrm{OLR}}_{T,i} \triangleq \max_{\mathbf{x} \in \mathbb{R}^d} \mathrm{Reg}^{\mathrm{OLR}}_{T,i}(\mathbf{x})$. Notably, although squared loss is exp-concave on bounded prediction ranges, the unbounded nature of the decision domain in this setting prevents a direct application of the Algorithm 1 discussed in the previous section. Before presenting our algorithm, we formalize a standard boundedness assumption regarding the feature vectors.

Assumption 4. *The feature vectors $\{\mathbf{h}_{t,i}\}_{t \in [T], i \in N}$ have bounded norms. Specifically, there exists an $H \geq 0$ such that $\|\mathbf{h}_{t,i}\|_2 \leq H$ holds for any $t \in [T]$ and $i \in [N]$.*

For each block z, agent i forms the covariance and label-feature statistics

$$C^0_i(z) = \sum_{\tau \in \mathcal{T}_z} \mathbf{h}_{\tau,i}\mathbf{h}_{\tau,i}^\top, \qquad Y^0_i(z) = \sum_{\tau \in \mathcal{T}_z} y_{\tau,i}\mathbf{h}_{\tau,i}.$$

Both statistics are gossiped during block $z+1$ and are available for use in block $z+2$. We use two ridge-regression copies, assigned to odd and even blocks, so that each copy's previous statistics have finished gossiping before its next prediction [Joulani et al., 2013]. Write $\mathcal{S}_z = \{s : 1 \leq s \leq z-2,\ s \equiv z \pmod 2\}$ for the earlier blocks assigned to the same copy; this set is empty for $z \leq 2$. We use $\mathcal{T}_z = \{(z-1)L+1, \ldots, \min\{zL, T\}\}$ to include a final incomplete block.

We now introduce Algorithm 2 for the D-OLR problem. Our approach builds upon the framework of block-based accelerated gossip, integrating it with the classic online ridge regression algorithm [Hoerl and Kennard, 1970, Vovk, 2001]. In contrast to the general exp-concave setting over bounded domains which necessitates the communication of both cumulative gradients and decision vectors, our algorithm exploits a fundamental structural property of D-OLR: the model update depends only on covariance and label-feature statistics, both of which can be aggregated locally and communicated through gossip. Specifically, Algorithm 2 proceeds over these blocks as follows:

- **Step 1: Local Ridge Prediction.** At each round $t \in \mathcal{T}_z$, agent i computes a ridge-regression estimate $\mathbf{x}_i(z,t)$ (Line 5) using the covariance and label-feature statistics aggregated via gossip from the earlier blocks $s \in \mathcal{S}_z$. The model is fixed during the block. Here $\Pi_{\mathbb{S}^d_+}$ denotes projection onto the positive semidefinite cone in Frobenius norm. This projection is needed because accelerated gossip uses signed combinations and need not preserve positive semidefiniteness; it ensures that the regularized matrix is bounded below by λI_d.
- **Step 2: Prediction and Local Feedback.** Then, agent i deploys the decision model $\mathbf{x}_i(z,t)$ and observes only its local ground-truth label $y_{t,i}$. With a slight abuse of notation, we define

Algorithm 2: Decentralized online ridge regression for agent $i \in [N]$

Input: gossip parameter θ, epoch length L, regularization parameter $\lambda > 0$, communication matrix P, clipping level Y;

for $z = 1, \dots, \lceil T/L \rceil$ **do**

  Set $k_Y = 0$;

  for $t = (z-1)L + 1, \dots, \min\{zL, T\}$ **do**

    Observe the local feature $\mathbf{h}_{t,i}$ and compute

$$\mathbf{x}_i(z,t) = \left(\lambda I_d + N \sum_{s \in \mathcal{S}_z} \Pi_{\mathbb{S}_+^d}(C_i(s))\right)^{-1} N \sum_{s \in \mathcal{S}_z} Y_i(s).$$

    Predict $\Pi_{[-Y,Y]}(\langle \mathbf{h}_{t,i}, \mathbf{x}_i(z,t) \rangle)$ and receive only the local label $y_{t,i}$;

    if $z \geq 2$ **then**

      Simultaneously update

$$C_i^{k_Y+1}(z-1) = (1+\theta) \sum_{j \in \mathcal{N}_i} P_{ij} C_j^{k_Y}(z-1) - \theta C_i^{k_Y-1}(z-1),$$

$$Y_i^{k_Y+1}(z-1) = (1+\theta) \sum_{j \in \mathcal{N}_i} P_{ij} Y_j^{k_Y}(z-1) - \theta Y_i^{k_Y-1}(z-1).$$

    Set $k_Y \leftarrow k_Y + 1$;

  if $z \geq 2$ *and* $k_Y = L$ **then**

    Set $C_i(z-1) = C_i^L(z-1)$ and $Y_i(z-1) = Y_i^L(z-1)$;

  Set $C_i^0(z) = C_i^{-1}(z) = \sum_{\tau \in \mathcal{T}_z} \mathbf{h}_{\tau,i} \mathbf{h}_{\tau,i}^\top$;

  Set $Y_i^0(z) = Y_i^{-1}(z) = \sum_{\tau \in \mathcal{T}_z} y_{\tau,i} \mathbf{h}_{\tau,i}$;

$\widehat{y}_{t,j}^i(z,t) \triangleq \Pi_{[-Y,Y]}(\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) \rangle)$ as the virtual evaluation of agent i's model on agent j's feature using Algorithm 2.

- **Step 3: Gossip of Covariance and Label-Feature Statistics.** While collecting new local labels during block z, the agents simultaneously perform accelerated gossip on both statistics compiled during the previous block $z-1$, applying the covariance update entrywise. After L gossip iterations, the local vector $Y_i^L(z-1)$ approximates the network average of the cumulative label-feature vectors $\frac{1}{N} \sum_{\tau \in \mathcal{T}_{z-1}} \sum_{j=1}^N y_{\tau,j} \mathbf{h}_{\tau,j}$ generated in block $z-1$. Likewise, $C_i^L(z-1)$ approximates $\frac{1}{N} \sum_{\tau \in \mathcal{T}_{z-1}} \sum_{j=1}^N \mathbf{h}_{\tau,j} \mathbf{h}_{\tau,j}^\top$. Both updates use the same neighbor-communication step, with $O(d^2)$ real numbers per message.

We remark again that in contrast to Algorithm 1 for the general exp-concave setting which requires gossiping both the cumulative gradients and the decision vectors, Algorithm 2 requires gossiping covariance and label-feature statistics, but not the decision vectors. The absence of a decision-gossip stage allows the procedure to operate with only a 1-block delay, rather than the 2-block delay inherent to Algorithm 1 as discussed before. Specifically, after agents collect their local covariance and label-feature statistics in block z, they execute the two gossip recursions simultaneously during block $z+1$ to approximate the network averages. These statistics are then incorporated into the same ridge-regression copy for use in block $z+2$. Thus, each copy uses all of its completed earlier blocks, and the interleaving costs only an absolute factor in regret.

The following theorem demonstrates that Algorithm 2 guarantees a logarithmic regret for D-OLR.

Theorem 3. *Suppose Assumption 4 holds, and let P be the communication matrix defined in Definition 1. In D-OLR, if every agent $i \in [N]$ executes Algorithm 2 with $\lambda > 0$, θ chosen according to Eq. (3), and*

$$L = \left\lceil \frac{\sqrt{2}}{(\sqrt{2}-1)\sqrt{1-\sigma_2(P)}} \log\left(\sqrt{14}\,T^3 N^{7/2}(1+H^2/\lambda)\right) \right\rceil, \tag{10}$$

then for any agent $i \in [N]$ and comparator $\mathbf{x} \in \mathbb{R}^d$, the regret is bounded by:

$$\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) = \mathcal{O}\left(\lambda\|\mathbf{x}\|_2^2 + \frac{Y^2H^2}{\lambda} + Y^2LNd\log\left(1+\frac{NTH^2}{\lambda d}\right)\right). \tag{11}$$

By substituting the choice of L from Eq. (10) into Eq. (11), we obtain a total regret of $\mathcal{O}\left(Nd(1-\sigma_2(P))^{-1/2}\log^2(T)\right)$. To our knowledge, Theorem 3 establishes the first logarithmic regret guarantee for D-OLR with locally observed features and labels.

4.1 Analysis of Algorithm 2

We briefly explain the main ideas behind the proof of Theorem 3; the complete proof is deferred to Appendix C. The analysis compares Algorithm 2 with a centralized online ridge-regression predictor that has access to all labels immediately. We make this comparison separately for the odd and even copies. Write

$$A(z) = \lambda I_d + \sum_{s\in\mathcal{S}_z}\sum_{j=1}^{N} C_j^0(s), \qquad b(z) = \sum_{s\in\mathcal{S}_z}\sum_{j=1}^{N} Y_j^0(s)$$

for the exact delayed covariance and label-feature statistics. For each block z and round $t \in \mathcal{T}_z$, define the reference predictor

$$\mathbf{x}^\star(z,t) = \left(A(z) + \sum_{\substack{\tau\in\mathcal{T}_z\\ \tau\le t}}\sum_{j=1}^{N}\mathbf{h}_{\tau,j}\mathbf{h}_{\tau,j}^\top\right)^{-1}\left(b(z) + \sum_{\substack{\tau\in\mathcal{T}_z\\ \tau\le t}}\sum_{j=1}^{N} y_{\tau,j}\mathbf{h}_{\tau,j}\right),$$

which corresponds to the centralized ridge-regression solution using all observations up to round t assigned to the same copy. It is used only for the analysis. The regret is then decomposed into two parts:

$$\begin{aligned}\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) = &\sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left(f_{t,j}(\mathbf{x}^\star(z,t)) - f_{t,j}(\mathbf{x})\right)\\ &+ \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left(\widehat{f}_{t,j}^{\,i}(z,t) - f_{t,j}(\mathbf{x}^\star(z,t))\right),\end{aligned} \tag{12}$$

where we define $f_{t,j}(\mathbf{x}) = \frac{1}{2}(\langle\mathbf{h}_{t,j},\mathbf{x}\rangle - y_{t,j})^2$ and $\widehat{f}_{t,j}^{\,i}(z,t) \triangleq \frac{1}{2}(\widehat{y}_{t,j}^{\,i}(z,t) - y_{t,j})^2$. The first term is precisely the sum of the regrets of the two centralized ridge-regression reference sequences. Applying the Be-The-Regularized-Leader lemma to each copy bounds this term by $\lambda\|\mathbf{x}\|_2^2$.

It remains to control the discrepancy between the decentralized predictor and the full-information ridge predictor. As the clipping operation guarantees that the squared loss is $2Y$-Lipschitz with respect to the raw predictions, the loss difference between two clipped predictors is at most $2Y$ times their prediction difference. This is why clipping is useful here: the models themselves can be unbounded, whereas the error estimate depends only on the bounded label range. Moreover, clipping the reference prediction cannot increase its squared loss. The complete scalar argument is given in Appendix C.

We introduce an auxiliary exact delayed ridge predictor $\mathbf{x}'(z,t) = A(z)^{-1}b(z)$. The analysis then separates a *delay error*, due to the absence of the current block's covariance and label-feature statistics, from a *gossip error*, due to imperfect averaging of completed blocks.

For the delay error, each copy is ordinary ridge regression over batches of at most LN examples. We retain the regularized-minimum increments in the centralized comparison and use the quadratic structure to control the loss of the fixed model throughout a batch. Clipping bounds each batch loss by

$2Y^2LN$. A matrix-inversion identity then bounds the excess over the regularized-minimum increment by $2Y^2LN$ times the increase in a log-determinant potential. Summing these increases over the odd and even copies gives $\mathcal{O}\big(Y^2LNd\log(1+NTH^2/(\lambda d))\big)$. This accounts jointly for the delay of the covariance and label-feature statistics; the intermediate calculations are included in the full proof.

For the gossip error, the covariance estimates are projected onto the positive semidefinite cone, so both the exact and approximate regularized matrices dominate λI_d. Non-expansiveness of this projection and the matrix-inverse identity $\widehat{A}^{-1} - A^{-1} = \widehat{A}^{-1}(A - \widehat{A})A^{-1}$ reduce the model difference to the covariance and label-feature gossip errors. Using the same accelerated-gossip estimate as in Lemma 1, with the block length in Eq. (10), their cumulative contribution to regret is at most $4Y^2H^2/\lambda$. Summing the centralized, delay, and gossip terms yields the claimed regret bound.

References

B. Awerbuch and R. D. Kleinberg. Adaptive routing with end-to-end feedback: Distributed learning and geometric approaches. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 45–53, 2004.

P. Bartlett, E. Hazan, and A. Rakhlin. Adaptive online gradient descent. *Advances in neural information processing systems*, 20, 2007.

J. C. Duchi, A. Agarwal, and M. J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic control*, 57(3):592–606, 2011.

A. Gyorgy and P. Joulani. Adapting to delays and data in adversarial multi-armed bandits. In *International Conference on Machine Learning*, pages 3988–3997. PMLR, 2021.

E. Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3-4): 157–325, 2016.

E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2):169–192, 2007.

A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970. doi: 10.1080/00401706.1970.10488634.

S. Hosseini, A. Chapman, and M. Mesbahi. Online distributed optimization via dual averaging. In *52nd IEEE Conference on Decision and Control*, pages 1484–1489. IEEE, 2013.

P. Joulani, A. Gyorgy, and C. Szepesvári. Online learning under delayed feedback. In *International conference on machine learning*, pages 1453–1461. PMLR, 2013.

D. Li, K. D. Wong, Y. H. Hu, and A. M. Sayeed. Detection, classification, and tracking of targets. *IEEE signal processing magazine*, 19(2):17–29, 2002.

J. Liu and A. S. Morse. Accelerated linear iterations for distributed averaging. *Annual Reviews in Control*, 35(2):160–165, 2011.

H. Luo, M. Zhang, and P. Zhao. Adaptive bandit convex optimization with heterogeneous curvature. In P.-L. Loh and M. Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 1576–1612. PMLR, 02–05 Jul 2022.

J. J. Mayo, H. Hadiji, and T. van Erven. Scale-free unconstrained online learning for curved losses. In P.-L. Loh and M. Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 4464–4497. PMLR, 02–05 Jul 2022.

A. Nedic and A. Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on automatic control*, 54(1):48–61, 2009.

F. Orabona. A modern introduction to online learning, 2026. URL https://arxiv.org/abs/1912.13213.

E. Ordentlich and T. M. Cover. The cost of achieving the best portfolio in hindsight. *Mathematics of Operations Research*, 23(4):960–982, 1998.

H. Qiu, E. Esposito, and M. Zhang. Exploiting curvature in online convex optimization with delayed feedback. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 50448–50479. PMLR, 2025.

O. Shorinwa, T. Halsted, J. Yu, and M. Schwager. Distributed optimization methods for multi-robot systems: Part 1—a tutorial [tutorial]. *IEEE Robotics & Automation Magazine*, 31(3):121–138, 2024.

V. Vovk. Competitive on-line statistics. *International Statistical Review*, 69(2):213–248, 2001.

Y. Wan. Black-box reductions for decentralized online convex optimization in changing environments. In *The Thirty Eighth Annual Conference on Learning Theory*, pages 5605–5631. PMLR, 2025.

Y. Wan, T. Wei, M. Song, and L. Zhang. Nearly optimal regret for decentralized online convex optimization. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 4862–4888. PMLR, 2024.

Y. Wan, T. Wei, B. Xue, M. Song, and L. Zhang. Optimal and efficient algorithms for decentralized online convex optimization. *Journal of Machine Learning Research*, 26(135):1–43, 2025.

L. Xiao and S. Boyd. Fast linear iterations for distributed averaging. *Systems & Control Letters*, 53(1): 65–78, 2004.

F. Yan, S. Sundaram, S. Vishwanathan, and Y. Qi. Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties. *IEEE Transactions on Knowledge and Data Engineering*, 25 (11):2483–2493, 2013.

H. Ye, L. Luo, Z. Zhou, and T. Zhang. Multi-consensus decentralized accelerated gradient descent. *Journal of machine learning research*, 24(306):1–50, 2023.

D. Yuan, A. Proutiere, G. Shi, et al. Multi-agent online optimization. *Foundations and Trends® in Optimization*, 7(2-3):81–263, 2024.

Y.-J. Zhang, P. Zhao, and M. Sugiyama. Non-stationary online learning for curved losses: Improved dynamic regret via mixability. *arXiv preprint arXiv:2506.10616*, 2025.

A Auxiliary Lemmas

In this section, we provide some auxiliary technical results that will be helpful throughout the paper.

The following lemma is a standard result standard be-the-leader lemma, see also [Gyorgy and Joulani, 2021, Theorem 3].

Lemma 3 (Be-The-Regularized-Leader)**.** *Let $\ell_1, \dots, \ell_T : \mathcal{X} \to \mathbb{R}$ be arbitrary functions, and let $\psi_0, \psi_1, \dots, \psi_T : \mathcal{X} \to \mathbb{R}$ satisfy $\psi_0 \equiv 0$. Define $F_t(\mathbf{x}) \triangleq \sum_{s=1}^{t} \ell_s(\mathbf{x}) + \psi_t(\mathbf{x})$ and $\mathbf{x}^\star(t) \in \arg\min_{\mathbf{x}\in\mathcal{X}} F_t(\mathbf{x}), \quad t \in [T]$. Then, for any $\mathbf{x} \in \mathcal{X}$,*

$$\sum_{t=1}^{T} \ell_t\big(\mathbf{x}^\star(t)\big) \le \sum_{t=1}^{T} \ell_t(\mathbf{x}) + \psi_T(\mathbf{x}) - \sum_{t=1}^{T} \Big(\psi_t\big(\mathbf{x}^\star(t)\big) - \psi_{t-1}\big(\mathbf{x}^\star(t)\big)\Big). \tag{13}$$

In particular, if

$$\psi_t(\mathbf{x}) - \psi_{t-1}(\mathbf{x}) \ge 0, \qquad \forall t \in [T],\ \forall \mathbf{x} \in \mathcal{X},$$

then

$$\sum_{t=1}^{T} \ell_t\big(\mathbf{x}^\star(t)\big) \le \sum_{t=1}^{T} \ell_t(\mathbf{x}) + \psi_T(\mathbf{x}).$$

Proof. Define

$$\tilde{\ell}_t(\mathbf{x}) \triangleq \ell_t(\mathbf{x}) + \psi_t(\mathbf{x}) - \psi_{t-1}(\mathbf{x}), \qquad t \in [T].$$

Then, for every $t \in [T]$,

$$\sum_{s=1}^{t} \tilde{\ell}_s(\mathbf{x}) = \sum_{s=1}^{t} \ell_s(\mathbf{x}) + \sum_{s=1}^{t} \big(\psi_s(\mathbf{x}) - \psi_{s-1}(\mathbf{x})\big) = \sum_{s=1}^{t} \ell_s(\mathbf{x}) + \psi_t(\mathbf{x}) = F_t(\mathbf{x}).$$

Hence,

$$\mathbf{x}^\star(t) \in \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{s=1}^{t} \tilde{\ell}_s(\mathbf{x}).$$

Applying the Be-The-Leader lemma to the sequence $\{\tilde{\ell}_t\}_{t=1}^{T}$, we obtain that for any $\mathbf{x} \in \mathcal{X}$,

$$\sum_{t=1}^{T} \tilde{\ell}_t\big(\mathbf{x}^\star(t)\big) \le \sum_{t=1}^{T} \tilde{\ell}_t(\mathbf{x}).$$

Expanding both sides yields

$$\sum_{t=1}^{T} \Big(\ell_t\big(\mathbf{x}^\star(t)\big) + \psi_t\big(\mathbf{x}^\star(t)\big) - \psi_{t-1}\big(\mathbf{x}^\star(t)\big)\Big) \le \sum_{t=1}^{T} \Big(\ell_t(\mathbf{x}) + \psi_t(\mathbf{x}) - \psi_{t-1}(\mathbf{x})\Big).$$

Since the right-hand side telescopes and $\psi_0 \equiv 0$, we have

$$\sum_{t=1}^{T} \big(\psi_t(\mathbf{x}) - \psi_{t-1}(\mathbf{x})\big) = \psi_T(\mathbf{x}).$$

Therefore,

$$\sum_{t=1}^{T} \ell_t\big(\mathbf{x}^\star(t)\big) + \sum_{t=1}^{T} \Big(\psi_t\big(\mathbf{x}^\star(t)\big) - \psi_{t-1}\big(\mathbf{x}^\star(t)\big)\Big) \le \sum_{t=1}^{T} \ell_t(\mathbf{x}) + \psi_T(\mathbf{x}).$$

Rearranging proves (13). If, in addition, $\psi_t(\mathbf{x}) - \psi_{t-1}(\mathbf{x}) \ge 0$, $\forall t \in [T],\ \forall \mathbf{x} \in \mathcal{X}$, then

$$\sum_{t=1}^{T} \Big(\psi_t\big(\mathbf{x}^\star(t)\big) - \psi_{t-1}\big(\mathbf{x}^\star(t)\big)\Big) \ge 0.$$

Dropping this nonnegative term gives

$$\sum_{t=1}^{T} \ell_t\big(\mathbf{x}^\star(t)\big) \le \sum_{t=1}^{T} \ell_t(\mathbf{x}) + \psi_T(\mathbf{x}),$$

which completes the proof. □

The following lemma is the quadratic bound of α-exp-concave functions.

Lemma 4 ([Hazan et al., 2007, Lemma 3]). *At each round $t \in [T]$, the loss function $f_{t,i}(\mathbf{x})$ of each agent $i \in [N]$ is α-exp-concave over $\mathcal{X}$, i.e., $f_{t,i}(\mathbf{y}) \geq f_{t,i}(\mathbf{x}) + \langle \nabla f_{t,i}(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + \frac{\beta}{2} \big(\langle \nabla f_{t,i}(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle \big)^2$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, where $\beta = \frac{1}{2} \min \left\{ \frac{1}{4GD}, \alpha \right\}$.*

Here we provide a proof of Lemma 1, which shows that after $\Theta((1-\sigma_2)^{-1/2} \log(TN))$ rounds of accelerated gossiping, the approximation error to the global average is $\mathcal{O}(1/(T^2N^2))$.

Lemma 1. *Let P be the communication matrix defined in Definition 1. Let*

$$\theta \triangleq \frac{1}{1 + \sqrt{1 - \sigma_2^2(P)}}, \qquad L \triangleq \left\lceil \frac{\sqrt{2} \log(T^2 \sqrt{14} N^{5/2})}{(\sqrt{2}-1)\sqrt{1-\sigma_2(P)}} \right\rceil. \tag{3}$$

Consider arbitrary vectors $\mathbf{u}_1, \ldots, \mathbf{u}_N \in \mathbb{R}^d$ satisfying $\|\mathbf{u}_i\|_2 \leq R, i \in [N]$. Let

$$U^0 = U^{-1} = \begin{bmatrix} \mathbf{u}_1^\top \\ \vdots \\ \mathbf{u}_N^\top \end{bmatrix}, \qquad \overline{\mathbf{u}} = \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i^0, \qquad \overline{U} = \begin{bmatrix} \overline{\mathbf{u}}^\top \\ \vdots \\ \overline{\mathbf{u}}^\top \end{bmatrix}.$$

Suppose accelerated gossip is run for L iterations with update $U^{k+1} = (1+\theta)PU^k - \theta U^{k-1}$ for $k \geq 0$. Then, for every agent $i \in [N]$, $\|\mathbf{u}_i^L - \overline{\mathbf{u}}\|_2 \leq \frac{2R}{T^2N^2}$.

Proof. By Proposition 1 of Ye et al. [2023], the accelerated gossip iterations satisfy[4]

$$\|U^L - \overline{U}\|_F \leq \sqrt{14} \left(1 - \left(1 - \frac{1}{\sqrt{2}}\right) \sqrt{1 - \sigma_2(P)}\right)^L \|U^0 - \overline{U}\|_F.$$

By the definition of L and the inequality $\log(1-x) \leq -x$ for $x \in (0,1)$, with $a = \left(1 - \frac{1}{\sqrt{2}}\right) \sqrt{1-\sigma_2(P)}$, we have $L \geq a^{-1} \log(T^2\sqrt{14}N^{5/2})$, and hence

$$\sqrt{14} \left(1 - \left(1 - \frac{1}{\sqrt{2}}\right) \sqrt{1 - \sigma_2(P)}\right)^L = \sqrt{14}(1-a)^L \leq \sqrt{14} e^{-aL} \leq \frac{1}{T^2 N^{5/2}}.$$

Moreover, since $\|\mathbf{u}_i\|_2 \leq R$ for all $i \in [N]$, we also have

$$\|\overline{\mathbf{u}}\|_2 = \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{u}_i \right\|_2 \leq \frac{1}{N} \sum_{i=1}^N \|\mathbf{u}_i\|_2 \leq R.$$

Hence, for every $i \in [N]$,

$$\|\mathbf{u}_i^0 - \overline{\mathbf{u}}\|_2 \leq \|\mathbf{u}_i^0\|_2 + \|\overline{\mathbf{u}}\|_2 = \|\mathbf{u}_i\|_2 + \|\overline{\mathbf{u}}\|_2 \leq 2R.$$

Therefore,

$$\|U^0 - \overline{U}\|_F = \left(\sum_{i=1}^N \|\mathbf{u}_i^0 - \overline{\mathbf{u}}\|_2^2 \right)^{1/2} \leq 2R\sqrt{N}.$$

Combining the previous inequalities gives

$$\|U^L - \overline{U}\|_F \leq \frac{1}{T^2 N^{5/2}} \cdot 2R\sqrt{N} = \frac{2R}{T^2 N^2}.$$

Since the Euclidean norm of each row is upper bounded by the Frobenius norm, we obtain

$$\|\mathbf{u}_i^L - \overline{\mathbf{u}}\|_2 \leq \|U^L - \overline{U}\|_F \leq \frac{2R}{T^2N^2}, \qquad i \in [N]. \tag{14}$$

This completes the proof. □

[4] In the proof of Proposition 1 of Ye et al. [2023], the upper-right block of the transition matrix should be $-\eta_w I$, rather than $-\eta_w W$. With this correction, the projection identities still hold, so the proof requires only this minor adjustment and the proposition remains valid.

B Omitted Details from Section 3

In this section, we provide the omitted details from Section 3. We begin by introducing two key lemmas for the proof of Theorem 1. We then combine the two lemmas to prove Theorem 1, and finally prove the lower bound in Theorem 2.

Lemma 5 reduces the decentralized regret of any agent to a sum of block-level surrogate losses. It shows that, after accelerated gossip, the regret can be controlled by linearized gradients plus an exp-concave quadratic correction term, up to lower-order consensus errors. The key role of this lemma is to connect the decentralized online exp-concave problem with the delayed ONS analysis used later.

Compared with Wan [2025, Theorem 1], which uses the uniform quadratic curvature of strongly convex losses, the reduction below retains the gradient-dependent rank-one correction of exp-concavity. The averaging and Jensen steps are standard; their role here is to preserve this correction in the expression passed to Lemma 2.

Lemma 5. *Under the same conditions as Algorithm 1. For any comparator $\mathbf{x} \in \mathcal{X}$, the regret of any agent $i \in [N]$ satisfies*

$$\mathrm{Reg}_{T,i}(\mathbf{x}) \le \sum_{z=1}^{T/L} \sum_{j=1}^{N} \left(\langle \mathbf{g}_j(z), \widehat{\mathbf{x}}_j(z) - \mathbf{x} \rangle - \frac{\beta}{2} \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x} \rangle^2 \right) + \mathcal{O}\left(\frac{GD}{NT} + \frac{\beta G^2 D^2}{NT} \right).$$

Proof. Before proving the lemma, we first record two consensus bounds implied by Lemma 1. For each block $z \in [T/L]$, recall

$$\overline{\mathbf{x}}(z) = \frac{1}{N} \sum_{j=1}^{N} \widehat{\mathbf{x}}_j(z), \qquad \overline{\mathbf{g}}(z) = \frac{1}{N} \sum_{t \in \mathcal{T}_z} \sum_{j=1}^{N} \nabla f_{t,j}\left(\mathbf{x}_j(z)\right).$$

By Assumption 2, the set $\mathcal{X}$ contains the origin and has diameter at most D. Since $\widehat{\mathbf{x}}_i(z) \in \mathcal{X}$, we have $\|\widehat{\mathbf{x}}_i(z)\|_2 \le D$. Hence, after L accelerated gossip steps in Algorithm 1, and using the non-expansiveness of projection onto a convex set [Orabona, 2026, Proposition 2.11], the played decision satisfies

$$\|\mathbf{x}_i(z) - \overline{\mathbf{x}}(z)\|_2 \le \left\|\mathbf{x}_i^L(z) - \overline{\mathbf{x}}(z)\right\|_2 \le \frac{2D}{T^2 N^2}. \tag{15}$$

Similarly, by Assumption 1, each local loss is G-Lipschitz over $\mathcal{X}$, and hence

$$\left\| \sum_{t \in T_z} \nabla f_{t,i}(\mathbf{x}_i(z)) \right\|_2 \le LG.$$

Applying Lemma 1 to the gradient-gossip process gives

$$\|\mathbf{g}_i(z) - \overline{\mathbf{g}}(z)\|_2 \le \frac{2LG}{T^2 N^2}. \tag{16}$$

For any $\mathbf{x} \in \mathcal{X}$, by Assumption 1, we have

$$\begin{aligned}
f_{t,j}\left(\mathbf{x}_i(z)\right) - f_{t,j}(\mathbf{x}) &\le f_{t,j}\left(\mathbf{x}_j(z)\right) - f_{t,j}(\mathbf{x}) + G\left\|\mathbf{x}_j(z) - \mathbf{x}_i(z)\right\|_2 \\
&\le \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x} \right\rangle - \frac{\beta}{2} \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x} \right\rangle^2 \\
&\quad + G\left\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z) + \overline{\mathbf{x}}(z) - \mathbf{x}_i(z)\right\|_2 \qquad \text{(Assumption 3 and Lemma 4)} \\
&\le \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \overline{\mathbf{x}}(z) - \mathbf{x} \right\rangle - \frac{\beta}{2} \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x} \right\rangle^2 \\
&\quad + \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \overline{\mathbf{x}}(z) \right\rangle + \frac{4GD}{T^2 N^2} \\
&\le \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \overline{\mathbf{x}}(z) - \mathbf{x} \right\rangle - \frac{\beta}{2} \left\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x} \right\rangle^2 \\
&\quad + G\left\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\right\|_2 + \frac{4GD}{T^2 N^2}
\end{aligned}$$

$$\leq \langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle - \frac{\beta}{2}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x}\rangle^2 + \frac{6GD}{T^2N^2}\,, \tag{17}$$

where the third and the last inequalities use Eq. (15). Summing over $z \in [T/L]$, $t \in \mathcal{T}_z$, and $j \in [N]$, we have

$$\begin{aligned}
&\sum_{z=1}^{T/L}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left(f_{t,j}\left(\mathbf{x}_i(z)\right) - f_{t,j}(\mathbf{x})\right) \\
&\leq \sum_{z=1}^{T/L}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left(\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle - \frac{\beta}{2}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x}\rangle^2\right) + \frac{6GD}{TN} \qquad \text{(Eq. (17))} \\
&= \sum_{z=1}^{T/L}\left(N\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle - \frac{\beta}{2}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x}\rangle^2\right) + \frac{6GD}{TN} \\
&\qquad\qquad (\overline{\mathbf{g}}(z) = \tfrac{1}{N}\textstyle\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\nabla f_{t,j}\left(\mathbf{x}_j(z)\right)) \\
&\leq \sum_{z=1}^{T/L}\left(N\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle - \frac{\beta}{2NL}\left(\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x}\rangle\right)^2\right) + \frac{6GD}{TN}, \qquad (18)
\end{aligned}$$

where the last inequality holds because of Jensen's inequality. Now we focus on

$$\begin{aligned}
&-\frac{\beta}{2NL}\left(\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \mathbf{x}\rangle\right)^2 \\
&= -\frac{\beta}{2NL}\left(\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\rangle + \sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\langle \nabla f_{t,j}\left(\mathbf{x}_j(z)\right), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle\right)^2 \\
&\leq -\frac{N\beta}{2L}\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle^2 + \beta LG^2D\sum_{j=1}^{N}\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\|_2 \\
&= -\frac{\beta}{2L}\sum_{j=1}^{N}\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle^2 + \beta LG^2D\sum_{j=1}^{N}\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\|_2 \\
&\leq -\frac{\beta}{2L}\sum_{j=1}^{N}\langle \mathbf{g}_j(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle^2 + \beta LG^2D\sum_{j=1}^{N}\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\|_2 + \beta GD^2\sum_{j=1}^{N}\|\mathbf{g}_j(z) - \overline{\mathbf{g}}(z)\|_2 \\
&\leq -\frac{\beta}{2L}\sum_{j=1}^{N}\langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x}\rangle^2 + 2\beta LG^2D\sum_{j=1}^{N}\|\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\|_2 + \beta GD^2\sum_{j=1}^{N}\|\mathbf{g}_j(z) - \overline{\mathbf{g}}(z)\|_2 \\
&\leq -\frac{\beta}{2L}\sum_{j=1}^{N}\langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x}\rangle^2 + \frac{6\beta G^2D^2L}{T^2N} \qquad (19)
\end{aligned}$$

Here, the first inequality is obtained by decomposing $\mathbf{x}_j(z) - \mathbf{x} = \left(\overline{\mathbf{x}}(z) - \mathbf{x}\right) + \left(\mathbf{x}_j(z) - \overline{\mathbf{x}}(z)\right)$, and applying $-(a+b)^2 \leq -a^2 + 2|a||b|$, Cauchy–Schwarz, Assumption 1, and Assumption 2. The equality follows from the fact that $\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle^2$ is independent of j. For the second inequality, we use the identity $|a^2 - b^2| = |a-b||a+b|$ with $a = \langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle$, and $b = \langle \mathbf{g}_j(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle$,together with Cauchy–Schwarz, Assumption 1, and Assumption 2. The third inequality is proved analogously by taking $a = \langle \mathbf{g}_j(z), \overline{\mathbf{x}}(z) - \mathbf{x}\rangle$ and $b = \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x}\rangle$.The last inequality follows from Eq. (15) and Eq. (16).

Combining Eq. (18) and Eq. (19), we have

$$\sum_{z=1}^{T/L}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left(f_{t,j}\left(\mathbf{x}_i(z)\right) - f_{t,j}(\mathbf{x})\right)$$

$$
\begin{aligned}
&\le \sum_{z=1}^{T/L}\left(N\left\langle \overline{\mathbf{g}}(z), \overline{\mathbf{x}}(z)-\mathbf{x}\right\rangle - \frac{\beta}{2L}\sum_{j=1}^{N}\left\langle \mathbf{g}_j(z), \mathbf{x}_j(z)-\mathbf{x}\right\rangle^2\right) + \frac{6GD}{TN} + \frac{6\beta G^2D^2}{TN} \\
&= \sum_{z=1}^{T/L}\sum_{j=1}^{N}\left(\left\langle \overline{\mathbf{g}}(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}\right\rangle - \frac{\beta}{2L}\left\langle \mathbf{g}_j(z), \mathbf{x}_j(z)-\mathbf{x}\right\rangle^2\right) + \frac{6GD}{TN} + \frac{6\beta G^2D^2}{TN} \quad \left(\overline{\mathbf{x}}(z)=\tfrac{1}{N}\textstyle\sum_{j=1}^{N}\hat{\mathbf{x}}_j(z)\right) \\
&\le \sum_{z=1}^{T/L}\sum_{j=1}^{N}\left(\left\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}\right\rangle - \frac{\beta}{2L}\left\langle \mathbf{g}_j(z), \mathbf{x}_j(z)-\mathbf{x}\right\rangle^2\right) + \frac{6GD}{TN} + \frac{6\beta G^2D^2}{TN} \\
&\quad + D\sum_{z=1}^{T/L}\sum_{j=1}^{N}\left\|\mathbf{g}_j(z)-\overline{\mathbf{g}}(z)\right\|_2 \\
&= \sum_{z=1}^{T/L}\sum_{j=1}^{N}\left(\left\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}\right\rangle - \frac{\beta}{2L}\left\langle \mathbf{g}_j(z), \mathbf{x}_j(z)-\mathbf{x}\right\rangle^2\right) + \mathcal{O}\left(\frac{GD}{NT}+\frac{\beta G^2D^2}{NT}\right),
\end{aligned}
$$

where the last inequality follows from Eq. (16). □

It remains to control the block-level surrogate regrets appearing in Lemma 5. This step is more delicate than in the strongly convex case because the linear term and the quadratic correction are evaluated at two different points: the preparatory decision $\widehat{\mathbf{x}}_j(z)$ and the executed decision $\mathbf{x}_j(z)$. To handle this mismatch, we compare the delayed FTRL iterate $\hat{\mathbf{x}}_j(z)$, which only uses feedback up to block $z-3$, with a current regularized leader that uses feedback up to block z. The comparison is carried out in the adaptive second-order norm induced by the accumulated gradient matrix.

Lemma 2. *Under the same conditions as Algorithm 1. For any $j\in[N]$ and every comparator $\mathbf{x}\in\mathcal{X}$,*

$$
\sum_{z=1}^{T/L}\left(\left\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}\right\rangle - \tfrac{\beta}{2L}\left\langle \mathbf{g}_j(z), \mathbf{x}_j(z)-\mathbf{x}\right\rangle^2\right) \le \mathcal{O}\left(\eta\|\mathbf{x}\|_2^2 + \frac{d(1+\beta GD)L\log\left(1+\frac{\beta G^2T}{\eta d}\right)}{\beta}\right).
$$

Proof. For all $j\in[N]$, define $F^\star_{z,j}(\mathbf{x}) \triangleq \sum_{s=1}^{z}\langle \mathbf{g}_j(s),\mathbf{x}\rangle + \psi^\star_{z,j}(\mathbf{x})$ and let $\mathbf{x}^\star_j(z)\in\arg\min_{\mathbf{x}\in\mathcal{X}} F^\star_{z,j}(\mathbf{x})$, where $\psi^\star_{z,j}(\mathbf{x}) \triangleq \frac{\beta}{2L}\sum_{s=1}^{z}\langle \mathbf{g}_j(s),\mathbf{x}_j(s)-\mathbf{x}\rangle^2 + \frac{\eta}{2}\|\mathbf{x}\|_2^2$. For all $j\in[N]$, we have

$$
\sum_{z=1}^{T/L}\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}\rangle = \sum_{z=1}^{T/L}\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}^\star_j(z)\rangle + \sum_{z=1}^{T/L}\langle \mathbf{g}_j(z), \mathbf{x}^\star_j(z)-\mathbf{x}\rangle\ . \tag{20}
$$

Using Lemma 3, we have

$$
\sum_{z=1}^{T/L}\langle \mathbf{g}_j(z), \mathbf{x}^\star_j(z)-\mathbf{x}\rangle \le \psi^\star_{T/L,j}(\mathbf{x}) - \sum_{z=1}^{T/L}\left(\psi^\star_{z,j}\left(\mathbf{x}^\star_j(z)\right) - \psi^\star_{z-1,j}\left(\mathbf{x}^\star_j(z)\right)\right) \le \psi^\star_{T/L,j}(\mathbf{x}), \tag{21}
$$

where the last inequality follows from $\psi^\star_{z,j}(\mathbf{x}) - \psi^\star_{z-1,j}(\mathbf{x}) = \frac{\beta}{2L}\langle \mathbf{g}_j(z),\mathbf{x}_j(z)-\mathbf{x}\rangle^2 \ge 0$.

It remains to bound

$$
\sum_{z=1}^{T/L}\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z)-\mathbf{x}^\star_j(z)\rangle.
$$

For notational convenience, for any integer $z\le 0$, we set $A_j(z)\triangleq \eta I_d$, For any integer $z\ge 1$, define

$$
A_j(z) \triangleq \eta I_d + \frac{\beta}{L}\sum_{s=1}^{z}\mathbf{g}_j(s)\mathbf{g}_j(s)^\top.
$$

For each $z\ge 1$, We define

$$
\hat{F}_{z,j}(\mathbf{x}) \triangleq \sum_{s=1}^{z-3}\left(\langle \mathbf{g}_j(s),\mathbf{x}\rangle + \frac{\beta}{2L}\langle \mathbf{g}_j(s),\mathbf{x}_j(s)-\mathbf{x}\rangle^2\right) + \frac{\eta}{2}\|\mathbf{x}\|_2^2
$$

and $\hat{\mathbf{x}}_j(z) \in \arg\min_{\mathbf{x}\in\mathcal{X}} \hat{F}_{z,j}(\mathbf{x})$. Since $F^\star_{z,j}$ is 1-strongly convex with respect to $\|\cdot\|_{A_j(z)}$, and $\mathbf{x}^\star_j(z)$ minimizes $F^\star_{z,j}$ over $\mathcal{X}$, we have

$$F^\star_{z,j}(\hat{\mathbf{x}}_j(z)) - F^\star_{z,j}(\mathbf{x}^\star_j(z)) \geq \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z)}. \tag{22}$$

Likewise, since $\hat{F}_{z,j}$ is 1-strongly convex with respect to $\|\cdot\|_{A_j(z-3)}$, and $\hat{\mathbf{x}}_j(z)$ minimizes $\hat{F}_{z,j}$ over $\mathcal{X}$, we have

$$\hat{F}_{z,j}(\mathbf{x}^\star_j(z)) - \hat{F}_{z,j}(\hat{\mathbf{x}}_j(z)) \geq \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z-3)}. \tag{23}$$

Adding Eq. (22) and Eq. (23) is equivalent to applying the stability lemma of Qiu et al. [2025, Lemma A.2] to these two regularized objectives. This is a standard strong-convexity argument, not a new ingredient; we include the derivation to make the treatment of the two decisions explicit.

Summing Eq. (22) and Eq. (23), and expanding the definitions of $F^\star_{z,j}$ and $\hat{F}_{z,j}$, yields

$$\begin{aligned}
&\frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z)} + \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z-3)} \\
&\leq \sum_{s=z-2}^{z} \langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\rangle \\
&\quad + \frac{\beta}{2L}\sum_{s=z-2}^{z}\Big(\langle \mathbf{g}_j(s), \mathbf{x}_j(s) - \hat{\mathbf{x}}_j(z)\rangle^2 - \langle \mathbf{g}_j(s), \mathbf{x}_j(s) - \mathbf{x}^\star_j(z)\rangle^2\Big).
\end{aligned}$$

Using $a^2 - b^2 = (a-b)(a+b)$, with $a = \langle \mathbf{g}_j(s), \mathbf{x}_j(s) - \hat{\mathbf{x}}_j(z)\rangle$ and $b = \langle \mathbf{g}_j(s), \mathbf{x}_j(s) - \mathbf{x}^\star_j(z)\rangle$ we obtain

$$\begin{aligned}
&\frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z)} + \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z-3)} \\
&\leq \sum_{s=z-2}^{z} \big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\rangle\big| \\
&\quad + \frac{\beta}{2L}\sum_{s=z-2}^{z} \big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\rangle\big| \cdot \big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) + \mathbf{x}^\star_j(z) - 2\mathbf{x}_j(s)\rangle\big|.
\end{aligned}$$

Using Assumption 1 and Assumption 2,

$$\big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) + \mathbf{x}^\star_j(z) - 2\mathbf{x}_j(s)\rangle\big| \leq \|\mathbf{g}_j(s)\|_2 \, \|\hat{\mathbf{x}}_j(z) + \mathbf{x}^\star_j(z) - 2\mathbf{x}_j(s)\|_2 \leq 2LGD,$$

and therefore

$$\frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z)} + \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z-3)} \leq (1+\beta GD)\sum_{s=z-2}^{z}\big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\rangle\big|.$$

Applying Cauchy–Schwarz in the pair of dual norms induced by $A_j(z)$,

$$\big|\langle \mathbf{g}_j(s), \hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\rangle\big| \leq \|\mathbf{g}_j(s)\|_{A_j(z)^{-1}} \|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|_{A_j(z)},$$

we get

$$\begin{aligned}
&\frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z)} + \frac{1}{2}\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|^2_{A_j(z-3)} \\
&\leq (1+\beta GD)\left(\sum_{s=z-2}^{z}\|\mathbf{g}_j(s)\|_{A_j(z)^{-1}}\right)\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|_{A_j(z)}.
\end{aligned}$$

Dropping the second nonnegative term on the left-hand side gives

$$\|\hat{\mathbf{x}}_j(z) - \mathbf{x}^\star_j(z)\|_{A_j(z)} \leq 2\,(1+\beta GD)\sum_{s=z-2}^{z}\|\mathbf{g}_j(s)\|_{A_j(z)^{-1}}.$$

Consequently,

$$\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z) - \mathbf{x}_j^\star(z) \rangle \le \|\mathbf{g}_j(z)\|_{A_j(z)^{-1}} \|\hat{\mathbf{x}}_j(z) - \mathbf{x}_j^\star(z)\|_{A_j(z)} \le 2\,(1+\beta GD)\,\|\mathbf{g}_j(z)\|_{A_j(z)^{-1}} \sum_{s=z-2}^{z} \|\mathbf{g}_j(s)\|_{A_j(z)^{-1}}. \tag{24}$$

Summing Eq. (24) over $z = 1, \dots, T/L$, we obtain

$$\begin{aligned}
\sum_{z=1}^{T/L} \langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z) - \mathbf{x}_j^\star(z) \rangle &\le 2\,(1+\beta GD) \sum_{z=1}^{T/L} \left(\|\mathbf{g}_j(z)\|_{A_j(z)^{-1}} \sum_{s=z-2}^{z} \|\mathbf{g}_j(s)\|_{A_j(z)^{-1}} \right) \\
&\le (1+\beta GD) \sum_{z=1}^{T/L} \left(3\|\mathbf{g}_j(z)\|^2_{A_j(z)^{-1}} + \sum_{s=z-2}^{z} \|\mathbf{g}_j(s)\|^2_{A_j(z)^{-1}} \right) \\
&\qquad \text{(AM-GM inequality)} \\
&\le (1+\beta GD) \sum_{z=1}^{T/L} \left(3\|\mathbf{g}_j(z)\|^2_{A_j(z)^{-1}} + \sum_{s=z-2}^{z} \|\mathbf{g}_j(s)\|^2_{A_j(s)^{-1}} \right) \quad (A_j(s) \preceq A_j(z)) \\
&\le 6\,(1+\beta GD) \sum_{z=1}^{T/L} \|\mathbf{g}_j(z)\|^2_{A_j(z)^{-1}} \\
&\le \frac{6d\,(1+\beta GD)\,L}{\beta} \log\left(1 + \frac{\beta T G^2}{\eta d}\right).
\end{aligned} \tag{25}$$

where the last inequality holds by using Orabona [2026, Lemma 7.29].

Substituting Eq. (21) and Eq. (25) in to Eq. (20), we have

$$\begin{aligned}
&\sum_{z=1}^{T/L} \left(\langle \mathbf{g}_j(z), \hat{\mathbf{x}}(z) - \mathbf{x} \rangle - \frac{\beta}{2L} \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x} \rangle^2 \right) \\
&\le \psi^\star_{T/L,j}(\mathbf{x}) - \sum_{z=1}^{T/L} \frac{\beta}{2L} \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x} \rangle^2 + \frac{6d\,(1+\beta GD)\,L}{\beta} \log\left(1 + \frac{\beta T G^2}{\eta d}\right) \\
&= \frac{\eta}{2}\|\mathbf{x}\|_2^2 + \frac{6d\,(1+\beta GD)\,L}{\beta} \log\left(1 + \frac{\beta T G^2}{\eta d}\right) \\
&= \mathcal{O}\left(\eta \|\mathbf{x}\|_2^2 + \frac{d\,(1+\beta GD)\,L}{\beta} \log\left(1 + \frac{\beta T G^2}{\eta d}\right) \right).
\end{aligned} \tag{26}$$

□

We now prove Theorem 1 by combining the regret decomposition with the delayed ONS bound.

Theorem 1. *Suppose that the sequence of loss functions $\{f_{t,i}\}_{t\in[T],i\in[N]}$ satisfy Assumption 1, 2, and 3. Let P be the communication matrix defined in Definition 1. If every agent $i \in [N]$ runs Algorithm 1 with $\beta = \frac{1}{2}\min\{\frac{1}{4GD}, \alpha\}$ and with parameters θ and L chosen in Eq. (3), then, for any agent $i \in [N]$, the regret is bounded by:*

$$\mathrm{Reg}_{T,i}(\mathbf{x}) = \mathcal{O}\left(\eta N \|\mathbf{x}\|_2^2 + \frac{Nd(1+\beta GD)L}{\beta} \log\left(1 + \frac{\beta G^2 T}{\eta d}\right) \right). \tag{6}$$

Proof. Fix any agent $i \in [N]$ and any comparator $\mathbf{x} \in \mathcal{X}$. By Lemma 5,

$$\mathrm{Reg}_{T,i} \le \sum_{z=1}^{T/L} \sum_{j=1}^{N} \left(\langle \mathbf{g}_j(z), \hat{\mathbf{x}}_j(z) - \mathbf{x} \rangle - \frac{\beta}{2L} \langle \mathbf{g}_j(z), \mathbf{x}_j(z) - \mathbf{x} \rangle^2 \right) + \mathcal{O}\left(\frac{GD}{NT} + \frac{\beta G^2 D^2}{NT} \right).$$

Applying Lemma 2 to each $j \in [N]$ gives

$$\mathrm{Reg}_{T,i}(\mathbf{x}) = \mathcal{O}\left(\eta N \|\mathbf{x}\|_2^2 + \frac{Nd(1+\beta GD)L}{\beta} \log\left(1 + \frac{\beta G^2 T}{\eta d}\right) \right).$$

This proves Eq. (6). Substituting the choice of L in Eq. (3) gives the stated dependence on the spectral gap. □

B.1 Omitted Details in Section 3.3

We now turn to the proof of the lower bound in Theorem 2. For completeness, we restate the lower-bound theorem and provide its proof.

Theorem 2. *Let $N = 4K + 2$ for some positive integer K, and let $\mathcal{X} = \Delta_d$.*[5] *There exists a graph $\mathcal{G} = ([N], E)$ with $P = I - \frac{1}{2(d_{\max}+1)} L_{\mathcal{G}}$ such that there exists a sequence of local convex α-exp-concave losses $\{f_{t,i} : \mathcal{X} \to \mathbb{R}\}_{i\in[N],t\in[T]}$ such that $\mathbb{E}[\mathrm{Reg}_{T,1}] \geq \Omega\Big(\frac{Nd}{\alpha\sqrt{(1-\sigma_2(P))}} \log(T/N)\Big)$.*

Proof. We prove the lower bound on a 1-connected cycle graph. By definition, we have $K = \frac{N-2}{4}$. Let

$$Z = \left\lfloor \frac{T-1}{K} \right\rfloor, \qquad c_i = iK, \quad i = 0, \ldots, Z, \qquad c_{Z+1} = T.$$

We divide the horizon into intervals

$$[c_0 + 1, c_1],\ [c_1 + 1, c_2],\ \ldots,\ [c_{Z-1} + 1, c_Z],$$

each of length K, and one remaining final interval $[c_Z + 1, c_{Z+1}]$.

The loss functions for each agent $n \in [N]$ at each $t \in [c_i + 1, c_{i+1}]$ are as follows. Specifically, we choose $h_0, \ldots, h_{Z-1}$ as the wealth function in portfolio optimization defined in [Ordentlich and Cover, 1998, Theorem 1] with horizon $n = Z$ but scaled by $\frac{1}{\alpha}$. Then, for each $i = 0, \ldots, Z-1$ and each $t \in [c_i + 1, c_{i+1}]$, define the local losses on the cycle as follows:

$$\begin{aligned} f_{t,1}(x) = \cdots = f_{t,K}(x) = 0, \\ f_{t,K+1}(x) = \cdots = f_{t,N-K+1}(x) = h_i(x), \\ f_{t,N-K+2}(x) = \cdots = f_{t,N}(x) = 0. \end{aligned}$$

For the final interval $t \in [c_Z + 1, c_{Z+1}]$, set $f_{t,j}(x) = 0$ for every $j \in [N]$. Direct calculation shows that each local loss function is convex and α-exp-concave, since each local loss is either one of the hard losses h_i or the zero function.

For every $i = 0, \ldots, Z-1$ and every $t \in [c_i + 1, c_{i+1}]$, the global loss is

$$f_t(x) = \sum_{j=1}^{N} f_{t,j}(x) = S h_i(x),$$

where $S = N - 2K + 1 = 2K + 3$. In terms of N, since $K = (N-2)/4$, we know that

$$SK = \frac{N-2}{4}\left(\frac{N-2}{2} + 3\right) = \frac{(N-2)(N+4)}{8} \geq \frac{N^2}{8},$$

for every $N \geq 4$.

Now consider agent 1. On the 1-connected cycle, every node $n \in \{K+1, \ldots, N-K+1\}$ is at graph distance at least K from agent 1. Since the length of $[c_i + 1, c_{i+1}]$ is exactly K, agent 1 cannot receive information generated by the current hidden loss h_i before making the decisions $x_1(c_i + 1), \ldots, x_1(c_{i+1})$.

Now we define the averaged decision of agent 1 during interval $[c_i + 1, c_{i+1}]$ as $\overline{x}_1(i) = \frac{1}{K}\sum_{t=c_i+1}^{c_{i+1}} x_1(t)$.

Since the decision set $\mathcal{X}$ is convex, we have $\overline{x}_1(i) \in \mathcal{X}$. Moreover, $\overline{x}_1(i)$ is a valid online decision against the centralized loss sequence $h_0, \ldots, h_{i-1}$ at round i, since $\overline{x}_1(i)$ does not have the knowledge of h_i. By convexity of h_i, we know that

$$h_i(\overline{x}_1(i)) = h_i\left(\frac{1}{K}\sum_{t=c_i+1}^{c_{i+1}} x_1(t)\right) \leq \frac{1}{K}\sum_{t=c_i+1}^{c_{i+1}} h_i(x_1(t)).$$

[5] To align with Assumption 2 used in our upper bound, one can shift the simplex Δ_d by $\frac{1}{d}\mathbf{1}$ to include the origin. The α-exp-concavity is invariant under such constant shifts.

Multiplying both sides by SK and summing over $i = 0, \ldots, Z-1$, we obtain that

$$SK \sum_{i=0}^{Z-1} h_i(\overline{x}_i) \le \sum_{t=c_i+1}^{c_{i+1}} \sum_{i=0}^{Z-1} S h_i(x_1(t)) = \sum_{t=c_i+1}^{c_{i+1}} \sum_{i=0}^{Z-1} f_t(x_1(t)) = \sum_{t=1}^{T} f_t(x_1(t)).$$

On the other hand, for any comparator $x \in \mathcal{X}$, we have

$$\sum_{t=1}^{T} f_t(x) = \sum_{i=0}^{Z-1} \sum_{t=c_i+1}^{c_{i+1}} S h_i(x) = SK \sum_{i=0}^{Z-1} h_i(x).$$

Thus $\min_{x\in\mathcal{K}} \sum_{t=1}^{T} f_t(x) = SK \min_{x\in\mathcal{K}} \sum_{i=0}^{Z-1} h_i(x)$. Combining the two displays, we get

$$\mathrm{Reg}_{T,1} = \sum_{t=1}^{T} f_t(x_1(t)) - \min_{x\in\mathcal{K}} \sum_{t=1}^{T} f_t(x) \ge SK \left(\sum_{i=0}^{Z-1} h_i(\overline{x}_i) - \min_{x\in\mathcal{K}} \sum_{i=0}^{Z-1} h_i(x) \right).$$

Taking expectation over the sampled losses and any internal randomness of the distributed algorithm, we know that the loss sequence $h_0, \ldots, h_{Z-1}$ defined in Ordentlich and Cover [1998],

$$\mathbb{E}[R_{T,1}] \ge SK \cdot \Omega\left(\frac{d}{\alpha} \log(Z) \right).$$

Since $SK \ge N^2/8$, this implies that $\mathbb{E}[R_{T,1}] \ge \Omega\left(\frac{N^2 d}{\alpha} \log(Z) \right) = \Omega(\frac{N^2 d}{\alpha} \log(T/N))$. Finally, for the 1-connected cycle graph with N, communication matrix $P = I - \frac{1}{2(d_{\max}+1)} L_{\mathcal{G}}$, [Wan et al., 2024, Lemma 6] shows that $\frac{\pi^2}{1-\sigma_2(P)} \le 4N^2$, implying that

$$\mathbb{E}[R_{T,1}] \ge \Omega\left(\frac{Nd}{\alpha\sqrt{1-\sigma_2(P)}} \log(T/N) \right).$$

This completes the proof. □

C Omitted Details from Section 4

In this section, we show the omitted details in Section 4. For completeness, we restate the theorem and provide its proof.

Theorem 3. *Suppose Assumption 4 holds, and let P be the communication matrix defined in Definition 1. In D-OLR, if every agent $i \in [N]$ executes Algorithm 2 with $\lambda > 0$, θ chosen according to Eq. (3), and*

$$L = \left\lceil \frac{\sqrt{2}}{(\sqrt{2}-1)\sqrt{1-\sigma_2(P)}} \log\left(\sqrt{14}\, T^3 N^{7/2} (1 + H^2/\lambda) \right) \right\rceil, \tag{10}$$

then for any agent $i \in [N]$ and comparator $\mathbf{x} \in \mathbb{R}^d$, the regret is bounded by:

$$\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) = \mathcal{O}\left(\lambda \|\mathbf{x}\|_2^2 + \frac{Y^2 H^2}{\lambda} + Y^2 L N d \log\left(1 + \frac{NTH^2}{\lambda d} \right) \right). \tag{11}$$

Proof. For each block z and round $t \in \mathcal{T}_z$, we define the reference decision $\mathbf{x}^\star(z,t)$ as the minimizer of

$$F_{z,t}^\star(\mathbf{x}) = \frac{\lambda}{2}\|\mathbf{x}\|_2^2 + \sum_{s\in\mathcal{S}_z} \sum_{\tau\in\mathcal{T}_s} \ell_\tau(\mathbf{x}) + \sum_{\substack{\tau\in\mathcal{T}_z \\ \tau\le t}} \ell_\tau(\mathbf{x}), \qquad \ell_t(\mathbf{x}) = \sum_{j=1}^{N} f_{t,j}(\mathbf{x}).$$

This is the reference predictor defined in the proof sketch: it represents the ridge-regression decision with full knowledge of all observations assigned to the same copy through the end of round t.

With the definition of $\mathbf{x}^\star(z,t)$, we decompose the regret as follows:

$$\begin{aligned}\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) &= \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[\widetilde{f}_{t,j}^{\,i}(z,t) - f_{t,j}(\mathbf{x})\right] \\ &= \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[f_{t,j}(\mathbf{x}^\star(z,t)) - f_{t,j}(\mathbf{x})\right] \\ &\quad + \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[\widetilde{f}_{t,j}^{\,i}(z,t) - f_{t,j}(\mathbf{x}^\star(z,t))\right], \end{aligned} \tag{27}$$

where $\widetilde{f}_{t,j}^{\,i}(z,t) \triangleq \frac{1}{2}(\widehat{y}_{t,j}^{\,i}(z,t) - y_{t,j})^2$ and $\widehat{y}_{t,j}^{\,i}(z,t) \triangleq \Pi_{[-Y,Y]}(\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t)\rangle)$.

Consider the first summation in Eq. (27). Within each copy, the reference decisions minimize the cumulative squared loss plus $\lambda\|\mathbf{x}\|_2^2/2$. Applying Lemma 3 separately to the two subsequences, with this regularizer after initialization and zero initial regularizer, gives

$$\sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[f_{t,j}(\mathbf{x}^\star(z,t)) - f_{t,j}(\mathbf{x})\right] \le \lambda\|\mathbf{x}\|_2^2. \tag{28}$$

Consequently,

$$\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) \le \lambda\|\mathbf{x}\|_2^2 + \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[\widetilde{f}_{t,j}^{\,i}(z,t) - f_{t,j}(\mathbf{x}^\star(z,t))\right]. \tag{29}$$

We next explain the role of clipping before controlling the delay and gossip terms. For notational convenience, denote

$$p_{t,j}^{i}(z,t) \triangleq \langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t)\rangle, \qquad p_{t,j}^{\star}(z,t) \triangleq \langle \mathbf{h}_{t,j}, \mathbf{x}^\star(z,t)\rangle.$$

For $|y| \le Y$, define the clipped scalar loss $\phi_y(u) \triangleq \frac{1}{2}(\Pi_{[-Y,Y]}(u) - y)^2$. Since $y \in [-Y,Y]$, projection onto $[-Y,Y]$ does not increase the squared distance to y, and hence

$$\phi_y(u) \le \frac{1}{2}(u-y)^2, \qquad \forall u \in \mathbb{R}.$$

Moreover, ϕ_y is $2Y$-Lipschitz: for all $u, v \in \mathbb{R}$,

$$\begin{aligned}|\phi_y(u) - \phi_y(v)| &= \frac{1}{2}\left|\Pi_{[-Y,Y]}(u) - \Pi_{[-Y,Y]}(v)\right|\left|\Pi_{[-Y,Y]}(u) + \Pi_{[-Y,Y]}(v) - 2y\right| \\ &\le \frac{1}{2}|u-v|\cdot 4Y = 2Y|u-v|.\end{aligned}$$

Therefore, we have

$$\begin{aligned}\widetilde{f}_{t,j}^{\,i}(z,t) - f_{t,j}(\mathbf{x}^\star(z,t)) &= \phi_{y_{t,j}}(p_{t,j}^{i}(z,t)) - \frac{1}{2}(p_{t,j}^{\star}(z,t) - y_{t,j})^2 \\ &\le \phi_{y_{t,j}}(p_{t,j}^{i}(z,t)) - \phi_{y_{t,j}}(p_{t,j}^{\star}(z,t)) \\ &\le 2Y|p_{t,j}^{i}(z,t) - p_{t,j}^{\star}(z,t)| \\ &= 2Y\left|\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) - \mathbf{x}^\star(z,t)\rangle\right|. \end{aligned} \tag{30}$$

Thus, the squared-loss comparison does not require a bound on either model or its raw prediction. We will use this Lipschitz property for the gossip error. For the delay of both sufficient statistics, we retain the quadratic information in the regularized objectives, as detailed next.

Define the auxiliary exact delayed predictor by

$$\mathbf{x}'(z,t) = A(z)^{-1}b(z), \qquad A(z) = \lambda I_d + \sum_{s\in\mathcal{S}_z}\sum_{j=1}^{N} C_j^0(s), \qquad b(z) = \sum_{s\in\mathcal{S}_z}\sum_{j=1}^{N} Y_j^0(s). \tag{31}$$

It uses exactly the blocks available to the corresponding copy, with exact averaging in place of gossip. In particular, it is fixed throughout block z. We separate the regret into the regret of this exact delayed predictor and its discrepancy from the decentralized predictor:

$$\begin{aligned}\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) = &\sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[\phi_{y_{t,j}}(\langle \mathbf{h}_{t,j}, \mathbf{x}'(z,t)\rangle) - f_{t,j}(\mathbf{x})\right] \\ &+ \sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left[\widetilde{f}_{t,j}^{\,i}(z,t) - \phi_{y_{t,j}}(\langle \mathbf{h}_{t,j}, \mathbf{x}'(z,t)\rangle)\right]. \end{aligned} \tag{32}$$

The first term accounts for the delay of the current block. By the clipping argument above, the second term is at most

$$2Y\sum_{z=1}^{\lceil T/L\rceil}\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\left|\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) - \mathbf{x}'(z,t)\rangle\right|. \tag{33}$$

We first consider the exact delayed predictor. Fix a block z and its assigned copy, and define the regularized objectives immediately before and after this block by

$$F_z^-(\mathbf{x}) = \frac{\lambda}{2}\|\mathbf{x}\|_2^2 + \sum_{s\in\mathcal{S}_z}\sum_{\tau\in\mathcal{T}_s}\ell_\tau(\mathbf{x}), \qquad F_z^+(\mathbf{x}) = F_z^-(\mathbf{x}) + \sum_{t\in\mathcal{T}_z}\ell_t(\mathbf{x}).$$

Let $M_z^- \triangleq \min_{\mathbf{x}} F_z^-(\mathbf{x})$ and $M_z^+ \triangleq \min_{\mathbf{x}} F_z^+(\mathbf{x})$. The minimizer of F_z^- is $\mathbf{x}'(z,t)$, whereas that of F_z^+ is the reference predictor at the last round of the block. Successive minimum increments telescope within each copy. Since each final minimum is at most the final objective evaluated at any comparator $\mathbf{x}$, we have

$$\sum_{z=1}^{\lceil T/L\rceil}(M_z^+ - M_z^-) \le \lambda\|\mathbf{x}\|_2^2 + \sum_{t=1}^{T}\sum_{j=1}^{N} f_{t,j}(\mathbf{x}). \tag{34}$$

This is the block form of the centralized regularized-leader comparison. Retaining these minimum increments lets us account for the delay in the covariance as well as in the labels.

To bound the loss within a block, stack its $N|\mathcal{T}_z|$ feature vectors into the rows of B_z, and its labels into $\mathbf{y}_z$. Completing the square around the minimizer $\mathbf{x}'(z,t)$ gives

$$\begin{aligned} M_z^+ - M_z^- &= \min_{\mathbf{v}}\left\{\frac{1}{2}\|\mathbf{v}\|_{A(z)}^2 + \frac{1}{2}\|B_z\mathbf{x}'(z,t) + B_z\mathbf{v} - \mathbf{y}_z\|_2^2\right\} \\ &= \frac{1}{2}(B_z\mathbf{x}'(z,t) - \mathbf{y}_z)^\top\left(I + B_zA(z)^{-1}B_z^\top\right)^{-1}(B_z\mathbf{x}'(z,t) - \mathbf{y}_z). \end{aligned}$$

The second equality follows by minimizing the quadratic and applying the matrix-inversion identity. Since the smallest eigenvalue of the inverse matrix is $1/(1 + \|B_zA(z)^{-1}B_z^\top\|_{\mathrm{op}})$, and clipping cannot increase the squared loss,

$$M_z^+ - M_z^- \ge \frac{\displaystyle\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\phi_{y_{t,j}}(\langle \mathbf{h}_{t,j}, \mathbf{x}'(z,t)\rangle)}{1 + \|B_zA(z)^{-1}B_z^\top\|_{\mathrm{op}}}.$$

Each clipped squared loss is at most $2Y^2$, and the block contains at most LN examples. Consequently,

$$\begin{aligned} &\sum_{t\in\mathcal{T}_z}\sum_{j=1}^{N}\phi_{y_{t,j}}(\langle \mathbf{h}_{t,j}, \mathbf{x}'(z,t)\rangle) - (M_z^+ - M_z^-) \\ &\le 2Y^2LN\frac{\|B_zA(z)^{-1}B_z^\top\|_{\mathrm{op}}}{1 + \|B_zA(z)^{-1}B_z^\top\|_{\mathrm{op}}} \\ &\le 2Y^2LN\log\left(\frac{\det(A(z) + B_z^\top B_z)}{\det A(z)}\right). \end{aligned} \tag{35}$$

For the last inequality, use the nonnegative eigenvalues of $B_z A(z)^{-1} B_z^\top$ and the determinant identity $\det(I + B_z A(z)^{-1} B_z^\top) = \det(A(z) + B_z^\top B_z)/\det A(z)$.

It remains to sum the log-determinant increases. Within each copy, $A(z) + B_z^\top B_z$ is precisely the covariance matrix before its next assigned block, so these increases telescope. The trace of either copy's final unregularized covariance is at most NTH^2. The arithmetic–geometric mean inequality for the eigenvalues therefore gives, for each copy, a total log-determinant increase at most $d \log(1 + NTH^2/(\lambda d))$. This is the elliptical-potential argument at block level; the factor LN accounts for the examples predicted before the copy receives its next statistics. Combining both copies with Eq. (34) and Eq. (35) yields

$$\sum_{z=1}^{\lceil T/L \rceil} \sum_{t \in \mathcal{T}_z} \sum_{j=1}^{N} \left[\phi_{y_{t,j}}(\langle \mathbf{h}_{t,j}, \mathbf{x}'(z,t) \rangle) - f_{t,j}(\mathbf{x}) \right] \le \lambda \|\mathbf{x}\|_2^2 + 4Y^2 LNd \log\left(1 + \frac{NTH^2}{\lambda d}\right). \tag{36}$$

We now consider the gossip term in Eq. (33). Recall that

$$\widehat{A}_i(z) = \lambda I_d + N \sum_{s \in \mathcal{S}_z} \Pi_{\mathbb{S}_+^d}(C_i(s)), \qquad \widehat{b}_i(z) = N \sum_{s \in \mathcal{S}_z} Y_i(s), \qquad \mathbf{x}_i(z,t) = \widehat{A}_i(z)^{-1} \widehat{b}_i(z).$$

Both $\widehat{A}_i(z)$ and $A(z)$ dominate λI_d. Subtracting the exact delayed predictor and applying the matrix-inverse identity gives

$$\begin{aligned} \mathbf{x}_i(z,t) - \mathbf{x}'(z,t) &= \widehat{A}_i(z)^{-1}(\widehat{b}_i(z) - b(z)) + \big(\widehat{A}_i(z)^{-1} - A(z)^{-1}\big) b(z) \\ &= \widehat{A}_i(z)^{-1}(\widehat{b}_i(z) - b(z)) - \widehat{A}_i(z)^{-1}(\widehat{A}_i(z) - A(z)) A(z)^{-1} b(z). \end{aligned}$$

Therefore, by Cauchy–Schwarz,

$$|\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) - \mathbf{x}'(z,t) \rangle| \le \|\mathbf{h}_{t,j}\|_2 \left(\frac{\|\widehat{b}_i(z) - b(z)\|_2}{\lambda} + \frac{\|\widehat{A}_i(z) - A(z)\|_{\mathrm{op}} \|b(z)\|_2}{\lambda^2} \right). \tag{37}$$

Compared with the label-statistic calculation, the second term accounts for gossiping the covariance matrix.

By Assumption 4, $|y_{\tau,i}| \le Y$ and $\|\mathbf{h}_{\tau,i}\|_2 \le H$. Hence,

$$\|Y_i^0(s)\|_2 \le \sum_{\tau \in \mathcal{T}_s} |y_{\tau,i}| \|\mathbf{h}_{\tau,i}\|_2 \le LYH, \qquad \|C_i^0(s)\|_F \le \sum_{\tau \in \mathcal{T}_s} \|\mathbf{h}_{\tau,i} \mathbf{h}_{\tau,i}^\top\|_F \le LH^2.$$

Apply the accelerated-gossip contraction used in the proof of Lemma 1, also to the vectorized covariance matrices. The choice of L in Eq. (10) makes the contraction factor at most $1/[T^3 N^{7/2}(1 + H^2/\lambda)]$. The initial distance to the average is at most $2\sqrt{N}$ times the displayed bound on an individual statistic. Thus, for every completed block s,

$$\begin{aligned} \left\| Y_i(s) - \frac{1}{N} \sum_{j=1}^{N} Y_j^0(s) \right\|_2 &\le \frac{2LYH}{T^3 N^3 (1 + H^2/\lambda)}, \\ \left\| C_i(s) - \frac{1}{N} \sum_{j=1}^{N} C_j^0(s) \right\|_F &\le \frac{2LH^2}{T^3 N^3 (1 + H^2/\lambda)}. \end{aligned}$$

The exact average covariance is positive semidefinite. Non-expansiveness of its metric projection in Frobenius norm therefore gives

$$\left\| \Pi_{\mathbb{S}_+^d}(C_i(s)) - \frac{1}{N} \sum_{j=1}^{N} C_j^0(s) \right\|_F \le \left\| C_i(s) - \frac{1}{N} \sum_{j=1}^{N} C_j^0(s) \right\|_F.$$

All blocks in $\mathcal{S}_z$ are complete, and $L|\mathcal{S}_z| \le T$. Summing the errors of completed blocks yields

$$\|\widehat{b}_i(z) - b(z)\|_2 \le N \sum_{s \in \mathcal{S}_z} \left\| Y_i(s) - \frac{1}{N} \sum_{j=1}^{N} Y_j^0(s) \right\|_2 \le \frac{2YH}{T^2 N^2 (1 + H^2/\lambda)},$$

$$\|\widehat{A}_i(z) - A(z)\|_{\mathrm{op}} \le N \sum_{s\in\mathcal{S}_z} \left\| \Pi_{\mathbb{S}^d_+}(C_i(s)) - \frac{1}{N}\sum_{j=1}^{N} C_j^0(s) \right\|_F \le \frac{2H^2}{T^2N^2(1+H^2/\lambda)}.$$

Moreover, $\|b(z)\|_2 \le NTYH$. Substituting these estimates into Eq. (37), and using $TN \ge 1$, gives

$$|\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) - \mathbf{x}'(z,t)\rangle| \le \frac{2YH}{\lambda TN}\|\mathbf{h}_{t,j}\|_2.$$

Consequently,

$$\begin{aligned} 2Y \sum_{z=1}^{\lceil T/L\rceil} \sum_{t\in\mathcal{T}_z} \sum_{j=1}^{N} |\langle \mathbf{h}_{t,j}, \mathbf{x}_i(z,t) - \mathbf{x}'(z,t)\rangle| &\le \frac{4Y^2H}{\lambda TN} \sum_{z=1}^{\lceil T/L\rceil} \sum_{t\in\mathcal{T}_z} \sum_{j=1}^{N} \|\mathbf{h}_{t,j}\|_2 \\ &\le \frac{4Y^2H^2}{\lambda}. \end{aligned} \tag{38}$$

The last inequality uses $\|\mathbf{h}_{t,j}\|_2 \le H$ and the fact that the blocks partition $[T]$.

Combining Eq. (32), Eq. (33), Eq. (36), and Eq. (38), we have

$$\mathrm{Reg}_{T,i}^{\mathrm{OLR}}(\mathbf{x}) \le \lambda\|\mathbf{x}\|_2^2 + \frac{4Y^2H^2}{\lambda} + 4Y^2LNd\log\left(1 + \frac{NTH^2}{\lambda d}\right), \tag{39}$$

which proves Theorem 3. □