

TIGHTER PERFORMANCE THEORY OF FEDEXPROX

Anonymous authors

Paper under double-blind review

ABSTRACT

We revisit FedExProx – a recently proposed distributed optimization method designed to enhance convergence properties of parallel proximal algorithms via extrapolation. In the process, we uncover a surprising flaw: its known theoretical guarantees on quadratic optimization tasks are no better than those offered by the vanilla Gradient Descent (GD) method. Motivated by this observation, we develop a novel analysis framework, establishing a tighter linear convergence rate for non-strongly convex quadratic problems. By incorporating both computation and communication costs, we demonstrate that FedExProx can indeed provably outperform GD, in stark contrast to the original analysis. Furthermore, we consider partial participation scenarios and analyze two adaptive extrapolation strategies – based on gradient diversity and Polyak stepsizes — again significantly outperforming previous results. Moving beyond quadratics, we extend the applicability of our analysis to general functions satisfying the Polyak-Łojasiewicz condition, outperforming the previous strongly convex analysis while operating under weaker assumptions. Backed by empirical results, our findings point to a new and stronger potential of FedExProx, paving the way for further exploration of the benefits of extrapolation in federated learning.

1 INTRODUCTION

Federated Learning (FL) (Konečný et al., 2016; McMahan et al., 2017) is a distributed machine learning paradigm where multiple edge devices collaboratively train a global model without the need for centralized data collection. Instead of transmitting raw data, each client computes updates based on their local datasets, and these are then aggregated at a central server to update the global model. This approach ensures that no sensitive information is shared with the orchestrating server or other clients, making it particularly useful in privacy-sensitive domains, such as healthcare (Rieke et al., 2020) and recommendation systems (Hard et al., 2018). Formally, FL solves the following optimization problem:

$$\min_{x \in \mathbb{R}^d} \left\{ f(x) := \frac{1}{n} \sum_{i=1}^n f_i(x) \right\}, \quad (1)$$

where $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ represents the empirical risk associated with the local dataset stored on client $i \in [n]$. Despite its advantages, FL introduces several nontrivial challenges, including *communication delays* caused by limited network bandwidth, and *partial participation* of clients due to random outages (Kairouz et al., 2021). Addressing these issues is essential to making FL efficient and reliable in real-world applications.

1.1 RELATED WORK

Let us briefly review the existing methods used to solve (1) in FL scenarios. For a detailed overview of the notation used, please refer to Table 1.

GD. The simplest and most naive approach is the *Gradient Descent* (GD) method (Nesterov, 2018), iterating $x_{k+1} = x_k - \gamma \nabla f(x_k) = x_k - \gamma \frac{1}{n} \sum_{i=1}^n \nabla f_i(x_k)$, where x_0 is a starting point and γ is the stepsize. Let us introduce the key assumptions typically employed in the analysis of GD in the convex world.

Assumption 1.1. The functions f_i are proper, closed and convex for all $i \in [n]$, and the function f attains a minimum at some (potentially non-unique) point $x_* \in \mathbb{R}^d$.

Assumption 1.2. The function f is L -smooth, i.e., $\|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\| \forall x, y \in \mathbb{R}^d$. Additionally, each $f_i, i \in [n]$, is differentiable and L_i -smooth, with $L_{\max} := \max_{i \in [n]} L_i$.

Under these assumptions, classical optimization theory guarantees that GD solves problem (1) and finds $\bar{x} \in \mathbb{R}^d$ such that $f(\bar{x}) - f(x_*) \leq \varepsilon$ in $\mathcal{O}(LR^2/\varepsilon)$ iterations, where $R^2 := \|x_* - x_0\|^2$. In an FL setting, finding a solution to the main problem requires each client to participate in every communication round, and thus to communicate with the server

$$\mathcal{O}\left(\frac{LR^2}{\varepsilon}\right) \quad (2)$$

times, transmitting the local gradient $\nabla f_i(x_k)$ at each iteration k . While this approach is conceptually simple, its direct application in distributed environments leads to several challenges related to communication overhead and scalability.

Partial Participation. In real-world FL scenarios, not all clients remain active for the entire duration of training. Instead, a subset of clients is chosen at each iteration, typically based on practical factors such as device availability (e.g., battery life or network conditions) or statistical considerations (e.g., data heterogeneity) (Kairouz et al., 2021; Tyurin & Richtárik, 2023).

Communication bottleneck. Another key problem in distributed training is the *communication bottleneck* (Ramesh et al., 2021; Kairouz et al., 2021). Since the overall performance of a distributed algorithm is the product of the number of communication rounds K needed to find a solution and the cost C of one such round, there exist two main strategies to addressing this issue: (i) minimizing C and/or (ii) minimizing K . Objective (i) is typically achieved via compression of the information transmitted between the clients and the server (Beznosikov et al., 2020; Gruntkowska et al., 2024). Objective (ii) can be accomplished by increasing the amount of local computation performed by clients between rounds, allowing for less frequent communication.

FedAvg and modern local methods. A common strategy for ensuring communication efficiency that falls into category (ii) is *local training*. The idea is simple: allow clients to do more work before transmitting the results to the server. If executed with care, this approach leads to each communication round providing more “informative” updates, ultimately resulting in more effective changes to the global model. One of the most popular methods in this class is Federated Averaging (FedAvg), introduced by McMahan et al. (2017). Within this framework, a subset of clients is selected in each round to perform local training using a gradient-based method of their choice (e.g., GD). This is followed by an averaging step on the server, where the local updates are aggregated. Despite its popularity, the theoretical properties of FedAvg are somewhat limited (Khaled et al., 2019; Koloskova et al., 2020). A major issue is *client drift*, which can lead to slower convergence (Karimireddy et al., 2020). This problem can be mitigated using modern drift correction techniques (Karimireddy et al., 2020; Gorbunov et al., 2021; Mishchenko et al., 2022). One notable advancement in this area is Fed-Exp, introduced by Jhunjunwala et al. (2023), which extends FedAvg by leveraging extrapolation to better control server updates.

FedProx and proximal methods. All methods discussed so far assume that clients have access to the gradients $\{\nabla f_i\}$ of local functions. A recent line of *proximal methods* (Bertsekas, 2011; Ryu & Boyd, 2014; Khaled & Jin, 2022; Richtárik et al., 2024) relies on a different, more powerful oracle – the *proximal operator* $\text{prox}_{\gamma f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, defined by

$$\text{prox}_{\gamma f}(x) = \arg \min_{z \in \mathbb{R}^d} \left\{ f(z) + \frac{1}{2\gamma} \|z - x\|^2 \right\}. \quad (3)$$

Here, f is a convex function and γ plays a role of a regularization parameter. It is important to note that computing $\text{prox}_{\gamma f}(\cdot)$ is an optimization problem on its own. However, the analyses of such methods essentially abstract away from the way of running the local operations, assuming that computing $\text{prox}_{\gamma f}(\cdot)$ is computationally cheap and can be performed using any solver, e.g., GD, Newton’s method or LBFGS (Nesterov, 2018; Liu & Nocedal, 1989). There exist numerous proximal methods addressing the single-node case. A classic approach is the Proximal Point Method (PPM) by Rockafellar (1976), which iterates $x_{k+1} = \text{prox}_{\gamma_k f}(x_k)$. Stochastic versions of PPM have also been studied (Ryu & Boyd, 2014; Bianchi, 2016; Patrascu & Necoara, 2018). Recently, interest in stochastic proximal methods has surged, with promising results from several works (Condat & Richtárik, 2023; Traoré et al., 2024; Sadiev et al., 2024; Combettes & Madariaga, 2024). In the context of distributed learning, one relevant work by Li et al. (2020) builds on the principles of FedAvg

to introduce its proximal variant – FedProx. The method replaces local gradient-based updates with proximal operator computations, leading to the update rule

$$x_{k+1} = \frac{1}{n} \sum_{i=1}^n \text{prox}_{\gamma f_i}(x_k). \quad (4)$$

By leveraging proximal operators, FedProx addresses some of the limitations of traditional GD in FL. However, simply averaging local updates at the server remains suboptimal and can be improved.

FedExProx. A recent advancement, and the main focus of this paper, is FedExProx, introduced by Li et al. (2024). This method improves on FedProx by employing the technique of *extrapolation*¹ (Combettes, 1997; Necoara et al., 2019). The algorithm operates under the *interpolation regime*:

Assumption 1.3. There exists $x_* \in \mathbb{R}^d$ such that $\nabla f_i(x_*) = 0$ for all $i \in [n]$. We denote the set of such minimizers by $\mathcal{X}_* := \{x \in \mathbb{R}^d : \nabla f_i(x) = 0\}$ and the projection of x onto \mathcal{X}_* by $\Pi(x)$.

This assumption is relatively mild and is often met in practice, e.g., when working with over-parameterized models (Ma et al., 2018). Instead of (4), FedExProx introduces the extrapolated update

$$x_{k+1} = x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^n \text{prox}_{\gamma f_i}(x_k) - x_k \right), \quad (5)$$

where α_k serves as the *extrapolation parameter*. The algorithm is analyzed by reformulating step (5) in terms of *Moreau envelopes* (Moreau, 1965): the update rule can be expressed equivalently as

$$x_{k+1} = x_k - \alpha_k \gamma \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k), \quad (6)$$

where the Moreau envelope $M_f^\gamma : \mathbb{R}^d \rightarrow \mathbb{R}$ of a convex function f is defined as

$$M_f^\gamma(x) = \min_{z \in \mathbb{R}^d} \left\{ f(z) + \frac{1}{2\gamma} \|z - x\|^2 \right\} \quad \forall x \in \mathbb{R}^d. \quad (7)$$

Notably, when $\alpha_k = 1$, (5) reduces to (4). However, this choice is not optimal, and choosing $\alpha_k > 1$ yields better complexity. If α_k is set to be constant across iterations of the algorithm, then using the optimal value $\alpha_k \equiv \alpha = \frac{1}{\gamma L_\gamma} > 1$ guarantees that FedExProx converges after

$$\mathcal{O} \left(\frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \right) \quad (8)$$

communication rounds in the convex case, where L_γ is the smoothness constant of the function $M^\gamma(x) := \frac{1}{n} \sum_{i=1}^n M_{f_i}^\gamma(x)$. Li et al. (2024) show that (8) is not larger than $\mathcal{O}(L_{\max}R^2/\varepsilon)$, and can be significantly smaller in many practical scenarios. Thus, FedExProx offers provably better iteration complexity compared to both FedProx and FedExP (Li et al., 2020; Jhunjunwala et al., 2023).

2 CONTRIBUTIONS

1. Although the iteration complexity (8) of FedExProx outperforms both FedProx and FedExP, we find that it is *not superior* to that of the simplest baseline, GD, on quadratic optimization tasks (see Remark 3.3). This raises a critical question regarding the utility of employing the proximal oracle in the first place.

2. Motivated by these pessimistic findings, we revisit the analysis of FedProx and prove a significantly more optimistic iteration complexity result. We establish a linear convergence rate for solving non-strongly convex quadratic distributed optimization problems of the form (1) (see Theorem 4.1) and demonstrate that our new result does indeed lead to the conclusion that FedExProx can significantly outperform GD. Our analysis assumes a realistic model that accounts for both computation and communication times – critical factors in real-world distributed optimization. We show that the total *time complexity* of FedExProx is never worse than that of GD, and can be strictly better when communication time dominates computation time (see Theorem 4.3), which is typically the case in FL scenarios. This stands in stark contrast to the previous analysis by Li et al. (2024), underscoring the superiority of our new findings.

¹Notably, the idea of extrapolation was explored by Jhunjunwala et al. (2023) in developing FedExP.

3. To account for the stochastic setting that involves partial participation of clients, we complement the above results with client sampling (see Theorems 5.1, 5.2, 7.5 and F.1).

4. Beyond constant extrapolation, we establish two novel results (Theorems 6.1 and F.1) that incorporate smoothness-adaptive strategies, based on gradient diversity (FedExProx-GraDS), and Polyak stepsize (FedExProx-StoPS), again significantly improving upon the result previously established by Li et al. (2024) in the quadratic case.

5. We extend the analysis beyond quadratics to functions satisfying the Polyak-Łojasiewicz condition, obtaining a linear convergence rate (see Theorems 7.2 and 7.5). In contrast, Li et al. (2024) derived a linear rate under strong convexity for the function f . Our approach is not only more general, as it relies on weaker assumptions, but also demonstrates improved dependence on problem-specific constants. Additionally, we establish a result in the PŁ setting that accounts for inexact computations of the proximal mappings (Theorem E.1).

6. The theoretical findings are validated with empirical experiments, which demonstrate the robustness and applicability of our framework.

3 NOT BETTER THAN GD ON QUADRATICS

Let us take a closer look at the complexity result (8). This complexity depends on the regularization parameter γ introduced in definition (3). Suppose for now that we solve (3) using GD. Then, the number of iterations needed to find $\text{prox}_{\gamma f_i}(x)$ with accuracy ε is

$$\mathcal{O}\left(\frac{L_i+1/\gamma}{1/\gamma} \log \frac{1}{\varepsilon}\right) = \mathcal{O}\left((\gamma L_i + 1) \log \frac{1}{\varepsilon}\right),$$

since the function $f_i(z) + \frac{1}{2\gamma} \|z - x\|^2$ is $(L_i + 1/\gamma)$ -smooth and $1/\gamma$ -strongly convex. Hence, γ controls the difficulty of calculating the proximal operator: the larger it is, the more difficult the problem. Therefore, to accelerate local computations, one would prefer to choose γ as small as possible. We can formalize this intuition by letting $\pi(\gamma)$ be the time per one iteration of FedExProx and making the following reasonable assumption:

Assumption 3.1. The time complexity $\pi(\gamma)$ of a FedExProx step is a non-decreasing function of γ .

Combining it with the iteration complexity result (8) established by Li et al. (2024), the total time required by FedExProx with $\alpha_k \equiv 1/\gamma L_\gamma$ to find \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ is

$$T(\gamma) := \pi(\gamma) \times \frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon}.$$

With this definition in place, we can now translate the main result of Li et al. (2024) (Theorem 1) to quadratic optimization problems, yielding the following pessimistic outcome:

Theorem 3.2. *Let Assumptions 1.3 and 3.1 hold. Consider solving a non-strongly convex quadratic optimization problem of the form (1), where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Then the assumptions of Theorem 1 by Li et al. (2024) hold, and*

$$T(\gamma) = \pi(\gamma) \times \frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \geq \pi(0) \times \frac{LR^2}{\varepsilon} \quad (9)$$

for all $\gamma > 0$. Moreover, when $\gamma \rightarrow 0$, then $\pi(\gamma) \times \frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \rightarrow \pi(0) \times \frac{LR^2}{\varepsilon}$, and FedExProx effectively reduces to GD.

Remark 3.3. In light of Theorem 3.2, GD performs no worse than FedExProx by Li et al. (2024), even if the time complexity $\pi(\gamma)$ of one FedExProx step does not depend on γ . Indeed, it holds that

$$\frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \geq \frac{LR^2}{\varepsilon}$$

for all $\gamma > 0$ (see the proof of Theorem 3.2).

These results, based on the original analysis of FedExProx, lead us to a rather disappointing conclusion: from a theoretical time complexity perspective, the optimal strategy appears to be using vanilla GD, disregarding proximal oracles entirely. The question that remains is: is this an inherent limitation of the method itself or merely a consequence of suboptimal analysis by Li et al. (2024)?

4 TIGHTER AND MORE OPTIMISTIC RESULTS

Are we truly constrained to GD, or is there room for improvement? It turns out that a tighter iteration complexity bound for FedExProx can be derived, revealing significantly more optimistic results. The following theorem presents a refined analysis, providing a new, improved complexity result for FedExProx and shedding light on its true capabilities.

Theorem 4.1. Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization problem (1) where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Under Assumption 1.3, FedExProx with $\alpha = 1/\gamma L_\gamma$ finds \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{L_\gamma}{\mu_\gamma^+} \log \frac{1}{\varepsilon}\right) \quad (10)$$

iterations, where L_γ is a smoothness constant of M^γ and μ_γ^+ is the smallest non-zero eigenvalue of the matrix $\nabla^2 M^\gamma$.

Remark 4.2. To the best of our knowledge, under the assumptions of Theorem 4.1, vanilla GD requires

$$\mathcal{O}\left(\frac{L}{\mu^+} \log \frac{1}{\varepsilon}\right) \quad (11)$$

iterations to solve the quadratic optimization problem, where μ^+ is the smallest non-zero eigenvalue of the matrix $\mathbf{A} = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$ (Richtárik & Takác, 2020).

Let us now demonstrate that our new result does indeed provide a tighter bound, leading to the conclusion that FedExProx can in fact outperform GD.

4.1 ONE STEP TIME COMPLEXITY $\pi(\gamma)$ IS $\mu + \tau(\gamma L_{\max} + 1)$.

As introduced in Section 1, the time per one global iteration of FedExProx has two main sources:

1. **Local computation:** In large-scale problems, each step (5) requires clients to compute $\text{prox}_{\gamma f_i}(x_k)$ iteratively. One of the simplest solvers is GD, which returns a solution of subproblem i after $\tilde{\mathcal{O}}(\gamma L_i + 1)$ local iterations² (see Section 3). If each gradient calculation takes τ seconds, the total time required for all clients to calculate $\text{prox}_{\gamma f_i}(x_k)$ is $\tilde{\mathcal{O}}(\tau \times (\gamma L_{\max} + 1))$, since the process is gated by the “slowest” client, associated with the problem with the highest smoothness constant.

2. **Communication:** Once the local computations are completed, clients must communicate their results before the server can execute the global step (5). We assume this communication takes μ seconds, which can be huge in FL environments (Kairouz et al., 2021).

Consequently, the total time per global iteration, $\pi(\gamma)$, is proportional to $\mu + \tau(\gamma L_{\max} + 1)$ (thus satisfying Assumption 3.1), and the total time complexity of FedExProx is given by

$$T_\mu(\gamma) := \tilde{\mathcal{O}}\left((\mu + \tau(\gamma L_{\max} + 1)) \times \frac{L_\gamma}{\mu_\gamma^+}\right). \quad (12)$$

Note that the total time complexity of GD is

$$T_{\text{GD}} := T_\mu(0) = \tilde{\mathcal{O}}\left((\mu + \tau) \times \frac{L}{\mu^+}\right).$$

Our next goal is to determine the value of γ that minimizes $T_\mu(\gamma)$, with the hope that this time the optimal choice does not result in $\gamma \rightarrow 0$. This is indeed the case whenever $\mu > \tau$.

Theorem 4.3. Consider the non-strongly convex quadratic optimization problem from Theorem 4.1. Up to a constant factor, the time complexity (12) is minimized by $\gamma \in \left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min\left\{\frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}\right\}\right]$ if $\frac{\mu}{\tau} \geq 2$ and by $\gamma \in \left[0, \max\left\{0, \min\left\{\frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}\right\}\right\}\right]$ if $\frac{\mu}{\tau} < 2$, and $T_\mu(\gamma) \leq T_{\text{GD}}$.

²Alternatively, an accelerated method could be employed, reducing the iteration count to $\tilde{\mathcal{O}}(\sqrt{\gamma L_i + 1})$. Our analysis can be extended to accommodate these faster solvers as well.

The theorem above establishes that the total time complexity of FedExProx is no worse than that of GD, and can be strictly better when communication time dominates computation time, highlighting the superiority of our new analysis. Notably, this was not the case for the previous analysis by Li et al. (2024) (Theorem 3.2). The potential for our complexity to significantly outperform that of GD is illustrated in the following simple example, where $T_\mu(\gamma) \ll T_{\text{GD}}$ for some $\gamma > 0$.

Example 4.4. Let the matrices $\{\mathbf{A}_i\}$ be diagonal, i.e., $\mathbf{A}_i = \text{diag}(a_{i1}, \dots, a_{id})$ for all $i \in [n]$, and $a_{ij} > 0$ for all $i \in [n], j \in [d]$. Then, according to (27) and (28), we get

$$\frac{L_\gamma}{\mu_\gamma^+} = \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} = \frac{\max_{j \in [d]} \sum_{i=1}^n \frac{\gamma a_{ij}}{1 + \gamma a_{ij}}}{\min_{j \in [d]} \sum_{i=1}^n \frac{\gamma a_{ij}}{1 + \gamma a_{ij}}},$$

where $\lambda_{\max}(\gamma)$ and $\lambda_{\min}^+(\gamma)$ are the largest and the smallest non-zero eigenvalues of the matrix $\mathbf{M} := \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})$, respectively (see Section C). Taking $\gamma = 1 / \min_{i \in [n], j \in [d]} a_{ij}$ and using the fact that $1 \geq \frac{x}{1+x} \geq \frac{1}{2}$ for all $x \geq 1$, we can conclude that $\frac{L_\gamma}{\mu_\gamma^+} \leq 2$. Thus,

$$T_\mu(\gamma) = \tilde{\mathcal{O}} \left(\mu + \tau \frac{\max_{i \in [n], j \in [d]} a_{ij}}{\min_{i \in [n], j \in [d]} a_{ij}} \right).$$

Suppose that communication is slow. Formally, let $\mu \geq \tau \max_{i \in [n], j \in [d]} a_{ij} / \min_{i \in [n], j \in [d]} a_{ij}$. Then $T_\mu(\gamma) = \tilde{\mathcal{O}}(\mu)$, while the total time complexity of GD is

$$T_{\text{GD}} = \tilde{\Omega} \left(\mu \times \frac{L}{\mu^+} \right) = \tilde{\Omega} \left(\mu \times \frac{\max_{j \in [d]} \sum_{i=1}^n a_{ij}}{\min_{j \in [d]} \sum_{i=1}^n a_{ij}} \right),$$

which is at least $\max_{j \in [d]} \sum_{i=1}^n a_{ij} / \min_{j \in [d]} \sum_{i=1}^n a_{ij}$ times worse!

A similar improvement can be observed in the general case. However, the derivation is significantly more complex, so to maintain clarity in the presentation, the example focuses on diagonal matrices.

5 PARTIAL PARTICIPATION

Thus far, we have concentrated on the full participation scenario. However, as outlined in Section 1.1, practical FL settings often involve only a subset of clients participating in each training round. To address this, we supplement our theory with a convergence result in the stochastic setting. For illustration, we consider nice sampling (see Section A.2), where at each iteration, a subset $\mathcal{S}_k \subseteq [n]$ of clients is selected uniformly at random from all subsets of size S . Although we use this sampling strategy as an example, other client selection methods can also be employed.

In this context, we can formulate the following complexity result.

Theorem 5.1. Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization problem (1) where $f_i(x) = \frac{1}{2} x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Under Assumption 1.3, FedExProx with nice sampling (Algorithm 2) with $\alpha = 1/\gamma L_{\gamma,S}$ finds \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ after

$$\mathcal{O} \left(\frac{L_{\gamma,S}}{\mu_\gamma^+} \log \frac{1}{\varepsilon} \right) \quad (13)$$

iterations, where $L_{\gamma,S} := \frac{n-S}{S(n-1)} \frac{L_{\max}}{1+\gamma L_{\max}} + \frac{n(S-1)}{S(n-1)} L_\gamma$, L_γ is the smoothness constant of M^γ and μ_γ^+ is the smallest non-zero eigenvalue of the matrix $\nabla^2 M^\gamma$.

Assuming the same time complexity model as in Section 4.1, the total time complexity of FedExProx with nice sampling is given by

$$T_\mu(\gamma, S) := \tilde{\mathcal{O}} \left((\mu + \tau (\gamma L_{\max} + 1)) \times \frac{L_{\gamma,S}}{\mu_\gamma^+} \right). \quad (14)$$

As it turns out, the optimal stepsize for the stochastic setting aligns with the one used in the deterministic scenario (see Theorem 4.3).

Theorem 5.2. *Up to a constant factor, the time complexity (14) is minimized by $\gamma \in \left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right]$ if $\frac{\mu}{\tau} \geq 2$ and by $\gamma \in \left[0, \max \left\{ 0, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right\} \right]$ if $\frac{\mu}{\tau} < 2$.*

Consequently, the conclusions from the previous section apply equally to the partial participation scenario: FedExProx performs at least as well as GD and can be strictly better when communication time exceeds computation time.

6 ADAPTIVITY

Next, we explore adaptive extrapolation strategies introduced by Horváth et al. (2022): *gradient diversity* (GraDS) and a variant of the classical *Polyak stepsize* (StoPS), both of which were later adapted for proximal methods by Li et al. (2024). As in the constant extrapolation case, we refine the analysis by Li et al. (2024) for quadratic problems. Although Theorem 6.1 refers to the full participation setting, the approach can be extended to the stochastic case (see Section F.2).

Theorem 6.1. *Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization problem (1), where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Let Assumption 1.3 hold and consider two adaptive extrapolation strategies:*

1. (FedExProx-GraDS) Set

$$\alpha_k = \alpha_k^{\text{GraDS}}(x_k) := \frac{\frac{1}{n} \sum_{i=1}^n \left\| \nabla M_{f_i}^\gamma(x_k) \right\|^2}{\left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2} \geq 1.$$

Then, the iterates of Algorithm 1 satisfy

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \min_{k=0, \dots, K-1} \alpha_k \gamma \frac{2+\gamma L_{\max}}{1+\gamma L_{\max}} \mu_\gamma^+ \right)^K \|x_0 - \Pi(x_0)\|^2. \quad (15)$$

2. (FedExProx-StoPS) Set

$$\alpha_k = \alpha_k^{\text{StoPS}}(x_k) := \frac{\frac{1}{n} \sum_{i=1}^n (M_{f_i}^\gamma(x_k) - \inf M_{f_i}^\gamma)}{\gamma \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2} \geq \frac{1}{2\gamma L_\gamma}.$$

Then, the iterates of Algorithm 1 satisfy

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{3}{2} \min_{k=0, \dots, K-1} \alpha_k \gamma \mu_\gamma^+ \right)^K \|x_0 - \Pi(x_0)\|^2. \quad (16)$$

Remark 6.2. Similar to the observations by Li et al. (2024), we see that, unlike in the constant extrapolation case (see Theorem 4.1), FedExProx with adaptive extrapolation benefits from *semi-adaptivity* to the smoothness constant. Specifically, it converges for any $\gamma > 0$, and since $\gamma \mu_\gamma^+$ is bounded above (see (28)), it suffices to choose a large enough γ to achieve the optimal performance.

Remark 6.3. Bounding $\min_{k=0, \dots, K-1} \alpha_k$ by $\frac{1}{2\gamma L_\gamma}$, inequality (16) implies that

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{3}{4} \frac{\mu_\gamma^+}{L_\gamma} \right)^K \|x_0 - \Pi(x_0)\|^2, \quad (17)$$

and hence Theorem 6.1 guarantees convergence of FedExProx-StoPS in

$$\mathcal{O} \left(\frac{L_\gamma}{\mu_\gamma^+} \log \frac{1}{\varepsilon} \right)$$

iterations, regardless of the choice of the stepsize γ . This matches the guarantee from Theorem 4.1, but without requiring prior knowledge of the optimal extrapolation parameter $\alpha = 1/\gamma L_\gamma$. The benefit comes with the trade-off of needing to know the minimum of the average of Moreau envelopes.

7 BETTER THEORY WITH POLYAK-ŁOJASIEWICZ CONDITION

The result from Section 4 can be extended to solving problem (1) for general functions that satisfy the Polyak-Łojasiewicz (PŁ) condition.

Assumption 7.1. The function M^γ satisfies PŁ-condition, i.e., there exists μ_γ^+ such that

$$\frac{1}{2} \|\nabla M^\gamma(x)\|^2 \geq \mu_\gamma^+ (M^\gamma(x) - M^\gamma(x_*))$$

for all $x \in \mathbb{R}^d$.

This assumption holds with $\mu_\gamma^+ \geq \mu^+ / (1 + \gamma L_{\max})$ if the function f satisfies PŁ condition with a constant μ^+ (see Lemma A.11). However, the choice $\mu_\gamma^+ = \mu^+ / (1 + \gamma L_{\max})$ is loose and leads to the problems described in Theorem 3.2. With this assumption in place, we can present the theorem.

Theorem 7.2. *Let Assumptions 1.1, 1.2, 1.3, and 7.1 hold. For all $\gamma > 0$, FedExProx (Algorithm 1) with $\alpha = 1/\gamma L_\gamma$ finds \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ in*

$$\mathcal{O}\left(\frac{L_\gamma}{\mu_\gamma^+} \log \frac{1}{\varepsilon}\right)$$

iterations, where L_γ is a smoothness constant of M^γ and μ_γ^+ is the PŁ constant.

Remark 7.3. In practice, obtaining an exact solution to the local problems is often infeasible. Instead, we can only calculate an inexact proximal operator that approximates the desired quantity within a certain accuracy budget. To accommodate this, we extend the above analysis to the scenario where the clients can only compute updates $\text{prox}_{\gamma f_i}^\delta(x)$ such that $\|\text{prox}_{\gamma f_i}^\delta(x) - \text{prox}_{\gamma f}(x)\|^2 \leq \delta$ (see Section E).

Remark 7.4. Li et al. (2024) also establish a linear rate for FedExProx. However, their result relies on the assumption that the function f is strongly convex, and hence that (1) has a unique solution. In contrast, our theorem is more general, as it allows for the possibility of multiple solutions.

In the partial participation scenario, we can establish the following result.

Theorem 7.5. *Let Assumptions 1.1, 1.2, 1.3, and 7.1 hold. For all $\gamma > 0$, FedExProx with nice sampling (Algorithm 2) and $\alpha = 1/\gamma L_\gamma$ finds \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ in*

$$\mathcal{O}\left(\frac{L_{\gamma,S}}{\mu_\gamma^+} \log \frac{1}{\varepsilon}\right)$$

iterations, where $L_{\gamma,S} := \left(\frac{n-S}{S(n-1)} \frac{L_{\max}}{1+\gamma L_{\max}} + \frac{n(S-1)}{S(n-1)} L_\gamma\right)$, L_γ is a smoothness constant of M^γ and μ_γ^+ is the PŁ constant.

Note that the complexities in Theorems 7.2 and 7.5 are entirely analogous to those in Theorems 4.1 and 5.1, with the only distinction being the substitution of the smallest non-zero eigenvalue of the function M^γ with the PŁ constant.

7.1 WHY DO WE GET A TIGHTER ANALYSIS?

To understand why the new analysis yields stronger guarantees, we need to examine the reasoning behind Theorem 3.2. When proving the result, we noticed that the main reason why (9) was true was the dependence of the complexity on the L_{\max} factor. At the same time, the complexity of GD depends only on L , and there exist many examples when $L_{\max} \gg L$. The question was: why does the original analysis of FedExProx involve L_{\max} , and can one circumvent this dependency? The reason behind it is the reliance of the proofs by Li et al. (2024) on Lemma A.9, which establishes the inequality $M^\gamma(x) - M^\gamma(x_*) \geq \frac{1}{1+\gamma L_{\max}} (f(x) - f(x_*))$ for all $x \in \mathbb{R}^d$. This inequality appears to be essential for deriving the convergence result $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ (where $\bar{x} \in \mathbb{R}^d$ is the output of FedExProx). Our main idea is to instead obtain convergence in terms of the distance between \bar{x} and x_* , i.e., to establish that $\mathbb{E}[\|\bar{x} - x_*\|^2] \leq \varepsilon$, and then translate it to $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ using the L -smoothness of the function f . In this way, one can avoid using Lemma A.9 when dealing with quadratic functions or problems satisfying the PŁ condition, ultimately obtaining much more favourable convergence guarantees.

The story behind our theoretical improvements is no less important than the improvements themselves, as it offers valuable insights and can be instructive for future research.

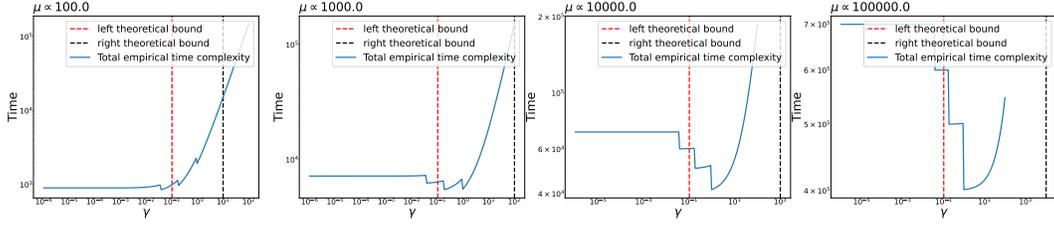
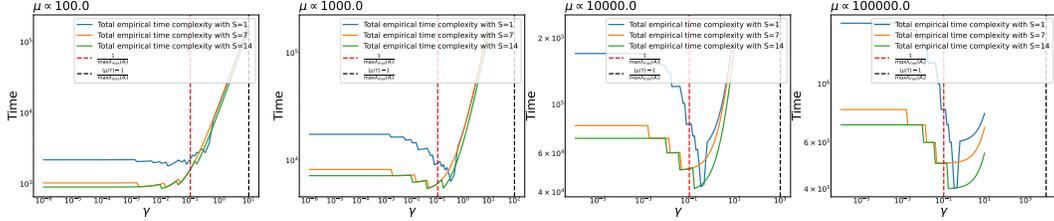


Figure 1: Empirical time complexities of FedExProx on a quadratic optimization task.

Figure 2: Empirical time complexities of FedExProx with partial client participation on a quadratic optimization task for $S \in \{1, 7, 14\}$ clients participating in each round.

8 EXPERIMENTS

To verify our theoretical results, we assess the empirical time complexity of FedExProx as a function of γ . We consider two types of optimization problems: synthetic quadratic optimization tasks and classification problems with a smooth hinge loss function. In both cases, local updates are computed using GD. Below, we first outline the methodology common to both experiments, followed by a detailed description of the specific settings and their corresponding results.

Total empirical time complexity. In accordance with the setup described in Section 4.1, we assume that one local iteration of GD takes τ seconds. Thus, the time needed by worker i to find $\text{prox}_{\gamma f_i}(x_k)$ at global iteration k is proportional to $\tau \times n_{ik}$, where n_{ik} is the number of GD iterations needed to find $\text{prox}_{\gamma f_i}(x_k)$ to a given small fixed accuracy. In the full participation case, the total empirical time complexity is given by

$$\sum_{k=0}^{K-1} (\mu + \tau \max_{i \in [n]} n_{ki}) \quad (18)$$

where K is the number of global iterations needed for FedExProx to converge to the desired accuracy and μ is the communication cost. Notice that we take the maximum over n_{ki} because the overall time is determined by the slowest client. One of experimental tasks is to verify that this empirical complexity aligns with the theoretical complexity derived in (14) (see Section 8.2). In experiments involving client sampling, we consider the complexity $\sum_{k=0}^{K-1} (\mu + \tau \max_{i \in S_k} n_{ki})$ instead, where S_k is the random subset of clients sampled at global iteration k .

8.1 EXPERIMENTS WITH QUADRATIC OPTIMIZATION PROBLEMS

In the first set of experiments, we consider quadratic optimization problems of the form $f(x) = \frac{1}{n} \sum_{i=1}^n \frac{1}{2} x^\top \mathbf{A}_i x$, where $\mathbf{A}_i \in \text{Sym}_+^7$, $i \in [n]$ are random positive semidefinite matrices with the smallest eigenvalue equal to zero. Each worker i computes the proximal mappings using GD with a stepsize of $1/\bar{L}_i$, where $\bar{L}_i = \lambda_{\max}(\mathbf{A}_i) + 1/\gamma$ is the smoothness constant of the local subproblem (3). The extrapolation parameter α is set to its optimal value from Theorem 4.1, using an explicitly computed value of L_γ . The data is distributed across $n = 14$ workers.

The results are presented in Figure 1. The dashed lines represent the theoretical bounds from Theorem 4.3, within which the optimal γ is expected to lie. One can see that when μ is relatively small ($\mu \propto 100$), the best choice of γ is near 0. However, as the communication cost μ increases, the best γ shifts to values greater than 0.1. A distinctive U-shape emerges, indicating the nontrivial optimal choice of γ . The results for partial participation from Figure 2 are qualitatively the same. These observations are fully consistent with our theoretical predictions.

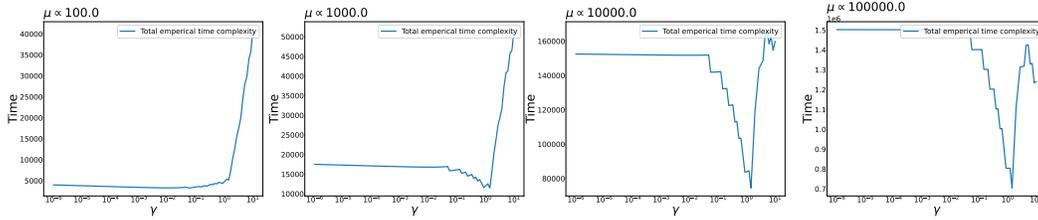


Figure 3: Empirical time complexities of FedExProx on a classification task with smooth hinge loss.

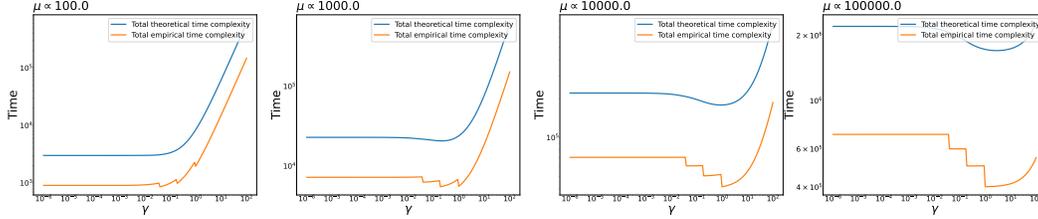


Figure 4: Comparison of theoretical time complexity (12) and empirical time complexity (18).

8.2 EXPERIMENTS WITH SMOOTH HINGE LOSS

We now turn to a classification problem with linearly separable data $\{(x_{ij}, y_{ij})\}_{i \in [n], j \in [m]}$, ensuring the interpolation regime. The function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as

$$f(w) = \frac{1}{n} \sum_{i=1}^n f_i(w) := \frac{1}{n} \sum_{i=1}^n \frac{1}{m} \sum_{j=1}^m \ell_{ij}(w),$$

where $(x_{ij}, y_{ij}) \in \mathbb{R}^d \times \{-1, 1\}$ for all $i \in [n], j \in [m]$ and ℓ_{ij} is the smooth hinge loss, defined by

$$\ell_{ij}(w) = \begin{cases} 0, & \text{if } y_{ij} w^T x_{ij} \geq 1, \\ \frac{1}{2}(1 - y_{ij} w^T x_{ij})^2, & \text{if } 0 < y_{ij} w^T x_{ij} < 1, \\ 1 - y_{ij} w^T x_{ij}, & \text{if } y_{ij} w^T x_{ij} \leq 0. \end{cases}$$

Each client $i \in [n]$ computes the local proximal mappings using GD with stepsize $1/\bar{L}_i$, where $\bar{L}_i = \frac{1}{m} \sum_{j=1}^m \|x_{ij}\|^2$ is an upper bound on the smoothness constant of the function f_i . The extrapolation parameter α is determined via grid search within the theoretical range given by Lemma A.7. We use $n = 4$ workers, each holding $m = 4$ datapoints, and the number of parameters is $d = 3$.

The results are presented in Figure 3. As in Section 8.1, these experiments also suggest a non-zero optimal choice of γ when communication is slow. Indeed, the values of γ minimizing the U-shapes in the last three plots are around 1. These experiments provide further support for our theory.

Comparing theoretical and empirical time complexities. Lastly, we compare the theoretical and empirical time complexities (12) and (18). Both are plotted on the same scale in Figure 4. We see that our theoretical model reflects the empirical dependencies (up to a multiplicative factor).

9 CONCLUSION

In this work, we revisit an empirically promising extrapolated parallel proximal method FedExProx. Motivated by the limitations of its original overly pessimistic analysis, we develop a novel analytical framework for non-strongly convex quadratic and PL cases. Our analysis significantly improves upon the previous results, offering insights into the true method’s performance. While our work focuses on the relatively simple case of quadratic problems, the experimental results indicate that the phenomena we describe extend beyond this setting. These findings suggest that the benefits of extrapolation in FL hold in broader optimization scenarios, motivating future research to rigorously establish such extensions. We believe that our work provides a foundational stepping stone toward better understanding the role of extrapolation in FL and can inspire further exploration to relax the assumptions made here, ultimately advancing the theory and practice of federated optimization.

REFERENCES

- 540
541
542 Amir Beck. *First-Order Methods in Optimization*. SIAM-Society for Industrial and Applied Math-
543 ematics, Philadelphia, PA, USA, 2017. ISBN 1611974984.
- 544 Dimitri P Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathemat-*
545 *ical programming*, 129(2):163–195, 2011.
- 546
547 Aleksandr Beznosikov, Samuel Horváth, Peter Richtárik, and Mher Safaryan. On biased compres-
548 sion for distributed learning. *arXiv preprint arXiv:2002.12410*, 2020.
- 549 P. Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Opti-*
550 *mization*, 26(4):2235–2260, 2016.
- 551
552 P. L. Combettes and J. I. Madariaga. Randomly activated proximal methods for nonsmooth convex
553 minimization. preprint arXiv:2403.10673, 2024.
- 554 Patrick L Combettes. Convex set theoretic image recovery by extrapolated iterations of parallel
555 subgradient projections. *IEEE Transactions on Image Processing*, 6(4):493–506, 1997.
- 556
557 L. Condat and P. Richtárik. RandProx: Primal-dual optimization algorithms with randomized prox-
558 imal updates. In *Proc. of Int. Conf. Learning Representations (ICLR)*, 2023.
- 559 Eduard Gorbunov, Filip Hanzely, and Peter Richtárik. Local SGD: Unified theory and new efficient
560 methods. In *International Conference on Artificial Intelligence and Statistics*, pp. 3556–3564.
561 PMLR, 2021.
- 562
563 Kaja Gruntkowska, Alexander Tyurin, and Peter Richtárik. Improving the worst-case bidirec-
564 tional communication complexity for nonconvex distributed optimization under function simi-
565 larity. *arXiv preprint arXiv:2402.06412*, 2024.
- 566
567 Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean
568 Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile
569 keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018.
- 570 Samuel Horváth, Konstantin Mishchenko, and Peter Richtárik. Adaptive learning rates for faster
571 stochastic gradient methods. *arXiv preprint arXiv:2208.05287*, 2022.
- 572 Divyansh Jhunjhunwala, Shiqiang Wang, and Gauri Joshi. FedExP: Speeding up federated averaging
573 via extrapolation. *arXiv preprint arXiv:2301.09604*, 2023.
- 574
575 Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin
576 Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Ad-
577 vances and open problems in federated learning. *Foundations and Trends® in Machine Learning*,
578 14(1–2):1–210, 2021.
- 579 Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-
580 gradient methods under the Polyak-Łojasiewicz condition. In *Machine Learning and Knowl-*
581 *edge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy,*
582 *September 19-23, 2016, Proceedings, Part I 16*, pp. 795–811. Springer, 2016.
- 583
584 Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and
585 Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In
586 *International conference on machine learning*, pp. 5132–5143. PMLR, 2020.
- 587 Ahmed Khaled and Chi Jin. Faster federated optimization under second-order similarity. *arXiv*
588 *preprint arXiv:2209.02257*, 2022.
- 589
590 Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. First analysis of local GD on hetero-
591 geneous data. *arXiv preprint arXiv:1909.04715*, 2019.
- 592
593 Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified
theory of decentralized SGD with changing topology and local updates. In *International Confer-*
ence on Machine Learning, pp. 5381–5393. PMLR, 2020.

- 594 Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and
595 Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv*
596 *preprint arXiv:1610.05492*, 2016.
- 597 Hanmin Li, Kirill Acharya, and Peter Richtarik. The power of extrapolation in federated learning.
598 *arXiv preprint arXiv:2405.13766*, 2024.
- 600 Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith.
601 Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Sys-*
602 *tems*, 2:429–450, 2020.
- 604 Dong C Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization.
605 *Mathematical programming*, 45(1):503–528, 1989.
- 606 Siyuan Ma, Raef Bassily, and Mikhail Belkin. The power of interpolation: Understanding the
607 effectiveness of SGD in modern over-parametrized learning. In *International Conference on*
608 *Machine Learning*, pp. 3325–3334. PMLR, 2018.
- 609 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
610 Communication-efficient learning of deep networks from decentralized data. In *Artificial intelli-*
611 *gence and statistics*, pp. 1273–1282. PMLR, 2017.
- 613 Konstantin Mishchenko, Grigory Malinovsky, Sebastian Stich, and Peter Richtárik. ProxSkip: Yes!
614 local gradient steps provably lead to communication acceleration! Finally! In *International*
615 *Conference on Machine Learning*, pp. 15750–15769. PMLR, 2022.
- 616 J.J. Moreau. Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de*
617 *France*, 93:273–299, 1965.
- 619 Ion Necoara, Peter Richtárik, and Andrei Patrascu. Randomized projection methods for convex
620 feasibility: Conditioning and convergence rates. *SIAM Journal on Optimization*, 29(4):2814–
621 2852, 2019.
- 622 Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018.
- 624 A. Patrascu and I. Necoara. Nonasymptotic convergence of stochastic proximal point methods for
625 constrained convex optimization. *Journal of Machine Learning Research*, 18:1–42, 2018.
- 626 Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen,
627 and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine*
628 *Learning*, pp. 8821–8831. PMLR, 2021.
- 630 Peter Richtárik and Martin Takáč. Parallel coordinate descent methods for big data optimization.
631 *Mathematical Programming*, 156:433–484, 2016.
- 632 Peter Richtárik and Martin Takáč. Stochastic reformulations of linear systems: algorithms and
633 convergence theory. *SIAM Journal on Matrix Analysis and Applications*, 41(2):487–524, 2020.
- 635 Peter Richtárik, Abdurakhmon Sadiev, and Yury Demidovich. A unified theory of stochastic proximal
636 point methods without smoothness. *arXiv preprint arXiv:2405.15941*, 2024.
- 637 Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon
638 Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. The future of digital
639 health with federated learning. *NPJ digital medicine*, 3(1):1–7, 2020.
- 640 R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM J. Control*
641 *Optim.*, 14(5):877–898, 1976.
- 642 R. Tyrrell Rockafellar and Roger J-B Wets. *Variational Analysis*. Grundlehren der mathematischen
643 *Wissenschaften*. Springer Berlin, Heidelberg, 1998.
- 644 Ernest K Ryu and Stephen Boyd. Stochastic proximal iteration: a non-asymptotic improvement
645 upon stochastic gradient descent. *Author website, early draft*, 2014.
- 646
- 647

648 Abdurakhmon Sadiev, Laurent Condat, and Peter Richtárik. Stochastic proximal point methods for
649 monotone inclusions under expected similarity. *arXiv preprint arXiv:2405.14255*, 2024.
650

651 Cheik Traoré, Vassilis Apidopoulos, Saverio Salzo, and Silvia Villa. Variance reduction techniques
652 for stochastic proximal point algorithms. *Journal of Optimization Theory and Applications*, pp.
653 1–30, 2024.

654 Alexander Tyurin and Peter Richtárik. A computation and communication efficient method for dis-
655 tributed nonconvex problems in the partial participation setting. *Advances in Neural Information*
656 *Processing Systems*, 2023.

657 Alexander Tyurin, Lukang Sun, Konstantin Burlachenko, and Peter Richtárik. Sharper rates
658 and flexible framework for nonconvex SGD with client and data sampling. *arXiv preprint*
659 *arXiv:2206.02275*, 2022.
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702	CONTENTS	
703		
704	1 Introduction	1
705	1.1 Related Work	1
706		
707		
708	2 Contributions	3
709		
710		
711	3 Not Better than GD on Quadratics	4
712		
713	4 Tighter and More Optimistic Results	5
714	4.1 One step time complexity $\pi(\gamma)$ is $\mu + \tau(\gamma L_{\max} + 1)$	5
715		
716	5 Partial Participation	6
717		
718		
719	6 Adaptivity	7
720		
721	7 Better Theory with Polyak-Łojasiewicz Condition	8
722	7.1 Why do we get a tighter analysis?	8
723		
724		
725	8 Experiments	9
726	8.1 Experiments with quadratic optimization problems	9
727	8.2 Experiments with smooth hinge loss	10
728		
729		
730	9 Conclusion	10
731		
732	A Background	16
733	A.1 Useful facts about Moreau envelope	16
734	A.2 Nice sampling	17
735	A.3 Note on other sampling strategies	18
736	A.4 Useful lemmas	18
737		
738		
739		
740	B FedExProx as (S)GD on Moreau Envelope Reformulation	19
741		
742		
743	C FedExProx for Quadratics	20
744	C.1 Properties of the Hessian	21
745	C.2 Proofs for Section 3	22
746	C.3 Tighter analysis	24
747	C.3.1 Full participation case	24
748	C.3.2 Partial participation case	27
749	C.4 Lemmas	30
750		
751		
752		
753	D FedExProx under PL Condition	33
754		
755	E FedExProx with Inexact Computations	35

756	F Adaptive Extrapolation	38
757		
758	F.1 Full participation case	38
759	F.2 Partial participation case	40
760		
761	G Notation	43
762		
763		
764		
765		
766		
767		
768		
769		
770		
771		
772		
773		
774		
775		
776		
777		
778		
779		
780		
781		
782		
783		
784		
785		
786		
787		
788		
789		
790		
791		
792		
793		
794		
795		
796		
797		
798		
799		
800		
801		
802		
803		
804		
805		
806		
807		
808		
809		

Appendix

A BACKGROUND

Before discussing the contributions of this work, we provide formal definitions and essential facts that will be repeatedly referenced in the proofs. A summary of the commonly used notation is provided in Table 1.

Definition A.1 (Proximal operator). The *proximal operator* $\text{prox}_{\gamma f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ of f is defined by

$$\text{prox}_{\gamma f}(x) = \arg \min_{z \in \mathbb{R}^d} \left\{ f(z) + \frac{1}{2\gamma} \|z - x\|^2 \right\},$$

where $\|\cdot\|$ is the standard Euclidean norm.

Definition A.2 (Moreau envelope). The *Moreau envelope* of an extended-real-valued function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d \cup \{\infty\}$ with stepsize $\gamma > 0$ is the function $M_f^\gamma : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$M_f^\gamma(x) = \min_{y \in \mathbb{R}^d} \left\{ f(y) + \frac{1}{2\gamma} \|y - x\|^2 \right\}$$

for all $x \in \mathbb{R}^d$.

A.1 USEFUL FACTS ABOUT MOREAU ENVELOPE

We will often rely on several useful properties of the Moreau envelopes, summarized below. In what follows, we denote

$$M^\gamma(x) := \frac{1}{n} \sum_{i=1}^n M_{f_i}^\gamma(x).$$

Lemma A.3 (Beck (2017) Theorem 6.60). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper, closed and convex. Then its Moreau envelope M_f^γ satisfies*

$$\nabla M_f^\gamma(x) = \frac{1}{\gamma}(x - \text{prox}_{\gamma f}(x))$$

for all $x \in \mathbb{R}^d$, for any $\gamma > 0$.

Lemma A.4 ((Li et al., 2024) Lemma 4). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be convex and L -smooth. Then M_f^γ is $\frac{L}{1+\gamma L}$ -smooth.*

Lemma A.5 ((Li et al., 2024) Lemma 5). *Let $f : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ be a proper, closed and convex function. Then, for all $\gamma > 0$, f and $M_f^\gamma(x)$ have the same set of minimizers and minimum.*

Lemma A.6 ((Beck, 2017) Theorem 6.55). *Let $f : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$ be a proper, closed and convex function. Then M_f^γ is also a convex function.*

Lemma A.7 ((Li et al., 2024) Lemma 7). *Let the functions $f_i : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$, $i \in [n]$, be proper, closed, convex and L_i -smooth. Then M^γ is convex and L_γ -smooth with*

$$\frac{1}{n^2} \sum_{i=1}^n \frac{L_i}{1 + \gamma L_i} \leq L_\gamma \leq \frac{1}{n} \sum_{i=1}^n \frac{L_i}{1 + \gamma L_i}.$$

Lemma A.8 ((Li et al., 2024) Lemma 8). *Let the functions $f_i : \mathbb{R}^d \mapsto \mathbb{R} \cup \{+\infty\}$, $i \in [n]$, be proper, closed and convex, and suppose that Assumption 1.3 holds. Then, for all $\gamma \in (0, \infty)$, the function $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ has the same set of minimizers and minimum as the function $M^\gamma(x)$.*

Lemma A.9. [Li et al. (2024) Lemma 10] *Let Assumptions 1.1, 1.2 and 1.3 hold. Then, for any $x_* \in \mathcal{X}_*$ and for all $x \in \mathbb{R}^d$, it holds that*

$$M_{f_i}^\gamma(x) - M_{f_i}^\gamma(x_*) \geq \frac{1}{1 + \gamma L_i} (f_i(x) - f_i(x_*)).$$

Consequently,

$$M^\gamma(x) - M^\gamma(x_*) \geq \frac{1}{1 + \gamma L_{\max}} (f(x) - f(x_*)).$$

The next lemma will play a role in the analysis of the method when assuming the Polyak-Łojasiewicz (PŁ) condition.

Assumption A.10 (Polyak-Łojasiewicz condition). The function f satisfies PŁ condition, i.e., there exists $\mu > 0$ such that

$$\frac{1}{2} \|\nabla f(x)\|^2 \geq \mu (f(x) - f(x_*))$$

for all $x \in \mathbb{R}^d$.

Lemma A.11. Let the function f satisfy Assumption A.10 with parameter μ^+ , and suppose that Assumptions 1.1, 1.2 and 1.3 hold. Then, Assumption 7.1 holds with parameter

$$\mu_\gamma^+ \geq \frac{\mu^+}{1 + \gamma L_{\max}}.$$

Proof. If f satisfies Assumption A.10, then under Assumption 1.1, Polyak-Łojasiewicz condition and inequality

$$f(x) - f(x_*) \geq \frac{\mu^+}{2} \|x - \Pi(x)\|^2 \quad \forall x \in \mathbb{R}^d, \quad (19)$$

are equivalent (Karimi et al., 2016), where $\Pi(x)$ is the projection of x onto the solution set \mathcal{X}_* of f and $x_* \in \mathcal{X}_*$. Then, using Lemma A.9, we get

$$M^\gamma(x) - M^\gamma(x_*) \geq \frac{1}{1 + \gamma L_{\max}} (f(x) - f(x_*)) \stackrel{(19)}{\geq} \frac{\mu^+}{2(1 + \gamma L_{\max})} \|x - \Pi(x)\|^2.$$

Due to Lemma A.8, M^γ and f share the set of minimizers. Thus, $\Pi(x)$ is the projection of x onto the solution set \mathcal{X}_* of M^γ . We can conclude that M^γ satisfies (19) (with M^γ and $\frac{\mu^+}{1 + \gamma L_{\max}}$ instead of f and μ^+), meaning that M^γ satisfies PŁ condition with parameter $\frac{\mu^+}{1 + \gamma L_{\max}}$. \square

A.2 NICE SAMPLING

Let us formally introduce the sampling strategy used. Fix a minibatch size $S \in [n]$ and let \mathcal{S} be a random subset of $[n]$ of size S , chosen uniformly at random from all $\binom{n}{S}$ subsets of $[n]$ of this size. Such a random set \mathcal{S} is known in the literature under the name *S-nice sampling* (Richtárik & Takáč, 2016).

In the proofs, we will rely on the following useful lemma:

Lemma A.12. Fix $\mathbf{B}_1, \dots, \mathbf{B}_n \in \mathbb{R}^{k \times l}$ and let \mathcal{S} be an S -nice sampling of the indices $[n]$. Then

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{1}{S} \sum_{i \in \mathcal{S}} \mathbf{B}_i \right)^\top \left(\frac{1}{S} \sum_{i \in \mathcal{S}} \mathbf{B}_i \right) \right] \\ &= \frac{n-S}{S(n-1)} \frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^\top \mathbf{B}_i + \frac{n(S-1)}{S(n-1)} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right). \end{aligned}$$

Proof. Define

$$\chi_i := \begin{cases} 1 & i \in \mathcal{S}, \\ 0 & i \notin \mathcal{S}. \end{cases}$$

Then

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{1}{S} \sum_{i \in \mathcal{S}} \mathbf{B}_i \right)^\top \left(\frac{1}{S} \sum_{i \in \mathcal{S}} \mathbf{B}_i \right) \right] \\ &= \mathbb{E} \left[\left(\frac{1}{S} \sum_{i=1}^n \chi_i \mathbf{B}_i \right)^\top \left(\frac{1}{S} \sum_{i=1}^n \chi_i \mathbf{B}_i \right) \right] \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{S^2} \mathbb{E} \left[\sum_{i=1}^n \chi_i^2 \mathbf{B}_i^\top \mathbf{B}_i + \sum_{i \neq j} \chi_i \chi_j \mathbf{B}_i^\top \mathbf{B}_j \right] \\
&= \frac{1}{S^2} \left(\sum_{i=1}^n \mathbb{E}[\chi_i] \mathbf{B}_i^\top \mathbf{B}_i + \sum_{i \neq j} \mathbb{E}[\chi_i \chi_j] \mathbf{B}_i^\top \mathbf{B}_j \right) \\
&= \frac{1}{S^2} \left(\frac{S}{n} \sum_{i=1}^n \mathbf{B}_i^\top \mathbf{B}_i + \frac{S(S-1)}{n(n-1)} \sum_{i \neq j} \mathbf{B}_i^\top \mathbf{B}_j \right) \\
&= \frac{1}{S} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^\top \mathbf{B}_i + \frac{S-1}{n(n-1)} \left(\left(\sum_{i=1}^n \mathbf{B}_i \right)^\top \left(\sum_{i=1}^n \mathbf{B}_i \right) - \sum_{i=1}^n \mathbf{B}_i^\top \mathbf{B}_i \right) \right) \\
&= \frac{n-S}{S(n-1)} \frac{1}{n} \sum_{i=1}^n \mathbf{B}_i^\top \mathbf{B}_i + \frac{n(S-1)}{S(n-1)} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right).
\end{aligned}$$

□

A.3 NOTE ON OTHER SAMPLING STRATEGIES

As discussed in Section 5, Algorithm 2 can handle virtually any (unbiased) sampling technique, with nice sampling used here as an illustrative example. For a comprehensive overview of other potential sampling strategies, we direct the reader to Tyurin et al. (2022), which offers a general framework for analyzing a broad range of sampling mechanisms.

A.4 USEFUL LEMMAS

Lemma A.13 ((Karimi et al., 2016) Theorem 2). *Assume that the function f is convex, L -smooth, and satisfies the Polyak-Łojasiewicz condition with parameter μ . Then*

$$f(x) - f(x_*) \geq \frac{\mu}{2} \|x - \Pi(x)\|^2,$$

for all $x \in \mathbb{R}^d$, where $\Pi(x)$ is the projection of x onto the solution set \mathcal{X}_* and $x_* \in \mathcal{X}_*$.

Lemma A.14 ((Nesterov, 2018)). *Let $D_f(x, y) := f(x) - f(y) - \langle \nabla f(y), x - y \rangle$ be the Bregman divergence of the function f . If f is convex, L -smooth and differentiable, then*

$$\frac{1}{L} \|\nabla f(x) - \nabla f(y)\|^2 \leq D_f(x, y) + D_f(y, x)$$

for any $x, y \in \mathbb{R}^d$.

Algorithm 1 FedExProx

1: **Parameters:** stepsize $\gamma > 0$, extrapolation parameter $\alpha_t > 0$, starting point $x_0 \in \mathbb{R}^d$
2: **for** $k = 0, 1, 2, \dots$ **do**
3: $x_{k+1} = x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^n \text{prox}_{\gamma f_i}(x_k) - x_k \right)$
4: **end for**

B FedExProx AS (S)GD ON MOREAU ENVELOPE REFORMULATION

The key idea behind the original analysis of FedExProx is to rewrite the step

$$x_{k+1} = x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^n \text{prox}_{\gamma f_i}(x_k) - x_k \right)$$

in terms of the Moreau envelopes of the local functions. Indeed, if the functions f_i are proper, closed and convex for all $i \in [n]$, by Lemma A.3, the update rule of Algorithm 1 can be rewritten as

$$x_{k+1} = x_k - \alpha_k \gamma \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k).$$

Analogously, in the partial participation case (Algorithm 2), we have

$$x_{k+1} = x_k - \alpha_k \gamma \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k).$$

It follows that Algorithms 1 and 2 are equivalent to GD and minibatch SGD for minimizing

$$M^\gamma(x) := \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x)$$

with stepsize $\alpha_k \gamma$.

1026 C FedExProx FOR QUADRATICS

1027 We begin by formally introducing the problem setup and the notation used throughout the proofs.

1028 We are interested in solving the quadratic optimization problem

$$1029 \min_{x \in \mathbb{R}^n} \left\{ f(x) := \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{2} x^\top \mathbf{A}_i x - b_i^\top x \right) \right\}, \quad (20)$$

1030 where $\mathbf{A}_i \in \text{Sym}_+^d := \{\mathbf{X} \in \mathbb{R}^{d \times d} \mid \mathbf{X} = \mathbf{X}^\top, \mathbf{X} \succeq 0\}$ and $b_i \in \mathbb{R}^d$, assuming the interpolation regime (Assumption 1.3), which in this case states that there exists x_* such that $\mathbf{A}_i x_* = b_i$ for all $i \in [n]$.

1031 We denote the set of minimizers by $\mathcal{X}_* := \{x \in \mathbb{R}^d : \frac{1}{n} \sum_{i=1}^n (\mathbf{A}_i x - b_i) = 0\}$ and let $\Pi(\cdot)$ be the projection onto \mathcal{X}_* :

$$1032 \Pi(x) = x - \left(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i \right)^\dagger \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{A}_i x - b_i) \right), \quad (21)$$

1033 where \dagger denotes the Moore-Penrose pseudoinverse.

1034 Using eigendecomposition, each matrix \mathbf{A}_i , $i \in [n]$ can be decomposed as $\mathbf{A}_i = \mathbf{Q}_i \mathbf{\Lambda}_i \mathbf{Q}_i^\top$, where $\mathbf{Q}_i = [q_{i,1}, \dots, q_{i,d}] \in \mathbb{R}^{d \times d}$ is the orthogonal matrix whose columns are eigenvectors of \mathbf{A}_i and $\mathbf{\Lambda}_i \in \mathbb{R}^{d \times d}$ is a diagonal matrix whose diagonal elements are the corresponding eigenvalues of \mathbf{A}_i , i.e., $\mathbf{\Lambda}_i = \text{diag}(\lambda_1(\mathbf{A}_i), \dots, \lambda_d(\mathbf{A}_i))$, where $0 = \lambda_1(\mathbf{A}_i) \leq \lambda_2(\mathbf{A}_i) \leq \dots \leq \lambda_d(\mathbf{A}_i)$. Let $\lambda_{\min}^+(\mathbf{A}_i)$ denote the smallest non-zero eigenvalue of \mathbf{A}_i and $\lambda_{\max}(\mathbf{A}_i)$ – the largest eigenvalue of \mathbf{A}_i .

1035 In general, the largest eigenvalue, smallest eigenvalue, and smallest positive eigenvalue of any matrix \mathbf{B} will be denoted by $\lambda_{\max}(\mathbf{B})$, $\lambda_{\min}(\mathbf{B})$ and $\lambda_{\min}^+(\mathbf{B})$, respectively.

1036 The proximal operators in this setting are

$$1037 \text{prox}_{\gamma f_i}(x) = (\gamma \mathbf{A}_i + \mathbf{I})^{-1} (x + \gamma b_i). \quad (22)$$

1038 Thus, by Lemma A.3,

$$1039 \nabla M_{f_i}^\gamma(x) = \frac{1}{\gamma} \left(x - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} (x + \gamma b_i) \right) = \frac{1}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) (x - x_*) \quad (23)$$

1040 for any $x_* \in \mathcal{X}_*$, where we use the fact that $\mathbf{A}_i x_* = b_i$. Furthermore, it can easily be shown that

$$1041 M_{f_i}^\gamma(x) = \frac{1}{2\gamma} \left(x^\top (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) x - 2\gamma b_i^\top (\gamma \mathbf{A}_i + \mathbf{I})^{-1} x - \gamma^2 b_i^\top (\gamma \mathbf{A}_i + \mathbf{I})^{-1} b_i \right) \\ 1042 = \frac{1}{2\gamma} \left((x - x_*)^\top (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) (x - x_*) - \gamma x_*^\top \mathbf{A}_i x_* \right), \quad (24)$$

1043 and hence

$$1044 M^\gamma(x) = \frac{1}{2} x^\top \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \right) x - \left(\frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1} b_i \right)^\top x \\ 1045 - \frac{\gamma}{2n} \sum_{i=1}^n b_i^\top (\gamma \mathbf{A}_i + \mathbf{I})^{-1} b_i \\ 1046 = \frac{1}{2} x^\top \mathbf{M} x - \left(\frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1} b_i \right)^\top x - \frac{\gamma}{2n} \sum_{i=1}^n b_i^\top (\gamma \mathbf{A}_i + \mathbf{I})^{-1} b_i \quad (25)$$

1047 where $\mathbf{M} := \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})$ is the Hessian matrix of the function M^γ . This can equivalently be written as

$$1048 M^\gamma(x) = \frac{1}{2} (x - x_*)^\top \mathbf{M} (x - x_*) - x_*^\top \left(\frac{1}{2n} \sum_{i=1}^n \mathbf{A}_i \right) x_* \quad (26)$$

1049 for any $x_* \in \mathcal{X}_*$. We shall denote by $\lambda_{\min}^+(\gamma)$ and $\lambda_{\max}(\gamma)$ the smallest non-zero eigenvalue and the largest eigenvalue of \mathbf{M} , respectively.

1080 C.1 PROPERTIES OF THE HESSIAN

1081 Next, we introduce several important properties of the matrix \mathbf{M} , which will be repeatedly used in
1082 the proofs that follow.

1083 In the main part of the paper, we stick to the notation L_γ and μ_γ^+ for easier comparison with the
1084 results by Li et al. (2024). However, in the proofs concerning the quadratic case, we adopt a more
1085 intuitive eigenvalue-based notation.

1086 **Fact 1.** Since $\mathbf{M} = \nabla^2 M^\gamma$, it holds that $L_\gamma = \lambda_{\max}(\gamma)$ and $\mu_\gamma^+ = \lambda_{\min}^+(\gamma)$.

1087 *Proof.* The result is an immediate consequence of (25). □

1088 **Fact 2.** Let $\lambda_{\max}(\gamma)$ and $\lambda_{\min}^+(\gamma)$ be the smallest non-zero eigenvalue and the largest eigenvalue of
1089 the matrix

$$\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}),$$

1090 respectively. Then

$$\lambda_{\max}(\gamma) = \max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x \quad (27)$$

1091 and

$$\lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in \ker(\mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x, \quad (28)$$

1092 where $\mathbf{A} := \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$ and $[b_j]_{jj}$ is the diagonal matrix with b_j as the j th entry.

1093 *Proof.* Using the eigendecomposition, we can write \mathbf{M} as

$$\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} \mathbf{Q}_i (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \mathbf{Q}_i^\top = \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top.$$

1094 The result follows from the identities

$$\lambda_{\max}(\gamma) = \max_{\|x\| \leq 1} x^\top \mathbf{M} x$$

1095 and

$$\lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in (\ker \mathbf{A})^\perp} x^\top \mathbf{M} x$$

1096 (see Lemma C.7). □

1097 **Fact 3.** For any $\gamma \geq 0$, it holds that $\lambda_{\max}(\gamma) \leq \lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}^+(\gamma) \leq \lambda_{\min}^+(\mathbf{A})$.

1098 *Proof.* From Fact 2, we have

$$\lambda_{\max}(\gamma) \stackrel{(27)}{=} \max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x,$$

1099 and similarly

$$\lambda_{\max}(\mathbf{A}) = \max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i [\lambda_j(\mathbf{A}_i)]_{jj} \mathbf{Q}_i^\top x.$$

1100 Since it always holds that $\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \leq \lambda_j(\mathbf{A}_i)$, we get

$$\lambda_{\max}(\gamma) \leq \lambda_{\max}(\mathbf{A}).$$

1101 The proof of the second inequality is analogous. □

C.2 PROOFS FOR SECTION 3

Theorem 3.2. *Let Assumptions 1.3 and 3.1 hold. Consider solving a non-strongly convex quadratic optimization problem of the form (1), where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Then the assumptions of Theorem 1 by Li et al. (2024) hold, and*

$$T(\gamma) = \pi(\gamma) \times \frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \geq \pi(0) \times \frac{LR^2}{\varepsilon} \quad (9)$$

for all $\gamma > 0$. Moreover, when $\gamma \rightarrow 0$, then $\pi(\gamma) \times \frac{L_\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \rightarrow \pi(0) \times \frac{LR^2}{\varepsilon}$, and FedExProx effectively reduces to GD.

Proof. We first prove that $L_\gamma(1 + \gamma L_{\max}) \geq L = \lambda_{\max}(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i)$ for all $\gamma \geq 0$. Since

$$\nabla^2 M_{f_i}(x) = \frac{1}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right),$$

we have

$$L_\gamma = \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) \right).$$

Therefore, using the fact that $L_{\max} = \max_i \lambda_{\max}(\mathbf{A}_i)$, we obtain

$$\begin{aligned} (1 + \gamma L_{\max})L_\gamma &= \left(1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i) \right) \frac{1}{\gamma} \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) \right) \\ &= \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) \right), \end{aligned} \quad (29)$$

and hence, it is sufficient to show that

$$\frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) \succeq \mathbf{A}_i. \quad (30)$$

Using the eigenvalue decomposition of \mathbf{A}_i , the expression can be rewritten as

$$\begin{aligned} \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) &= \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{Q}_i \mathbf{\Lambda}_i \mathbf{Q}_i^\top + \mathbf{I})^{-1} \right) \\ &= \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \mathbf{Q}_i \left(\mathbf{I} - (\gamma \mathbf{\Lambda}_i + \mathbf{I})^{-1} \right) \mathbf{Q}_i^\top. \end{aligned}$$

Now, letting $[\mathbf{B}]_j$ be the j th diagonal element of matrix \mathbf{B} , we have

$$\begin{aligned} \left[\frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{\Lambda}_i + \mathbf{I})^{-1} \right) \right]_j &= \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left[\mathbf{I} - (\gamma \mathbf{\Lambda}_i + \mathbf{I})^{-1} \right]_j \\ &= \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \left(1 - \frac{1}{\gamma [\mathbf{\Lambda}_i]_j + 1} \right) = \frac{1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)}{\gamma} \frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \geq \lambda_j(\mathbf{A}_i), \end{aligned}$$

and hence

$$\frac{(1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i))}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{\Lambda}_i + \mathbf{I})^{-1} \right) \succeq \mathbf{A}_i.$$

Multiplying by \mathbf{Q}_i from the left and by \mathbf{Q}_i^\top from the right, we see that (30) holds, and we can conclude that

$$(1 + \gamma L_{\max})L_\gamma \geq \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i \right) = L$$

for all $\gamma \geq 0$. Inequality (9) follows from Assumption 3.1.

We now turn to the second part of the theorem and show that $\pi(\gamma) \times \frac{L\gamma(1+\gamma L_{\max})R^2}{\varepsilon} \rightarrow \pi(0) \times \frac{LR^2}{\varepsilon}$ as $\gamma \rightarrow 0$. The idea behind the proof is that for small enough γ , it holds that

$$\begin{aligned} (1 + \gamma L_{\max})L_\gamma &\stackrel{(29)}{=} \left(1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)\right) \frac{1}{\gamma} \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}\right) \right) \\ &\approx \frac{1}{\gamma} \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \gamma \mathbf{A}_i \right) = \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i \right) = L. \end{aligned}$$

More formally, we have

$$\mathbf{I} - (\mathbf{I} + \gamma \mathbf{A}_i)^{-1} = \gamma \mathbf{A}_i - \gamma^2 \mathbf{A}_i^2 + \gamma^3 \mathbf{A}_i^3 + \dots = \sum_{k=1}^{\infty} (-1)^{k-1} \gamma^k \mathbf{A}_i^k,$$

so that

$$(1 + \gamma L_{\max})L_\gamma = \left(\frac{1}{\gamma} + \max_i \lambda_{\max}(\mathbf{A}_i) \right) \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^{\infty} (-1)^{k-1} \gamma^k \mathbf{A}_i^k \right).$$

By applying the dominated convergence theorem to exchange the limit operation and summation, we obtain

$$\begin{aligned} \lim_{\gamma \rightarrow 0} (1 + \gamma L_{\max})L_\gamma &= \lim_{\gamma \rightarrow 0} \left(\left(1 + \gamma \max_i \lambda_{\max}(\mathbf{A}_i)\right) \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^{\infty} (-1)^{k-1} \gamma^{k-1} \mathbf{A}_i^k \right) \right) \\ &= \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \sum_{k=1}^{\infty} \lim_{\gamma \rightarrow 0} (-1)^{k-1} \gamma^{k-1} \mathbf{A}_i^k \right) = \lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i \right) = L, \end{aligned}$$

as required.

Lastly, letting $\gamma \rightarrow 0$, we have $M_{f_i}^\gamma(x) \rightarrow f_i(x)$ for all x , and consequently $M^\gamma(x) \rightarrow f(x)$ (see, e.g., (Rockafellar & Wets, 1998)). Since Algorithm 1 is equivalent to GD for minimizing $M^\gamma(x)$ with stepsize $\alpha\gamma = 1/L_\gamma$, it follows that its iterates converge to those of vanilla GD with stepsize $1/L$. \square

1242 C.3 TIGHTER ANALYSIS

1243 C.3.1 FULL PARTICIPATION CASE

1244 **Iteration complexity.** We first establish the convergence rate of the algorithm.

1245 **Theorem C.1.** *Let Assumption 1.3 hold and let x_k be the iterates of FedExProx (Algorithm 1) applied*
 1246 *to problem (20). Then*

$$1247 \quad \|x_K - \Pi(x_K)\|^2 \leq (1 - \alpha\gamma(2 - \alpha\gamma\lambda_{\max}(\gamma))\lambda_{\min}^+(\gamma))^K \|x_0 - \Pi(x_0)\|^2,$$

1248 where $\lambda_{\max}(\gamma)$ and $\lambda_{\min}^+(\gamma)$ are the largest and the smallest positive eigenvalues of the matrix
 1249 $\mathbf{M} := \frac{1}{\gamma}(\mathbf{I} - \frac{1}{n}\sum_{i=1}^n(\gamma\mathbf{A}_i + \mathbf{I})^{-1})$, respectively.

1250 The optimal choice of stepsize γ and extrapolation parameter α is $\alpha\gamma = \frac{1}{\lambda_{\max}(\gamma)}$, in which case the
 1251 rate becomes

$$1252 \quad \|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{\lambda_{\min}^+(\gamma)}{\lambda_{\max}(\gamma)}\right)^K \|x_0 - \Pi(x_0)\|^2.$$

1253 Hence, the number of iterations needed to reach an ϵ -solution is

$$1254 \quad K \geq \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log\left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon}\right).$$

1255 *Proof of Theorem C.1.* The update rule of FedExProx can be written as

$$1256 \quad x_{k+1} = x_k + \alpha \left(\frac{1}{S} \sum_{i=1}^S \text{prox}_{\gamma f_i}(x_k) - x_k \right)$$

$$1257 \quad \stackrel{(22)}{=} x_k + \alpha \left(\frac{1}{S} \sum_{i=1}^S (\gamma\mathbf{A}_i + \mathbf{I})^{-1}(x_k + \gamma b_i) - x_k \right)$$

$$1258 \quad = x_k + \alpha \left(\frac{1}{S} \sum_{i=1}^S ((\gamma\mathbf{A}_i + \mathbf{I})^{-1} - \mathbf{I}) \right) x_k - \alpha \left(\frac{1}{S} \sum_{i=1}^S ((\gamma\mathbf{A}_i + \mathbf{I})^{-1} - \mathbf{I}) \right) \Pi(x_k),$$

1259 where we use the fact that $\mathbf{A}_i x_* = b_i$ for all $x_* \in \mathcal{X}_*$. Therefore, using symmetry of \mathbf{A}_i

$$1260 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 \leq \|x_{k+1} - \Pi(x_k)\|^2$$

$$1261 \quad = \left\| \left((1 - \alpha)\mathbf{I} + \alpha \frac{1}{n} \sum_{i=1}^n (\gamma\mathbf{A}_i + \mathbf{I})^{-1} \right) (x_k - \Pi(x_k)) \right\|^2$$

$$1262 \quad = \left\| \left(\mathbf{I} - \alpha \underbrace{\left(\frac{1}{n} \sum_{i=1}^n (\mathbf{I} - (\gamma\mathbf{A}_i + \mathbf{I})^{-1}) \right)}_{:=\gamma\mathbf{M}} \right) (x_k - \Pi(x_k)) \right\|^2$$

$$1263 \quad = (x_k - \Pi(x_k))^T (\mathbf{I} - \alpha\gamma\mathbf{M})^T (\mathbf{I} - \alpha\gamma\mathbf{M}) (x_k - \Pi(x_k))$$

$$1264 \quad = (x_k - \Pi(x_k))^T (\mathbf{I} - 2\alpha\gamma\mathbf{M} + \alpha^2\gamma^2\mathbf{M}^2) (x_k - \Pi(x_k)).$$

1265 Since $\mathbf{M}^2 \leq \lambda_{\max}(\gamma)\mathbf{M}$, it follows that

$$1266 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 \leq (x_k - \Pi(x_k))^T (\mathbf{I} - 2\alpha\gamma\mathbf{M} + \lambda_{\max}(\gamma)\alpha^2\gamma^2\mathbf{M}) (x_k - \Pi(x_k))$$

$$1267 \quad = \|x_k - \Pi(x_k)\|^2 - \alpha\gamma(2 - \lambda_{\max}(\gamma)\alpha\gamma)(x_k - \Pi(x_k))^T \mathbf{M} (x_k - \Pi(x_k))$$

$$1268 \quad \leq \|x_k - \Pi(x_k)\|^2 - \alpha\gamma(2 - \lambda_{\max}(\gamma)\alpha\gamma)\lambda_{\min}^+(\gamma) \|x_k - \Pi(x_k)\|^2$$

1269 for any $\alpha\gamma$ such that $2 - \lambda_{\max}(\gamma)\alpha\gamma \geq 0$, where the last inequality follows from the fact that
 1270 $x_k - \Pi(x_k) \in \text{range}(\mathbf{M})$ (see Lemma C.6). Consequently,

$$1271 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 \leq (1 - \alpha\gamma(2 - \alpha\gamma\lambda_{\max}(\gamma))\lambda_{\min}^+(\gamma)) \|x_k - \Pi(x_k)\|^2.$$

It remains to unroll the recurrence and substitute the optimal step size $\alpha\gamma = \frac{1}{\lambda_{\max}(\gamma)}$ to obtain the rate

$$\|x_{k+1} - \Pi(x_{k+1})\|^2 \leq \left(1 - \frac{\lambda_{\min}^+(\gamma)}{\lambda_{\max}(\gamma)}\right) \|x_k - \Pi(x_k)\|^2.$$

□

Theorem 4.1. Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization problem (1) where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Under Assumption 1.3, FedExProx with $\alpha = 1/\gamma L_\gamma$ finds \bar{x} such that $\mathbb{E}[f(\bar{x})] - f(x_*) \leq \varepsilon$ after

$$\mathcal{O}\left(\frac{L_\gamma}{\mu_\gamma^+} \log \frac{1}{\varepsilon}\right) \quad (10)$$

iterations, where L_γ is a smoothness constant of M^γ and μ_γ^+ is the smallest non-zero eigenvalue of the matrix $\nabla^2 M^\gamma$.

Proof. The result follows directly from Theorem C.1, smoothness of f , and Fact 1 by substituting $\lambda_{\max}(\gamma) = \lambda_{\max}(\nabla^2 M^\gamma) = L_\gamma$ and $\lambda_{\min}^+(\gamma) = \lambda_{\min}^+(\nabla^2 M^\gamma) = \mu_\gamma^+$. □

Remark C.2. Recall from the proof of Fact 2 that

$$\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top,$$

where $\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \in \left[\frac{\lambda_j(\mathbf{A}_i)}{2}, \lambda_j(\mathbf{A}_i) \right]$ if $\gamma \leq \frac{1}{\lambda_j(\mathbf{A}_i)}$. Consequently, for $\gamma \leq \frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$, we have

$$\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top \preceq \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i [\lambda_j(\mathbf{A}_i)]_{jj} \mathbf{Q}_i^\top = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i = \mathbf{A},$$

and analogously, $\mathbf{M} \succeq \frac{1}{2} \mathbf{A}$. Using (27) and (28) gives

$$\frac{1}{2} \lambda_{\max}(\mathbf{A}) = \frac{1}{2} \max_{\|x\| \leq 1} x^\top \mathbf{A} x \leq \lambda_{\max}(\gamma) = \max_{\|x\| \leq 1} x^\top \mathbf{M} x \leq \max_{\|x\| \leq 1} x^\top \mathbf{A} x = \lambda_{\max}(\mathbf{A}),$$

and similarly

$$\frac{1}{2} \lambda_{\min}^+(\mathbf{A}) \leq \lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in (\ker \mathbf{A})^\perp} x^\top \mathbf{M} x \leq \lambda_{\min}^+(\mathbf{A}).$$

Noting that $L_\gamma = \lambda_{\max}(\gamma)$, $\mu_\gamma^+ = \lambda_{\min}^+(\gamma)$, $L = \lambda_{\max}(\mathbf{A})$ and $\mu^+ = \lambda_{\min}^+(\mathbf{A})$, it follows that the rates (10) of FedExProx and (32) of GD coincide up to a constant factor when γ is sufficiently small. This should come as no surprise, as Theorem 3.2 shows that FedExProx effectively reduces to GD as $\gamma \rightarrow 0$.

Time complexity. Let us now prove the result from Section 4.1. We assume that each worker i computes the proximal operator using an iterative method, with the running time proportional to the condition number of the subproblem. Given that $\lambda_{\min}(\mathbf{A}_i) = 0$ for all $i \in [n]$ (otherwise the problem becomes trivial), the time required for all clients to compute $\text{prox}_{\gamma f_i}(x_k)$ at each global iteration k is

$$\tau \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)\right)$$

seconds, where τ is the time per one iteration of solving the subproblem. Then, it takes μ seconds to aggregate the updates at the server and move on to the next step. Since, according to Theorem C.1, the number of iterations needed to reach an ε -solution is

$$K = \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\varepsilon} \right),$$

the total time required to solve the global problem is

$$T_\mu(\gamma) := \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\varepsilon} \right) \times \left(\mu + \tau \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)\right) \right). \quad (31)$$

Theorem 4.3. Consider the non-strongly convex quadratic optimization problem from Theorem 4.1. Up to a constant factor, the time complexity (12) is minimized by $\gamma \in \left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right]$ if $\frac{\mu}{\tau} \geq 2$ and by $\gamma \in \left[0, \max \left\{ 0, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right\} \right]$ if $\frac{\mu}{\tau} < 2$, and $T_\mu(\gamma) \leq T_{\text{GD}}$.

Proof. Recall that $\lambda_{\max}(\gamma)$ and $\lambda_{\min}^+(\gamma)$ are the eigenvalues of the matrix

$$\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \stackrel{(2)}{=} \frac{1}{n} \sum_{i=1}^n \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top.$$

Since if $\gamma \geq \frac{1}{\lambda_j(\mathbf{A}_i)}$ and $\lambda_j(\mathbf{A}_i) > 0$, then $\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \in [\frac{1}{2}, 1)$, and if $\gamma \leq \frac{1}{\lambda_j(\mathbf{A}_i)}$, then $\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \in \left[\frac{\gamma \lambda_j(\mathbf{A}_i)}{2}, \gamma \lambda_j(\mathbf{A}_i) \right)$, it follows that

$$\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} = \begin{cases} \Theta\left(\frac{1}{\gamma}\right), & \gamma \geq \frac{1}{\lambda_j(\mathbf{A}_i)}, \\ \Theta(\lambda_j(\mathbf{A}_i)), & \gamma < \frac{1}{\lambda_j(\mathbf{A}_i)}. \end{cases} \quad (32)$$

Moreover, the identities (27) and (28) from Fact 2 tell us that

$$\lambda_{\max}(\gamma) = \max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x$$

and

$$\lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in \ker(\mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x,$$

where $\mathbf{A} := \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$. Due to (27), (28) and (32), the function $T_\mu(\gamma)$ is approximately constant for all $\gamma \leq \frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$. If $\gamma \geq \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$ then $\tau(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)) \geq \mu$, and consequently,

$$\begin{aligned} T_\mu(\gamma) &= \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right) \times \left(\mu + \tau \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) \right) \\ &\leq 2\tau \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right) \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) := 2\bar{g}(\gamma). \end{aligned}$$

Hence, $\bar{g}(\gamma) \leq T_\mu(\gamma) \leq 2\bar{g}(\gamma)$, and the term to be minimized is

$$\bar{T}_\mu(\gamma) := \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right).$$

Using (27) and (28), this can be written as

$$\bar{T}_\mu(\gamma) = \frac{\max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \lambda_j(\mathbf{A}_i) \right]_{jj} \mathbf{Q}_i^\top x}{\min_{\|x\|=1, x \in \ker(\mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x},$$

where the numerator is non-decreasing and the denominator is non-increasing as a function of γ . It follows that $\bar{g}(\gamma)$ is non-decreasing for $\gamma \geq \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$. Hence, when $\frac{\mu}{\tau} \geq 2$, the optimal (up to a multiplicative factor) γ belongs to the interval

$$\left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \right]. \quad (33)$$

When $\frac{\mu}{\tau} < 2$, the time-complexity is non-decreasing for $\gamma \geq \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$ and approximately constant otherwise, so the choice

$$\gamma \in \left[0, \max \left\{ 0, \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \right\} \right] \quad (34)$$

Algorithm 2 FedExProx with partial participation

-
- 1: **Parameters:** stepsize $\gamma > 0$, extrapolation parameter $\alpha_k > 0$, starting point $x_0 \in \mathbb{R}^d$, batch size S
 - 2: **for** $k = 0, 1, 2, \dots$ **do**
 - 3: Sample a minibatch $\mathcal{S}_k \subseteq [n]$ uniformly from all subsets of cardinality S
 - 4: $x_{k+1} = x_k + \alpha_k \left(\frac{1}{S} \sum_{i \in \mathcal{S}_k} \text{prox}_{\gamma f_i}(x_k) - x_k \right)$
 - 5: **end for**
-

is optimal (again, up to a constant factor).

Now, suppose that $\gamma \geq \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}$. In this case, the diagonal entries of the matrix $\left[\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj}$ are either to 0 (when $\lambda_j(\mathbf{A}_i) = 0$) or lie within the interval $[\frac{1}{2}, 1)$. Therefore, the ratio

$$\frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} = \frac{\max_{\|x\| \leq 1} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x}{\min_{\|x\|=1, x \in \ker(\mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x}$$

is approximately constant in γ . Consequently, the time

$$T_\mu(\gamma) := \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right) \times \left(\mu + \tau \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) \right)$$

is an increasing function of γ for $\gamma \geq \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}$, and the optimal range for γ can be tightened from the previously established bounds (33) and (34) to

$$\left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right]$$

when $\frac{\mu}{\tau} \geq 2$ and

$$\gamma \in \left[0, \max \left\{ 0, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right\} \right]$$

when $\frac{\mu}{\tau} < 2$.

□

C.3.2 PARTIAL PARTICIPATION CASE

Iteration complexity. We again start with establishing the iteration complexity of the algorithm.

Theorem C.3. *Let Assumption 1.3 hold and let x_k be the iterates of minibatch FedExProx (Algorithm 2) applied to problem (20). Then*

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq (1 - \alpha\gamma (2 - \alpha\gamma L_{\gamma,S}) \lambda_{\min}^+(\gamma))^K \|x_0 - \Pi(x_0)\|^2,$$

where $\lambda_{\min}^+(\gamma)$ is the smallest positive eigenvalue of the matrix $\mathbf{M} := \frac{1}{\gamma} (\mathbf{I} - \frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1})$

and $L_{\gamma,S} := \frac{n-S}{S(n-1)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} + \frac{n(S-1)}{S(n-1)} \lambda_{\max}(\gamma)$.

The optimal choice of stepsize γ and extrapolation parameter α is $\alpha\gamma = \frac{1}{L_{\gamma,S}}$, in which case the rate becomes

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \frac{\lambda_{\min}^+(\gamma)}{L_{\gamma,S}} \right)^K \|x_0 - \Pi(x_0)\|^2.$$

Hence, the number of iterations needed to reach an ϵ -solution is

$$K \geq \frac{L_{\gamma,S}}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right). \quad (35)$$

1458 *Remark C.4.* Similar to the proof of Theorem 4.1, the result of Theorem 5.1 follows directly from
 1459 Theorem C.3, smoothness of f , and Fact 1 by noting that $\lambda_{\max}(\gamma) = \lambda_{\max}(\nabla^2 M\gamma) = L_\gamma$ and
 1460 $\lambda_{\min}^+(\gamma) = \lambda_{\min}^+(\nabla^2 M\gamma) = \mu_\gamma^+$.
 1461

1462 *Proof of Theorem C.3.* The update rule of FedExProx can be written as
 1463

$$\begin{aligned}
 1464 \quad x_{k+1} &= x_k + \alpha \left(\frac{1}{S} \sum_{i \in \mathcal{S}_k} \text{prox}_{\gamma f_i}(x_k) - x_k \right) \\
 1465 &= x_k + \alpha \left(\frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} (x_k + \gamma b_i) - x_k \right) \\
 1466 &= x_k + \alpha \left(\frac{1}{S} \sum_{i \in \mathcal{S}_k} ((\gamma \mathbf{A}_i + \mathbf{I})^{-1} - \mathbf{I}) x_k - \alpha \left(\frac{1}{S} \sum_{i \in \mathcal{S}_k} ((\gamma \mathbf{A}_i + \mathbf{I})^{-1} - \mathbf{I}) \Pi(x_k), \right. \right. \\
 1467 & \\
 1468 & \\
 1469 & \\
 1470 & \\
 1471 & \\
 1472 &
 \end{aligned}$$

1473 where we use the fact that $\mathbf{A}_i x_* = b_i$ for all $x_* \in \mathcal{X}_*$. Therefore
 1474

$$\begin{aligned}
 1475 \quad & \|x_{k+1} - \Pi(x_{k+1})\|^2 \\
 1476 & \leq \|x_{k+1} - \Pi(x_k)\|^2 \\
 1477 & = \left\| \left((1-\alpha)\mathbf{I} + \alpha \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) (x_k - \Pi(x_k)) \right\|^2 \\
 1478 & = \|x_k - \Pi(x_k)\|^2 - 2\alpha (x_k - \Pi(x_k))^T \left(\mathbf{I} - \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) (x_k - \Pi(x_k)) \\
 1479 & \quad + \alpha^2 (x_k - \Pi(x_k))^T \left(\mathbf{I} - \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right)^T \left(\mathbf{I} - \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) (x_k - \Pi(x_k)) \\
 1480 & = \|x_k - \Pi(x_k)\|^2 - 2\alpha\gamma (x_k - \Pi(x_k))^T \mathbf{M}_k (x_k - \Pi(x_k)) \\
 1481 & \quad + \alpha^2 \gamma^2 (x_k - \Pi(x_k))^T \mathbf{M}_k^T \mathbf{M}_k (x_k - \Pi(x_k)), \tag{36} \\
 1482 & \\
 1483 & \\
 1484 & \\
 1485 & \\
 1486 & \\
 1487 & \\
 1488 & \\
 1489 &
 \end{aligned}$$

1490 where $\mathbf{M}_k := \frac{1}{\gamma} \left(\mathbf{I} - \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right)$. Next, using the fact that $\mathbf{M}^2 \preceq \lambda_{\max}(\gamma)\mathbf{M}$ and
 1491 applying Lemma A.12 with $\mathbf{B}_i = \frac{1}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right)$, the expectation of $\mathbf{M}_k^T \mathbf{M}_k$ is
 1492

$$\begin{aligned}
 1493 \quad \mathbb{E} [\mathbf{M}_k^T \mathbf{M}_k] &= \frac{1}{\gamma^2} \frac{n-S}{S(n-1)} \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})^\top (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) + \frac{n(S-1)}{S(n-1)} \mathbf{M}^\top \mathbf{M} \\
 1494 &= \frac{1}{\gamma^2} \frac{n-S}{S(n-1)} \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})^2 + \frac{n(S-1)}{S(n-1)} \mathbf{M}^2 \\
 1495 &\leq \frac{1}{\gamma} \frac{n-S}{S(n-1)} \max_{i \in [n]} \lambda_{\max}(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \\
 1496 &\quad + \frac{n(S-1)}{S(n-1)} \lambda_{\max}(\gamma)\mathbf{M} \\
 1497 &= \frac{1}{\gamma} \frac{n-S}{S(n-1)} \max_{i \in [n]} \left(\frac{\gamma \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \lambda_{\max}(\mathbf{A}_i)} \right) \mathbf{M} + \frac{n(S-1)}{S(n-1)} \lambda_{\max}(\gamma)\mathbf{M} \\
 1498 &= \underbrace{\left(\frac{n-S}{S(n-1)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} + \frac{n(S-1)}{S(n-1)} \lambda_{\max}(\gamma) \right)}_{:=L_{\gamma,S}} \mathbf{M}. \tag{37} \\
 1499 & \\
 1500 & \\
 1501 & \\
 1502 & \\
 1503 & \\
 1504 & \\
 1505 & \\
 1506 & \\
 1507 & \\
 1508 & \\
 1509 &
 \end{aligned}$$

1510 Hence, taking expectation conditioned on x_k in (36) gives
 1511

$$\mathbb{E}_k \left[\|x_{k+1} - \Pi(x_{k+1})\|^2 \right]$$

$$\begin{aligned}
&\leq \|x_k - \Pi(x_k)\|^2 + \alpha^2 \gamma^2 (x_k - \Pi(x_k))^T \mathbb{E}_k [\mathbf{M}_k^T \mathbf{M}_k] (x_k - \Pi(x_k)) \\
&\quad - 2\alpha\gamma (x_k - \Pi(x_k))^T \mathbb{E}_k [\mathbf{M}_k] (x_k - \Pi(x_k)) \\
&\stackrel{(37)}{=} \|x_k - \Pi(x_k)\|^2 + \alpha^2 \gamma^2 L_{\gamma,S} (x_k - \Pi(x_k))^T \mathbf{M} (x_k - \Pi(x_k)) \\
&\quad - 2\alpha\gamma (x_k - \Pi(x_k))^T \mathbf{M} (x_k - \Pi(x_k)) \\
&= \|x_k - \Pi(x_k)\|^2 - \alpha\gamma (2 - \alpha\gamma L_{\gamma,S}) (x_k - \Pi(x_k))^T \mathbf{M} (x_k - \Pi(x_k)) \\
&\leq \|x_k - \Pi(x_k)\|^2 - \alpha\gamma (2 - \alpha\gamma L_{\gamma,S}) \lambda_{\min}^+(\gamma) (x_k - \Pi(x_k))^T (x_k - \Pi(x_k)),
\end{aligned}$$

where the last inequality follows from the fact that $x_k - \Pi(x_k) \in \text{range}(\mathbf{M})$ (see Lemma C.6). Taking expectation again,

$$\mathbb{E} \left[\|x_{k+1} - \Pi(x_{k+1})\|^2 \right] \leq (1 - \alpha\gamma (2 - \alpha\gamma L_{\gamma,S}) \lambda_{\min}^+(\gamma)) \mathbb{E} \left[\|x_k - \Pi(x_k)\|^2 \right].$$

It remains to unroll the recurrence. \square

Time complexity. Let us consider the same setup as in Section C.3.1, i.e., at each iteration k , each client computes a proximal operator $\text{prox}_{\gamma f_i}(x_k)$ in at most

$$\tau \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right)$$

seconds, where τ is the time per one iteration of solving the subproblem, and it takes μ seconds to aggregate the updates at the server and move on to the next step. According to Theorem C.3, the number of iterations needed to reach an ϵ -solution is

$$K = \frac{L_{\gamma,S}}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right),$$

where $L_{\gamma,S} := \frac{n-S}{S(n-1)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} + \frac{n(S-1)}{S(n-1)} \lambda_{\max}(\gamma)$. Hence, the total time required to solve the global problem is at most

$$T_\mu(\gamma, S) := \frac{L_{\gamma,S}}{\lambda_{\min}^+(\gamma)} \log \left(\frac{\|x_0 - \Pi(x_0)\|^2}{\epsilon} \right) \times \left(\mu + \tau \times \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) \right). \quad (38)$$

Theorem 5.2. *Up to a constant factor, the time complexity (14) is minimized by $\gamma \in \left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right]$ if $\frac{\mu}{\tau} \geq 2$ and by $\gamma \in \left[0, \max \left\{ 0, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right\} \right]$ if $\frac{\mu}{\tau} < 2$.*

Proof. Fix some $S \in [n]$. Then, we are interested in minimizing

$$\begin{aligned}
\bar{T}_\mu(\gamma) &= \frac{n-S}{S(n-1)} \frac{1}{\lambda_{\min}^+(\gamma)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \underbrace{\left(\mu + \tau \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) \right)}_{:= \bar{g}_1(\gamma)} \\
&\quad + \frac{n(S-1)}{S(n-1)} \frac{\lambda_{\max}(\gamma)}{\lambda_{\min}^+(\gamma)} \underbrace{\left(\mu + \tau \left(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \right) \right)}_{:= \bar{g}_2(\gamma)} \\
&:= \frac{n-S}{S(n-1)} \bar{g}_1(\gamma) + \frac{n(S-1)}{S(n-1)} \bar{g}_2(\gamma).
\end{aligned}$$

We know from Theorem 4.3 that the second term of this expression, $\bar{g}_2(\gamma)$, is minimized by

$$\left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right]$$

when $\frac{\mu}{\tau} \geq 2$ and

$$\gamma \in \left[0, \max \left\{ 0, \min \left\{ \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)} \right\} \right\} \right]$$

when $\frac{\mu}{\tau} < 2$.

The same argument can be applied to the term $\bar{g}_1(\gamma)$. To be more precise, due to (27), (28), and (32), the function $\bar{g}_1(\gamma)$ is approximately constant for all $\gamma \leq \frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$. For $\gamma > \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$, we have $\frac{\mu}{\tau} + 1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) < 2(1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i))$. Consequently, letting $\bar{g}(\gamma) := 2\tau \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{\lambda_{\min}^+(\gamma)}$, $\bar{g}_1(\gamma)$ can be bounded as

$$\frac{1}{2} \bar{g}(\gamma) < \bar{g}_1(\gamma) < \bar{g}(\gamma).$$

Since $\lambda_{\min}^+(\gamma)$ is a non-increasing function of γ , $\bar{g}(\gamma)$ is non-decreasing. It follows that for $\gamma > \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}$, $\bar{g}_1(\gamma)$ is non-decreasing in γ (up to a constant factor), and we arrive at the same conclusions as in the case of $\bar{g}_2(\gamma)$: $\bar{g}_1(\gamma)$ is minimized by $\gamma \in \left[\frac{1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}, \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \right]$ if $\frac{\mu}{\tau} \geq 2$ and by $\gamma \in \left[0, \frac{\frac{\mu}{\tau} - 1}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \right]$ if $\frac{\mu}{\tau} < 2$.

On the other hand, by following a similar argument as in the proof of Theorem 4.3, the expression

$$\frac{1}{\lambda_{\min}^+(\gamma)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} = \frac{\frac{\gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}}{\min_{\|x\|=1, x \in \ker(\mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x}$$

is approximately constant for $\gamma \geq \frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}$. This implies that $\bar{g}_1(\gamma)$ is increasing, and hence the intervals can be bounded above by $\frac{1}{\min_{i \in [n]} \lambda_{\min}^+(\mathbf{A}_i)}$.

The conclusion follows from the fact that the minimum of a convex combination of $\bar{g}_1(\gamma)$ and $\bar{g}_2(\gamma)$ must lie within the same interval. \square

Remark C.5. The only term in (35) and (38) that depends on S is

$$L_{\gamma, S} := \frac{n - S}{S(n - 1)} \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} + \frac{n(S - 1)}{S(n - 1)} \lambda_{\max}(\gamma).$$

Recall that $\lambda_{\max}(\gamma)$ is the largest eigenvalue of the matrix $\mathbf{M} = \frac{1}{n} \sum_{i=1}^n \frac{1}{2\gamma} (\mathbf{I} - (\mathbf{I} + \gamma \mathbf{A}_i)^{-1})$, so

$$\lambda_{\max}(\gamma) \leq \frac{1}{\gamma} \frac{1}{n} \sum_{i=1}^n \frac{\lambda_{\max}(\mathbf{A}_i) \gamma}{\lambda_{\max}(\mathbf{A}_i) \gamma + 1} \leq \frac{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{\max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i) \gamma + 1}.$$

As a result, since $\frac{n-S}{S(n-1)}$ is decreasing, and $\frac{n(S-1)}{S(n-1)}$ is increasing in S , both the iteration and time complexities are increasing functions of S . This underscores the advantage of involving a larger number of clients in the training process.

C.4 LEMMAS

Lemma C.6. *Let x_k be the iterates of SGD applied to the problem $\min_x \{f(x) := \frac{1}{n} \sum_{i=1}^n (\frac{1}{2} x^\top \mathbf{B}_i x + c_i^\top x + d_i)\}$, where the matrices \mathbf{B}_i are symmetric. Then $x_k - \Pi(x_k) \in \text{range}(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i)$ for all k .*

Proof. By definition of Π , we have

$$\Pi(x) = x - \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\dagger \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{B}_i x + c_i) \right).$$

Hence, using the identity $\mathbf{A}^\top(\mathbf{A}\mathbf{A}^\top)^\dagger = \mathbf{A}^\dagger$,

$$\begin{aligned} x_k - \Pi(x_k) &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\dagger \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{B}_i x_k + c_i) \right) \\ &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\top \left(\left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right) \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \right)^\top \right)^\dagger \left(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i \mathbb{E}[x_k] - c_i \right), \end{aligned}$$

and so, by symmetry of \mathbf{B}_i , we get $x_k - \Pi(x_k) \in \text{range}(\frac{1}{n} \sum_{i=1}^n \mathbf{B}_i)$. \square

Lemma C.7. Let $\mathbf{M} := \frac{1}{n} \sum_{i=1}^n \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})$, where $\gamma > 0$ and $\mathbf{A}_i \in \text{Sym}_+^d$ for all $i \in [n]$. The smallest positive eigenvalue of \mathbf{M} is given by

$$\lambda_{\min}^+(\gamma) := \min_{\|x\|=1, x \in (\ker \mathbf{A})^\perp} \frac{1}{n} \sum_{i=1}^n x^\top \mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top x,$$

where $\mathbf{A} = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$ and $[b_j]_{jj}$ denotes a diagonal matrix with b_j as the j th entry.

Proof. First, observe that the matrices $\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}$ are symmetric, and their eigenvalues are given by

$$1 - \frac{1}{1 + \gamma \lambda_j(\mathbf{A}_i)} = \frac{\gamma \lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)}.$$

Consequently, \mathbf{M} is a sum of symmetric positive definite matrices and is therefore also symmetric positive definite. We now claim that

$$\lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in (\ker \mathbf{M})^\perp} x^\top \mathbf{M} x.$$

First, choosing x to be a multiple of the eigenvector of \mathbf{M} corresponding to $\lambda_{\min}^+(\gamma)$, we see that

$$\lambda_{\min}^+(\gamma) \geq \min_{\|x\|=1, x \in (\ker \mathbf{M})^\perp} x^\top \mathbf{M} x.$$

To establish the reverse inequality, let $\{e_i\}$ be an orthonormal eigenbasis of \mathbf{M} and let x be such that $\|x\| = 1$. Then we can write $x = \sum_{i=1}^d \alpha_i e_i$, and since $x \in (\ker \mathbf{M})^\perp$, all coefficients corresponding to an eigenvalue 0 vanish. Thus

$$x^\top \mathbf{M} x = \sum_{i=1}^d \alpha_i^2 \lambda_i \geq \lambda_{\min}^+(\gamma) \sum_{i=1}^d \alpha_i^2 = \lambda_{\min}^+(\gamma) \|x\|^2 = \lambda_{\min}^+(\gamma).$$

This proves that

$$\lambda_{\min}^+(\gamma) = \min_{\|x\|=1, x \in (\ker \mathbf{M})^\perp} x^\top \mathbf{M} x.$$

Now, each matrix $\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}$ can be decomposed as $\mathbf{Q}_i \left[\frac{\lambda_j(\mathbf{A}_i)}{1 + \gamma \lambda_j(\mathbf{A}_i)} \right]_{jj} \mathbf{Q}_i^\top$. Therefore what remains to be proven is that $\ker(\mathbf{M}) = \ker(\mathbf{A})$. To this end, take $x \in \ker(\mathbf{M})$. Then $\mathbf{M}x = 0$, and since $\gamma > 0$, we have

$$\frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1} x = x. \quad (39)$$

Next, observe that

$$\lambda_{\max} \left(\frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) \leq \frac{1}{n} \sum_{i=1}^n \lambda_{\max} ((\gamma \mathbf{A}_i + \mathbf{I})^{-1}) = \frac{1}{n} \sum_{i=1}^n \frac{1}{1 + \gamma \lambda_{\min}(\mathbf{A}_i)}.$$

Now consider two cases.

1674 First, if there exists $j \in [n]$ such that $\lambda_{\min}(\mathbf{A}_j) > 0$, the above upper bound is strictly less than 1.
 1675 This implies that there exists no nonzero x that satisfies equation (39), and so $\ker(\mathbf{M}) = \{0\}$. But
 1676 $\lambda_{\min}(\mathbf{A}_j) > 0$ also implies that

$$1677 \lambda_{\min} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{A}_i \right) \geq \frac{1}{n} \sum_{i=1}^n \lambda_{\min}(\mathbf{A}_i) > 0.$$

1680 As a result, $\ker(\mathbf{M}) = \{0\} = \ker(\mathbf{A})$.

1682 Now, let us suppose that $\lambda_{\min}(\mathbf{A}_i) = 0$ for all $i \in [n]$. In this case, $\lambda_{\max}((\gamma\mathbf{A}_i + \mathbf{I})^{-1}) = 1$ for
 1683 all $i \in [n]$. Since all matrices $(\gamma\mathbf{A}_i + \mathbf{I})^{-1}$ are symmetric positive definite with maximum eigenvalue
 1684 equal to 1, it follows that equation (39) holds if and only if

$$1686 (\gamma\mathbf{A}_i + \mathbf{I})^{-1}x = x$$

1687 for all $i \in [n]$. This is equivalent to $\gamma\mathbf{A}_i x = 0$ for all $i \in [n]$. Since $\gamma > 0$ and \mathbf{A}_i are symmetric
 1688 positive definite, we can, in turn, equivalently express it as $\mathbf{A}x = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i x = 0$.

1689 Consequently, (39) holds (i.e., $x \in \ker(\mathbf{M})$) if and only if $x \in \ker(\mathbf{A})$ and hence $\ker(\mathbf{A}) = \ker(\mathbf{M})$.

1691 The final expression in the statement of the Lemma follows from eigendecomposition (see Fact 2).
 1692 □

D FedExProx UNDER PL CONDITION

Theorem D.1. *Let Assumptions 1.1, 1.2, 1.3, and 7.1 hold, and assume that $\alpha\gamma \leq \frac{2}{L_{\gamma,S}}$. Then the iterates of Algorithm 2 satisfy*

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \alpha\gamma(2 - \alpha\gamma L_{\gamma,S}) \frac{\mu_\gamma^+}{2} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right]$$

and hence

$$\mathbb{E} [f(x_K) - f(\Pi(x_K))] \leq \frac{L}{2} \left(1 - \alpha\gamma(2 - \alpha\gamma L_{\gamma,S}) \frac{\mu_\gamma^+}{2} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right],$$

where $L_{\gamma,S} := \left(\frac{n-S}{S(n-1)} \frac{L_{\max}}{1+\gamma L_{\max}} + \frac{n(S-1)}{S(n-1)} L_\gamma \right)$. For the optimal choice of stepsize and extrapolation parameter $\alpha\gamma = 1/L_{\gamma,S}$, these rates become

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \frac{\mu_\gamma^+}{2L_{\gamma,S}} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right]$$

and

$$\mathbb{E} [f(x_K) - f(\Pi(x_K))] \leq \frac{L}{2} \left(1 - \frac{\mu_\gamma^+}{2L_{\gamma,S}} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right].$$

Remark D.2. The above theorem proves Theorem 7.5. Furthermore, by setting $S = n$, we recover the result from Theorem 7.2.

Proof. The proof closely follows the proof of Theorem 3 from Li et al. (2024). Recall from Section B that Algorithm 2 is equivalent to SGD for minimizing $M^\gamma(x) := \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x)$ with stepsize $\alpha\gamma$, and its updates can be written as

$$x_{k+1} = x_k - \alpha\gamma \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k).$$

Then, for any $x_* \in \mathcal{X}_*$

$$\begin{aligned} & \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] \\ &= \|x_k - x_*\|^2 - 2\alpha\gamma \left\langle x_k - x_*, \mathbb{E}_k \left[\frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k) \right] \right\rangle + \alpha^2 \gamma^2 \mathbb{E}_k \left[\left\| \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \right] \\ &\stackrel{\text{(A.8)}}{=} \|x_k - x_*\|^2 - 2\alpha\gamma \langle x_k - x_*, \nabla M^\gamma(x_k) - \nabla M^\gamma(x_*) \rangle + \alpha^2 \gamma^2 \mathbb{E}_k \left[\left\| \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \right] \\ &= \|x_k - x_*\|^2 - 2\alpha\gamma (D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k)) + \alpha^2 \gamma^2 \mathbb{E}_k \left[\left\| \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \right], \quad (40) \end{aligned}$$

where $D_{M^\gamma}(x, y) := M^\gamma(x) - M^\gamma(y) - \langle \nabla M^\gamma(y), x - y \rangle$. Next, applying Lemma A.12 with $\mathbf{B}_i = \nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*) \in \mathbb{R}^{d \times 1}$, the last term in the above inequality can be written as

$$\begin{aligned} & \mathbb{E}_k \left[\left\| \frac{1}{S} \sum_{i \in \mathcal{S}_k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \right] \stackrel{\text{(A.5)}}{=} \mathbb{E}_k \left[\left\| \frac{1}{S} \sum_{i \in \mathcal{S}_k} (\nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*)) \right\|^2 \right] \\ & \stackrel{\text{(A.12)}}{=} \frac{n-S}{S(n-1)} \frac{1}{n} \sum_{i=1}^n \left\| \nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*) \right\|^2 \\ & \quad + \frac{n(S-1)}{S(n-1)} \left\| \frac{1}{n} \sum_{i=1}^n (\nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*)) \right\|^2. \end{aligned}$$

Looking at the first term of the inequality above, using convexity and smoothness of the functions $M_{f_i}^\gamma$ (Lemmas A.4 and A.6), we have

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \left\| \nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*) \right\|^2 &\stackrel{\text{(A.14)}}{\leq} \frac{1}{n} \sum_{i=1}^n \frac{L_i}{1 + \gamma L_i} \left(D_{M_{f_i}^\gamma}(x_k, x_*) + D_{M_{f_i}^\gamma}(x_*, x_k) \right) \\ &\leq \frac{L_{\max}}{1 + \gamma L_{\max}} \frac{1}{n} \sum_{i=1}^n \left(D_{M_{f_i}^\gamma}(x_k, x_*) + D_{M_{f_i}^\gamma}(x_*, x_k) \right) \\ &= \frac{L_{\max}}{1 + \gamma L_{\max}} \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right). \end{aligned}$$

Next, since by Lemma A.7, the function M^γ is convex and smooth, the second term can be bounded as

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \left(\nabla M_{f_i}^\gamma(x_k) - \nabla M_{f_i}^\gamma(x_*) \right) \right\|^2 &= \left\| \nabla M^\gamma(x_k) - \nabla M^\gamma(x_*) \right\|^2 \\ &\stackrel{\text{(A.14)}}{\leq} L_\gamma \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right). \end{aligned}$$

Applying these bounds in (40) gives

$$\begin{aligned} &\mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] \\ &\leq \|x_k - x_*\|^2 - 2\alpha\gamma \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right) \\ &\quad + \alpha^2 \gamma^2 \underbrace{\left(\frac{n-S}{S(n-1)} \frac{L_{\max}}{1 + \gamma L_{\max}} + \frac{n(S-1)}{S(n-1)} L_\gamma \right)}_{:=L_{\gamma,S}} \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right) \\ &= \|x_k - x_*\|^2 - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right). \end{aligned}$$

By Lemma A.7, M^γ is convex and smooth, so $D_{M^\gamma}(x_*, x_k) \geq 0$, and by Assumption 7.1 and Lemma A.13, we have

$$D_{M^\gamma}(x_k, x_*) = M^\gamma(x_k) - M^\gamma(x_*) \geq \frac{\mu_\gamma^+}{2} \|x_k - \Pi(x_k)\|^2.$$

Therefore, for $\alpha\gamma L_{\gamma,S} \leq 2$

$$\begin{aligned} \mathbb{E}_k \left[\|x_{k+1} - x_*\|^2 \right] &\leq \|x_k - x_*\|^2 - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \left(D_{M^\gamma}(x_k, x_*) + D_{M^\gamma}(x_*, x_k) \right) \\ &\leq \|x_k - x_*\|^2 - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \frac{\mu_\gamma^+}{2} \|x_k - \Pi(x_k)\|^2. \end{aligned}$$

Taking expectation and letting $x_* = \Pi(x_k)$ gives

$$\begin{aligned} \mathbb{E} \left[\|x_{k+1} - \Pi(x_{k+1})\|^2 \right] &\leq \mathbb{E} \left[\|x_{k+1} - \Pi(x_k)\|^2 \right] \\ &\leq \mathbb{E} \left[\|x_k - \Pi(x_k)\|^2 \right] - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \frac{\mu_\gamma^+}{2} \mathbb{E} \left[\|x_k - \Pi(x_k)\|^2 \right]. \end{aligned}$$

Unrolling the recurrence, we obtain the first result

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \frac{\mu_\gamma^+}{2} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right].$$

Lastly, using L -smoothness of f , it follows that

$$\begin{aligned} \mathbb{E} [f(x_K) - f(\Pi(x_K))] &\leq \frac{L}{2} \mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \\ &\leq \frac{L}{2} \left(1 - \alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right) \frac{\mu_\gamma^+}{2} \right)^K \mathbb{E} \left[\|x_0 - \Pi(x_0)\|^2 \right]. \end{aligned}$$

Substituting $\alpha\gamma = 1/L_{\gamma,S}$, which minimizes the expression $\alpha\gamma \left(2 - \alpha\gamma L_{\gamma,S} \right)$, finishes the proof. \square

1836 E FedExProx WITH INEXACT COMPUTATIONS

1837
1838 In practice, solving (3) exactly is often infeasible, and we can only find a vector $\text{prox}_{\gamma f}^{\delta}(x)$ such that

$$1839 \quad \|\text{prox}_{\gamma f}^{\delta}(x) - \text{prox}_{\gamma f}(x)\|^2 \leq \delta,$$

1840 where δ represents the accuracy of the approximate solution to (3). As a result, we can only calculate
1841 an inexact gradient of the Moreau envelope, defined as

$$1842 \quad \nabla M_f^{\gamma, \delta}(x) := \frac{1}{\gamma}(x - \text{prox}_{\gamma f}^{\delta}(x)).$$

1843 One can easily show that

$$1844 \quad \|\nabla M_f^{\gamma, \delta}(x) - \nabla M_f^{\gamma}(x)\|^2 \leq \frac{\delta}{\gamma^2}.$$

1845 With these inexact updates, Algorithm 1 iterates

$$1846 \quad x_{k+1} = x_k + \alpha_k \left(\frac{1}{n} \sum_{i=1}^n \text{prox}_{\gamma f_i}^{\delta}(x_k) - x_k \right) = x_k - \alpha_k \gamma \underbrace{\frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^{\gamma, \delta}(x_k)}_{:= \nabla M^{\gamma, \delta}(x_k)}, \quad (41)$$

1847 where

$$1848 \quad \begin{aligned} \|\nabla M^{\gamma, \delta}(x) - \nabla M^{\gamma}(x)\|^2 &= \left\| \frac{1}{n} \sum_{i=1}^n \left(\nabla M_{f_i}^{\gamma, \delta}(x) - \nabla M_{f_i}^{\gamma}(x) \right) \right\|^2 \\ &= \frac{1}{\gamma^2} \left\| \frac{1}{n} \sum_{i=1}^n \left(\text{prox}_{\gamma f}^{\delta}(x) - \text{prox}_{\gamma f}(x) \right) \right\|^2 \\ &\leq \frac{1}{\gamma^2} \frac{1}{n} \sum_{i=1}^n \|\text{prox}_{\gamma f}^{\delta}(x) - \text{prox}_{\gamma f}(x)\|^2 \leq \frac{\delta}{\gamma^2}. \end{aligned} \quad (42)$$

1849 With this in mind, we proceed to our analysis.

1850 **Theorem E.1.** Consider the inexact FedExProx method in the full participation setting defined
1851 in (41). Let Assumptions 1.1, 1.2, 1.3, and 7.1 hold, and set $\alpha\gamma = \frac{1}{2L\gamma}$. Then, the iterates of
1852 the algorithm satisfy

$$1853 \quad \|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{\mu_{\gamma}^+}{8L\gamma}\right)^K \|x_0 - \Pi(x_0)\|^2 + \frac{20\delta}{(\mu_{\gamma}^+)^2 \gamma^2},$$

1854 and hence

$$1855 \quad f(x_K) - f(\Pi(x_K)) \leq \frac{L}{2} \left(1 - \frac{\mu_{\gamma}^+}{8L\gamma}\right)^K \|x_0 - \Pi(x_0)\|^2 + \frac{10L\delta}{(\mu_{\gamma}^+)^2 \gamma^2}. \quad (43)$$

1856 *Remark E.2.* Note that the price one has to pay for inexactness is minimal. To illustrate this, suppose
1857 that clients solve the local problems using GD. Then, taking

$$1858 \quad \delta = \frac{(\mu_{\gamma}^+)^2 \gamma^2 \varepsilon}{20L},$$

1859 we ensure that the last term in (43) is small:

$$1860 \quad \frac{10L\delta}{(\mu_{\gamma}^+)^2 \gamma^2} \leq \frac{\varepsilon}{2}.$$

1861 Since problem (3) is strongly convex, the local complexity is proportional to

$$1862 \quad \mathcal{O}\left(\kappa \times \log \frac{1}{\delta}\right) = \mathcal{O}\left(\kappa \times \log \left(\frac{20L}{(\mu_{\gamma}^+)^2 \gamma^2 \varepsilon}\right)\right),$$

1863 where $\kappa \gg 1$ is the condition number. In practice, this logarithmic factor can be ignored and treated
1864 as a constant.

1890 *Proof of Theorem E.1.* The updates can be written as

$$1891 \quad x_{k+1} = x_k - \alpha\gamma \nabla M^{\gamma, \delta}(x_k).$$

1893 Thus

$$1894 \quad \begin{aligned} 1895 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \|x_{k+1} - \Pi(x_k)\|^2 \\ 1896 \quad &= \|x_k - \Pi(x_k)\|^2 - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^{\gamma, \delta}(x_k) \rangle + \alpha^2\gamma^2 \|\nabla M^{\gamma, \delta}(x_k)\|^2 \\ 1897 \quad &= \|x_k - \Pi(x_k)\|^2 - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^\gamma(x_k) \rangle \\ 1898 \quad &\quad - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k) \rangle \\ 1899 \quad &\quad + \alpha^2\gamma^2 \|\nabla M^\gamma(x_k) + (\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k))\|^2. \end{aligned}$$

1902 Using Jensen's inequality, we get

$$1903 \quad \begin{aligned} 1904 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \|x_k - \Pi(x_k)\|^2 - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^\gamma(x_k) \rangle \\ 1905 \quad &\quad - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k) \rangle \\ 1906 \quad &\quad + 2\alpha^2\gamma^2 \|\nabla M^\gamma(x_k)\|^2 + 2\alpha^2\gamma^2 \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2. \end{aligned} \quad (44)$$

1908 Following the same steps as in the proof of Theorem D.1 gives

$$1909 \quad \begin{aligned} 1910 \quad \|x_k - \Pi(x_k)\|^2 - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^\gamma(x_k) \rangle + 2\alpha^2\gamma^2 \|\nabla M^\gamma(x_k)\|^2 \\ 1911 \quad \leq (1 - \alpha\gamma(1 - \alpha\gamma L_\gamma)\mu_\gamma^+) \|x_k - \Pi(x_k)\|^2 \\ 1912 \quad \leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{2}\right) \|x_k - \Pi(x_k)\|^2 \end{aligned}$$

1914 for $\alpha\gamma \leq \frac{1}{2L_\gamma}$, and hence

$$1915 \quad \begin{aligned} 1916 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 &\stackrel{(44)}{\leq} \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{2}\right) \|x_k - \Pi(x_k)\|^2 \\ 1917 \quad &\quad - 2\alpha\gamma \langle x_k - \Pi(x_k), \nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k) \rangle \\ 1918 \quad &\quad + 2\alpha^2\gamma^2 \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2. \end{aligned}$$

1922 Young's inequality yields

$$1923 \quad \begin{aligned} 1924 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{2}\right) \|x_k - \Pi(x_k)\|^2 + \alpha\gamma\lambda \|x_k - \Pi(x_k)\|^2 \\ 1925 \quad &\quad + \frac{\alpha\gamma}{\lambda} \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2 + 2\alpha^2\gamma^2 \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2, \end{aligned}$$

1928 where $\lambda > 0$ is a free parameter. Taking $\lambda = \frac{\mu_\gamma^+}{4}$, we obtain

$$1929 \quad \begin{aligned} 1930 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{4}\right) \|x_k - \Pi(x_k)\|^2 + \frac{4\alpha\gamma}{\mu_\gamma^+} \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2 \\ 1931 \quad &\quad + 2\alpha^2\gamma^2 \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2 \\ 1932 \quad &\leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{4}\right) \|x_k - \Pi(x_k)\|^2 + \frac{5\alpha\gamma}{\mu_\gamma^+} \|\nabla M^{\gamma, \delta}(x_k) - \nabla M^\gamma(x_k)\|^2, \end{aligned}$$

1936 where the last inequality due to $\alpha\gamma \leq \frac{1}{2L_\gamma}$ and $L_\gamma \geq \mu_\gamma^+$. Lastly, due to (42), we have

$$1937 \quad \|x_{k+1} - \Pi(x_{k+1})\|^2 \leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{4}\right) \|x_k - \Pi(x_k)\|^2 + \frac{5\alpha\delta}{\mu_\gamma^+ \gamma}.$$

1941 Unrolling the recursion, one can show that

$$1942 \quad \|x_K - \Pi(x_K)\|^2 \leq \left(1 - \alpha\gamma \frac{\mu_\gamma^+}{4}\right)^K \|x_0 - \Pi(x_0)\|^2 + \frac{20\delta}{(\mu_\gamma^+)^2 \gamma^2},$$

1944 and hence

$$\begin{aligned} 1945 & \\ 1946 & f(x_K) - f(\Pi(x_K)) \leq \frac{L}{2} \|x_K - \Pi(x_K)\|^2 \\ 1947 & \\ 1948 & \leq \frac{L}{2} \left(1 - \alpha\gamma\frac{\mu_\gamma^+}{4}\right)^K \|x_0 - \Pi(x_0)\|^2 + \frac{10L\delta}{(\mu_\gamma^+)^2 \gamma^2}. \\ 1949 & \\ 1950 & \end{aligned}$$

1951 \square

1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997

F ADAPTIVE EXTRAPOLATION

The optimal extrapolation parameter values in Theorems 4.1, 5.1, 7.2 and 7.5 depend on L_γ (or $L_{\gamma,S}$). This dependency can be avoided by employing *adaptive* extrapolation strategies (Horváth et al., 2022). Specifically, we analyze methods based on gradient diversity (FedExProx-GraDS) and the Polyak stepsize (FedExProx-StoPS).

F.1 FULL PARTICIPATION CASE

Theorem 6.1. Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization problem (1), where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Let Assumption 1.3 hold and consider two adaptive extrapolation strategies:

1. (FedExProx-GraDS) Set

$$\alpha_k = \alpha_k^{\text{GraDS}}(x_k) := \frac{\frac{1}{n} \sum_{i=1}^n \|\nabla M_{f_i}^\gamma(x_k)\|^2}{\left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2} \geq 1.$$

Then, the iterates of Algorithm 1 satisfy

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \min_{k=0, \dots, K-1} \alpha_k \gamma \frac{2+\gamma L_{\max}}{1+\gamma L_{\max}} \mu_\gamma^+\right)^K \|x_0 - \Pi(x_0)\|^2. \quad (15)$$

2. (FedExProx-StoPS) Set

$$\alpha_k = \alpha_k^{\text{StoPS}}(x_k) := \frac{\frac{1}{n} \sum_{i=1}^n (M_{f_i}^\gamma(x_k) - \inf M_{f_i}^\gamma)}{\gamma \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2} \geq \frac{1}{2\gamma L_\gamma}.$$

Then, the iterates of Algorithm 1 satisfy

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{3}{2} \min_{k=0, \dots, K-1} \alpha_k \gamma \mu_\gamma^+\right)^K \|x_0 - \Pi(x_0)\|^2. \quad (16)$$

Proof. We start with the standard decomposition

$$\begin{aligned} \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \|x_{k+1} - \Pi(x_k)\|^2 \\ &= \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma \left\langle \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k), x_k - \Pi(x_k) \right\rangle \\ &\quad + \alpha_k^2 \gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\ &\stackrel{(23)}{=} \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\ &\quad + \alpha_k^2 \gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2. \end{aligned} \quad (45)$$

Now, let us consider the two adaptive extrapolation strategies.

1. (FedExProx-GraDS) In this case, the last term of (45) can be rewritten as

$$\begin{aligned} &\alpha_k^2 \gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\ &= \alpha_k \gamma^2 \frac{1}{n} \sum_{i=1}^n \left\| \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\ &\stackrel{(23)}{=} \alpha_k \gamma^2 \frac{1}{n} \sum_{i=1}^n \left\| \frac{1}{\gamma} \left(\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right) (x_k - \Pi(x_k)) \right\|^2 \end{aligned}$$

$$\begin{aligned}
&= \alpha_k (x_k - \Pi(x_k))^\top \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})^\top (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \right) (x_k - \Pi(x_k)) \\
&\leq \alpha_k \max_{i \in [n]} \lambda_{\max} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) (x_k - \Pi(x_k))^\top \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \right) (x_k - \Pi(x_k)) \\
&= \alpha_k \gamma \frac{\gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)).
\end{aligned}$$

Substituting this in (45) gives

$$\begin{aligned}
&\|x_{k+1} - \Pi(x_{k+1})\|^2 \\
&\leq \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
&\quad + \alpha_k \gamma \frac{\gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
&= \|x_k - \Pi(x_k)\|^2 - \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
&\leq \left(1 - \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \lambda_{\min}^+(\gamma) \right) \|x_k - \Pi(x_k)\|^2,
\end{aligned}$$

where the last inequality follows from the fact that $x_k - \Pi(x_k) \in \text{range}(\mathbf{M})$ (see Lemma C.6). Applying this bound iteratively for $k = 0, \dots, K-1$ gives

$$\begin{aligned}
&\|x_K - \Pi(x_K)\|^2 \\
&\leq \left(1 - \gamma \min_{k=0, \dots, K-1} \alpha_k \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \lambda_{\min}^+(\gamma) \right)^K \|x_0 - \Pi(x_0)\|^2.
\end{aligned}$$

2. (FedExProx-StoPS) First, note that L_γ -smoothness of M^γ implies that

$$M^\gamma(x_k) - \inf M^\gamma \geq \frac{1}{2L_\gamma} \|\nabla M^\gamma(x_k)\|^2,$$

and hence $\alpha_{k,S}(x_k) \geq \frac{1}{2\gamma L_\gamma}$.

Now, the last term of (45) is

$$\begin{aligned}
\alpha_k^2 \gamma^2 \left\| \frac{1}{n} \sum_{i=1}^n \nabla M_{f_i}^\gamma(x_k) \right\|^2 &= \alpha_k \gamma \frac{1}{n} \sum_{i=1}^n \left(M_{f_i}^\gamma(x_k) - \inf M_{f_i}^\gamma \right) \\
&= \alpha_k \gamma (M^\gamma(x_k) - M^\gamma(\Pi(x_k))) \\
&\stackrel{(26)}{=} \frac{1}{2} \alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)).
\end{aligned}$$

Therefore, using Lemma C.6, we get

$$\begin{aligned}
\|x_{k+1} - \Pi(x_{k+1})\|^2 &\stackrel{(45)}{\leq} \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
&\quad + \frac{1}{2} \alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
&\leq \left(1 - \frac{3}{2} \alpha_k \gamma \lambda_{\min}^+(\gamma) \right) \|x_k - \Pi(x_k)\|^2.
\end{aligned}$$

Applying this bound iteratively for $k = 0, \dots, K-1$ gives

$$\|x_K - \Pi(x_K)\|^2 \leq \left(1 - \frac{3}{2} \min_{k=0, \dots, K-1} \alpha_k \gamma \lambda_{\min}^+(\gamma) \right)^K \|x_0 - \Pi(x_0)\|^2.$$

2105

□

2106 F.2 PARTIAL PARTICIPATION CASE
2107

2108 **Theorem F.1.** Fix any $\gamma > 0$ and consider solving non-strongly convex quadratic optimization
2109 problem (1), where $f_i(x) = \frac{1}{2}x^\top \mathbf{A}_i x - b_i^\top x$ for all $i \in [n]$, with $\mathbf{A}_i \in \text{Sym}_+^d$ and $b_i \in \mathbb{R}^d$. Let
2110 Assumption 1.3 hold and consider two adaptive extrapolation strategies:

2111 1. (FedExProx-GraDS) Set

$$2112 \alpha_k = \alpha_{k,S}^{\text{GraDS}}(x_k, \mathcal{S}_k) := \frac{\frac{1}{S} \sum_{i \in \mathcal{S}^k} \left\| \nabla M_{f_i}^\gamma(x_k) \right\|^2}{\left\| \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2} \geq 1. \quad 2113$$

2114 Then, the iterates of Algorithm 2 satisfy

$$2115 \mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \inf \alpha_k \gamma \frac{2 + \gamma L_{\max}}{1 + \gamma L_{\max}} \lambda_{\min}^+(\gamma) \right)^K \|x_0 - \Pi(x_0)\|^2, \quad (46)$$

2116 where

$$2117 \inf \alpha_k := \inf_{x \in \mathbb{R}^d, \mathcal{S} \subseteq [n], |\mathcal{S}|=S} \alpha_{k,S}^{\text{GraDS}}(x, \mathcal{S}).$$

2118 2. (FedExProx-StoPS) Set

$$2119 \alpha_k = \alpha_{k,S}^{\text{StoPS}}(x_k, \mathcal{S}_k) := \frac{\frac{1}{S} \sum_{i \in \mathcal{S}^k} \left(M_{f_i}^\gamma(x_k) - \inf M_{f_i}^\gamma \right)}{\gamma \left\| \frac{1}{n} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2}. \quad 2120$$

2121 Then, the iterates of Algorithm 2 satisfy

$$2122 \mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \frac{3}{2} \inf \alpha_k \gamma \lambda_{\min}^+(\gamma) \right)^K \|x_0 - \Pi(x_0)\|^2, \quad (47)$$

2123 where

$$2124 \inf \alpha_k := \inf_{x \in \mathbb{R}^d, \mathcal{S} \subseteq [n], |\mathcal{S}|=S} \alpha_{k,S}^{\text{StoPS}}(x, \mathcal{S}).$$

2125 **Remark F.2.** In the single node setting ($S = 1$),

$$2126 \inf_{x \in \mathbb{R}^d, \mathcal{S} \subseteq [n], |\mathcal{S}|=1} \alpha_{k,1}^{\text{GraDS}}(x, \mathcal{S}) = 1. \quad 2127$$

2128 However, as more clients participate, the extrapolation parameter may exceed 1, resulting in im-
2129 proved performance of FedExProx-GraDS.

2130 **Remark F.3.** In the single node setting ($S = 1$), Lemma A.4 implies that

$$2131 \inf_{x \in \mathbb{R}^d, \mathcal{S} \subseteq [n], |\mathcal{S}|=1} \alpha_{k,1}^{\text{StoPS}}(x, \mathcal{S}) = \frac{1}{2} \frac{1 + \gamma L_{\max}}{\gamma L_{\max}}. \quad 2132$$

2133 As the number of participating clients increases, the extrapolation parameter may exceed this bound,
2134 leading to better performance of FedExProx-StoPS.

2135 **Proof of Theorem F.1.** As in the full batch case, we start with the decomposition

$$2136 \begin{aligned} 2137 \|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \|x_{k+1} - \Pi(x_k)\|^2 \\ 2138 &= \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma \left\langle \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k), x_k - \Pi(x_k) \right\rangle \\ 2139 &\quad + \alpha_k^2 \gamma^2 \left\| \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \end{aligned} \quad 2140$$

$$\begin{aligned}
& \stackrel{(23)}{=} \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)) \\
& \quad + \alpha_k^2 \gamma^2 \left\| \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2, \tag{48}
\end{aligned}$$

where $\mathbf{M}_k := \frac{1}{\gamma} (\mathbf{I} - \frac{1}{S} \sum_{i \in \mathcal{S}^k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1})$. Now, let us consider the two adaptive extrapolation strategies.

1. (FedExProx-GraDS) In this case, the last term of (48) can be rewritten as

$$\begin{aligned}
& \alpha_k^2 \gamma^2 \left\| \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\
& = \alpha_k \gamma^2 \frac{1}{S} \sum_{i \in \mathcal{S}^k} \left\| \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\
& \stackrel{(23)}{=} \alpha_k \gamma^2 \frac{1}{S} \sum_{i \in \mathcal{S}^k} \left\| \frac{1}{\gamma} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) (x_k - \Pi(x_k)) \right\|^2 \\
& = \alpha_k (x_k - \Pi(x_k))^\top \left(\frac{1}{S} \sum_{i \in \mathcal{S}^k} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1})^\top (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \right) (x_k - \Pi(x_k)) \\
& \leq \alpha_k \max_{i \in [n]} \lambda_{\max} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) (x_k - \Pi(x_k))^\top \left(\frac{1}{S} \sum_{i \in \mathcal{S}^k} (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) \right) (x_k - \Pi(x_k)) \\
& = \alpha_k \gamma \frac{\gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)).
\end{aligned}$$

Substituting this in (48) gives

$$\begin{aligned}
& \|x_{k+1} - \Pi(x_{k+1})\|^2 \\
& \leq \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)) \\
& \quad + \alpha_k \gamma \frac{\gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)) \\
& = \|x_k - \Pi(x_k)\|^2 - \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)) \\
& \leq \|x_k - \Pi(x_k)\|^2 - \inf \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M}_k (x_k - \Pi(x_k)),
\end{aligned}$$

where we view α_k as a function of x and \mathcal{S} . Taking expectation conditioned on x_k , we obtain

$$\begin{aligned}
& \mathbb{E}_k \left[\|x_{k+1} - \Pi(x_{k+1})\|^2 \right] \\
& \leq \|x_k - \Pi(x_k)\|^2 - \inf \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbb{E}_k [\mathbf{M}_k] (x_k - \Pi(x_k)) \\
& = \|x_k - \Pi(x_k)\|^2 - \inf \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} (x_k - \Pi(x_k))^\top \mathbf{M} (x_k - \Pi(x_k)) \\
& \leq \left(1 - \inf \alpha_k \gamma \frac{2 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)}{1 + \gamma \max_{i \in [n]} \lambda_{\max}(\mathbf{A}_i)} \lambda_{\min}^+(\gamma) \right) \|x_k - \Pi(x_k)\|^2,
\end{aligned}$$

where the last inequality follows from the fact that $x_k - \Pi(x_k) \in \text{range}(\mathbf{M})$ (see Lemma C.6). Taking full expectation and applying this bound iteratively for $k = 0, \dots, K-1$ gives the final result.

2. (FedExProx-StoPS) In this case, the last term of (48) is

$$\begin{aligned}
& \alpha_k^2 \gamma^2 \left\| \frac{1}{S} \sum_{i \in \mathcal{S}^k} \nabla M_{f_i}^\gamma(x_k) \right\|^2 \\
&= \alpha_k \gamma \frac{1}{S} \sum_{i \in \mathcal{S}^k} \left(M_{f_i}^\gamma(x_k) - \inf M_{f_i}^\gamma \right) \\
&\stackrel{(24)}{=} \frac{1}{2} \alpha_k \frac{1}{S} \sum_{i \in \mathcal{S}^k} (x_k - \Pi(x_k))^T (\mathbf{I} - (\gamma \mathbf{A}_i + \mathbf{I})^{-1}) (x_k - \Pi(x_k)) \\
&= \frac{1}{2} \alpha_k \gamma (x_k - \Pi(x_k))^T \mathbf{M}_k (x_k - \Pi(x_k)).
\end{aligned}$$

Therefore, (48) gives

$$\begin{aligned}
\|x_{k+1} - \Pi(x_{k+1})\|^2 &\leq \|x_k - \Pi(x_k)\|^2 - 2\alpha_k \gamma (x_k - \Pi(x_k))^T \mathbf{M}_k (x_k - \Pi(x_k)) \\
&\quad + \frac{1}{2} \alpha_k \gamma (x_k - \Pi(x_k))^T \mathbf{M}_k (x_k - \Pi(x_k)) \\
&\leq \|x_k - \Pi(x_k)\|^2 - \frac{3}{2} \inf \alpha_k \gamma (x_k - \Pi(x_k))^T \mathbf{M}_k (x_k - \Pi(x_k)).
\end{aligned}$$

Taking full expectation and applying Lemma C.6, we obtain

$$\mathbb{E} \left[\|x_{k+1} - \Pi(x_{k+1})\|^2 \right] \leq \left(1 - \frac{3}{2} \inf \alpha_k \gamma \lambda_{\min}^+(\gamma) \right) \|x_k - \Pi(x_k)\|^2,$$

and hence

$$\mathbb{E} \left[\|x_K - \Pi(x_K)\|^2 \right] \leq \left(1 - \frac{3}{2} \inf \alpha_k \gamma \lambda_{\min}^+(\gamma) \right)^K \|x_0 - \Pi(x_0)\|^2.$$

□

G NOTATION

Notation	
$\ \cdot\ $	Standard Euclidean norm
$\langle \cdot, \cdot \rangle$	Standard Euclidean inner product
$[k]$	$:= \{1, \dots, k\}$
$\mathbb{E}_k[\cdot]$	Expectation conditioned on the first k iterations
n	Number of workers/nodes/clients/devices
d	Dimensionality of the problem
γ	Stepsize
α_k	Extrapolation parameter
L, L_i, L_{\max}	Smoothness constants (Assumption 1.2)
$\text{prox}_{\gamma f}(x)$	Proximal operator of the function f with parameter γ (see (3))
M_f^γ	Moreau envelope of the function f with parameter γ (see (7))
M^γ	$:= \frac{1}{n} \sum_{i=1}^n M_{f_i}^\gamma$
L_γ	Smoothness constant of M^γ
μ_γ^+	The smallest non-zero eigenvalue of the matrix $\nabla^2 M^\gamma$ (see Theorems 4.1, 5.1) / PL constant (see Theorems 7.2, 7.5)
$\pi(\gamma)$	Total time per global iteration of Algorithm 1 / 2
$T_\mu(\gamma)$	Time complexity of Algorithm 1
$T_\mu(\gamma, S)$	Time complexity of Algorithm 2
\mathcal{X}_*	$:= \{x \in \mathbb{R}^d : \nabla f_i(x) = 0\}$
$\Pi(x)$	Projection of x onto the solution set \mathcal{X}_*
R^2	$:= \ x_* - x_0\ ^2$
$D_f(x, y)$	Bregman divergence between x and y associated with a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$
\mathbf{I}	Identity matrix
\mathbf{A}_i	Local data matrix stored on worker i , with spectral decomposition $\mathbf{A}_i = \mathbf{Q}_i \mathbf{\Lambda}_i \mathbf{Q}_i^\top$
\mathbf{A}	$:= \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$
Sym_+^d	$:= \{\mathbf{X} \in \mathbb{R}^{d \times d} \mid \mathbf{X} = \mathbf{X}^\top, \mathbf{X} \succeq 0\}$ - the set of symmetric positive semidefinite matrices
$\ker \mathbf{A}$	Kernel of a matrix \mathbf{A}
\mathbf{M}_k	$:= \frac{1}{\gamma} \left(\mathbf{I} - \frac{1}{S} \sum_{i \in S_k} (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right)$
\mathbf{M}	$:= \frac{1}{\gamma} \left(\mathbf{I} - \frac{1}{n} \sum_{i=1}^n (\gamma \mathbf{A}_i + \mathbf{I})^{-1} \right)$
$\lambda_{\max}(\mathbf{B})$	The largest eigenvalue of matrix \mathbf{B}
$\lambda_{\min}(\mathbf{B})$	The smallest eigenvalue of matrix \mathbf{B}
$\lambda_{\min}^+(\mathbf{B})$	The smallest positive eigenvalue of matrix \mathbf{B}
$[\mathbf{B}]_j$	The j th diagonal element of a diagonal matrix \mathbf{B}
$[b_j]_{jj}$	A diagonal matrix with b_j as the j th entry

Table 1: Frequently used notation.