

A Tight Convergence Analysis for Stochastic Gradient Descent with Delayed Updates

Yossi Arjevani

Ohad Shamir

Nathan Srebro

Weizmann Institute of Science

TTI Chicago

Rehovot 7610001, Israel

Chicago, IL 60637

{yossi.arjevani, ohad.shamir}@weizmann.ac.il

nati@ttic.edu

Abstract

We provide tight finite-time convergence bounds for gradient descent and stochastic gradient descent on quadratic functions, when the gradients are delayed and reflect iterates from τ rounds ago. First, we show that without stochastic noise, delays strongly affect the attainable optimization error: In fact, the error can be as bad as non-delayed gradient descent ran on only $1/\tau$ of the gradients. In sharp contrast, we quantify how stochastic noise makes the effect of delays negligible, improving on previous work which only showed this phenomenon asymptotically or for much smaller delays. Also, in the context of distributed optimization, the results indicate that the performance of gradient descent with delays is competitive with synchronous approaches such as mini-batching. Our results are based on a novel technique for analyzing convergence of optimization algorithms using generating functions.

h.OC] 26 Jun 201

1 Introduction

Gradient-based optimization methods are widely used in machine learning and other large-scale applications, due to their simplicity and scalability. However, in their standard formulation, they are also strongly synchronous and iterative in nature: In each iteration, the update step is based on the gradient at the current iterate, and we need to wait for this computation to finish before moving to the next iterate. For example, to minimize some function F , plain stochastic gradient descent initializes at some point \mathbf{w}_0 , and computes iterates of the form

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta(\nabla F(\mathbf{w}_k) + \xi_k), \quad (1)$$

where $\nabla F(\mathbf{w}_k)$ is the gradient of F at \mathbf{w}_k , η is the step size and ξ_1, ξ_2, \dots are independent zero-mean noise terms. Unfortunately, in several important applications, a direct implementation of this is too costly. For example, consider a setting where we wish to optimize a function F using a distributed platform, consisting of several machines with shared memory. We can certainly implement gradient descent, by letting one of the machines compute the gradient at each iteration, but this is clearly wasteful, since just one machine is non-idle at any given time. Thus, it is highly desirable to use methods which parallelize the computation. One approach is to employ *mini-batch gradient* methods, which parallelize the computation of the stochastic gradient, and their analysis is relatively well understood (e.g. [6, 5, 19, 24]). However, these methods are still generally iterative and synchronous in nature, and hence can suffer from problems such as having to wait for the slowest machine at each iteration.

A second and popular approach is to utilize *asynchronous* gradient methods. With these methods, each update step is not necessarily based just on the gradient of the current iterate, but possibly on the gradients of earlier iterates (often called *stale updates*). For example, when optimizing a function using several machines, each machine might read the current iterate from a shared parameter server, compute

the gradient at that iterate, and then update the parameters, even though other machines might have performed other updates to the parameters in the meantime. Although such asynchronous methods often work well in practice, analyzing them is much trickier than synchronous methods.

In our work, we focus on arguably the simplest possible variant of these methods, where we perform plain stochastic gradient descent on a convex function F on \mathbb{R}^d , with a fixed delay of $\tau > 0$ in the gradient computation:

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta(\nabla F(\mathbf{w}_{k-\tau}) + \xi_k), \quad (2)$$

where we assume that $\mathbf{w}_0 = \mathbf{w}_1 = \dots = \mathbf{w}_\tau$. Compared to Eq. (1), we see that the gradient is computed with respect to $\mathbf{w}_{k-\tau}$ rather than \mathbf{w}_k . Already in this simple formulation, the precise effect of the delay on the convergence rate is not completely clear. For example, for a given number of iterations k , how large can τ be before we might expect a significant deterioration in the accuracy? And under what conditions? Although there exist some prior results in this direction (which we survey in the related work section below), these questions have remained largely open.

In this paper, we aim at providing a tight, finite-time convergence analysis for stochastic gradient descent with delays, focusing on the simple case where F is a convex quadratic function. Although a quadratic assumption is non-trivial, it arises naturally in problems such as least squares, and is an important case study since all smooth and convex function are locally quadratic close to their minimum (hence, our results should still hold in a local sense). In future work, we hope to show that our results are also applicable more generally.

First, we consider the case of *deterministic* delayed gradient descent (DGD, defined in Eq. (2) with $\xi_k = \mathbf{0}$). Assuming the step size η is chosen appropriately, we prove that

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq 5\mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \exp\left(-\frac{\lambda(k+1)}{10\mu(\tau+1)}\right)$$

after k iterations, over the class of λ -strongly convex μ -smooth quadratic functions with a minimum at \mathbf{w}^* , and

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq \frac{17\mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2 (\tau+1)}{k+1}$$

over the class of μ -smooth convex quadratic functions with minimum at \mathbf{w}^* . In terms of iteration complexity, the number of iterations k required to achieve a fixed optimization error of at most ϵ in the strongly convex and the convex cases is therefore

$$\mathcal{O}\left(\tau + \kappa \ln\left(\frac{\mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2}{\epsilon}\right)\right) \quad \text{and} \quad \mathcal{O}\left(\tau + \frac{\mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2}{\epsilon}\right) \quad (3)$$

respectively, where $\kappa := \mu/\lambda$ is the so-called condition number¹. When τ is a bounded constant, these bounds match the known iteration complexity of standard gradient descent without delays [17]. However, as τ increases, both bounds deteriorate linearly with τ . Notably, in our setting of delayed gradients, this

¹ Following standard convention, we use here the \mathcal{O} -notation to hide constants, and tilde $\tilde{\mathcal{O}}$ -notation to hide constants and factors polylogarithmic in the problem parameters.

implies that DGD is no better than a trivial algorithm, which performs a single gradient step, and then waits for τ rounds till the delayed gradient is received, before performing the next step (thus, the algorithm is equivalent to non-delayed gradient descent with k/τ gradient steps, resulting in the same linear deterioration of the iteration complexity with τ).

Despite these seemingly weak guarantees, we show that they are in fact tight in terms of τ , by proving that this linear dependence on τ is unavoidable with standard gradient-based methods (including gradient descent). The dependence on the other problem parameters in our lower bounds is a bit weaker than our upper bounds, but can be matched by an *accelerated* gradient descent procedure (see Sec. 3 for more details).

In the second part of our paper, we consider the case of *stochastic* delayed gradient descent (SDGD, defined in (2)). Assuming ξ_k satisfies $E[k\xi_k k^2] \leq \sigma^2$ and that the step size η is appropriately tuned, we prove that

$$E[F(\mathbf{w}_k) - F(\mathbf{w}^*)] \leq O\left(\frac{\sigma^2}{\lambda k} + \mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right)\right). \quad (4)$$

for λ -strongly convex, μ -smooth quadratic functions with minimum at \mathbf{w}^* , and

$$E[F(\mathbf{w}_k) - F(\mathbf{w}^*)] \leq O\left(\frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2 \sigma^2}{\lambda k} + \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2 \mu}{k}\right) \quad (5)$$

for μ -smooth convex quadratic functions. In terms of iteration complexity, these correspond to

$$O\left(\frac{\sigma^2}{\lambda \rho} + \tau \kappa \ln\left(\frac{\lambda \|\mathbf{w}_0 - \mathbf{w}^*\|^2}{\rho}\right)\right) \quad \text{and} \quad O\left(\frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2 \sigma^2}{\lambda \rho} + \tau \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2 \mu}{\rho}\right) \quad (6)$$

in the strongly convex and convex cases respectively, where again $\kappa := \mu/\lambda$. As in the deterministic case, when τ is a bounded constant, these bounds match the known iteration complexity bounds for standard gradient descent without delays [3, 20]. Moreover, these bounds match the bounds for the deterministic case in Eq. (3) when $\sigma^2 = 0$ (i.e. zero noise), as they should. However, in sharp contrast to the deterministic case, the dependence on τ in Eq. (6) is quite different: The delay τ only appears in second-order terms (as $\rho \rightarrow 0$), and its influence becomes negligible when ρ is small enough. The same effect can be seen in Eq. (4) and Eq. (5): Once the number of iterations k is large enough, the first term in both bounds dominates, and τ no longer plays a role. More specifically:

- In the strongly convex case, the effect of the delay becomes negligible once the target accuracy ρ is sufficiently smaller than $O(\sigma^2/(\mu\tau))$, or when the number of iterations k is sufficiently larger than $\Omega(\tau\mu/\lambda)$. In other words, assuming the condition number μ/λ is bounded, we can have the delay τ nearly as large as the total number of iterations k (up to log-factors), without significant deterioration in the convergence rate. Note that this is a mild requirement, since if $\tau \geq k$, the algorithm receives no gradients and makes no updates.
- In the convex case, the effect of the delay becomes negligible once the target accuracy ρ is sufficiently smaller than $O(\sigma^2/(\mu\tau))$, or when the number of iterations k is sufficiently larger than $\Omega(\tau\mu/\lambda)$.

$\mathbf{w}^*k\mu\tau/\sigma^2$). Compared to the strongly convex case, here the regime is the same in terms of ρ , but the regime in terms of k is more restrictive: We need k to scale quadratically (rather than linearly) with τ .

Thus, the maximal delay τ with no performance deterioration is order of \sqrt{k} .

Finally, it is interesting to compare our bounds to those of *mini-batch* stochastic gradient descent (SGD), which can be seen as a synchronous gradient-based method to cope with delays, especially in distributed optimization and learning problems [6, 5, 1]. In mini-batch SGD, each update step is performed only after accumulating and averaging a mini-batch of b stochastic gradients, all with respect to the same point:

$$k \in \{0, b, 2b, \dots\}, \mathbf{w}_{k+b} = \mathbf{w}_k - \eta \cdot \frac{1}{b} \sum_{i=1}^b (\nabla F(\mathbf{w}_k) + \xi_{k+i}),$$

Although the algorithm makes an update only every b stochastic gradient computations, the averaging reduces the stochastic noise, and helps speed up convergence. Moreover, this can be seen as a particular type of algorithm with delayed updates (with the delay correspond to b), as we use $\nabla F(\mathbf{w}_k)$ to compute iterate \mathbf{w}_{k+b} . The important difference is that it is an inherently synchronous method, that waits for all b stochastic gradients to be computed before performing an update step. Remarkably, the bounds we proved above for delayed SGD are essentially identical to those known for mini-batch SGD, with the delay τ replaced by the mini-batch size b (at least in the convex case where mini-batch SGD has been more thoroughly analyzed). This indicates that an asynchronous method like delayed SGD can potentially match the performance of synchronous methods like mini-batch SGD, even without requiring synchronization – an important practical advantage.

Analyzing gradient descent with delays is notoriously tricky, due to the dependence of the updates on iterates produced many iterations ago. The technique we introduce for deriving our upper bounds is primarily based on *generating functions*, and might be useful for studying other optimization algorithms. We discuss this approach more thoroughly in Section 2. The rest of the paper is devoted mostly to presenting the formal theorems and an explanation of how they are derived (with technical details relegated to the supplementary material).

Related Work

There is a huge literature on asynchronous versions of gradient-based methods (see for example the seminal book [2]), including treating the effect of delay. However, most of these do not consider the setting we study here. For example, there has been much recent interest in asynchronous algorithms, in a model where there is a delay in updating individual *coordinates* in a shared parameter vector (e.g., the Hogwild! algorithm of [18], or more recently [14, 13]). Of course, this is a different model than ours, where the updates use a full gradient vector. Other works (such as [21]) focus on a setting where different agents in a network can perform local communication, which is again a different model than ours. Yet other works focus on sharp but asymptotic results, and do not provide guarantees after a fixed number k of iterations (e.g., [4]).

Moving closer to our setting, [15] showed convergence for delayed gradient descent, with the result shown in an adversarial online learning setting, for general convex functions, and this bound is known to be implying an τ/k convergence rate for convex functions. A similar bound on average regret has been optimal [10]. These results differ from our setting, in that they consider possibly non-smooth functions, in which the dependence on k is no better than $1/\sqrt{k}$ even without delays and no noise, and where the delay τ

always plays a significant role. In contrast, we focus here on smooth functions, where rates better than $1/\sqrt{k}$ are possible, and where the effect of τ is more subtle. In [7], the authors study a setting very similar to ours in the deterministic case, and manage to prove a linear convergence rate, but for a less standard algorithm, different than the one we study here (with iterates of the form $\mathbf{w}_{t+1} = \mathbf{w}_{t-\tau} - \nabla F(\mathbf{w}_{t-\tau})$).

Perhaps the works closest to ours are [1, 8], which study stochastic gradient descent with delayed gradients. Moreover, they consider a setting more general than ours, where the delay at each iteration is any integer up to τ (rather than fixed τ), and the functions are not necessarily quadratic. On the flip side, their bounds are significantly weaker. For example, for smooth convex functions and an appropriate step size, [1, Corollary 1] show a bound of

$$O\left(\frac{\sigma^2 + \tau + 1}{\sqrt{k} + \sigma^2 k}\right).$$

in terms of k, τ, σ . Note that this bound is vacuous in the deterministic or near-deterministic case (where $\sigma^2 \approx 0$), and is weaker than our bounds. With a different choice of the step size, it is possible to get a non-vacuous bound even if $\sigma^2 \rightarrow 0$, but the dependence on τ becomes even stronger. [8] improve the bound

to $O\left(\frac{\sigma^2 + \tau + 1}{\sqrt{k} + k}\right)$ and $O\left(\frac{\sigma^2 + \tau + 1}{k + k^2}\right)$.

On the convex and strongly convex case respectively. Even if $\sigma^2 = 0$, the iteration complexity is $O(\tau^2/\rho)$ and

(τ^2/ρ) , and implies a quadratic dependence on τ (whereas in our bounds the scaling is linear). When σ^2 is positive, the effect of delay on the bound is negligible only up to $O(\sqrt{k})$ (in contrast to $O(\sqrt{k})$ or even $O(k)$ in our bounds). We note that there are several other works which study a similar setting (such as [22]), but do not result in bounds which improve on the above.

Finally, we note that [12] attempt to show that for stochastic gradient descent with delayed updates, the dependence on the delay τ is negligible after sufficiently many iterations. Unfortunately, as pointed out in [1], the analysis contains a bug which make the results invalid.

2 Framework and the Generating Functions Approach

Throughout, we will assume that F is a convex quadratic function specified by

$$F(\mathbf{w}) = \frac{1}{2} \mathbf{w}^T A \mathbf{w} + \mathbf{b}^T \mathbf{w} + c \tag{7}$$

where $A \in \mathbb{R}^{d \times d}$ is a positive semi-definite matrix whose eigenvalues a_1, \dots, a_d are in $[0, \mu]$ (where μ is the smoothness parameter), $\mathbf{b} \in \mathbb{R}^d$ and $c \in \mathbb{R}$. To make the optimization problem meaningful, we further assume that F is bounded from below, which implies that it has some minimizer $\mathbf{w}^* \in \mathbb{R}^d$ at which the gradient vanishes (for completeness, we provide a proof in Lemma 3 in the supplementary material).

Letting $\mathbf{e}_k = \mathbf{w}_k - \mathbf{w}^*$, it is easily verified that

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) = \frac{1}{2} \|\nabla F(\mathbf{w}_k)\|^2 = \frac{1}{2} \|\nabla F \mathbf{e}_k\|^2, \tag{8}$$

so our goal will be to analyze the dynamics of \mathbf{e}_k .

To explain our technique, consider the iterates of DGD on the function F , which can be written as $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \nabla F(\mathbf{w}_{k-\tau}) = \mathbf{w}_k - \eta (A \mathbf{w}_{k-\tau} + \mathbf{b})$. Since $\nabla F(\mathbf{w}^*) = 0$, we have $\mathbf{w}^* = \mathbf{w}^* - \eta (A \mathbf{w}^* + \mathbf{b})$, by which it follows that the error term $\mathbf{e}_k = \mathbf{w}_k - \mathbf{w}^*$, satisfies the recursion $\mathbf{e}_{k+1} = \mathbf{e}_k - \eta A \mathbf{e}_{k-\tau}$, and (by definition of the algorithm) $\mathbf{e}_0 = \mathbf{e}_1 = \dots = \mathbf{e}_\tau$. By some simple arguments, our analysis then boils down to bounding the elements of the scalar-valued version of this sequence, namely

$$\begin{aligned} e_{k+1} &= e_k - \eta a_k e_{k-\tau} \\ e_k &= e_k - \alpha e_{k-\tau}, \quad k \geq \tau \end{aligned} \tag{9}$$

for some integer $\tau \geq 0$ and non-negative real number $\alpha \geq 0$. To analyze this sequence, we rely on tools from the area of generating functions, which have proven very effective in studying growth rates of sequences in many areas of mathematics. We now turn to briefly describe these functions and our approach (for general surveys on generating functions, see [25, 9, 23], to name a few).

Generally speaking, generating functions are formal power series associated with infinite sequences of numbers. Concretely, given a sequence (b_k) of numbers in a ring R , we define the corresponding generating function as a formal power series in z , defined as $\sum_{k=0}^{\infty} b_k z^k$. The set of all formal power series in z over R is denoted by $R[[z]]$. Moreover, given two power series defined by sequences (a_k) and (c_k) , we can define their addition as the power series corresponding to $(a_k + c_k)$, and their multiplication as the coefficients of the Cauchy product of the power series, namely $(\sum_{k=0}^{\infty} a_k z^k)(\sum_{k=0}^{\infty} c_k z^k) = \sum_{k=0}^{\infty} (\sum_{l=0}^k a_l c_{k-l}) z^k$.

In particular, over the reals, $\mathbb{R}[[z]]$ endowed with addition and multiplication is a commutative ring, and the set of matrices with elements in $\mathbb{R}[[z]]$ (with the standard addition and multiplication operations) forms a matrix algebra, denoted by $M(\mathbb{R}[[z]])$. We will often use the fact that any matrix, whose entries

are power series with scalar coefficients, can also be written as a power series with matrix-valued coefficients: More formally, $M(\mathbb{R}[[z]])$ is naturally identified with the ring of formal power series with real matrix coefficients $M(\mathbb{R})[[z]]$. To extract the coefficients of a given $M(z) \in M(\mathbb{R}[[z]])$, we shall use the conventional bracket notation $[z^k]M(z)$, defined to be a matrix whose entries are the k 'th coefficients of the respective formal power series.

Returning to Eq. (9), we write (b_k) as a formal power series denoted by $f(z)$, and proceed as follows,

$$\begin{aligned} f(z) &= \sum_{k=0}^{\infty} b_k z^k + \sum_{k=\tau+1}^{\infty} (b_{k-1} - \alpha b_{k-\tau-1}) z^k = \sum_{k=0}^{\tau} b_k z^k + \sum_{k=\tau+1}^{\infty} b_{k-1} z^k - \alpha \sum_{k=\tau+1}^{\infty} b_{k-\tau-1} z^k \\ &= \sum_{k=0}^{\tau} b_k z^k + z \left(f(z) - \sum_{k=0}^{\tau-1} b_k z^k \right) - \alpha z^{\tau+1} f(z) = b_0 + (z - \alpha z^{\tau+1}) f(z). \end{aligned} \tag{10}$$

Denoting

$$\pi_{\alpha}(z) := 1 - z + \alpha z^{\tau+1}$$

and rearranging terms gives

$$f(z) = \frac{b_0}{\pi_{\alpha}(z)} \implies b_k = [z^k] f(z) = [z^k] \frac{b_0}{\pi_{\alpha}(z)} \tag{11}$$

(by a well-known fact, $\pi_{\alpha}(z)$ is invertible in $\mathbb{R}[[z]]$, as its constant term 1 is trivially invertible in \mathbb{R} – see surveys mentioned above). We now see that the problem of bounding the coefficients (b_k) is reduced to that of estimating the coefficients of the rational function $1/\pi_{\alpha}(z)$, written as a power series. Note that for the analogous problem where the elements of the sequence are vectors $(b_k) \in \mathbb{R}^n$ and the factor α is replaced by αA for some square matrix A , the same derivation as above yields

$$\sum_{k=0}^{\infty} b_k z^k = (I - z + \alpha A z^{\tau+1})^{-1} b_0$$

(likewise, $I - z + Az$ is invertible in $M(\mathbb{R})[[z]]$ as its constant term I is invertible in $M(\mathbb{R})$).

To estimate the coefficients of $1/\pi_{\alpha}(z)$, we form its corresponding partial fraction decomposition. First, we note that as a polynomial of degree $\tau+1$, $\pi_{\alpha}(z)$ has $\tau+1$ roots $\zeta_1, \dots, \zeta_{\tau+1}$ (possibly complex-valued, and all non-zero since $\pi_{\alpha}(0) = 1$ for any $\alpha \in \mathbb{R}$). Assuming α is chosen so that all the roots are distinct (equivalently, $\pi'_{\alpha}(\zeta_i) \neq 0, \forall i \in \{1, \dots, \tau+1\}$), we have by a standard derivation

$$\frac{1}{\pi_{\alpha}(z)} = \sum_{i=1}^{\tau+1} \frac{1}{\pi'_{\alpha}(\zeta_i)(z - \zeta_i)} = \sum_{i=1}^{\tau+1} \frac{1}{\pi'_{\alpha}(\zeta_i)\zeta_i} \frac{1}{1 - \frac{z}{\zeta_i}} = \sum_{i=1}^{\tau+1} \frac{1}{\pi'_{\alpha}(\zeta_i)\zeta_i} \sum_{k=0}^{\infty} \left(\frac{z}{\zeta_i}\right)^k$$

Thus,

$$z^k A \left(\frac{1}{\pi_{\alpha}(z)} \right) = \sum_{i=1}^{\tau+1} \frac{1}{\pi'_{\alpha}(\zeta_i)\zeta_i^{k+1}}. \tag{12}$$

To bound the magnitude of $\frac{1}{\zeta_i}$ and $\frac{1}{\pi'_{\alpha}(\zeta_i)}$, we invoke the following lemma, whose proof (in the supplementary material) relies on some tools from complex analysis:

Lemma 1. Let $\alpha \in (0, 1/20(\tau + 1)]$, and assume $|\zeta_1| \leq |\zeta_2| \leq \dots \leq |\zeta_{\tau+1}|$, then

1. ζ_1 is a real scalar satisfying $1/\zeta_1 \leq 1 - \alpha$, and for $|\zeta_i| > 1, |1/\zeta_i| \leq 1 - \frac{\alpha}{2(\tau+1)}$.
2. $|\pi_\alpha(z)| > 1/4$, for any $i \in [\tau + 1]$.

With this lemma at hand, we have

$$|z^k| \left| \frac{1}{\pi_\alpha(z)} \right| \leq 2(1 - \alpha)^{k+1} + 2\tau \left(1 - \frac{\alpha}{2(\tau+1)} \right)^{k+1} \leq 2(1 - \alpha)^{k+1} \left(1 + \tau \exp\left(-\frac{k+1}{\tau+1}\right) \right),$$

where the last inequality is due to Lemma 5 (provided in the supplementary material). Moreover, one can use elementary arguments to show that $|[z^k]1/\pi_\alpha(z)| \leq 1$, as long as $\alpha \in [0, 1/\tau]$ (see

Lemma 2 in the supplementary material). Overall, for any $\tau \geq 0$, we have

$$\begin{cases} |z^k| \left| \frac{1}{\pi_\alpha(z)} \right| \leq 1 & 0 \leq k \leq (\tau + 1) \ln(2(\tau + 1)) - 1 \\ |z^k| \left| \frac{1}{\pi_\alpha(z)} \right| \leq 3(1 - \alpha)^{k+1} & k \geq (\tau + 1) \ln(2(\tau + 1)), \end{cases} \quad (13)$$

which, using Eq. (11), gives the

desired bounds on the elements (b_k) defined in Eq. (9).

3 Deterministic Delayed Gradient Descent

We start by analyzing the convergence of DGD for λ -strongly convex and μ -smooth quadratic functions, where the eigenvalues of A are assumed to lie in $[\lambda, \mu]$ for some $\mu \geq \lambda > 0$.

Following the same line of the derivation as in Eq. (10), we obtain $\mathbf{e}(z) = (I - Iz + \eta Az^{\tau+1})^{-1} \mathbf{e}_0$.

Letting $[d] := \{1, 2, \dots, d\}$, it follows that for any $k \geq (\tau + 1) \ln(2(\tau + 1))$,

$$\begin{aligned} \|e_k\| &= \left\| [z^k] \left((I - Iz + \eta Az^{\tau+1})^{-1} \mathbf{e}_0 \right) \right\| \stackrel{(a)}{=} \left\| [z^k] \left((I - Iz + \eta Az^{\tau+1})^{-1} \right) \mathbf{e}_0 \right\| \\ &\stackrel{(b)}{\leq} \max_{i \in [d]} \left| [z^k] \frac{1}{\pi_\alpha(z)} \right| \|\mathbf{e}_0\| \stackrel{(c)}{\leq} \eta \alpha^{-k+1} \|\mathbf{e}_0\| \stackrel{(d)}{\leq} 3(1 - \eta\lambda)^{k+1} \|\mathbf{e}_0\| \end{aligned} \quad (14)$$

where (a) follows by the linearity, (that reveals that the spectral norm of a matrix polynomial equals the absolute value of the same polynomial in one of its eigenvalues), (b) by eigendecomposition of A , and (d) by

the fact that $a_i \geq \lambda$ for all i . Moreover, by Eq. (8) and the fact that all eigenvalues of A are at most μ , we arrive at the following bound:

Theorem 1. For any delay $\tau \geq 0$ and $k \geq (\tau + 1)\ln(2(\tau + 1))$, running DGD with step size $\eta \in (0, 1/(20\mu(\tau + 1)))$ on a μ -smooth, λ -strongly convex quadratic function yields

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq 5\mu(1 - \eta\lambda)^{2(k+1)} \|\mathbf{w}_0 - \mathbf{w}^*\|^2.$$

In particular, setting $\eta = \Omega(1/\mu\tau)$, we get that

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq 5\mu \|\mathbf{w}_0 - \mathbf{w}^*\|^2 \exp\left(-\Omega\left(\frac{k\lambda}{\mu\tau}\right)\right).$$

Note that the assumption that $k \geq (\tau + 1)\ln(2(\tau + 1))$ is very mild, since if $k \leq \tau$ then the algorithm trivially makes no updates after k rounds.

We now turn to analyze the case of μ -smooth convex quadratic functions, where the eigenvalues of the matrix A are assumed to lie in $[0, \mu]$. Following the same derivation as in Ineq. 14 and using Ineq. 8, we have for any $k \geq (\tau + 1)\ln(2(\tau + 1))$ and $\eta \in (0, 1/(20\mu(\tau + 1)))$,

$$\begin{aligned} F(\mathbf{w}_k) - F(\mathbf{w}^*) &= \frac{1}{2} \|\sqrt{A} \mathbf{e}_k\|^2 = \frac{1}{2} \|\sqrt{A}|z^k| (I - Iz + \eta Az^{\tau+1})^{-1} \mathbf{e}_0\|^2 \\ &\stackrel{(a)}{\leq} \frac{1}{2} \left(3 \max_{i \in [d]} \sqrt{a_i} (1 - \eta a_i)^{k+1} \right)^2 \|\mathbf{e}_0\|^2 \stackrel{(b)}{\leq} \frac{9}{4e\eta(k+1)} \|\mathbf{e}_0\|^2 \end{aligned} \quad (15)$$

where $e = 2.718\dots$ is Euler's number, (a) is by the fact that the spectral norm of a matrix polynomial equals the absolute value of the same polynomial in one of its eigenvalues, and (b) is by the fact that $\sqrt{a_i} (1 - \eta a_i)^{k+1} \leq 1/\sqrt{2e\eta(k+1)}$ for any $i \in [d]$ (see Lemma 7 in the supplementary material). We have thus arrived at the following bound for the convex case:

Theorem 2. For any delay $\tau \geq 0$ and $k \geq (\tau + 1)\ln(2(\tau + 1))$, running DGD with step size $\eta \in (0, 1/(20\mu(\tau + 1)))$ on a μ -smooth convex quadratic function yields

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq \frac{9}{4e\eta(k+1)} \|\mathbf{w}_0 - \mathbf{w}^*\|^2.$$

In particular, if we set $\eta = \Omega(1/\mu\tau)$, we get that

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \leq \mathcal{O}\left(\frac{\mu\tau \|\mathbf{w}_0 - \mathbf{w}^*\|^2}{k}\right).$$

As discussed in the introduction, the theorems above imply that a delay of τ increases the iteration complexity by a factor of τ . We now show lower bounds which imply that this linear dependence on τ is unavoidable, for a large family of gradient-based algorithms (of which gradient descent is just a special

case). Specifically, we will consider any iterative algorithm producing iterates $\mathbf{w}_0, \mathbf{w}_1, \dots$ which satisfies the following:

$$\mathbf{w}_0 = \dots = \mathbf{w}_\tau = \mathbf{0} \text{ and } \forall k \geq \tau, \mathbf{w}_{k+1} \in \text{span}\{\nabla F(\mathbf{w}_0), \nabla F(\mathbf{w}_1), \dots, \nabla F(\mathbf{w}_{k-\tau})\}. \quad (16)$$

This is a standard assumption in proving optimization lower bounds (see [17]), and is satisfied by most standard gradient-based methods, and in particular our DGD algorithm. We also note that this algorithmic assumption can be relaxed at the cost of a more involved proof, similar to [16, 26] in the non-delayed case.

Theorem 3. *Consider any algorithm satisfying Eq. (16). Then the following holds for any $k \geq \tau + 1$ and sufficiently large dimensionality d :*

- *There exists a μ -smooth, λ -strongly convex function F over \mathbb{R}^d , such that*

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \geq \frac{\lambda}{4} \exp\left(-\frac{5k}{(\sqrt{\mu/\lambda} + 1)(\tau + 1)}\right) \|\mathbf{w}_0 - \mathbf{w}^*\|^2.$$

- *There exists a μ -smooth, convex quadratic function F over \mathbb{R}^d , such that*

$$F(\mathbf{w}_k) - F(\mathbf{w}^*) \geq \frac{\mu(\tau + 1)^2 \|\mathbf{w}_0 - \mathbf{w}^*\|^2}{4L^2}.$$

The proof of the theorem is very similar to standard optimization lower bounds for gradient-based methods without delays (e.g. [17, 11]), and is presented in the supplementary material. In fact, our main contribution is to recognize that the proof technique easily extends to incorporate delays.

In terms of iteration complexity, these bounds correspond to $\frac{2(\tau + \sqrt{\mu/\lambda} \ln(\lambda) \|\mathbf{w}_0 - \mathbf{w}^*\|^2/\epsilon)}{\epsilon}$ in the strongly convex case, and $\frac{2(\tau + \sqrt{\mu/\lambda} \|\mathbf{w}_0 - \mathbf{w}^*\|^2/\epsilon)}{\epsilon}$ in the convex case, which show that the linear dependence on τ is inevitable. The dependence on the other problem parameters is somewhat better than in our upper bounds, but this is not just an artifact of the analysis: In our delayed setting, the lower bounds can be matched by running *accelerated* gradient descent (AGD) [17], where each time we perform an accelerated gradient descent step, and then stay idle for τ iterations till we get the gradient of the current point. Overall, we perform k/τ accelerated gradient steps, and can apply the standard analysis of AGD to get an iteration complexity which is τ times the iteration complexity of AGD without delays. These match the lower bounds above up to constants. We believe it is possible to prove a similar upper bound for AGD performing an update with a delayed gradient at every iteration (like our DGD procedure), but the analysis is more challenging than for plain gradient descent, and we leave it to future work.

4 Stochastic Delayed Gradient Descent

In this section, we study the case of noisy (stochastic) gradient updates, and the SDGD algorithm, in which the influence of the delay is quite different than in the noiseless case. Instantiating SDGD for quadratic $F(\mathbf{w})$ (defined in (7)) results in the following update rule

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \nabla F(\mathbf{w}_{k-\tau} + \varrho_k) = \mathbf{w}_k - \eta (A\mathbf{w}_{k-\tau} + \mathbf{b} + \xi_k), \quad (17)$$

where $\xi_k, k \geq 0$ are independent zero-mean noise terms satisfying $\mathbb{E}[k\xi_k k^2] \leq \sigma^2$. As before, in terms of the error term $\mathbf{e}_k = \mathbf{w}_k - \mathbf{w}^*$, Eq. (17) reads as $\mathbf{e}_{k+1} = \mathbf{e}_k - \eta A\mathbf{e}_{k-\tau} - \eta \xi_k$. Given a realization of (ξ_k) , we denote its associated formal power series by $\mathbf{e}(z) = \sum_{k \geq 0} \mathbf{e}_k z^k$. By an analysis similar to before, we get that the formal power series of the error terms (\mathbf{e}_k) satisfies

$$\mathbf{e}(z) = (I - Iz + \eta Az^{\tau+1})^{-1}(\mathbf{e}_0 - \eta g(z)).$$

We can now bound the error terms by extracting the corresponding coefficients of $\mathbf{e}(z)$. Letting $D := (I - Iz + \eta Az^{\tau+1})^{-1}$, we have for any $k \geq (\tau + 1)\ln(2(\tau + 1))$

$$\begin{aligned} \mathbb{E}\|F(\mathbf{w}_k) - F(\mathbf{w}^*)\| &= \mathbb{E}\left[\left\|\sqrt{\lambda}\mathbf{e}_k\right\|^2\right] = \mathbb{E}\left[\left\|\sqrt{\lambda}[z^k](D(\mathbf{e}_0 - \eta g(z)))\right\|^2\right] \\ &\stackrel{(a)}{=} \mathbb{E}\left[\left\|\sqrt{\lambda}[z^k]D\mathbf{e}_0\right\|^2\right] + \eta^2 \mathbb{E}\left[\left\|\sqrt{\lambda}[z^k](Dg(z))\right\|^2\right] \\ &\stackrel{(b)}{=} \mathbb{E}\left[\left\|\sqrt{\lambda}[z^k]D\mathbf{e}_0\right\|^2\right] + \eta^2 \mathbb{E}\left[\left\|\sqrt{\lambda}\sum_{i=0}^k ([z^i]D)\xi_{k-i}\right\|^2\right] \\ &\stackrel{(c)}{\leq} \mathbb{E}\left[\left\|\sqrt{\lambda}[z^k]D\right\|^2\right]\|\mathbf{e}_0\|^2 + \eta^2 \sigma^2 \sum_{i=0}^k \mathbb{E}\left[\left\|\sqrt{\lambda}[z^i]D\right\|^2\right], \end{aligned} \quad (18)$$

where (a) follows by the linearity of the bracket operation $[z^k]$ and the assumption that $\mathbb{E}[\xi_k] = 0$ for all k (hence $\mathbb{E}[g(z)] = 0$), (b) follows by the Cauchy product for formal power series, and (c) by the hypothesis that ξ_k are independent and satisfy $\mathbb{E}[k\xi_k k^2] \leq \sigma^2$ for all k . We then upper bound both terms, building on Ineq. 13 (see the supplementary material for a full derivation), resulting in the following theorem:

Theorem 4. *Assuming the step η satisfies holds for SDGD:*

$$\eta \in \left(0, \frac{1}{20\mu(\tau+1)}\right), \text{ and } k \geq (\tau + 1)\ln(2(\tau + 1)), \text{ the following}$$

- For λ -strongly convex, μ -smooth quadratic convex functions, $\mathbb{E}[F(\mathbf{w}_k) - F(\mathbf{w}^*)]$ is at most

$$\frac{1}{2\mu} \exp(-2\eta\lambda(k+1))\|\mathbf{w}_0 - \mathbf{w}^*\|^2 + \frac{\eta^2\sigma^2}{2} \left(\mu(\tau+1)\ln(2(\tau+1)) + \frac{1 + e + \ln(\frac{1}{\eta\lambda})}{e\eta} \right).$$

In particular, by tuning η appropriately,

$$\mathbb{E}(F(\mathbf{w}_k) - F(\mathbf{w}^*)) \leq \mathcal{O}\left(\frac{\sigma^2}{\lambda k} + \mu \|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right)\right).$$

- For μ -smooth quadratic convex functions, $\mathbb{E}[F(\mathbf{w}_k) - F(\mathbf{w}^*)]$ is at most

$$\frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2}{4\mu(k+1)} + \eta^2 \sigma^2 \left(\mu(\tau+1) \ln(2(\tau+1)) + \frac{9}{2e\eta} (1 + \ln(k+1)) \right).$$

In particular, by tuning η appropriately,

$$\mathbb{E}(F(\mathbf{w}_k) - F(\mathbf{w}^*)) \leq \mathcal{O}\left(\frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2}{\sqrt{k}} + \frac{\|\mathbf{w}_0 - \mathbf{w}^*\|^2}{k}\right).$$

As discussed in the introduction in detail, the theorem implies that the effect of τ is negligible once k is sufficiently large.

Acknowledgements

This work was supported by NSF-BSF award 1718970.

References

- [1] Alekh Agarwal and John C Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*, pages 873–881, 2011.
- [2] Dimitri P Bertsekas and John N Tsitsiklis. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.
- [3] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- [4] Sorathan Chaturapruek, John C Duchi, and Christopher Ré. Asynchronous stochastic convex optimization: the noise is in the noise and sgd don't care. In *Advances in Neural Information Processing Systems*, pages 1531–1539, 2015.
- [5] Andrew Cotter, Ohad Shamir, Nati Srebro, and Karthik Sridharan. Better mini-batch algorithms via accelerated gradient methods. In *Advances in neural information processing systems*, pages 1647–1655, 2011.
- [6] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13(Jan):165–202, 2012.
- [7] Hamid Reza Feyzmahdavian, Arda Aytikin, and Mikael Johansson. A delayed proximal gradient method with linear convergence rate. In *Machine Learning for Signal Processing (MLSP), 2014 IEEE International Workshop on*, pages 1–6. IEEE, 2014.

- [8] Hamid Reza Feyzmahdavian, Arda Aytekin, and Mikael Johansson. An asynchronous mini-batch algorithm for regularized stochastic optimization. *IEEE Transactions on Automatic Control*, 61(12):3740–3754, 2016.
- [9] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. cambridge University press, 2009.
- [10] Pooria Joulani, Andras Gyorgy, and Csaba Szepesvári. Online learning under delayed feedback. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1453–1461, 2013.
- [11] Guanghui Lan and Yi Zhou. An optimal randomized incremental gradient method. *arXiv preprint arXiv:1507.02000*, 2015.
- [12] John Langford, Alex J Smola, and Martin Zinkevich. Slow learners are fast. *Advances in Neural Information Processing Systems*, 22:2331–2339, 2009.
- [13] Rémi Leblond, Fabian Pederegosa, and Simon Lacoste-Julien. Improved asynchronous parallel optimization analysis for stochastic incremental methods. *arXiv preprint arXiv:1801.03749*, 2018.
- [14] Horia Mania, Xinghao Pan, Dimitris Papailiopoulos, Benjamin Recht, Kannan Ramchandran, and Michael I Jordan. Perturbed iterate analysis for asynchronous stochastic optimization. *arXiv preprint arXiv:1507.06970*, 2015.
- [15] A Nedic, Dimitri P Bertsekas, and Vivek S Borkar. Distributed asynchronous incremental subgradient methods. *Studies in Computational Mathematics*, 8(C):381–407, 2001.
- [16] AS Nemirovsky and DB Yudin. Problem complexity and method efficiency in optimization. 1983. *Willey-Interscience, New York*, 1983.
- [17] Yurii Nesterov. *Introductory lectures on convex optimization*, volume 87. Springer Science & Business Media, 2004.
- [18] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in neural information processing systems*, pages 693–701, 2011.
- [19] Ohad Shamir and Nathan Srebro. Distributed stochastic optimization and learning. In *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pages 850–857. IEEE, 2014.
- [20] Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International Conference on Machine Learning*, pages 71–79, 2013.
- [21] Benjamin Sirb and Xiaojing Ye. Decentralized consensus algorithm with delayed and stochastic gradients. *arXiv preprint arXiv:1604.05649*, 2016.

- [22] Suvrit Sra, Adams Wei Yu, Mu Li, and Alexander J Smola. Adadelay: Delay adaptive distributed stochastic convex optimization. *arXiv preprint arXiv:1508.05003*, 2015.
- [23] Richard P Stanley. Enumerative combinatorics. vol. i, the wadsworth & brooks/cole mathematics series, wadsworth & brooks, 1986.
- [24] Martin Takáč, Avleen Singh Bijral, Peter Richtárik, and Nati Srebro. Mini-batch primal and dual methods for svms. In *ICML (3)*, pages 1022–1030, 2013.
- [25] Herbert S Wilf. *generatingfunctionology*. AK Peters/CRC Press, 2005.
- [26] Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. In *Advances in neural information processing systems*, pages 3639–3647, 2016.

A Proof of Lemma 1

Recall that $z + \alpha z^{\tau+1}$, and its roots, denoted by ζ_i , are ordered such that $|\zeta_1| \leq |\zeta_2| \leq \dots$. In order to bound from above the magnitude of $1/\zeta_i$, we analyze a related polynomial

$p_\alpha(z) = z^{\tau+1}\pi_\alpha(1/z)$ which takes the following explicit form

$$p_\alpha(z) = z^{\tau+1} - z^\tau + \alpha = (z - 1)z^\tau + \alpha.$$

The roots of p_α are precisely $1/\zeta_i$ (note that, $\pi_\alpha(0) = 1 = 0^0$, hence $\zeta_i = 0^0, i \in \{1, \dots, \tau + 1\}$). Thus, bounding from above (below) the magnitude of the roots of $p_\alpha(z)$ gives an upper (lower) bound for $|1/\zeta_i|$.

We first establish that for any $\alpha \in (0, \frac{1}{2(\tau+1)})$ has a real-valued root in $(\frac{1}{2(\tau+1)}, 1 - \alpha)$. Indeed, for any such α , we have on the one hand,

$$p_\alpha(1 - \alpha) = -\alpha(1 - \alpha)^\tau + \alpha = \alpha(1 - (1 - \alpha)^\tau) \geq 0,$$

and on the other hand (using the fact that $(1 - 1/2x)^x \geq 1/2$ for all $x \geq 1$),

$$\begin{aligned} p_\alpha\left(\frac{1}{2(\tau+1)}\right) &= \frac{1}{2(\tau+1)} \left(\frac{1}{2(\tau+1)}\right)^\tau + \alpha \\ &= \frac{1}{2(\tau+1)} \left(\frac{1}{2(\tau+1)}\right)^{\tau+1} + \alpha \\ &\leq \frac{1}{2(\tau+1)} \left(\frac{1}{2}\right)^{\tau+1} + \alpha < \frac{1}{20(\tau+1)} + \alpha \leq 0 \end{aligned} \tag{19}$$

so by continuity of p_α , we get that a real-valued root exists in $(\frac{1}{2(\tau+1)}, 1 - \alpha)$.

Next, we show that τ non-dominant roots of p_α are of absolute value smaller than $\frac{1}{2(\tau+1)}$. To this end, we invoke Rouché's theorem, which states that for any two holomorphic functions f, g in some region $K \subseteq \mathbb{C}$ with closed contour ∂K , iff for any $z \in \partial K$, then f and $f + g$ have the same number of zeros (counted with multiplicity) inside. In particular, choosing $f(z) = -z^\tau$, $g(z) = z^{\tau+1} + \alpha$ and $K = \{z : |z| \leq R\}$, it follows that if $|z^{\tau+1} + \alpha| < |-z^\tau|$ for all z such that $|z| = R$, then $f + g$

(which equals our polynomial p_α) has the same number of zeros as $f = -z^\tau$ inside K (namely, exactly τ).

However, since p_α is a degree $\tau + 1$ polynomial, it has exactly $\tau + 1$ roots, so the only root of absolute value larger than R is the real-valued one we found earlier. It remains to verify the condition $|z^{\tau+1} + \alpha| < |-z^\tau|$ for all z such that $|z| = R$.

For that, it is sufficient to show that $|z^{\tau+1}| + \alpha < |z^\tau|$ for all such z , or equivalently, $R + \alpha < R$.

$$R^\tau = \left(1 - \frac{3}{2(\tau+1)}\right)^\tau = \left(1 - \frac{3}{2(\tau+1)}\right)^\tau - \left(1 - \frac{3}{2(\tau+1)}\right)^{\tau+1} + \left(1 - \frac{3}{2(\tau+1)}\right)^{\tau+1}$$

$$= \frac{3}{2(\tau+1)} \left(1 - \frac{3}{2(\tau+1)}\right)^\tau + R^{\tau+1}.$$

By the inequality $1 - 1/(x+1) \geq \exp(-1/x)$ (see Lemma 4 below), we have

$$\frac{3}{2(\tau+1)} \geq \exp\left(\frac{-1}{2/3\tau - 1/3}\right) \implies \left(1 - \frac{3}{2(\tau+1)}\right)^\tau \geq \exp\left(\frac{-\tau}{2/3\tau - 1/3}\right).$$

It is straightforward to verify that

$$\frac{1}{2/3\tau - 1/3} \geq -3, \tau \geq 1,$$

$$R^\tau \geq \frac{3}{2(\tau+1)} \left(1 - \frac{3}{2(\tau+1)}\right)^\tau + R^{\tau+1}$$

$$\geq \frac{3}{2e^3(\tau+1)} + R^{\tau+1}$$

$$\geq \frac{1}{20(\tau+1)} + R^{\tau+1} \geq \alpha + R^{\tau+1},$$

implying that

where in the last inequality we used the assumption that $\zeta_i \in \left(0, \frac{1}{20(\tau+1)}\right)$. As mentioned earlier, the roots of p_α are exactly the reciprocals of the roots of π_α , therefore we conclude

$$\left|\frac{1}{\zeta_i}\right| \leq 1 - \frac{3}{2(\tau+1)}, i \in [\tau] \tag{20}$$

We now turn to bound $|\pi'_\alpha(\zeta_i)|$ from above. By definition, any root of π_α satisfies $\zeta_i^{\tau+1} - \zeta_i + 1 = 0$. Thus, $\alpha\zeta_i^\tau = \frac{1-\zeta_i}{\zeta_i}$ (note that as mentioned in the first part of the proof, $\zeta_i = 06$). This, in turn, gives

$$\pi'_\alpha(\zeta_i) = \alpha(\tau+1)\zeta_i^{\tau-1} - 1 = \frac{(\tau+1)(\zeta_i - 1)}{\zeta_i} - 1$$

$$= \frac{(\tau+1)(\zeta_i - 1) - \zeta_i}{\zeta_i}$$

$$= \frac{(\tau+1)\zeta_i - (\tau+1) - \zeta_i}{\zeta_i}$$

$$= \frac{\tau\zeta_i - (\tau+1)}{\zeta_i} = \tau - \frac{(\tau+1)}{\zeta_i} = (\tau+1) \left(\frac{\tau}{\tau+1} - \frac{1}{\zeta_i}\right) \tag{21}$$

In the previous parts of the proof, we showed that the distance from any root of p_α to the contour $\{z \mid |z| = 1 - 1/(\tau+1)\}$ is bounded from below by $\frac{1}{2(\tau+1)}$ (Ineq. 19 and Ineq. 20), therefore

$$|\pi'_\alpha(z)| = (\tau+1) \left| 1 - \frac{1}{\tau+1} - \frac{1}{z} \right| \geq \frac{\tau+1}{2(\tau+1)} = \frac{1}{2}, \quad i=1, \dots, \tau+1$$

thus concluding the proof.

B Technical Lemmas

Lemma 2. For any $\alpha \in [0, 1/\tau]$ and $k \geq 0$, it holds that $|[z^k]1/\pi_\alpha(z)| \leq 1$.

Proof. Recall that by Eq. (11), $b_k = \frac{1}{\tau+1} + \alpha \sum_{i=k-\tau}^k b_i$. Therefore, suffices it to prove that (b_k) (defined in 9) with $b_0 = 1$ and $\alpha \in [0, 1/\tau]$, satisfies $|b_k| \leq 1$ for any $k \geq 0$.

For the sake of simplicity, we slightly extend (b_k) to the negative indices by defining $b_{-\tau} = b_{-\tau+1} = \dots = b_{-1} = 1$. We proceed by full induction. The base case holds trivially by the definition of the initial conditions of b_k . For the induction step, suppose that $|b_0|, \dots, |b_k| \leq 1$. We have $b_{k+1} = b_k - \alpha b_{k-\tau}$, and therefore

$$b_{k+1} = (1-\alpha)b_k + \alpha(b_k - b_{k-\tau}) = (1-\alpha)b_k + \alpha \sum_{i=k-\tau}^k (b_{i+1} - b_i)$$

Using the recurrence relation again, this equals

$$(1-\alpha)b_k + \alpha \sum_{i=k-\tau}^k (-\alpha b_{i-\tau}) = (1-\alpha)b_k - \alpha \left(\alpha \sum_{i=k-2\tau}^k b_i \right)$$

By the induction hypothesis, this equals $(1-\alpha)b_k + \alpha r_k$, where $|r_k| \leq \alpha\tau \leq 1$. Thus, b_{k+1} is a weighted average of b_k and r_k which are both in $[-1, +1]$ by the induction hypothesis and the above, implying that we must have $b_{k+1} \in [-1, +1]$ as well. Thus, proving the induction step.

□

Lemma 3. Let $F(w) = \frac{1}{2} w^T A w + b^T w$, $A \in \mathbb{R}^{d \times d}$, $b \in \mathbb{R}^d$ be a convex quadratic function defined over \mathbb{R}^d . If F is bounded from below, then F has a minimizer at which the gradient vanishes.

Proof. Since F is convex and twice differentiable, A is positive semidefinite. In particular, we have $\mathbb{R}^d = \ker(A) \oplus \text{im}(A)$ (namely, the direct sum of the null space and the image space of A). Thus, \mathbf{b} can be expressed as a sum of two orthogonal vectors $\mathbf{b} = \mathbf{b}_\perp + \mathbf{b}^-$, where $\mathbf{b}_\perp \in \ker(A)$ and $\mathbf{b}^- \in \text{im}(A)$. For any $\alpha \in \mathbb{R}$, we have

$$F(\alpha \mathbf{b}_\perp) = \frac{1}{2}((\mathbf{b}_\perp)^\top A \mathbf{b}_\perp) \alpha^2 + \alpha \mathbf{b}^\top \mathbf{b}_\perp = \alpha k \mathbf{b}^\perp k^2.$$

By the hypothesis, F is bounded from below, hence \mathbf{b}_\perp must vanish (otherwise we can take $\alpha \rightarrow -\infty$ and make F as negative as we wish). In particular, $\mathbf{b} = \mathbf{b}^- \in \text{im}(A)$. Let $\mathbf{y} \in \mathbb{R}^d$ be such that $A\mathbf{y} = \mathbf{b}$, then $\nabla F(-\mathbf{y}) = A(-\mathbf{y}) + \mathbf{b} = 0$. Lastly, F is convex, therefore $-\mathbf{y}$ must be a (global) minimizer, thus concluding the proof.

□

Lemma 4. For any $x > 0$, it holds that $1 - 1/(x + 1) \geq \exp(-1/x)$.

Proof. Since $(\ln(1 + x))' = 1/(1 + x) > 0$ for any $x > -1$, it follows by the mean-value theorem that for any $x > 0$

$$\ln(1 + x) = \ln(1 + x) - \ln(1) = \frac{1}{1 + \xi} x$$

for some $\xi \in (0, x)$, hence $\ln(1 + x) \leq x$ for any $x > 0$. In particular, for any $x > 0$ we have

$$\ln\left(1 + \frac{1}{x}\right) \leq \frac{1}{x} \implies \ln\left(\frac{x}{x+1}\right) \geq -\frac{1}{x}$$

Taking the exponent of both sides yields the desired lower bound. □

Lemma 5. Let $\tau \geq 0$. If $\alpha \in (0, 1/[20(\tau + 1)])$ then

$$\left(\frac{1 - \frac{3}{2(\tau+1)}}{1 - \alpha}\right)^{k+1} \leq \exp\left(-\frac{k+1}{\tau+1}\right).$$

In particular, for $k \geq (\tau + 1)\ln(2(\tau + 1)) - 1$, we have

$$1 + \tau \left(\frac{1 - \frac{3}{2(\tau+1)}}{1 - \alpha}\right)^{k+1} \leq 3/2$$

(22)

Proof.

$$\left(\frac{1 - \frac{3}{2(\tau+1)}}{1 - \alpha}\right)^{k+1} = \left(\frac{1 - \alpha + \alpha - \frac{3}{2(\tau+1)}}{1 - \alpha}\right)^{k+1} = \left(1 + \frac{\alpha - \frac{3}{2(\tau+1)}}{1 - \alpha}\right)^{k+1} \leq \left(1 + \alpha - \frac{3}{2(\tau+1)}\right)^{k+1},$$

where the latter inequality follows from that fact that $\alpha < \frac{1}{20(\tau+1)} < \frac{3}{2(\tau+1)}$. Now,

$$\begin{aligned} \left(1 + \alpha - \frac{3}{2(\tau+1)}\right)^{k+1} &\leq \exp\left((k+1)\left(\alpha - \frac{3}{2(\tau+1)}\right)\right) \leq \exp\left((k+1)\left(\frac{1}{20(\tau+1)} - \frac{3}{2(\tau+1)}\right)\right) \\ &= \exp\left(\frac{k+1}{\tau+1}\left(\frac{3}{2} - \frac{1}{20}\right)\right) \leq \exp\left(-\frac{k+1}{\tau+1}\right) \end{aligned}$$

Lastly, to derive Ineq. 22, we have

$$1 + \tau \left(\frac{1 - \frac{1}{2(\tau+1)}}{1 - \frac{1}{2}} \right)^{k+1} \leq 1 + (\tau + 1) \exp\left(-\frac{k+1}{\tau+1}\right) = 1 + \exp\left(\ln(\tau+1) - \frac{k+1}{\tau+1}\right) \leq 1 + 1/2$$

where the last inequality by the assumption $k \geq (\tau + 1)\ln(2(\tau + 1)) - 1$.

□

C Proof of Thm. 3

The proof technique is based on a construction, first presented in [17, Section 2.1.2], which has been proven effective in various settings of optimization since then.

First, we address the strongly convex case. Given $\mu > \lambda > 0$, we consider the following function (devised by [11]):

$$F(\mathbf{w}) := \frac{\mu(\kappa - 1)}{4} \left(\frac{1}{2} \langle A\mathbf{w}, \mathbf{w} \rangle - \langle \mathbf{e}_1, \mathbf{w} \rangle \right) + \frac{\lambda}{2} \|\mathbf{w}\|^2 \quad (23)$$

where $\kappa = \mu/\lambda$ as before, \mathbf{e}_1 denotes the first unit vector, and A is a $d \times d$ matrix defined as follows

$$A = \begin{pmatrix} 2 & -1 & 0 & 0 & \dots & 0 & 0 & 0 \\ -1 & 2 & -1 & 0 & \dots & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & \dots & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & \dots & 0 & -1 & \frac{\sqrt{\kappa+1}}{\sqrt{\kappa-1}} \end{pmatrix}. \quad (24)$$

It can be easily verified that F is μ -smooth and λ -strongly convex function. Moreover, by [11, Lemma 8], it follows that the minimizer of f is $\mathbf{w}^* = (q, q^2, \dots, q^d)$ where $q = (\sqrt{\kappa} - 1)/(\sqrt{\kappa} + 1)$. In particular, if $\mathbf{w} \in \mathbb{R}^d$ is a vector whose all non-zero entries are located in the first m coordinates, where m is such that $d \geq m/2 + \log(1/2)/\log(q^2)$, then

$$\frac{\|\mathbf{w} - \mathbf{w}^*\|^2}{\|\mathbf{w}\|^2} \geq \frac{\sum_{i=1}^m \frac{1}{q^{2i}} \|\mathbf{w}_i\|^2}{\sum_{i=1}^m \|\mathbf{w}_i\|^2} = q^{2(m+1)} \frac{1 - q^{2(d-m-1)}}{1 - q^{2d}} \geq \frac{1}{2} q^{2(m+1)} \geq \frac{1}{2} \exp\left(-\frac{4(m+1)}{\sqrt{\kappa}-1}\right)$$

where the last two inequalities follow from [11, Lemma 9.b] and Lemma 4, respectively. Therefore, by bookkeeping which entries of the iterates are non-zero, we can bound from below the distance to the minimizer.

To this end, we will need the following lemma which, based on the tridiagonal structure of the Hessian of F , determines the non-zero entries:

Lemma 6. Let $F : \mathbb{R}^d \rightarrow \mathbb{R}$ be a convex quadratic function specified as follows $F(\mathbf{w}) = \frac{\lambda}{2} \mathbf{w}^T A \mathbf{w} + d \langle \mathbf{c}, \mathbf{w} \rangle$, where A is a tridiagonal matrix and c, d are real scalars. Assuming that the iterates produced by a given optimization algorithm satisfy $\mathbf{w}_0 = \dots = \mathbf{w}_\tau = \mathbf{0}$ and

$$\forall k \geq \tau, \mathbf{w}_{k+1} \in \text{span}\{\nabla F(\mathbf{w}_0), \nabla F(\mathbf{w}_1), \dots, \nabla F(\mathbf{w}_{k-\tau})\},$$

then $\mathbf{w}_k \in \text{span}\{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{\lfloor k/(\tau+1) \rfloor}\}$ for all $k \geq 0$ (where \mathbf{q}_0 denotes the vector of all zeros, and \mathbf{q}_i denote the i 'th standard unit vector).

Proof. First, note that, given a vector $\mathbf{w} \in \mathbb{R}^d$, such that $\mathbf{w} \in \text{span}\{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_m\}$ for some $m \geq 0$, we have

$$\nabla F(\mathbf{w}) = cA\mathbf{w} + d\mathbf{q}_1.$$

Since the entries of \mathbf{w} are all zero start from the $m + 1$ coordinate, $cA\mathbf{w}$ is a linear combination of the first m columns of A . Being a A tridiagonal matrix, it follows that all the entries of $cA\mathbf{w}$ are zero, except for its first $m + 1$ coordinates, that is, $cA\mathbf{w} \in \text{span}\{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{m+1}\}$. Together, $\nabla F(\mathbf{w}) = cA\mathbf{w} + d\mathbf{q}_1 \in \text{span}\{\mathbf{q}_1, \dots, \mathbf{q}_{m+1}\}$.

We proceed by full induction. For $k = 0, \dots, \tau$, the claim holds trivially. Now, assume the claim holds for all $i \leq k$, where $k \geq \tau$, we show that the claim holds for $k + 1$. By the induction hypothesis,

$\mathbf{w}_i \in \text{span}\{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{\lfloor i/(\tau+1) \rfloor}\}$ for all $i \leq k$. Therefore, by the first part of the proof, we have, $\nabla F(\mathbf{w}_i) \in \text{span}\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{\lfloor i/(\tau+1) \rfloor + 1}\}$ for all $i \leq k$, by which we conclude that $\text{span}\{\nabla F(\mathbf{w}_0), \nabla F(\mathbf{w}_1), \dots, \nabla F(\mathbf{w}_{k-\tau})\} \subseteq \text{span}\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{\lfloor (k-\tau)/(\tau+1) \rfloor + 1}\}$. Thus, by the linear span assumption, it follows that

$$\mathbf{w}_{k+1} \in \text{span}\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{\lfloor (k-\tau)/(\tau+1) \rfloor + 1}\}. \quad (26)$$

Observing that,

$$\lfloor (k-\tau)/(\tau+1) \rfloor + 1 = \lfloor (k-\tau)/(\tau+1) + 1 \rfloor = \lfloor (k+1)/(\tau+1) \rfloor,$$

concludes the proof. □

Overall, by Lemma 6, the k 'th iterate w_k , has all its entries zero, except for (possibly) the first $\lfloor k/(\tau+1) \rfloor$ first coordinates. By Ineq. 25, for any $k \geq \tau$, we then have

$$\|\mathbf{w}_k\|_2 \geq \frac{1}{2} \exp\left(-\frac{\lfloor k/(\tau+1) \rfloor + 1}{\sqrt{\kappa} - 1}\right) \geq \frac{1}{2} \exp\left(-\frac{\lfloor k/(\tau+1) \rfloor + 1}{\sqrt{\kappa} - 1}\right).$$

For the convex case, we use a construction (devised by [17]) similar to that of the strongly convex case. Let $\mu > 0$ be fixed and consider the following function

$$f_k(\mathbf{w}) = \frac{\mu}{2} (A_k \mathbf{w}, \mathbf{w}) - \langle \mathbf{e}_1, \mathbf{w} \rangle,$$

where A_k is a $d \times d$ matrix defined as follows

$$A = \begin{matrix} \tau-2 & -1 & 0 & 0 & \dots & 0 & 0 & 0 & & \\ & 2 & -1 & 0 & \dots & 0 & 0 & 0 & & \\ & \dots & \dots & \dots & \dots & \dots & \dots & \dots & & \\ & 0 & 0 & 0 & \dots & -1 & 2 & -1 & & \\ & 0 & 0 & 0 & \dots & 0 & -1 & 2 & 0_{k,d-k} & \\ \tau \times \tau & & & & & 0_{d-k,k} & & & \tau \times \tau & \end{matrix}$$

where $0_{m,n}$ is an $m \times n$ zero matrix. Given an iteration number k such that $\lfloor \frac{d-1}{\tau} \rfloor + 1 \leq k \leq \frac{d-1}{\tau}(\tau+1)$, we take our function $f(w) = \frac{1}{2} \|w\|_2^2 + \dots$. Using Lemma 6, the only (possibly) non-zero entries of the k 'th iterate w_k are the first $\lfloor k/(\tau+1) \rfloor$ coordinates. Thus, following the same lines of proof as in [17, Theorem 2.1.6] yields

$$\frac{F(w_k) - F(w^*)}{\|w_k\|^2} \geq \frac{3L}{32(\lfloor k/(\tau+1) \rfloor + 1)^2} \geq \frac{L(\tau+1)}{45k^2}.$$

D Proof of Thm. 4

We will first state and prove the following auxiliary lemma:

Lemma 7. *The following holds for any $\eta > 0$:*

- For any $k \geq 1$,

$$\max_{\{a: 0 < a < 1/\eta\}} a(1-\eta a)^k \leq \frac{1}{e\eta},$$

where $e = 2.718\dots$ is Euler's number.

In

$$\sum_{i=0}^k \max_{\{a: 0 < a < 1/\eta\}} a(1-\eta a)^i$$

particular, $\eta a)^{2(i+1)} \leq$

$\frac{1}{e\eta} H_k \leq \frac{1}{2\eta} (1 + \ln(k+1))$, where H_k denotes the

- If, in addition, we assume that $a > \lambda$ for some constant $\lambda > 0$, then

$$\sum_{i=0}^k \max_{\{a: \lambda < a < 1/\eta\}} a(1-\eta a)^{2(i+1)} \leq \frac{1 + e + \ln(\frac{1}{\eta\lambda})}{e\eta}$$

Proof. By the well-known inequality $1 + x \leq \exp(x)$, $x \in \mathbb{R}$, and since for the domain over which we optimize it holds that $1 - \eta a > 0$, we have for any $k \geq 1$

$$a(1 - \eta a)^k \leq a \exp(-\eta a k).$$

Let us denote the latter by $\psi(a) := a \exp(-\eta a k)$, and derive for it the desired upper bound.

Taking the derivative of ψ and setting to zero, gives

$$(1 - \eta k a) \exp(-\eta a k) = 0.$$

Therefore, the only stationary point of ψ is $a^* = \frac{1}{\eta k}$. Since ψ' is positive for $a < a^*$ and negative for $a > a^*$, it follows that a^* is a global maximum, at which the value of ψ is $\frac{1}{\eta k} \exp(-1)$, concluding the first part of the proof.

Now, let $\lambda > 0$. Since, the only maximizer of $a(1 - \eta a)^k$ is at $a = \frac{1}{\eta k}$, if $\lambda \geq \frac{1}{2\eta(i+1)}$, or equivalently $\lambda \geq \frac{1}{2\eta}$, then $\max_{\{a : \lambda < a < 1/\eta\}} a(1 - \eta a)^{2(i+1)} \leq \lambda(1 - \eta \lambda)^{2(i+1)}$. Therefore,

$$\begin{aligned} & \sum_{\lambda < a < 1/\eta} \max_{a(1 - \eta a)^{2(i+1)}} \leq \sum_{i=1}^{\lfloor 2\eta\lambda - 1 \rfloor} \max_{a(1 - \eta a)^{2(i+1)}} \\ & + \sum_{i=\lfloor 2\eta\lambda - 1 \rfloor}^k \max_{a(1 - \eta a)^{2(i+1)}} \\ & \leq \sum_{i=1}^{\lfloor 2\eta\lambda - 1 \rfloor} \lambda(1 - \eta \lambda)^{2(i+1)} + \sum_{i=\lfloor 2\eta\lambda - 1 \rfloor}^k \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \sum_{i=1}^{\lfloor 2\eta\lambda - 1 \rfloor} (1 - \eta \lambda)^{2(i+1)} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - (1 - \eta \lambda)^{2(\lfloor 2\eta\lambda - 1 \rfloor + 1)}}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - (1 - \eta \lambda)^{2\eta\lambda}}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - \exp(-2\eta\lambda)}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - \exp(-2\eta\lambda)}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - \exp(-2\eta\lambda)}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \\ & \leq \frac{1}{\eta} \frac{1 - \exp(-2\eta\lambda)}{1 - (1 - \eta \lambda)^2} + \frac{1}{\eta k} \exp(-1) \end{aligned}$$

□

We now turn to prove Thm. 4 itself. By Ineq. 18 we have

$$2\mathbb{E}[F(\mathbf{w}_k) - F(\mathbf{w}^*)] \leq \|\sqrt{A}[z^k]D\|^2 \|\mathbf{e}_0\|^2 + \eta^2 \sigma^2 \sum_{i=0}^k \left\| \sqrt{A}[z^i]D \right\|^2. \quad (27)$$

We will bound each of the terms above separately. Assuming $\eta \in \left(0, \frac{1}{20\mu(\tau+1)}\right)$ we have by Ineq. 14 and Ineq. 13,

$$\begin{aligned} \left\| \sqrt{A}[z^k]D \right\|^2 &= \left\| \sqrt{A}[z^k] \left((I - I_z + \eta A z^{\tau+1})^{-1} \right) \right\|^2 \\ &\leq \max_{i \in [d]} \left| \sqrt{a_i}[z^k] \frac{1}{\pi_{\eta a_i}(z)} \right|^2 \\ &\leq \begin{cases} \max_{i \in [d]} a_i & 0 \leq k \leq (\tau+1) \ln(2(\tau+1)) - 1 \\ 9 \max_{i \in [d]} a_i (1 - \alpha)^{2(k+1)} & k \geq (\tau+1) \ln(2(\tau+1)). \end{cases} \end{aligned} \quad (28)$$

Thus, for the first term, assuming $k \geq (\tau+1) \ln(2(\tau+1))$, we have

$$\begin{aligned} \overline{k} \sqrt{A}[z^k]D k_2 &\leq 9 \max_{i \in [d]} a_i (1 - \eta a_i)^{2(k+1)} \leq 9\mu \max_{i \in [d]} (1 - \eta a_i)^{2(k+1)} \\ &\leq 9\mu \exp(-2\eta\lambda(k+1)). \end{aligned} \quad (29)$$

Bounding the second term in Ineq. 27 is somewhat more involved and requires separating into the two regimes stated in Ineq. 28:

$$\begin{aligned} \sum_{i=0}^k \left\| \sqrt{A}[z^i]D \right\|^2 &\leq \sum_{i=0}^{(\tau+1) \ln(2(\tau+1)) - 1} \left\| \sqrt{A}[z^i]D \right\|^2 + \sum_{i=(\tau+1) \ln(2(\tau+1))}^k \left\| \sqrt{A}[z^i]D \right\|^2 \\ &\leq \mu(\tau+1) \ln(2(\tau+1)) + 9 \sum_{i=0}^k \max_{i \in [d]} a_i (1 - \eta a_i)^{2(i+1)} \end{aligned} \quad (30)$$

We proceed by considering the strongly convex case and the convex case separately. For the strongly convex case we have by Lemma 7

$$\begin{aligned} \sum_{i=0}^k \left\| \sqrt{A}[z^i]D \right\|^2 &\leq \mu(\tau+1) \ln(2(\tau+1)) + 9 \sum_{i=0}^k \max_{i \in [d]} a_i (1 - \eta a_i)^{2(i+1)} \\ &\leq \mu(\tau+1) \ln(2(\tau+1)) + \frac{1 + e + \ln\left(\frac{1}{\eta\lambda}\right)}{\eta}. \end{aligned}$$

Together with Ineq. 27 and Ineq. 29, this implies that for $k \geq (\tau+1) \ln(2(\tau+1))$,

$$\begin{aligned} 2\mathbb{E}[F(\mathbf{w}_k) - F(\mathbf{w}^*)] &\leq \|\sqrt{A}[z^k]D\|^2 \|\mathbf{e}_0\|^2 + \eta^2 \sigma^2 \sum_{i=0}^k \left\| \sqrt{A}[z^i]D \right\|^2 \\ &\leq 9\mu \exp(-2\eta\lambda(k+1)) \|\mathbf{e}_0\|^2 + \eta^2 \sigma^2 \left(\mu(\tau+1) \ln(2(\tau+1)) + \frac{1 + e + \ln\left(\frac{1}{\eta\lambda}\right)}{\eta} \right), \end{aligned}$$

resulting in the first bound stated in the theorem. To get the second bound, we show how to optimally tune the step size η (up to log factors). Ignoring the log factors, the bound above is

$$\mathbb{E}(F(\mathbf{w}_k) - F(\mathbf{w}^*)) \leq \mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp(-2\eta\lambda k) + \eta^2\sigma^2\left(\mu\tau + \frac{1}{\eta}\right)\right).$$

Moreover, since we assume that $\eta \leq O(1/\mu\tau)$, we get that $\mu\tau$ is dominated (up to constants) by $1/\eta$, so we can simplify the above to

$$\mathbb{E}(F(\mathbf{w}_k) - F(\mathbf{w}^*)) \leq \mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp(-2\eta\lambda k) + \eta\sigma^2\right). \quad (31)$$

We now consider three cases:

If $0 \leq \frac{\ln(\lambda\mu\|\mathbf{e}_0\|^2 k/\sigma^2)}{2\lambda k} \leq \frac{1}{20\mu\tau}$, we can pick $\eta = \frac{\ln(\lambda\mu\|\mathbf{e}_0\|^2 k/\sigma^2)}{2\lambda k}$, and get that Eq. (31) is

$$\mathcal{O}\left(\frac{\sigma^2}{\lambda k}\right) = \mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right) + \frac{\sigma^2}{\lambda k}\right)$$

- If $\frac{\ln(\lambda\mu\|\mathbf{e}_0\|^2 k/\sigma^2)}{2\lambda k} < 0$, it follows that $\|\mathbf{e}_0\|^2 \leq \frac{\sigma^2}{\lambda}$. In that case, we pick $\eta = 0$, and get that Eq. (31) is

$$\mathcal{O}\left(\mu\|\mathbf{e}_0\|^2\right) \leq \mathcal{O}\left(\frac{\sigma^2}{\lambda k}\right) = \mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right) + \frac{\sigma^2}{\lambda k}\right).$$

- If $\frac{\ln(\lambda\mu\|\mathbf{e}_0\|^2 k/\sigma^2)}{2\lambda k} > \frac{1}{20\mu\tau}$, we pick $\eta = \frac{1}{20\mu\tau}$, and get that Eq. (31) is

$$\mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right) + \frac{\sigma^2}{\mu\tau}\right) \leq \mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right) + \frac{\sigma^2}{\lambda k}\right).$$

Collecting the three cases above, we get a bound of

$$\mathcal{O}\left(\mu\|\mathbf{e}_0\|^2 \exp\left(-\frac{\lambda k}{10\mu\tau}\right) + \frac{\sigma^2}{\lambda k}\right)$$

as required.

For the convex case, we have by Ineq. 27, Ineq. 28 and Lemma 7, that for $k \geq (\tau + 1)\ln(2(\tau + 1))$

$$\mathbb{E}(F(\mathbf{w}_k) - F(\mathbf{w}^*)) \leq \frac{9}{4c\eta(k+1)}\|\mathbf{e}_0\|^2 + \frac{\eta^2\sigma^2}{2}\left(\mu(\tau+1)\ln(2(\tau+1)) + \frac{9}{2c\eta}(1 + \ln(k+1))\right),$$

resulting in the third bound in the theorem. To get the fourth bound, we now show how to optimally tune the step size η (up to log factors). Ignoring the log factors, the bound above is

$$\mathcal{O}\left(\frac{\|\mathbf{e}_0\|^2}{\eta k} + \eta^2\sigma^2\left(\mu\tau + \frac{1}{\eta}\right)\right).$$

As in the strongly convex case, since we assume $\eta \leq O(1/(\mu\tau))$, we can simplify the above to

$$O\left(\frac{\|e_0\|^2}{\eta k} + \eta \sigma^2\right).$$

We now consider two cases:

If $\frac{\|e_0\|}{\sigma\sqrt{k}} < \frac{1}{20\sqrt{\mu_T}}$, we choose $\eta = \frac{\|e_0\|}{\sigma\sqrt{k}}$, and get

$$O\left(\frac{\|e_0\|\sigma}{\sqrt{k}}\right) = O\left(\frac{\|e_0\|^2\mu_T}{k} + \frac{\|e_0\|\sigma}{\sqrt{k}}\right).$$

If $\frac{\|e_0\|}{\sigma\sqrt{k}} > \frac{1}{20\sqrt{\mu_T}}$, we choose $\eta = \frac{1}{20\sqrt{\mu_T}}$, and get

$$O\left(\frac{\|e_0\|^2\mu_T}{k} + \frac{\sigma^2}{\mu_T}\right) \leq O\left(\frac{\|e_0\|^2\mu_T}{k} + \frac{\|e_0\|\sigma}{\sqrt{k}}\right).$$

Collecting the two cases above, we get a bound of

$$O\left(\frac{\|e_0\|^2\mu_T}{k} + \frac{\|e_0\|\sigma}{\sqrt{k}}\right)$$

as required.