

Tradeoffs Between Convergence Rate and Noise Amplification for Momentum-Based Accelerated Optimization Algorithms

Hesameddin Mohammadi , Member, IEEE, Meisam Razaviyayn , Member, IEEE, and Mihailo R. Jovanović , Fellow, IEEE

Abstract—In this article, we study momentum-based first-order optimization algorithms in which the iterations utilize information from the two previous steps and are subject to an additive white noise. This setup uses noise to account for uncertainty in either gradient evaluation or iteration updates, and it includes Polyak’s heavy-ball and Nesterov’s accelerated methods as special cases. For strongly convex quadratic problems, we use the steady-state variance of the error in the optimization variable to quantify noise amplification and identify fundamental stochastic performance tradeoffs. Our approach utilizes the Jury stability criterion to provide a novel geometric characterization of conditions for linear convergence, and it reveals the relation between the noise amplification and convergence rate as well as their dependence on the condition number and the constant algorithmic parameters. This geometric insight leads to simple alternative proofs of standard convergence results and allows us to establish “uncertainty principle” of strongly convex optimization: for the two-step momentum method with linear convergence rate, the lower bound on the product between the settling time and noise amplification scales quadratically with the condition number. Our analysis also identifies a key difference between the gradient and iterate noise models: while the amplification of gradient noise can be made arbitrarily small by sufficiently decelerating the algorithm, the best achievable variance for the iterate noise model increases linearly with the settling time in the decelerating regime. Finally, we introduce two parameterized families of algorithms that strike a balance between noise amplification and settling time while preserving orderwise Pareto optimality for both noise models.

Received 7 January 2024; revised 18 June 2024; accepted 2 August 2024. Date of publication 3 September 2024; date of current version 30 January 2025. This work was supported by the National Science Foundation under Grant ECCS 1708906 and Grant ECCS 1809833. Recommended by Associate Editor T. Faulwasser. (Corresponding author: Mihailo R. Jovanović.)

Hesameddin Mohammadi and Mihailo R. Jovanović are with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: hesamedm@usc.edu; mihailo@usc.edu).

Meisam Razaviyayn is with the Daniel J. Epstein Department of Industrial and Systems Engineering, University of Southern California, Los Angeles, CA 90089 USA (e-mail: razaviya@usc.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TAC.2024.3453656>, provided by the authors.

Digital Object Identifier 10.1109/TAC.2024.3453656

Index Terms—Convergence rate, convex optimization, first-order algorithms, heavy-ball method, Nesterov’s accelerated algorithm, noise amplification, performance tradeoffs, settling time.

I. INTRODUCTION

ACCELERATED first-order algorithms [1], [2], [3] are often used for solving large-scale optimization problems [4], [5] because of their scalability, fast convergence, and low per-iteration complexity. Their convergence properties [6], [7], [8], [9], [10], [11], [12], [13], [14], worst-case performance guarantees with respect to function variations [15], [16], [17], and extensions to nonsmooth composite optimization problems [18], [19], [20] have been carefully studied. However, stochastic performance in the presence of noise has received less attention [21], [22], [23], [24], [25], [26]. Prior studies indicate that inaccuracies in the computation of gradient values can adversely impact the convergence rate of accelerated methods and that gradient descent may have advantages relative to its accelerated variants in noisy environments [27], [28], [29], [30], [31]. In contrast to gradient descent, accelerated algorithms can also exhibit *undesirable transient behavior* [32], [33], [34]; for convex quadratic problems, the nonnormal dynamic modes in accelerated algorithms induce large transient responses of the error in the optimization variable [34].

Analyzing the performance of accelerated algorithms with additive white noise that arises from uncertainty in gradient evaluation dates back to Polyak [21]. In this reference, Polyak established the optimal linear convergence rate for strongly convex quadratic problems and used time-varying parameters to obtain convergence in the error variance at a sublinear rate with an improved constant factor compared to gradient descent. Under strong convexity, noisy algorithms with constant parameters converge at a linear rate to a stationary distribution in Wasserstein distance; the convergence rate along with bounds on transient behavior and steady-state variance were obtained in [35]. The existence of tradeoffs between the convergence rate and the steady-state variance was also demonstrated.

Acceleration in a sublinear regime can also be achieved for smooth strongly convex problems with diminishing stepsize [36] and averaging can be used to prevent the accumulation of gradient noise by accelerated algorithms [37]. In [38], it has been further shown that the iteration complexity of any first-order

noisy algorithm for strongly convex problems is subject to a fundamental lower bound that consists of bias and variance error terms that decay to zero at linear and sublinear rates, respectively. To achieve this lower bound, a generic accelerated stochastic approximation framework was developed in [38]; this framework can be specialized to obtain optimal or nearly optimal methods. In addition, Aybat et al. [39] propose a multistage algorithm based on properly adjusting the stepsize to strike a balance between noise amplification and convergence rate. Therein, the proposed Nesterov-like algorithm with explicit choice of parameters does not require knowledge of noise variance and achieves Pareto optimality.

For standard accelerated methods with constant parameters, control-theoretic tools were utilized in [40] and [41] to study the steady-state variance of the error in the optimization variable for smooth strongly convex problems. In particular, for the parameters that optimize convergence rates for quadratic problems, tight upper and lower bounds on the noise amplification of gradient descent, heavy-ball method, and Nesterov's accelerated algorithm were developed [40]. These bounds are expressed in terms of the condition number κ and the problem dimension n , and they demonstrate opposite trends relative to the settling time: *for a fixed problem dimension n , accelerated algorithms increase noise amplification by a factor of $\Theta(\sqrt{\kappa})$ relative to gradient descent.* Similar result also holds for heavy-ball and Nesterov's algorithms with parameters that provide convergence rate $\rho \leq 1 - c/\sqrt{\kappa}$ with $c > 0$ [40]. Furthermore, for strongly convex optimization problems with a condition number κ , tight and attainable upper bounds for noise amplification of gradient descent and Nesterov's accelerated method were provided [40].

In this article, we extend the results of [40] to the class of first-order algorithms with three constant parameters in which the iterations involve information from the two previous steps. This class includes heavy-ball and Nesterov's accelerated algorithms as special cases and we examine its stochastic performance for strongly convex quadratic problems. Our results are complementary to [42], which evaluates stochastic performance in both the objective and the iterate errors, and to a recent work [41] that studies the steady-state variance of the error associated with the point at which the gradient is evaluated. Van Scoy and Lessard [41] combined the theory with computational experiments to demonstrate that a parameterized family of heavy-ball-like methods with reduced stepsize provides Pareto-optimal algorithms for simultaneous optimization of convergence rate and amplification of gradient noise. In contrast to [41], we establish analytical lower bounds on the product of the settling time and the steady-state variance of the error in the optimization variable that hold for any constant stabilizing parameters and for both gradient and iterate noise models. Our lower bounds scale with the square of the condition number and reveal a fundamental limitation of this class of algorithms.

In another related work [43], the tradeoff between the convergence rate and risk of suboptimality for the class of two-step momentum algorithms was studied. Can and Gurbuzbalaban [43] characterized the convergence rate and steady-state variance and proposed a systematic and computationally tractable approach based on solutions of semidefinite programming problems to

achieve tradeoffs in the risk-averse setting for strongly convex problems. The impact of parameters on the convergence rate and steady-state variance for momentum-based algorithms with extensions to nonconvex problems has also been studied [44]. Therein, the authors utilized second-order Taylor series expansion in the stepsize to reveal nonintuitive trends for the effect of momentum parameters on the stationary variance. In the context of differential privacy, theoretical bounds along with numerical observations were used to quantify noise amplification of noisy accelerated algorithms [45]. In addition, Kuru et al. [45] obtained optimal noise-robust heavy-ball algorithm and proposed multi-stage variants of accelerated algorithms that attenuate noise in the gradient while enjoying an improved convergence rate.

In addition to considering noise in gradient evaluation, we study the stochastic performance of algorithms when noise is directly added to the iterates (rather than the gradient). For this iterate noise model, we establish an alternative lower bound on the noise amplification. This bound scales linearly with the settling time and is orderwise tight for settling times that are larger than that of gradient descent with the standard stepsize. In this decelerated regime, we identify a key difference between the two noise models: while the impact of gradient uncertainties on variance amplification can be made arbitrarily small by decelerating the two-step momentum algorithm, the best achievable variance for the iterate noise model increases linearly with the settling time in the decelerated regime.

Our results build upon a simple, yet powerful geometric viewpoint, which clarifies the relation between condition number, convergence rate, and algorithmic parameters for strongly convex quadratic problems. This viewpoint allows us to present alternative proofs for the following:

- 1) the optimal convergence rate of the two-step momentum algorithm, which recovers Nesterov's fundamental lower bound on the convergence rate [46] for *finite-dimensional problems* [47];
- 2) the optimal rates achieved by standard gradient descent, heavy-ball method, and Nesterov's accelerated algorithm [9].

In addition, it enables a novel geometric characterization of noise amplification in terms of stability margins and it allows us to precisely quantify tradeoffs between convergence rate and robustness to noise.

We also introduce two parameterized families of algorithms that are structurally similar to the heavy-ball and Nesterov's accelerated algorithms. These algorithms utilize continuous transformations from gradient descent to the corresponding accelerated algorithm (with the optimal convergence rate) via a homotopy path, and they can be used to provide additional insight into the tradeoff between convergence rate and noise amplification. We prove that these parameterized families are orderwise (in terms of the condition number) Pareto-optimal for simultaneous minimization of settling time and noise amplification. Another family of algorithms that facilitates similar tradeoff was proposed in [11], and it includes the fastest known algorithm for the class of smooth strongly convex problems. We also utilize negative momentum parameters to decelerate a heavy-ball-like family of algorithms relative to gradient descent

with the optimal stepsize. For both noise models, our parameterized family yields orderwise optimal algorithms and it allows us to further highlight the key difference between them in the decelerated regime.

In contrast to conjugate gradient methods that are exceedingly sensitive to poor-conditioning and noise, momentum-based first-order optimization algorithms are flexible enough to be deployed to complex optimization landscapes and environments; and to benefit from recent extensive hardware optimization and parallelization in modern platforms that utilize GPUs. In spite of broad applicability of these algorithms with constant parameters, a clear understanding of fundamental tradeoffs between variance amplification and convergence rate is not available in the existing literature. This article addresses this issue by the following:

- 1) providing a novel geometric insight into the dependence of convergence rate and variance amplification on the algorithmic parameters;
- 2) identifying the product between the variance amplification J and the settling time T_s , $J \times T_s$, as an important “conserved quantity” of two-step momentum algorithms;
- 3) establishing *tight bounds* on $J \times T_s$ in terms of the square of the condition number.

The rest of this article is organized as follows. In Section II, we provide preliminaries and background material, and in Section III, we summarize our key contributions. In Section IV, we introduce the tools and ideas that enable our analysis. In particular, we utilize the Jury stability criterion to provide a novel geometric characterization of stability and ρ -linear convergence and exploit this insight to derive simple alternative proofs of standard convergence results and quantify fundamental stochastic performance tradeoffs. In Section V, we introduce two parameterized families of algorithms that allow us to constructively tradeoff settling time and noise amplification. In Section VI, we provide proofs of our main results, and finally, Section VII concludes this article.

II. PRELIMINARIES AND BACKGROUND

For the unconstrained optimization problem

$$\underset{x}{\text{minimize}} \quad f(x) \quad (1)$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is a strongly convex function with a Lipschitz continuous gradient ∇f , we consider noisy momentum-based first-order algorithms that use information from the two previous steps to update the optimization variable

$$\begin{aligned} x^{t+2} &= x^{t+1} + \beta(x^{t+1} - x^t) \\ &\quad - \alpha \nabla f(x^{t+1} + \gamma(x^{t+1} - x^t)) + \sigma_w w^t. \end{aligned} \quad (2)$$

Here, t is the iteration index, α is the stepsize, β and γ are momentum parameters, σ_w is the noise magnitude, and w^t is an additive white noise with zero mean and identity covariance

$$\mathbb{E}[w^t] = 0, \quad \mathbb{E}[w^t(w^\tau)^T] = I\delta(t - \tau)$$

where δ is the Kronecker delta and \mathbb{E} is the expected value operator. In this article, we consider two noise models.

- 1) *Iterate noise* ($\sigma_w = \sigma$): Models uncertainty in computing the iterates of (2), where σ denotes the stepsize-independent noise magnitude.
- 2) *Gradient noise* ($\sigma_w = \alpha\sigma$): Models uncertainty in the gradient evaluation. In this case, the stepsize α directly impacts magnitude of the additive noise.

Iterate noise model captures scenarios with uncertainties in optimization variables because of roundoff, quantization, and communication errors. This model has also been used to improve generalization and robustness in machine learning [48]. On the other hand, the second noise model accounts for gradient computation error or scenarios in which the gradient is estimated from noisy measurements [49]. Also, noise may be intentionally added to the gradient for privacy reasons [50].

Remark 1: An alternative noise model with $\sigma_w = \sqrt{\alpha}\sigma$ has been used to escape local minima in stochastic gradient descent [51] and to provide nonasymptotic guarantees in nonconvex learning [52], [53]. This model arises from a discretization of the continuous-time Langevin diffusion dynamics [52], and for strongly convex quadratic problems, our framework can be used to examine acceleration/robustness tradeoffs. For algorithms that are faster than the standard gradient descent, this model has orderwise identical performance bounds as the other two models and the only difference arises in decelerated regime. We omit details for brevity.

Special cases of (2) include noisy gradient descent ($\beta = \gamma = 0$), Polyak’s heavy-ball method ($\gamma = 0$), and Nesterov’s accelerated algorithm ($\gamma = \beta$). In the absence of noise (i.e., for $\sigma = 0$), the parameters (α, β, γ) can be selected such that the iterates converge linearly to the globally optimal solution [8]. For the family of smooth strongly convex problems, the parameters that yield the fastest known linear convergence rate were provided in [12].

A. Linear Dynamics for Quadratic Problems

Let \mathcal{Q}_m^L denote the class of m -strongly convex L -smooth quadratic functions

$$f(x) = \frac{1}{2}x^T Q x - q^T x \quad (3)$$

with the condition number $\kappa := L/m$, where q is a vector and $Q = Q^T \succ 0$ is the Hessian matrix with eigenvalues

$$L = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n = m > 0.$$

For the quadratic objective function in (3), we can use a linear time-invariant (LTI) state-space model to describe the *two-step momentum algorithm* (2) with constant parameters

$$\begin{aligned} \psi^{t+1} &= A\psi^t + Bw^t \\ z^t &= C\psi^t \end{aligned} \quad (4a)$$

where ψ^t is the state, $z^t := x^t - x^*$ is the performance output that determines the error to the optimal solution x^* , and w^t is the white stochastic input. In particular, choosing $\psi^t := [(x^t - x^*)^T \ (x^{t+1} - x^*)^T]^T$ yields

TABLE I

SETTLING TIMES $T_s := 1/(1 - \rho)$ [9, PROPOSITION 1] ALONG WITH THE CORRESPONDING NOISE AMPLIFICATION BOUNDS IN (10) [40, TH. 4] FOR THE PARAMETERS THAT OPTIMIZE THE LINEAR CONVERGENCE RATE ρ FOR STRONGLY CONVEX QUADRATIC FUNCTION $f \in \mathcal{Q}_m^L$ WITH THE CONDITION NUMBER $\kappa := L/m$

method	fastest parameters (α, β, γ)	T_s	J_{\min}/σ_w^2	J_{\max}/σ_w^2
Gradient	$(2/(L + m), 0, 0)$	$(\kappa + 1)/2$	$\Theta(\kappa) + n$	$n\Theta(\kappa)$
Heavy-ball	$(4/(\sqrt{L} + \sqrt{m})^2, (1 - 2/(\sqrt{\kappa} + 1))^2, 0)$	$(\sqrt{\kappa} + 1)/2$	$\Theta(\kappa\sqrt{\kappa}) + n\Theta(\sqrt{\kappa})$	$n\Theta(\kappa\sqrt{\kappa})$
Nesterov	$(4/(3L + m), 1 - 4/(\sqrt{3\kappa + 1} + 2), \beta)$	$\sqrt{3\kappa + 1}/2$	$\Theta(\kappa\sqrt{\kappa}) + n$	$n\Theta(\kappa\sqrt{\kappa})$

Here, n is the dimension of x , and σ_w^2 is the variance of the White noise.

$$A = \begin{bmatrix} 0 & I \\ -\beta I + \gamma\alpha Q & (1 + \beta)I - (1 + \gamma)\alpha Q \end{bmatrix}$$

$$B^T = [0 \quad \sigma_w I], \quad C = [I \quad 0]. \quad (4b)$$

B. Convergence Rate

An algorithm is stable if in the absence of noise (i.e., $\sigma_w = 0$), the state converges linearly with some rate $\rho < 1$

$$\|\psi^t\|_2 \leq c_t \rho^t \|\psi^0\|_2 \text{ for all } t \geq 1 \quad (5)$$

for all $f \in \mathcal{Q}_m^L$ and all initial conditions ψ^0 , where $c_t > 0$ grows at most polynomially with t . For LTI system (4a), the spectral radius $\rho(A)$ determines the best achievable convergence rate. In addition

$$T_s := 1/(1 - \rho) \quad (6)$$

provides the first-order approximation in ϵ to the *settling time*, i.e., the number of iterations required to reach a given desired accuracy ϵ ; see Appendix A. For the class \mathcal{Q}_m^L of high-dimensional functions (i.e., for $n \gtrsim T_s$), Nesterov established the fundamental lower bound on the settling time (convergence rate) of any first-order algorithm [8]

$$T_s \geq (\sqrt{\kappa} + 1)/2. \quad (7)$$

This lower bound is sharp and it is achieved by the heavy-ball method with the parameters provided in Table I [9]. We note that polynomial factors c_t may appear because of nonmonotonic transient responses for nonnormal dynamics [34], [43]. In addition, the restriction $n \gtrsim T_s$ in (7) can be lifted for the class of two-step momentum algorithms with constant parameters [47]. Robust control techniques have also been used to extend this result to algorithms that involve more than two previous steps [54].

C. Noise Amplification

For LTI system (4a) driven by an additive white noise w^t , $\mathbb{E}(\psi^{t+1}) = A \mathbb{E}(\psi^t)$. Thus, $\mathbb{E}(\psi^t) = A^t \mathbb{E}(\psi^0)$, and for any stabilizing parameters (α, β, γ) , the iterates reach a statistical steady state with $\lim_{t \rightarrow \infty} \mathbb{E}(\psi^t) = 0$ and a variance that can be computed from the solution of the algebraic Lyapunov equation [40], [55]. We call the steady-state variance of the error in the optimization variable noise (or variance) amplification

$$J := \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=0}^t \mathbb{E}(\|x^k - x^*\|_2^2). \quad (8)$$

In addition to the algorithmic parameters (α, β, γ) , the entire spectrum $\{\lambda_i \mid i = 1, \dots, n\}$ of the Hessian matrix Q impacts the noise amplification J of algorithm (2) [40].

Remark 2: An alternative performance metric that examines the steady-state variance of $y^t - x^*$ was considered in [41], where $y^t := x^t + \gamma(x^t - x^{t-1})$ is the point at which the gradient is evaluated in (2). For all $\gamma \geq 0$, we have $J_x \leq J_y \leq (1 + 2\gamma)^2 J_x$, where the subscripts x and y denote the noise amplification in terms of the error in x^t and y^t . Thus, these performance metrics are within a constant factor of each other for bounded values of non-negative momentum parameter γ .

D. Parameters That Optimize Convergence Rate

For special instances of the two-step momentum algorithm (2) applied to strongly convex quadratic problems, namely gradient descent (gd), heavy-ball method (hb), and Nesterov's accelerated algorithm (na), the parameters that yield the fastest convergence rates were established in [9], [21]. These parameters along with the corresponding rates and the noise amplification bounds are provided in Table I. The convergence rates are determined by the spectral radius of the corresponding A -matrices and the noise amplification bounds are computed by examining the solution to the algebraic Lyapunov equation and determining the functions $f \in \mathcal{Q}_m^L$ for which the steady-state variance is maximized/minimized [40, Proposition 1]. Since the optimal rate for the heavy-ball method meets the fundamental lower bound (7), this choice of parameters also optimizes the convergence rate of (2) for $f \in \mathcal{Q}_m^L$.

For the optimal parameters provided in Table I, there is a $\Theta(\sqrt{\kappa})$ improvement in settling times of the heavy-ball and Nesterov's accelerated algorithms relative to gradient descent

$$T_s = \begin{cases} \Theta(\kappa) & \text{gd} \\ \Theta(\sqrt{\kappa}) & \text{hb, na} \end{cases} \quad (9)$$

where $a = \Theta(b)$ means that a lies within constant factors of b as $b \rightarrow \infty$. This improvement makes accelerated algorithms popular for problems with large condition number κ .

While convergence rate is only affected by the largest and smallest eigenvalues of Q , the entire spectrum of Q influences the noise amplification J . On the other hand, the largest and smallest values of J over the function class \mathcal{Q}_m^L

$$J_{\max} := \max_{f \in \mathcal{Q}_m^L} J, \quad J_{\min} := \min_{f \in \mathcal{Q}_m^L} J \quad (10)$$

depend only on the noise magnitude σ_w , the algorithmic parameters (α, β, γ) , the problem dimension n , and the extreme eigenvalues m and L of Q .

For the parameters that optimize convergence rates, tight upper and lower bounds on the noise amplification were developed in [40, Th. 4]. These bounds are expressed in terms of the condition number κ and the problem dimension n , and they demonstrate opposite trends relative to the settling time. In particular, for gradient descent

$$J_{\max} = \sigma_w^2 n \Theta(\kappa), \quad J_{\min} = \sigma_w^2 (\Theta(\kappa) + n) \quad (11a)$$

and for accelerated algorithms

$$J_{\max} = \sigma_w^2 n \Theta(\kappa \sqrt{\kappa})$$

$$J_{\min} = \begin{cases} \sigma_w^2 (\Theta(\kappa \sqrt{\kappa}) + n \Theta(\sqrt{\kappa})) & \text{hb} \\ \sigma_w^2 (\Theta(\kappa \sqrt{\kappa}) + n) & \text{na.} \end{cases} \quad (11b)$$

Thus, for fixed problem dimension n and noise magnitude σ_w , accelerated algorithms increase noise amplification by a factor of $\Theta(\sqrt{\kappa})$ relative to gradient descent for the parameters that optimize convergence rates. While similar result also holds for heavy-ball and Nesterov's algorithms with arbitrary values of parameters α and β that provide settling time $T_s \leq c\sqrt{\kappa}$ with $c > 0$ [40, Th. 8], in this article, we establish fundamental tradeoffs between noise amplification and settling time for the class of the two-step momentum algorithms (2) with arbitrary stabilizing values of constant parameters (α, β, γ) .

III. SUMMARY OF MAIN RESULTS

In this section, we summarize our key contributions regarding robustness/convergence tradeoff for noisy two-step momentum algorithm (2). In addition, our geometric characterization of stability and ρ -linear convergence allows us to provide alternative proofs of standard convergence results and quantify fundamental performance tradeoffs. The proofs of results presented here can be found in Section VI.

A. Bounded Noise Amplification for Stabilizing Parameters

For a discrete-time LTI system with a convergence rate ρ , the distance of the eigenvalues to the unit circle is larger than $1 - \rho$. We use this stability margin to establish an upper bound on the noise amplification J of the two-step momentum method (2) for any stabilizing parameters (α, β, γ) .

Theorem 1 (Upper bounds): Let the parameters (α, β, γ) be such that the two-step momentum algorithm (2) converges linearly with the rate $\rho = 1 - 1/T_s$ for all $f \in \mathcal{Q}_m^L$. Then

$$J \leq \frac{\sigma_w^2 (1 + \rho^2)}{(1 + \rho)^3} n T_s^3 \quad (12a)$$

where n is the problem size. Furthermore, for the gradient noise model ($\sigma_w = \alpha\sigma$)

$$J \leq \frac{\sigma^2 (1 + \rho)(1 + \rho^2)}{L^2} n T_s^3. \quad (12b)$$

For $\rho < 1$, both upper bounds in (12) scale with nT_s^3 and they are exact for the heavy-ball method with the parameters that optimize the convergence rate provided by Table I. However, these bounds are not tight for all stabilizing parameters; e.g., applying (12a) to gradient descent with the optimal stepsize $\alpha = 2/(L + m)$ yields $J \leq \sigma_w^2 n \Theta(\kappa^3)$, which is off by a factor of κ^2 ; cf., Table I. The bound in (12b) is obtained by combining (12a) with $\alpha L \leq (1 + \rho)^2$, which follows from the conditions for ρ -linear convergence in Section IV. For entropic risk, bounds on noise amplification were derived in [43]. In contrast, Theorem 1 uses a geometric viewpoint to capture an explicit, exact cubic dependence of J_{\max} on the settling time.

B. Tradeoff Between Settling Time and Noise Amplification

For any stabilizing constant parameters (α, β, γ) in the two-step momentum algorithm (2), we next establish lower bounds on the largest and the smallest noise amplification J_{\max} and J_{\min} over the class of functions \mathcal{Q}_m^L , as defined in (10), in terms of the settling time T_s .

Theorem 2 (Reciprocal lower bounds): Let the parameters (α, β, γ) be such that the two-step momentum algorithm (2) converges linearly with the rate $\rho = 1 - 1/T_s$ for all $f \in \mathcal{Q}_m^L$. Then, J_{\max} and J_{\min} in (10) satisfy

$$J_{\max} \geq \sigma_w^2 \left((n-1) \frac{\kappa^2}{64} + \frac{\sqrt{\kappa} + 1}{2} \right) T_s^{-1} \quad (13a)$$

$$J_{\min} \geq \sigma_w^2 \left(\frac{\kappa^2}{64} + (n-1) \frac{\sqrt{\kappa} + 1}{2} \right) T_s^{-1}. \quad (13b)$$

Furthermore, for the gradient noise model ($\sigma_w = \alpha\sigma$), we have

$$J_{\max} \geq \frac{\sigma^2}{L^2} \left((n-1) \frac{\kappa^2}{4} + \max \left\{ \frac{\kappa^2}{T_s^3}, \frac{1}{4} \right\} \right) T_s^{-1} \quad (14a)$$

$$J_{\min} \geq \frac{\sigma^2}{L^2} \left(\frac{\kappa^2}{4} + (n-1) \max \left\{ \frac{\kappa^2}{T_s^3}, \frac{1}{4} \right\} \right) T_s^{-1}. \quad (14b)$$

For both noise models, the condition number κ restricts the performance of the two-step momentum algorithm with constant parameters: *for a fixed problem size n , all four lower bounds in Theorem 2 demonstrate the quadratic dependence on the condition number κ for both $J_{\max} \times T_s$ and $J_{\min} \times T_s$.* Relative to the dominant term in κ , the problem dimension n appears in a multiplicative fashion for the lower bounds on J_{\max} and in an additive fashion for the lower bounds on J_{\min} . We note that the fundamental lower bound on T_s in (7) holds for large problem dimension n for any first-order algorithm. In contrast, Theorem 2 holds for arbitrary n for the class of two-step momentum algorithms with constant parameters.

Theorem 3 (Linear lower bounds): Let the parameters (α, β, γ) be such that the two-step momentum algorithm (2) achieves the convergence rate $\rho = \rho(A) = 1 - 1/T_s$, where the

matrix A is given by (4). Then, J_{\max} and J_{\min} satisfy

$$J_{\max} \geq \sigma_w^2 \left((n-1) \frac{T_s}{2(1+\rho)^2} + 1 \right) \quad (15a)$$

$$J_{\min} \geq \sigma_w^2 \left(\frac{T_s}{2(1+\rho)^2} + (n-1) \right). \quad (15b)$$

We observe that the lower bounds on J_{\max} and J_{\min} in Theorem 3 grow linearly with T_s .

C. Accuracy of Lower Bounds

In this subsection, we establish upper bounds on J_{\max} and J_{\min} for a parameterized family of heavy-ball-like algorithms in terms of the settling time T_s . By comparison to the lower bounds established in Theorems 2 and 3, we prove that for any settling time T_s these bounds are *orderwise tight* in κ .

Theorem 4 (Upper bounds): For the class of strongly convex quadratic functions \mathcal{Q}_m^L with the condition number $\kappa = L/m$, let the scalar ρ be such that the fundamental lower bound $T_s = 1/(1-\rho) \geq (\sqrt{\kappa} + 1)/2$ given by (7) holds. Then, the two-step momentum algorithm (2) with parameters

$$\alpha = \frac{(1+\rho)(1+\beta/\rho)}{L}, \quad \beta = \rho \frac{\kappa - (1+\rho)/(1-\rho)}{\kappa + (1+\rho)/(1-\rho)} \quad (16)$$

and $\gamma = 0$, converges linearly with the rate ρ . In addition, J_{\max} and J_{\min} in (10) satisfy

$$J_{\max} \leq \begin{cases} \sigma_w^2 n \kappa ((\kappa + 1)/2) T_s^{-1}, & \text{if } T_s \leq \tau \\ \sigma_w^2 n T_s, & \text{if } T_s \geq \tau \end{cases} \quad (17a)$$

$$J_{\min} \leq \begin{cases} \sigma_w^2 \kappa (\kappa + n - 1) T_s^{-1}, & \text{if } T_s \leq \tau \\ 2\sigma_w^2 (1 + (n-2)/\kappa) T_s, & \text{if } T_s \geq \tau \end{cases} \quad (17b)$$

where $\tau := (\kappa + 1)/2$. Furthermore, for the gradient noise model ($\sigma_w = \alpha\sigma$), we have

$$J_{\max} \leq \sigma^2 n \kappa ((\kappa + 1)/L^2) T_s^{-1} \quad (18a)$$

$$J_{\min} \leq \sigma^2 2\kappa ((\kappa + 4n - 7)/L^2) T_s^{-1}. \quad (18b)$$

Theorem 4 provides upper bounds on J_{\max} and J_{\min} for a family of heavy-ball-like algorithms ($\gamma = 0$) parameterized by the settling time T_s . We note that the condition $T_s \geq (\kappa + 1)/2$ in Theorem 4 corresponds to nonpositive values of the momentum parameter $\beta \leq 0$. For the iterate noise model with $T_s \leq (\kappa + 1)/2$ and for the gradient noise model with any settling time, the upper bounds in Theorem 4 scale as κ^2 for both $J_{\max} \times T_s$ and $J_{\min} \times T_s$. This scaling matches the scaling of the corresponding lower bounds in Theorem 2. Thus, for the gradient noise model, the upper and lower bounds are orderwise tight (in κ) for any settling time. On the other hand, for the iterate noise model, the lower bounds in Theorem 2 are tight only in the accelerated regime $T_s \leq (\kappa + 1)/2$. For this noise model, in the nonaccelerated regime $T_s \geq (\kappa + 1)/2$, the alternative lower bounds established in Theorem 3 are tight as they orderwise match the upper bounds in Theorem 4.

Remark 3: Since \mathcal{Q}_m^L is a subset of the class of m -strongly convex functions with L -Lipschitz continuous gradients, the

fundamental lower bounds on $J_{\max} \times T_s$ established in Theorem 2 carry over to this broader class of problems. Thus, the restriction imposed by the condition number on the tradeoff between settling time and noise amplification goes beyond \mathcal{Q}_m^L and holds for general strongly convex problems.

Remark 4: The upper bounds in Theorem 4 are obtained for a particular choice of constant parameters. Thus, they also provide upper bounds on the best achievable noise amplification bounds $J_{\max}^* := \min_{\alpha, \beta, \gamma} J_{\max}$ and $J_{\min}^* := \min_{\alpha, \beta, \gamma} J_{\min}$ for a settling time T_s ; see Fig. 1.

IV. CONVERGENCE AND NOISE AMPLIFICATION: GEOMETRIC CHARACTERIZATION

In this section, we examine the relation between the convergence rate and noise amplification of the two-step momentum algorithm (2) for strongly convex quadratic problems. In particular, the eigenvalue decomposition of the Hessian matrix Q allows us to bring the dynamics into n decoupled second-order systems parameterized by the eigenvalues of Q and the algorithmic parameters (α, β, γ) . We utilize the Jury stability criterion to provide a novel geometric characterization of stability and ρ -linear convergence and exploit this insight to derive alternative proofs of standard convergence results and quantify fundamental performance tradeoffs.

A. Modal Decomposition

We utilize the eigenvalue decomposition of the Hessian matrix $Q = Q^T \succ 0$, $Q = V\Lambda V^T$, where Λ is the diagonal matrix of the eigenvalues and V is the orthogonal matrix of the corresponding eigenvectors. The change of variables $\hat{x}^t := V^T(x^t - x^*)$ and $\hat{w}^t := V^T w^t$ allows us to bring (4) into n decoupled second-order subsystems

$$\hat{\psi}_i^{t+1} = \hat{A}_i \hat{\psi}_i^t + \hat{B}_i \hat{w}_i^t, \quad \hat{z}_i^t = \hat{C}_i \hat{\psi}_i^t \quad (19a)$$

where \hat{w}_i^t is the i th component of the vector $\hat{w}^t \in \mathbb{R}^n$, $\hat{\psi}_i^t = [\hat{x}_i^t \ \hat{x}_i^{t+1}]^T$, $\hat{B}_i = [0 \ \sigma_w]^T$, $\hat{C}_i = [1 \ 0]$,

$$\hat{A}_i = \hat{A}(\lambda_i) := \begin{bmatrix} 0 & 1 \\ -a(\lambda_i) & -b(\lambda_i) \end{bmatrix} \quad (19b)$$

and

$$a(\lambda) := \beta - \gamma\alpha\lambda, \quad b(\lambda) := (1 + \gamma)\alpha\lambda - (1 + \beta). \quad (19c)$$

B. Conditions for Linear Convergence

For the class of strongly convex quadratic functions \mathcal{Q}_m^L , the best convergence rate ρ is determined by the largest spectral radius of the matrices $\hat{A}(\lambda)$ in (19) for $\lambda \in [m, L]$

$$\rho = \max_{\lambda \in [m, L]} \rho(\hat{A}(\lambda)). \quad (20)$$

For the heavy-ball and Nesterov's accelerated methods, analytical expressions for $\rho(\hat{A}(\lambda))$ were developed and algorithmic parameters that optimize convergence rate were obtained in [9]. Unfortunately, these expressions do not provide insight into the relation between convergence rates and noise amplification. In this article, we ask the dual question:

Iterate noise $\sigma_w = 1$	$J \leq \frac{1+\rho^2}{(1+\rho)^3} n T_s^3$ (12a) \circ
	$J_{\max}^* \leq \begin{cases} \frac{1}{2} n \kappa (\kappa + 1) T_s^{-1} & \text{if } T_s \leq (\kappa + 1)/2 \\ n T_s & \text{if } T_s \geq (\kappa + 1)/2 \end{cases}$ (17a) \blacksquare
	$J_{\max}^* \geq \max \left\{ \left(\frac{1}{64} (n-1) \kappa^2 + \frac{\sqrt{\kappa+1}}{2} \right) T_s^{-1}, \frac{1}{8} (n-1) T_s + 1 \right\}$ (13a) \blacktriangle
	$J_{\max}^* \geq \max \left\{ \left(\frac{1}{64} \kappa^2 + (n-1) \frac{\sqrt{\kappa+1}}{2} \right) T_s^{-1}, \frac{1}{8} T_s + n - 1 \right\}$ (15a) \blacktriangle
	$J_{\min}^* \leq \begin{cases} \kappa (\kappa + n - 1) T_s^{-1} & \text{if } T_s \leq (\kappa + 1)/2 \\ 2 \left(1 + \frac{n-2}{\kappa} \right) T_s & \text{if } T_s \geq (\kappa + 1)/2 \end{cases}$ (17b) \square
$J_{\min}^* \geq \max \left\{ \left(\frac{1}{64} \kappa^2 + (n-1) \frac{\sqrt{\kappa+1}}{2} \right) T_s^{-1}, \frac{1}{8} T_s + n - 1 \right\}$ (13b) \blacktriangle	
$J_{\min}^* \geq \max \left\{ \left(\frac{1}{64} \kappa^2 + (n-1) \frac{\sqrt{\kappa+1}}{2} \right) T_s^{-1}, \frac{1}{8} T_s + n - 1 \right\}$ (15b) \blacktriangle	
Gradient noise $\sigma_w = \alpha$	$J \leq \frac{(1+\rho)(1+\rho^2)}{L^2} n T_s^3$ (12b) \circ
	$J_{\max}^* \leq \frac{1}{L^2} n \kappa (\kappa + 1) T_s^{-1}$ (18a) \blacksquare
	$J_{\max}^* \geq \frac{1}{L^2} \left(\frac{1}{4} (n-1) \kappa^2 + \max \{ \kappa^2 / T_s^3, 1/4 \} \right) T_s^{-1}$ (14a) \blacktriangle
	$J_{\min}^* \leq \frac{2}{L^2} \kappa (\kappa + 4n - 7) T_s^{-1}$ (18b) \square
	$J_{\min}^* \geq \frac{1}{L^2} \left(\frac{1}{4} \kappa^2 + (n-1) \max \{ \kappa^2 / T_s^3, 1/4 \} \right) T_s^{-1}$ (14b) \blacktriangle




Fig. 1. Summary of the results established in Theorems 1–4 for $\sigma^2 = 1$. The top and bottom rows correspond to the iterate and gradient noise models, respectively, and they illustrate 1) $J_{\max}^* := \min_{\alpha, \beta, \gamma} \max_f J$ and $J_{\min}^* := \min_{\alpha, \beta, \gamma} \min_f J$ subject to a settling time T_s for $f \in \mathcal{Q}_m^L$ (black curves); and 2) their corresponding upper (maroon curves) and lower (red curves) bounds in terms of the condition number $\kappa = L/m$, problem size n , and settling time T_s . The upper bounds on J established in Theorem 1 are marked by blue curves. The dark shaded region and its union with the light shaded region respectively correspond to all possible pairs $(T_s, \max_f J)$ and $(T_s, \min_f J)$ for $f \in \mathcal{Q}_m^L$ and any stabilizing parameters (α, β, γ) .

For a fixed convergence rate ρ , what is the largest condition number κ that can be handled by the two-step momentum algorithm (2) with constant parameters?

We note that the matrices $\hat{A}(\lambda)$ share the same structure as

$$M = \begin{bmatrix} 0 & 1 \\ -a & -b \end{bmatrix} \quad (21a)$$

with the real scalars a and b and that the characteristic polynomial associated with the matrix M is given by

$$F(z) := \det(zI - M) = z^2 + bz + a. \quad (21b)$$

We next utilize the Jury stability criterion [56, ch.. 4-3] to provide conditions for stability of the matrix M given by (21a).

Lemma 1: For the matrix $M \in \mathbb{R}^{2 \times 2}$ given by (21a)

$$\rho(M) < 1 \iff (b, a) \in \Delta \quad (22a)$$

where the stability set

$$\Delta := \{(b, a) \mid |b| - 1 < a < 1\} \quad (22b)$$

is an open triangle in the (b, a) -plane with vertices

$$X = (-2, 1), \quad Y = (2, 1), \quad Z = (0, -1). \quad (22c)$$

Proof: The characteristic polynomial $F(z)$ associated with the matrix M is given by (21b) and the Jury stability criterion [56,

ch. 4-3] provides necessary and sufficient conditions for stability, $|a| < 1$, $F(\pm 1) = 1 \pm b + a > 0$. The condition $a > -1$ is ensured by the positivity of $F(\pm 1)$. \blacksquare

For any $\rho > 0$, the spectral radius $\rho(M)$ of the matrix M is smaller than ρ if and only if $\rho(M/\rho)$ is smaller than 1. This observation in conjunction with Lemma 1 allow us to obtain necessary and sufficient conditions for stability with the linear convergence rate ρ of the two-step momentum algorithm (2).

Lemma 2: For any positive scalar $\rho < 1$ and the matrix $M \in \mathbb{R}^{2 \times 2}$ given by (21a), we have

$$\rho(M) \leq \rho \iff (b, a) \in \Delta_\rho \quad (23a)$$

where the ρ -linear convergence set

$$\Delta_\rho := \{(b, a) \mid \rho(|b| - \rho) \leq a \leq \rho^2\} \quad (23b)$$

is a closed triangle in the (b, a) -plane with vertices

$$X_\rho = (-2\rho, \rho^2), \quad Y_\rho = (2\rho, \rho^2), \quad Z_\rho = (0, -\rho^2). \quad (23c)$$

Proof: See Appendix C in [57]. \blacksquare

Fig. 2 illustrates the stability and the ρ -linear convergence sets Δ and Δ_ρ . We note that for any $\rho \in (0, 1)$, we have $\Delta_\rho \subset \Delta$. This can be verified by observing that the vertices (X_ρ, Y_ρ, Z_ρ) of Δ_ρ all lie in Δ .




Fig. 2. Stability set Δ (the open, cyan triangle) in (22b) and the ρ -linear convergence set Δ_ρ (the closed, yellow triangle) in (23b) along with the corresponding vertices. For (b, a) (black dot) associated with the matrix M in (21a), the distances (d, h, l) in (30) are marked by black lines.

Remark 5: The characterization of the sets Δ and Δ_ρ in both continuous and discrete-time settings along with extensions to higher order systems has been discussed in the literature; e.g., see [58] and [59]. In particular, it has been shown that Δ_ρ is the convex hull of the coefficients associated with the special polynomials $(z - \rho)^2$, $(z + \rho)^2$, and $(z - \rho)(z + \rho)$.

For the two-step momentum algorithm (2), the functions $a(\lambda)$ and $b(\lambda)$ given by (19c) satisfy the affine relation

$$(1 + \gamma)a(\lambda) + \gamma b(\lambda) = \beta - \gamma. \quad (24)$$

This fact in conjunction with Lemmas 1 and 2 allow us to derive conditions for stability and the convergence rate. A similar approach for polynomials of arbitrary degree has been taken in [59], where the authors analyzed affine constraints on the coefficients in optimizing the convergence rate. For second-order polynomials, we note that the rate of convergence can also be directly characterized as a function of $(a(\lambda), b(\lambda))$. This approach was utilized in [43, Lemma 3.1].

Lemma 3: The two-step momentum algorithm (2) with constant parameters (α, β, γ) is stable for all functions $f \in \mathcal{Q}_m^L$ if and only if the following equivalent conditions hold:

- 1) $(b(\lambda), a(\lambda)) \in \Delta$ for all $\lambda \in [m, L]$;
- 2) $(b(\lambda), a(\lambda)) \in \Delta$ for $\lambda \in \{m, L\}$.

Furthermore, the linear convergence rate $\rho < 1$ is achieved for all functions $f \in \mathcal{Q}_m^L$ if and only if the following equivalent conditions hold:

- 1) $(b(\lambda), a(\lambda)) \in \Delta_\rho$ for all $\lambda \in [m, L]$;
- 2) $(b(\lambda), a(\lambda)) \in \Delta_\rho$ for $\lambda \in \{m, L\}$.

Here, $(b(\lambda), a(\lambda))$ is given by (19c), and the stability and ρ -linear convergence triangles Δ and Δ_ρ are given by (22b) and (23b), respectively.

Proof: The conditions in 1) follow from combining (20) with Lemma 1 (for stability) and Lemma 2 (for ρ -linear convergence). The conditions in 2) follow from the facts that Δ and Δ_ρ are convex sets and that $(b(\lambda), a(\lambda))$ is a line segment with end points corresponding to $\lambda = m$ and $\lambda = L$. ■

Lemma 3 exploits the affine relation (24) between $a(\lambda)$ and $b(\lambda)$ and the convexity of the sets Δ and Δ_ρ to establish necessary and sufficient conditions for stability and ρ -linear convergence: *the inclusion of the end points of the line segment $(b(\lambda), a(\lambda))$ associated with the extreme eigenvalues m and L of the matrix Q in the corresponding triangle.* A similar approach was taken in [41, Appendix A.1], where the affine nature of

the conditions resulting from the Jury stability criterion with respect to λ was used to conclude that $\rho(\hat{A}(\lambda))$ is a quasi-convex function of λ and show that the extreme points m and L determine $\rho(A)$. In contrast, we exploit the triangular shapes of the stability and ρ -linear convergence sets Δ and Δ_ρ and utilize this geometric insight to identify the parameters that optimize the convergence rate and to establish tradeoffs between noise amplification and convergence rate. The following corollary is immediate.

Corollary 1: Let the two-step momentum algorithm (2) with constant parameters (α, β, γ) minimize a function $f \in \mathcal{Q}_m^L$ with a linear rate $\rho < 1$. Then, the convergence rate ρ is achieved for all functions $f \in \mathcal{Q}_m^L$.

Proof: Lemma 3 implies that only the extreme eigenvalues m and L of Q determine ρ . Since all functions $f \in \mathcal{Q}_m^L$ share the same extreme eigenvalues, this completes the proof. ■

For the two-step momentum algorithm (2) with constant parameters, Lemma 3 leads to a simple alternative proof for the fundamental lower bound (7) on the settling time established by Nesterov. Our proof utilizes the fact that for any point $(b(\lambda), a(\lambda)) \in \Delta_\rho$, the horizontal signed distance to the edge XZ of the stability triangle Δ satisfies

$$d(\lambda) := a(\lambda) + b(\lambda) + 1 = \alpha\lambda \quad (25)$$

where a and b are given by (19c); see Fig. 2.

Proposition 1: Let the two-step momentum algorithm (2) with constant parameters (α, β, γ) achieve the convergence rate $\rho < 1$ for all functions $f \in \mathcal{Q}_m^L$. Then, lower bound (7) on the settling time holds and it is achieved by the heavy-ball method with the parameters provided in Table I.

Proof: Let $d(m) = \alpha m$ and $d(L) = \alpha L$ denote the values of the function $d(\lambda)$ associated with the points $(b(m), a(m))$ and $(b(L), a(L))$, where (b, a) and d are given by (19c) and (25), respectively. Lemma 3 implies that $(b(L), a(L))$ and $(b(m), a(m))$ lie in the ρ -linear convergence triangle Δ_ρ . Thus

$$d_{\max}/d_{\min} \geq d(L)/d(m) = \kappa \quad (26)$$

where d_{\max} and d_{\min} are the largest and smallest values of d among all points $(b, a) \in \Delta_\rho$. From the shape of Δ_ρ , we conclude that d_{\max} and d_{\min} correspond to the vertices Y_ρ and X_ρ of Δ_ρ given by (23c); see Fig. 2. Thus

$$d_{\max} = d_{Y_\rho} = 1 + \rho^2 + 2\rho = (1 + \rho)^2 \quad (27a)$$

$$d_{\min} = d_{X_\rho} = 1 + \rho^2 - 2\rho = (1 - \rho)^2. \quad (27b)$$

Combining (26) with (27) yields

$$\kappa = \frac{d(L)}{d(m)} \leq \frac{d_{\max}}{d_{\min}} = \frac{(1 + \rho)^2}{(1 - \rho)^2}. \quad (28)$$

Rearranging terms in (28) gives lower bound (7). ■

To provide additional insight, we next examine the implications of Lemma 3 for gradient descent, Polyak's heavy-ball, and Nesterov's accelerated algorithms. In all three cases, our dual approach recovers the optimal convergence rates provided in Table I. From the affine relation (24), it follows that $(b(\lambda), a(\lambda))$ with $\lambda \in [m, L]$ for




Fig. 3. For a fixed ρ -linear convergence triangle Δ_ρ (yellow), dashed blue lines mark the line segments $(b(\lambda), a(\lambda))$ with $\lambda \in [m, L]$ for gradient descent, Polyak's heavy-ball, and Nesterov's accelerated methods as particular instances of the two-step momentum algorithm (2) with constant parameters. The solid blue line segments correspond to the parameters for which the algorithm achieves rate ρ for the largest possible condition number given by (29).

- 1) gradient descent ($\beta = \gamma = 0$), is a horizontal line segment parameterized by $a(\lambda) = 0$;
- 2) heavy-ball method ($\gamma = 0$), is a horizontal line segment parameterized by $a(\lambda) = \beta$;
- 3) Nesterov's accelerated method ($\beta = \gamma$), is a line segment parameterized by $a(\lambda) = -\beta b(\lambda)/(1 + \beta)$.

These observations are illustrated in Fig. 3, and as we show in the proof of Lemma 3, to obtain the largest possible condition number for which the convergence rate ρ is feasible for each algorithm, one needs to find the largest ratio $d(L)/d(m) = \kappa$ among all possible orientations for the line segment $(b(\lambda), a(\lambda))$ with $\lambda \in [m, L]$ to lie within Δ_ρ .

- 1) For gradient descent, the largest ratio $d(L)/d(m)$ corresponds to the intersections of the horizontal axis and the edges $Y_\rho Z_\rho$ and $X_\rho Z_\rho$ of the triangle Δ_ρ , which are given by $(\rho, 0)$ and $(-\rho, 0)$, respectively. Thus

$$\kappa = d(L)/d(m) \leq (1 + \rho)/(1 - \rho). \quad (29a)$$

Rearranging terms in (29a) yields a lower bound on the settling time for gradient descent $1/(1 - \rho) \geq (\kappa + 1)/2$. This lower bound is tight as it can be achieved by choosing the parameters in Table I, which place $(b(\lambda), a(\lambda))$ to $(\rho, 0)$ and $(-\rho, 0)$ for $\lambda = L$ and $\lambda = m$, respectively.

- 2) For the heavy-ball method, the optimal rate is recovered by designing the parameters (α, β) such that the vertices X_ρ and Y_ρ belong to the horizontal line segment $(b(\lambda), a(\lambda))$

$$\kappa = d(L)/d(m) \leq (1 + \rho)^2/(1 - \rho)^2. \quad (29b)$$

By choosing $d(L) = d_{Y_\rho}$ and $d(m) = d_{X_\rho}$, we recover the optimal parameters provided in Table I and achieve the fundamental lower bound (7) on the convergence rate.

- 3) For Nesterov's accelerated method, the largest ratio $d(L)/d(m)$ corresponds to the line segment $X_\rho X'_\rho$ that passes through the origin, where $X'_\rho = (2\rho/3, -\rho^2/3)$ lies on the edge $Y_\rho Z_\rho$; see Appendix C in [57]. Thus

$$\kappa = d(L)/d(m) \leq (1 + \rho)(3 - \rho)/(3(1 - \rho)^2). \quad (29c)$$

Rearranging terms in this inequality provides a lower bound on the settling time $1/(1 - \rho) \geq \sqrt{3\kappa + 1}/2$. This lower bound is tight and it can be achieved with the parameters provided in Table I, which place $(b(L), a(L))$ to X'_ρ and $(b(m), a(m))$ to X_ρ .

Fig. 3 illustrates the optimal orientations discussed previously.

C. Noise Amplification

To quantify the noise amplification of the two-step momentum algorithm (2), we utilize an alternative characterization of the stability triangle Δ . As illustrated in Fig. 2, let d and l denote the horizontal signed distances of the point (a, b) to the edges XZ and YZ

$$\begin{aligned} d(\lambda) &:= a(\lambda) + b(\lambda) + 1 \\ l(\lambda) &:= a(\lambda) - b(\lambda) + 1 \end{aligned} \quad (30a)$$

and let h denote its vertical signed distance to the edge XY

$$h(\lambda) := 1 - a(\lambda). \quad (30b)$$

Then, the equivalence condition

$$(b, a) \in \Delta \iff h, d, l > 0 \quad (31)$$

follows from the definition of the set Δ in (22b).

While analytical expressions for J in terms of the algorithmic parameters have been derived in the literature (e.g., see [60, Th. 1] for gradient descent, heavy-ball, and Nesterov's accelerated algorithms, and [39, Appendix A] and [43, Lemma A.1] for the general case), the novelty of Theorem 5 lies in expressing J in terms of the reciprocals of the distances $d(\lambda_i)$, $h(\lambda_i)$, and $l(\lambda_i)$ of the point $(b(\lambda_i), a(\lambda_i))$ to the edges of the stability triangle for the noisy two-step momentum algorithm (2). This geometric insight facilitates proofs of our main results. The proof of Theorem 5 is straightforward and it is omitted for brevity.

Theorem 5: For a strongly convex quadratic objective function $f \in \mathcal{Q}_m^L$ with the Hessian matrix Q , the steady-state variance of $x^t - x^*$ for the two-step momentum algorithm (2) with any stabilizing parameters (α, β, γ) is determined by

$$J = \sum_{i=1}^n \frac{\sigma_w^2}{2h(\lambda_i)} \left(\frac{1}{l(\lambda_i)} + \frac{1}{d(\lambda_i)} \right) =: \sum_{i=1}^n \hat{J}(\lambda_i).$$

Here, $\hat{J}(\lambda_i)$ denotes the modal contribution of the i th eigenvalue λ_i of Q to the steady-state variance, (d, h, l) are defined in (30), and (a, b) are given by (19c).

Theorem 5 demonstrates that J depends on the entire spectrum of the Hessian matrix Q and not only on its extreme eigenvalues m and L , which determine the convergence rate. Since for any $f \in \mathcal{Q}_m^L$ the extreme eigenvalues of Q are fixed at m and L , we have

$$\begin{aligned} J_{\max} &= \hat{J}(m) + \hat{J}(L) + (n - 2)\hat{J}_{\max} \\ J_{\min} &= \hat{J}(m) + \hat{J}(L) + (n - 2)\hat{J}_{\min} \end{aligned} \quad (32)$$

where $\hat{J}_{\max} := \max_{\lambda \in [m, L]} \hat{J}(\lambda)$, $\hat{J}_{\min} := \min_{\lambda \in [m, L]} \hat{J}(\lambda)$. We use (32) to determine explicit bounds on J_{\max} and J_{\min} in terms of the condition number and the settling time.

V. DESIGNING ORDERWISE PARETO-OPTIMAL ALGORITHMS WITH ADJUSTABLE PARAMETERS

We now utilize the geometric insight developed in Section IV to design algorithmic parameters that tradeoff settling time and noise amplification. In particular, we introduce two parameterized families of heavy-ball-like ($\gamma = 0$) and Nesterov-like ($\gamma = \beta$) algorithms that provide *continuous transformations* from gradient descent to the corresponding accelerated algorithm (with the optimal convergence rate) via a homotopy path parameterized by the settling time T_s . For both the iterate and gradient noise models, we establish an orderwise tight scaling $\Theta(\kappa^2)$ for $J_{\max} \times T_s$ and $J_{\min} \times T_s$ in accelerated regime (i.e., when T_s is smaller than the settling time of gradient descent with the optimal stepsize, $(\kappa + 1)/2$). This is a direct extension of [40, Th. 4], which studied gradient descent and its accelerated variants for the parameters that optimize the corresponding convergence rates.

We also examine performance tradeoffs for the parameterized family of heavy-ball-like algorithms with negative momentum parameter $\beta < 0$. This decelerated regime corresponds to settling times larger than $(\kappa + 1)/2$ and it captures a key difference between the two noise models: *for $T_s \geq (\kappa + 1)/2$, J_{\max} and J_{\min} grow linearly with the settling time T_s for the iterate noise model and they remain inversely proportional to T_s for the gradient noise model.* Comparison with the lower bounds in Theorems 2 and 3 shows that the parameterized family of heavy-ball-like methods yields orderwise optimal (in κ and T_s) J_{\max} and J_{\min} for both noise models. The results presented here prove all upper bounds in Theorems 4 and 3.

A. Parameterized Family of Heavy-Ball-Like Methods

For the two-step momentum algorithm (2) with $\gamma = 0$, the line segment $(b(\lambda), a(\lambda))$ parameterized by $\lambda \in [m, L]$ is parallel to the b -axis in the (b, a) -plane and it satisfies $a(\lambda) = \beta$. As described in Section IV, gradient descent and heavy-ball methods with the optimal parameters provided in Table I are obtained for $\beta = 0$ and $\beta = \rho^2$, respectively, and the corresponding end points $(b(m), a(m))$ and $(b(L), a(L))$ lie at the edges $X_\rho Z_\rho$ and $Y_\rho Z_\rho$ of the ρ -linear convergence triangle Δ_ρ . Inspired by this observation, we propose a family of parameters for which $\beta = c\rho^2$, for some scalar $c \in [-1, 1]$, and determine the stepsize α such that the aforementioned end points lie at $X_\rho Z_\rho$ and $Y_\rho Z_\rho$

$$\alpha = (1 + \rho)(1 + c\rho)/L, \quad \beta = c\rho^2, \quad \gamma = 0. \quad (33)$$

This yields a continuous transformation between the standard heavy-ball method ($c = 1$) and gradient descent ($c = 0$) for a fixed condition number κ . In addition, the momentum parameter β in (33) becomes negative for $c < 0$; see Fig. 3 for an illustration. In Lemma 4, we provide expressions for the scalar c as well as for \hat{J}_{\max} and \hat{J}_{\min} defined in (32) in terms of the condition number κ and the convergence rate ρ .

Lemma 4: For the class of functions \mathcal{Q}_m^L with the condition number $\kappa = L/m$, let the scalar ρ be such that

$$T_s = 1/(1 - \rho) \geq (\sqrt{\kappa} + 1)/2.$$

Then, the two-step momentum algorithm (2) with parameters (33) achieves the convergence rate ρ , and the largest and smallest values \hat{J}_{\max} and \hat{J}_{\min} of $\hat{J}(\lambda)$ satisfy

$$\hat{J}_{\max} = \hat{J}(m) = \hat{J}(L) = \frac{\sigma_w^2 (\kappa + 1)}{2(1 - c\rho^2)(1 + \rho)(1 + c\rho)}$$

$$\hat{J}_{\min} = \hat{J}(\hat{\lambda}) = \frac{\sigma_w^2}{(1 + c\rho^2)(1 - c\rho^2)}$$

where $\hat{\lambda} := (m + L)/2$ and the scalar c is given by

$$c := \frac{\kappa - (1 + \rho)/(1 - \rho)}{\rho(\kappa + (1 + \rho)/(1 - \rho))} \in [-1, 1]. \quad (34)$$

Proof: See Appendix D in [57]. \blacksquare

The parameters in (33) with c given by (34) are equivalent to the parameters presented in Theorem 4. Lemma 4 in conjunction with (32) allow us to compute J_{\max} and J_{\min} .

Corollary 2: The parameterized family of heavy-ball-like methods (33) satisfies $J_{\max} = n\hat{J}(m) = n\hat{J}(L)$ and $J_{\min} = 2\hat{J}(m) + (n - 2)\hat{J}(\hat{\lambda})$, where $\hat{J}(m)$ and $\hat{J}(\hat{\lambda})$ are given in Lemma 4, and J_{\max} and J_{\min} defined in (10) are the largest and smallest values of J when the algorithm is applied to $f \in \mathcal{Q}_m^L$ with the condition number $\kappa = L/m$.

Proposition 2 uses the expressions in Corollary 2 to establish orderwise tight upper and lower bounds on J_{\max} and J_{\min} for the parameterized family of heavy-ball-like algorithms (33). Our upper and lower bounds are within constant factors of each other and they are expressed in terms of the problem size n , condition number κ , and settling time T_s .

Proposition 2: For the parameterized family of heavy-ball-like methods (33), J_{\max} and J_{\min} in (10) satisfy

$$J_{\max} \times T_s = \sigma_w^2 p_{1c}(\rho) n \kappa (\kappa + 1) \quad (35a)$$

$$J_{\min} \times T_s = \sigma_w^2 \kappa (2 p_{1c}(\rho) (\kappa + 1) + (n - 2) p_{2c}(\rho)). \quad (35b)$$

Furthermore, for the gradient noise model ($\sigma_w = \alpha\sigma$)

$$J_{\max} \times T_s = \sigma^2 p_{3c}(\rho) n \kappa (\kappa + 1) \quad (36a)$$

$$J_{\min} \times T_s = \sigma^2 \kappa (2 p_{3c}(\rho) (\kappa + 1) + (n - 2) p_{4c}(\rho)) \quad (36b)$$

where

$$p_{1c}(\rho) := q_c(\rho)/(2(1 + \rho)^2(1 + c\rho^2))$$

$$p_{2c}(\rho) := q_c(\rho)/((1 + \rho)(1 + c\rho^2)(1 + c\rho))$$

$$p_{3c}(\rho) := q_c(\rho)/(2L^2)$$

$$p_{4c}(\rho) := q_c(\rho)q_{-c}(\rho)(1 + \rho)/L^2$$

$$q_c(\rho) := (1 - c\rho)/(1 - c\rho^2). \quad (37)$$

In addition, for $c \in [0, 1]$, $p_{1c}(\rho) \in [1/64, 1/2]$ and $p_{2c}(\rho) \in [1/16, 1]$; and for $c \in [-1, 1]$, $p_{3c}(\rho) \in [1/(4L^2), 1/L^2]$ and $p_{4c}(\rho) \in [1/(4L^2), 4/L^2]$.

Proof: See Appendix D [57]. ■

Proposition 3: For the parameterized family of heavy-ball-like methods (33) with $c \in [-1, 0]$, J_{\max} and J_{\min} in (10) satisfy

$$J_{\max} = \sigma_w^2 p_{5c}(\rho) n (1 + 1/\kappa) T_s \quad (38a)$$

$$J_{\min} = \sigma_w^2 (2 p_{5c}(\rho) (1 + 1/\kappa) + p_{6c}(\rho)(n - 2)/\kappa) T_s \quad (38b)$$

where $p_{5c}(\rho) := 1/(2(1 + |c|\rho)(1 + |c|\rho^2)) \in [1/8, 1/2]$ and $p_{6c}(\rho) := 2(1 + \rho)p_{5c}(\rho)q_{-c}(\rho) \in [1/8, 2]$.

Proof: See Appendix D in [57]. ■

The upper bounds in Theorems 4 and 3 follow from Propositions 2 and 3, respectively. Since these upper bounds have the same scaling as the corresponding lower bounds in Theorems 2 and 3 that hold for all stabilizing parameters (α, β, γ) , this demonstrates tightness of lower bounds for all settling times and for both noise models.

B. Parameterized Family of Nesterov-Like Methods

For the two-step momentum algorithm (2) with $\gamma = \beta$, the line segment $(b(\lambda), a(\lambda))$ parameterized by $\lambda \in [m, L]$ passes through the origin. As described in Section IV, gradient descent and Nesterov's method with the optimal parameters provided in Table I are obtained for $a = 0$ and $a = -(\rho/2)b$, respectively, and the corresponding end points $(b(m), a(m))$ and $(b(L), a(L))$ lie on the edges $X_\rho Z_\rho$ and $Y_\rho Z_\rho$ of the ρ -linear convergence triangle Δ_ρ . To provide a continuous transformation between these two standard algorithms, we introduce a parameter $c \in [0, 1/2]$, and let the line segment satisfy $a(\lambda) = -c\rho b(\lambda)$ and take its end points at the edges $X_\rho Z_\rho$ and $Y_\rho Z_\rho$; see Fig. 3 for an illustration. This can be accomplished with the following choice of parameters:

$$\begin{aligned} \alpha &= (1 + \rho)(1 + c - c\rho)/(L(1 + c)) \\ \gamma &= \beta = c\rho^2/((\alpha L - 1)(1 + c)). \end{aligned} \quad (39)$$

For the parameterized family of Nesterov-like algorithms (39), Proposition 4 establishes the settling time and characterizes the dependence of $J_{\min} \times T_s$ and $J_{\max} \times T_s$ on the condition number κ and the problem size n .

Proposition 4: For the class of functions \mathcal{Q}_m^L with the condition number $\kappa = L/m$, let the scalar ρ be such that

$$T_s = 1/(1 - \rho) \in [(\sqrt{3\kappa + 1})/2, (\kappa + 1)/2].$$

The two-step momentum algorithm (2) with parameters (39) achieves the convergence rate ρ and satisfies

$$J_{\max} \times T_s \geq \sigma_w^2 ((n - 1)\kappa(\kappa + 1)/32 + \sqrt{3\kappa + 1})/2$$

$$J_{\max} \times T_s \leq 6 \sigma_w^2 n \kappa (3\kappa + 1)$$

$$J_{\min} \times T_s \geq \sigma_w^2 (\kappa(\kappa + 1)/32 + (n - 1)\sqrt{3\kappa + 1})/2$$

$$J_{\min} \times T_s \leq \sigma_w^2 (6 \kappa (3\kappa + 1) + (n - 1)(\kappa + 1))/2$$

where J_{\max} and J_{\min} are the largest and smallest values that J can take when the algorithm is applied to $f \in \mathcal{Q}_m^L$ with the condition number $\kappa = L/m$, and the scalar $c \in [0, 1/2]$ is the




Fig. 4. Triangle Δ_ρ (yellow) and the line segments $(b(\lambda), a(\lambda))$ with $\lambda \in [m, L]$ (blue) for gradient descent with reduced stepsize (40) and heavy-ball-like method (41), which place the end point $(b(m), a(m))$ at X_ρ and the end point $(b(L), a(L))$ at $(2c'\rho, \rho^2)$ on the edge $X_\rho Y_\rho$, where $c' := \kappa(1 - \rho)^2/\rho - (1 + \rho^2)/\rho$ ranges over the interval $[-1, 1]$.

solution to the quadratic equation

$$\kappa(1 - \rho)(1 - c\rho - c^2(1 + \rho)) = (1 + \rho)(1 - c\rho - c^2(1 - \rho)).$$

Proof: See Appendix D in [57]. ■

Since α in (39) satisfies $\alpha L \in [1, 3]$, comparing the upper bounds in Proposition 4 with the lower bounds in Theorem 2 shows that, for settling times $T_s \leq (\kappa + 1)/2$, (39) achieves orderwise optimal J_{\max} and J_{\min} for both the iterate ($\sigma_w = \sigma$) and gradient ($\sigma_w = \alpha\sigma$) noise models.

C. Impact of Reducing the Stepsize

When the only source of uncertainty is a noisy gradient, i.e., $\sigma_w = \alpha\sigma$, one can attempt to reduce the noise amplification J by decreasing the stepsize α at the expense of increasing the settling time $T_s = 1/(1 - \rho)$ [13], [24], [39], [41]. In particular, for gradient descent, α can be reduced from its optimal value $2/(L + m)$ by keeping $(b(m), a(m))$ at $(-\rho, 0)$ and moving the point $(b(L), a(L))$ from $(\rho, 0)$ toward $(-\rho, 0)$ along the horizontal axis; see Fig. 4. This can be accomplished with

$$\alpha = (1 + c\rho)/L, \quad \gamma = \beta = 0 \quad (40)$$

for some $c \in [-1, 1]$ parameterizing $(b(L), a(L)) = (c\rho, 0)$. In this case, the settling time satisfies $T_s = (\kappa + c)/(c + 1) \in [(\kappa + 1/2), \infty)$ and similar arguments to those presented in the proof of Lemma 4 can be used to obtain

$$\begin{aligned} \hat{J}_{\max} &= \hat{J}(m) = \sigma^2 \kappa^2 (1 - \rho)/L^2 \\ \hat{J}_{\min} &= \begin{cases} \hat{J}(L) & = \sigma^2 \alpha^2 / (1 - c^2 \rho^2) & c \leq 0 \\ \hat{J}(1/\alpha) & = \sigma^2 \alpha^2 & c \geq 0. \end{cases} \end{aligned}$$

For a fixed n , the stepsize in (40) yields a $\Theta(\kappa^2)$ scaling for both $J_{\max} \times T_s$ and $J_{\min} \times T_s$ for all $c \in [-1, 1]$. Thus, gradient descent with reduced stepsize orderwise matches the lower bounds in Theorem 2. An integral quadratic constraints (IQC)-based approach [40, Lemma 1] was utilized in [41, Th. 13] to show that α in (40) also yields the previously discussed convergence rate and worst-case noise amplification for one-point m -strongly convex L -smooth functions.

Remark 6: Any desired settling time $T_s = 1/(1 - \rho) \in [(\sqrt{\kappa} + 1)/2, \infty)$ can be achieved by the heavy-ball-like method with reduced stepsize

$$\alpha = (1 - \rho)^2/m, \quad \beta = \rho^2, \quad \gamma = 0. \quad (41)$$




Fig. 5. Triangle Δ_ρ (yellow) and the line segments $(b(\lambda), a(\lambda))$ with $\lambda \in [m, L]$ (blue) for the Nesterov-like method with reduced stepsize (40), which places the end point $(b(m), a(m))$ at X_ρ and the end point $(b(L), a(L))$ at $(-2c\rho, c\rho^2)$, where $c := \kappa - 1/(1 - \rho)^2$ ranges over the interval $[0, 1]$.

This choice yields $J_{\max} = \sigma^2 n \kappa^2 (1 - \rho^4) / (L^2 (1 + \rho)^4)$ for the gradient noise model $\sigma_w = \alpha \sigma$ [41, Th. 9]; see Fig. 4. In addition, by considering the error in $y^t = x^t + \gamma(x^t - x^{t-1})$ as the performance metric, it was stated and numerically verified in [41] that the choice of parameters (41) yields Pareto-optimal algorithms for simultaneously optimizing J_{\max} and ρ for the gradient noise model $\sigma_w = \alpha \sigma$. We note that the settling time $T_s = \Theta(\kappa)$ of gradient descent with standard stepsizes ($\alpha = 1/L$ or $2/(m + L)$) can be achieved via (41) by reducing α to $O(1/(\kappa L))$. In contrast, the parameterized family of heavy-ball-like methods (33) is orderwise Pareto-optimal (cf., Theorems 2 and 3) while maintaining $\alpha \in [1/L, 4/L]$.

Remark 7: Any desired settling time $T_s = 1/(1 - \rho) \in [\sqrt{\kappa}, \infty)$ can be achieved by the Nesterov-like method with reduced stepsize

$$\alpha = (1 - \rho)^2 / m, \quad \beta = \gamma = \frac{\rho}{2 - \rho}. \quad (42)$$

This choice makes the line segment $(b(\lambda), a(\lambda))$ for $\lambda \in [m, L]$ pass through the origin with the endpoint $(b(m), a(m))$ at the vertex $X_\rho = (-2\rho, \rho^2)$ of the ρ -exponential stability triangle Δ_ρ , and it yields $(b(L), a(L)) = (-2c\rho, c\rho^2)$, where $c := \kappa - T_s^2$ is a nonnegative constant; see Fig. 5. For the gradient noise model $\sigma_w = \alpha \sigma$, these parameters lead to

$$\hat{J}_{\max} = \hat{J}(m) = \frac{\sigma^2 \kappa^2 (1 + \rho^2)}{L^2 (1 + \rho)^3 T_s} \approx \frac{\sigma^2 \kappa^2}{4L^2 T_s}. \quad (43)$$

We note that, as $\rho \rightarrow 1$, the largest modal contribution to noise amplification \hat{J}_{\max} becomes inversely proportional to the settling time T_s . The family of parameters in (42) were utilized in [39] to propose the multistage accelerated stochastic gradient (M-ASG) algorithm as a means to systematically tradeoff convergence rate and noise amplification. For strongly convex problems, this algorithm optimally reduces the error in function values, thereby matching the fundamental lower bound established in [61]. In particular, at every stage $k \in \{0, 1, \dots\}$, M-ASG performs a specific restart that balances the initial condition followed by N_k Nesterov-like-iterations with

$$\alpha_k := \frac{1}{4^k L}, \quad \beta_k = \gamma_k := \frac{1 - \sqrt{\alpha_k m}}{1 + \sqrt{\alpha_k m}} \quad (44)$$

where N_k is proportional to $2^k / \kappa$, m is the parameter of strong convexity, and L is the Lipschitz constant. When restricted to strongly convex quadratic problems, the parameters in (44) are




Fig. 6. Exponential stability triangles Δ_{ρ_k} along with the line segments $(b_k(\lambda), a_k(\lambda))$ with $\lambda \in [m, L]$ associated with the first three stages $k = 0, 1, 2$ in the M-ASG algorithm. The parabola $b^2 = 4a$ determines the locus of vertices X_ρ and Y_ρ .

identical to those in (42) with the convergence rate $\rho_k = 1 - 1/(2^k \sqrt{\kappa})$; see Fig. 6. It is straightforward to show that while M-ASG reduces the largest contribution to noise amplification $\hat{J}(m)$ to half by going to the next stage, it also doubles the settling time. Finally, contrasting (43) with the lower bounds in (51a) and (14a), established in Theorem 3, allows us to conclude that M-ASG preserves $J \times T_s$ near the Pareto-optimal curve at each stage while achieving the optimal iteration complexity [61] by successively reducing the stepsize to half of its previous value and utilizing a suitable iteration count N_k .

VI. PROOFS OF THEOREMS 1–4

A. Proof of Theorem 1

From Theorem 5, it follows that we can use upper bounds on $\hat{J}(\lambda)$ over $\lambda \in [m, L]$ to establish an upper bound on J . Since the algorithm achieves the convergence rate ρ , combining (20) and Lemma 2 yield $(b(\lambda), a(\lambda)) \in \Delta_\rho$ for all $\lambda \in [m, L]$. As we demonstrate in Appendix B, the function \hat{J} is convex in (b, a) over the stability triangle Δ . In addition, $\Delta_\rho \subset \Delta$ is the convex hull of the points X_ρ, Y_ρ, Z_ρ in the (b, a) -plane. Since the maximum of a convex function over the convex hull of a finite set of points is attained at one of these points, \hat{J} attains its maximum over Δ_ρ at X_ρ, Y_ρ , or Z_ρ .

Using the definition of X_ρ, Y_ρ , and Z_ρ in (23c), the affine relations (30), and the analytical expression for \hat{J} in Theorem 5, it follows that the maximum occurs at vertices X_ρ and Y_ρ

$$\hat{J}_{\max} := \max_{\lambda \in [m, L]} \hat{J}(\lambda) = \sigma_w^2 (1 + \rho^2) / ((1 - \rho)^3 (1 + \rho)^3)$$

where we use $d_{X_\rho} = l_{Y_\rho} = (1 - \rho)^2$, $l_{X_\rho} = d_{Y_\rho} = (1 + \rho)^2$, and $h_{X_\rho} = h_{Y_\rho} = 1 - \rho^2$. Combining the aforementioned identity with Theorem 5 completes the proof of (12a).

We use an argument similar to the proof of Proposition 1 to prove (12b). In particular, since $(b(L), a(L)) \in \Delta_\rho$, we have

$$\alpha L = d(L) \leq d_{\max} = (1 + \rho)^2$$

where d given by (25) is the horizontal signed distance to the edge XZ of the stability triangle Δ . On the other hand, d_{\max} is the largest value that d can take among all points $(b, a) \in \Delta_\rho$ and it corresponds to the vertex Y_ρ ; see (27a). Combining this inequality with $\sigma_w = \alpha \sigma$ and (12a) completes the proof of Theorem 1.




Fig. 7. Line \mathcal{L} (blue, dashed) and the intersection point G , along with the distances d_1 , h_1 , d_G , and h_G as introduced in the proof of Theorem 2.

B. Proof of Theorem 2

Using the expression $J = \sum_i \hat{J}(\lambda_i)$ established in Theorem 5, we have the decomposition

$$J = \hat{J}(m) + \sum_{i=1}^{n-1} \hat{J}(\lambda_i). \quad (45)$$

To prove the lower bounds (13b) and (14b) on J_{\min} , we establish a lower bound on $\hat{J}(m) \times T_s$ that scales quadratically with κ , and a general lower bound on $\hat{J}(\lambda) \times T_s$.

1) Case $\sigma_w = \sigma$: The proof of (13b) utilizes the inequalities

$$\hat{J}(m) \times T_s \geq \sigma_w^2 \kappa^2 / (2(1 + \rho)^5) \quad (46a)$$

$$\hat{J}(\lambda) \times T_s \geq \sigma_w^2 (\sqrt{\kappa} + 1) / 2. \quad (46b)$$

We first prove (46a). Our approach builds on the proof of Proposition 1. In particular, $d(\lambda) = \alpha\lambda$ for the point $(b(\lambda), a(\lambda))$, where d and (b, a) are defined in (30) and (19c), respectively. Thus, $d(m) = d(L)/\kappa$. Furthermore, Lemma 3 implies $(b(\lambda), a(\lambda)) \in \Delta_\rho$ for $\lambda \in [m, L]$. Thus, the trivial inequality $d(L) \leq d_{\max}$ leads to

$$d(m) \leq d_{\max}/\kappa = (1 + \rho)^2/\kappa \quad (47)$$

where $d_{\max} = (1 + \rho)^2$ is the largest value that d can take among all points $(b, a) \in \Delta_\rho$; see (27a). We now use Theorem 5 to write

$$\hat{J}(\lambda) = \frac{\sigma_w^2 (d(\lambda) + l(\lambda))}{2d(\lambda)h(\lambda)l(\lambda)} \geq \frac{\sigma_w^2}{2d(\lambda)h(\lambda)}. \quad (48)$$

Next, we lower bound the right-hand side of (48). Let \mathcal{L} be the line that passes through $(b(\lambda), a(\lambda))$, which is parallel to the edge XZ of the stability triangle Δ , and let G be the intersection of \mathcal{L} and the edge $X_\rho Z_\rho$ of the ρ -stability triangle Δ_ρ ; see Fig. 7 for an illustration. It is easy to verify that

$$h_G \geq h(\lambda), \quad d_G = d(\lambda) \quad (49a)$$

where h_G and d_G correspond to the values of h and d associated with the point G . In addition, since G lies on the edge $X_\rho Z_\rho$, h_G and d_G satisfy the affine relation

$$h_G = 1 - \rho + d_G \rho / (1 - \rho). \quad (49b)$$

This follows from the equation of the line $X_\rho Z_\rho$ in the (b, a) -plane and from the definitions of d and h in (30). Furthermore, combining (49a) and (49b) implies

$$h(\lambda)(1 - \rho) \leq h_G(1 - \rho) = (1 - \rho)^2 + \rho d(\lambda). \quad (50a)$$

For $\lambda = m$, we can further write

$$2d(m) \left((1 - \rho)^2 + \rho d(m) \right) \leq 2(1 + \rho)^5 / \kappa^2 \quad (50b)$$

where the inequality is obtained from (28) and (47). Combining (48), (50a), and (50b) completes the proof of (46a).

Next, we prove the general lower bound in (46b). As we demonstrate in Appendix B, the modal contribution \hat{J} to the noise amplification is a convex function of (b, a) which takes its minimum $\hat{J}_{\min} = \sigma_w^2$ over the stability triangle Δ at the origin $b = a = 0$. Combining this fact with the lower bound in (7) on ρ completes the proof of (46b).

Finally, we can obtain the lower bound (13b) on J_{\min} by combining (45) and (46).

2) Case $\sigma_w = \alpha\sigma$: The proof of (14b) utilizes

$$\hat{J}(\lambda) \times T_s \geq \sigma^2 / (2\lambda^2(1 + \rho)) \quad (51a)$$

$$\hat{J}(\lambda) \times T_s \geq \sigma^2 (1 - \rho)^3 \kappa^2 / L^2. \quad (51b)$$

In particular, (14b) follows from using (51a) for $\lambda = m$ and taking the maximum of (51a) and (51b) for the other eigenvalues to bound the expression for J in Theorem 5.

We first prove (51a). By combining (48) and (50a), we obtain

$$\frac{\hat{J}(\lambda)}{1 - \rho} \geq \frac{\alpha^2 \sigma^2}{2d(\lambda) \left((1 - \rho)^2 + \rho d(\lambda) \right)}. \quad (52)$$

Since $d(\lambda) \geq d_{\min} := (1 - \rho)^2$, where d_{\min} is the smallest value of d over Δ_ρ , [cf. (27b)], we can write

$$\frac{\alpha^2 \sigma^2}{(1 - \rho)^2 + \rho d(\lambda)} \geq \frac{\alpha^2 \sigma^2}{d(\lambda)(1 + \rho)} = \frac{\sigma^2 d(\lambda)}{\lambda^2 (1 + \rho)}. \quad (53)$$

Combining (52) and (53) completes the proof of (51a).

To prove (51b), we use $d(\lambda) \geq d_{\min} := (1 - \rho)^2$ and $d(m) = \alpha m$, to obtain $\alpha \geq (1 - \rho)^2 \kappa / L$. Combining this inequality with $\hat{J}_{\min} = \sigma_w^2 = \alpha^2 \sigma^2$ yields (51b). Finally, we obtain the lower bound (14b) on J_{\min} by combining (45) and (51).

To obtain the lower bounds (13a) and (14a) on J_{\max} , we consider a quadratic function for which the Hessian has $n - 1$ eigenvalues at $\lambda = m$ and one eigenvalue at $\lambda = L$. For such a function, we can use Theorem 5 to write

$$J_{\max} \geq J = (n - 1)\hat{J}(m) + \hat{J}(L). \quad (54)$$

3) Case $\sigma_w = \sigma$: To prove (13a), we use inequalities in (46a) and (46b) to bound $\hat{J}(m)/(1 - \rho)$ and $\hat{J}(L)/(1 - \rho)$ in (54), respectively.

4) Case $\sigma_w = \alpha\sigma$: To prove (14a), we use inequality in (51a) with $\lambda = m$ to lower bound $\hat{J}(m)/(1 - \rho)$, and combine (51a) and (51b) to lower bound $\hat{J}(L)/(1 - \rho)$ in (54).

C. Proof of Theorem 3

The following proposition allows us to prove the lower bounds in Theorem 3.

Proposition 5: Let $\rho = \rho(A) = 1 - 1/T_s$ be the convergence rate of the two-step momentum algorithm (2). Then, the largest and smallest modal contributions to noise amplification given by (32) satisfy $\hat{J}_{\max} \geq \sigma_w^2 / (2(1 + \rho)^2) T_s$ and $\hat{J}_{\min} \geq \sigma_w^2$.

Proof: The inequality $J_{\min} \geq \sigma_w^2$ follows from the fact that \hat{J} , as a function of (b, a) , takes its minimum value at the

origin; see Appendix B. The proof for \hat{J}_{\max} utilizes the fact that for any constant parameters (α, β, γ) and fixed condition number, the spectral radius $\rho(A)$ corresponds to the smallest ρ -linear convergence triangle Δ_ρ that contains the line segment $(b(\lambda), a(\lambda))$ for $\lambda \in [m, L]$. Thus, at least one of the end points $(b(m), a(m))$ or $(b(L), a(L))$ will be on the boundary of the triangle $\Delta_{\rho(A)}$. Combining this with the fact that $d(m) \leq d(L)$, it follows that at least one of the following holds:

$$\begin{aligned} (b(m), a(m)) &\in X_\rho Z_\rho \text{ or } X_\rho Y_\rho \\ (b(L), a(L)) &\in Y_\rho Z_\rho \text{ or } X_\rho Y_\rho. \end{aligned}$$

Together with the concrete values of vertices (22c) in terms of ρ , this yields

$$1 - \rho \geq \min \{h(m), h(L), l(L)/(1 + \rho), d(m)/(1 + \rho)\}. \quad (55)$$

Also, using Theorem 5 and noting that the maximum values that $h(\lambda)$, $d(\lambda)$, and $l(\lambda)$ can take among Δ_ρ are given by $1 + \rho^2$, $(1 + \rho)^2$, and $(1 + \rho)^2$, respectively, we can write

$$\begin{aligned} \hat{J}(m) &\geq \frac{\sigma_w^2}{2h(m)d(m)} \\ &\geq \max \left\{ \frac{\sigma_w^2}{2h(m)(1 + \rho)^2}, \frac{\sigma_w^2}{2d(m)(1 + \rho^2)} \right\} \\ \hat{J}(L) &\geq \frac{\sigma_w^2}{2h(L)l(L)} \\ &\geq \max \left\{ \frac{\sigma_w^2}{2h(L)(1 + \rho)^2}, \frac{\sigma_w^2}{2l(L)(1 + \rho^2)} \right\}. \quad (56) \end{aligned}$$

Finally, by the convexity of \hat{J} (see Appendix B), we have $\hat{J}_{\max} \geq \max\{\hat{J}(m), \hat{J}(L)\}$. Combining this with (55) and (56) completes the proof. ■

The lower bounds in Theorem 3 follow from combining Proposition 5 with the expression for J in Theorem 5.

D. Proof of Theorem 4

As described in Section V, the parameters in Theorem 4 are obtained by placing the end points of the horizontal line segment $(b(\lambda), a(\lambda))$ parameterized by $\lambda \in [m, L]$ at the edges $X_\rho Z_\rho$ and $Y_\rho Z_\rho$ of the ρ -linear convergence triangle Δ_ρ . These parameters can be equivalently represented by (33) where the scalar c given in Lemma 4 satisfies $c \in [0, 1]$ if and only if $T_s \leq (\kappa + 1)/2$ and it satisfies $c \in [-1, 0]$ if and only if $T_s \geq (\kappa + 1)/2$. The proof of Theorem 4 follows from combining Lemma 4 and Propositions 2 and 3.

VII. CONCLUSION

We have examined the amplification of stochastic disturbances for a class of two-step momentum algorithms in which the iterates are perturbed by an additive white noise that arises from uncertainties in gradient evaluation or in computing the iterates. For both noise models, we establish lower bounds on the product of the settling time and the smallest/largest steady-state

variance of the error in the optimization variable. These bounds scale with κ^2 for all stabilizing parameters, which reveals a fundamental limitation imposed by the condition number κ in designing algorithms that tradeoff noise amplification and convergence rate. In addition, we provide a novel geometric viewpoint of stability and ρ -linear convergence. This viewpoint brings insight into the relation between noise amplification, convergence rate, and algorithmic parameters. It also allows us to

- 1) take an alternative approach to optimizing convergence rates for standard algorithms;
- 2) identify key similarities and differences between the iterate and gradient noise models;
- 3) introduce parameterized families of algorithms for which the parameters can be continuously adjusted to tradeoff noise amplification and settling time.

By utilizing positive and negative momentum parameters in accelerated and decelerated regimes, respectively, we demonstrate that a parameterized family of the heavy-ball-like algorithms can achieve orderwise Pareto optimality for all settling times and both noise models.

Our ongoing work focuses on extending these results to algorithms with more complex structures including update strategies that utilize information from more than the last two iterates and time-varying algorithmic parameters [62]. It is also of interest to identify fundamental performance limitations of stochastic gradient descent algorithms in which both additive and multiplicative stochastic disturbances exist [63], [64].

APPENDIX

A. Settling Time

If ρ denotes the linear convergence rate, $T_s = 1/(1 - \rho)$ quantifies the *settling time*. The inequality in (5) shows that $c\rho^t \leq \epsilon$ provides a sufficient condition for reaching the accuracy level ϵ with $\|\psi^t\|_2/\|\psi^0\|_2 \leq \epsilon$. Taking the logarithm of $c\rho^t \leq \epsilon$ and using the first-order Taylor series approximation $\log(1 - x) \approx -x$ around $x = 0$ yields a sufficient condition on the number of iterations t for an algorithm to reach ϵ -accuracy

$$t \geq \log(\epsilon/c)/\log(1 - 1/T_s) \approx T_s \log(c/\epsilon).$$

B. Convexity of Modal Contribution \hat{J} to Noise Amplification

By Theorem 5, we have

$$\frac{\hat{J}}{\sigma_w^2} = \frac{d + l}{2dh} = \frac{1}{2hd} + \frac{1}{2hl}$$

where we have dropped the dependence on λ for simplicity. The functions $1/(2hd)$ and $1/(2hl)$ are both convex over the positive orthant $d, h, l > 0$. Thus, \hat{J} is convex with respect to (d, h, l) . In addition, since d, h , and l are all affine functions of a and b , we can use the equivalence relation in (31) to conclude that \hat{J} is also convex in (b, a) over the stability triangle Δ . Finally, since $b(\lambda)$ and $a(\lambda)$ are affine in λ , it follows that for any stabilizing parameters, \hat{J} is also convex with respect to λ over the interval $[m, L]$.

Convexity of \hat{J} allows us to use first-order conditions to find its minimizer. In particular, since for $\sigma_w = 1$

$$\begin{aligned}\frac{\partial \hat{J}}{\partial d} &= -\frac{1}{2hd^2}, & \frac{\partial \hat{J}}{\partial l} &= -\frac{1}{2hl^2}, & \frac{\partial \hat{J}}{\partial h} &= -\frac{l+d}{2h^2dl} \\ \frac{\partial d}{\partial a} &= \frac{\partial l}{\partial a} = -\frac{\partial h}{\partial a} = \frac{\partial d}{\partial b} = -\frac{\partial l}{\partial b} = 1, & \frac{\partial h}{\partial b} &= 0\end{aligned}$$

it is easy to verify that $\partial \hat{J} / \partial a = \partial \hat{J} / \partial b = 0$ at $a = b = 0$. Thus, \hat{J} takes its minimum $\hat{J}_{\min} = \sigma_w^2$ over the stability triangle Δ at $a = b = 0$, which corresponds to $d = h = l = 1$.

C. Proofs of Section IV

1) Proof of Lemma 2: We start by noting that $\rho(M) \leq \rho$ if and only if $\rho(M') \leq 1$ where $M' := M/\rho$. The characteristic polynomial associated with M' , $F_{\rho}(z) = z^2 + (b/\rho)z + a/\rho^2$, allows us to use similar arguments to those presented in the proof of Lemma 1 to show that

$$\rho(M') \leq 1 \iff (b/\rho, a/\rho^2) \in \Delta_1 \quad (57)$$

where $\Delta_1 := \{(b, a) \mid |b| - 1 \leq a \leq 1\}$ is the closure of the set Δ in (22b). Finally, the condition on the right-hand side of (57) is equivalent to $(b, a) \in \Delta_{\rho}$, where Δ_{ρ} is given by (23b).

REFERENCES

- [1] B. T. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Comput. Math. Math. Phys.*, vol. 4, no. 5, pp. 1–17, 1964.
- [2] Y. Nesterov, "A method for solving the convex programming problem with convergence rate $O(1/k^2)$," in *Proc. Dokl. Akad. Nauk SSSR*, 1983, vol. 27, pp. 543–547.
- [3] Y. Nesterov, "Gradient methods for minimizing composite objective functions," *Math. Prog.*, vol. 140, no. 1, pp. 125–161, 2013.
- [4] L. Bottou and Y. Le Cun, "On-line learning for very large data sets," *Appl. Stochastic Models Bus. Ind.*, vol. 21, no. 2, pp. 137–151, 2005.
- [5] M. Hong, M. Razaviyayn, Z.-Q. Luo, and J.-S. Pang, "A unified algorithmic framework for block-structured optimization involving Big Data: With applications in machine learning and signal processing," *IEEE Signal Process. Mag.*, vol. 33, no. 1, pp. 57–77, Jan. 2016.
- [6] A. Badithela and P. Seiler, "Analysis of the heavy-ball algorithm using integral quadratic constraints," in *Proc. IEEE Amer. Control Conf.*, 2019, pp. 4081–4085.
- [7] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 1139–1147.
- [8] Y. Nesterov, *Lectures on Convex Optimization* (Springer Optimization and Its Applications), vol. 137. Cham, Switzerland: Springer, 2018.
- [9] L. Lessard, B. Recht, and A. Packard, "Analysis and design of optimization algorithms via integral quadratic constraints," *SIAM J. Optim.*, vol. 26, no. 1, pp. 57–95, 2016.
- [10] B. Hu and L. Lessard, "Dissipativity theory for Nesterov's accelerated method," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1549–1557.
- [11] S. Cyrus, B. Hu, B. Van Scoy, and L. Lessard, "A robust accelerated optimization algorithm for strongly convex functions," in *Proc. Amer. Control Conf.*, 2018, pp. 1376–1381.
- [12] B. Van Scoy, R. A. Freeman, and K. M. Lynch, "The fastest known globally convergent first-order method for minimizing strongly convex functions," *IEEE Control Syst. Lett.*, vol. 2, no. 1, pp. 49–54, Jan. 2018.
- [13] M. Fazlyab, A. Ribeiro, M. Morari, and V. M. Preciado, "Analysis of optimization algorithms via integral quadratic constraints: Nonstrongly convex problems," *SIAM J. Optim.*, vol. 28, no. 3, pp. 2654–2689, 2018.
- [14] C. W. Scherer, C. Ebenbauer, and T. Holicki, "Optimization algorithm synthesis based on integral quadratic constraints: A tutorial," in *Proc. 62th IEEE Conf. Decis. Control*, 2023, pp. 2995–3002.
- [15] Y. Drori and M. Teboulle, "Performance of first-order methods for smooth convex minimization: A novel approach," *Math. Program.*, vol. 145, no. 1, pp. 451–482, 2014.
- [16] D. Kim and J. A. Fessler, "Generalizing the optimized gradient method for smooth convex minimization," *SIAM J. Optim.*, vol. 28, no. 2, pp. 1920–1950, 2018.
- [17] A. B. Taylor, "Convex interpolation and performance estimation of first-order methods for convex optimization," Ph.D. dissertation, Inst. Inf. Commun. Technol., Electron. Appl. Math., Universite Catholique de Louvain, Ottignies-Louvain-la-Neuve, Belgium, 2017.
- [18] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [19] Y. Ouyang, Y. Chen, G. Lan, and E. Pasilio Jr., "An accelerated linearized alternating direction method of multipliers," *SIAM J. Imag. Sci.*, vol. 8, no. 1, pp. 644–681, 2015.
- [20] J. Bai, W. W. Hager, and H. Zhang, "An inexact accelerated stochastic ADMM for separable convex optimization," *Comput. Optim. Appl.*, vol. 81, no. 2, pp. 479–518, 2022.
- [21] B. T. Polyak, "Comparison of the convergence rates for single-step and multi-step optimization algorithms in the presence of noise," *Eng. Cybern.*, vol. 15, no. 1, pp. 6–10, 1977.
- [22] Y. Bengio, "Gradient-based optimization of hyperparameters," *Neural Comput.*, vol. 12, no. 8, pp. 1889–1900, 2000.
- [23] D. Maclaurin, D. Duvenaud, and R. Adams, "Gradient-based hyperparameter optimization through reversible learning," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2113–2122.
- [24] K. Yuan, B. Ying, and A. H. Sayed, "On the influence of momentum acceleration on online learning," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 6602–6667, 2016.
- [25] A. Beirami, M. Razaviyayn, S. Shahrapour, and V. Tarokh, "On optimal generalizability in parametric learning," in *Proc. Neural Inf. Process.*, 2017, pp. 3458–3468.
- [26] S. Michalowsky, C. Scherer, and C. Ebenbauer, "Robust and structure exploiting optimisation algorithms: An integral quadratic constraint approach," *Int. J. Control*, vol. 94, no. 11, pp. 2956–2979, 2021.
- [27] Z.-Q. Luo and P. Tseng, "Error bounds and convergence analysis of feasible descent methods: A general approach," *Ann. Oper. Res.*, vol. 46, no. 1, pp. 157–178, 1993.
- [28] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, pp. 400–407, 1951.
- [29] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM J. Optim.*, vol. 19, no. 4, pp. 1574–1609, 2009.
- [30] O. Devolder, "Exactness, inexactness and stochasticity in first-order methods for large-scale convex optimization," Ph.D. dissertation, Inst. Inf. Commun. Technol., Electron. Appl. Math., University of Louvain, Louvain-la-Neuve, Belgium, 2013.
- [31] P. Dvurechensky and A. Gasnikov, "Stochastic intermediate gradient method for convex problems with stochastic inexact oracle," *J. Optim. Theory Appl.*, vol. 171, no. 1, pp. 121–145, 2016.
- [32] B. O'Donoghue and E. Candes, "Adaptive restart for accelerated gradient schemes," *Found. Comput. Math.*, vol. 15, pp. 715–732, 2015.
- [33] B. T. Polyak and G. V. Smirnov, "Transient response in matrix discrete-time linear systems," *Autom. Remote Control*, vol. 80, no. 9, pp. 1645–1652, 2019.
- [34] H. Mohammadi, S. Samuelson, and M. R. Jovanović, "Transient growth of accelerated optimization algorithms," *IEEE Trans. Automat. Control*, vol. 68, no. 3, pp. 1823–1830, Mar. 2023.
- [35] B. Can, M. Gurbuzbalaban, and L. Zhu, "Accelerated linear convergence of stochastic momentum methods in Wasserstein distances," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 891–901.
- [36] O. Devolder, "Stochastic first order methods in smooth convex optimization," CORE Discussion, 2011, Paper 2011/70.
- [37] M. B. Cohen, J. Diakonikolas, and L. Orecchia, "On acceleration with noise-corrupted gradients," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 1019–1028.
- [38] S. Ghadimi and G. Lan, "Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization, II: Shrinking procedures and optimal algorithms," *SIAM J. Opt.*, vol. 23, no. 4, pp. 2061–2089, 2013.
- [39] N. S. Aybat, A. Fallah, M. Gurbuzbalaban, and A. Ozdaglar, "A universally optimal multistage accelerated stochastic gradient method," in *Proc. Neural Inf. Process.*, 2019.
- [40] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Robustness of accelerated first-order algorithms for strongly convex optimization problems," *IEEE Trans. Automat. Control*, vol. 66, no. 6, pp. 2480–2495, Jun. 2021.

- [41] B. Van Scoy and L. Lessard, "The speed-robustness trade-off for first-order methods with additive gradient noise," 2021, *arXiv:2109.05059*.
- [42] N. S. Aybat, A. Fallah, M. Gurbuzbalaban, and A. Ozdaglar, "Robust accelerated gradient methods for smooth strongly convex functions," *SIAM J. Opt.*, vol. 30, no. 1, pp. 717–751, 2020.
- [43] B. Can and M. Gurbuzbalaban, "Entropic risk-averse generalized momentum methods," 2022, *arXiv:2204.11292*.
- [44] I. Gitman, H. Lang, P. Zhang, and L. Xiao, "Understanding the role of momentum in stochastic gradient methods," in *Proc. Neural Inf. Process.*, 2019, vol. 32.
- [45] N. Kuru, S. Ilker Birbil, M. Gurbuzbalaban, and S. Yildirim, "Differentially private accelerated optimization algorithms," *SIAM J. Opt.*, vol. 32, no. 2, pp. 795–821, 2022.
- [46] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, vol. 87. Berlin, Germany: Springer, 2004.
- [47] Y. Arjevani, S. Shalev-Shwartz, and O. Shamir, "On lower and upper bounds in smooth and strongly convex optimization," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 4303–4353, 2016.
- [48] Z. He, A. S. Rakin, and D. Fan, "Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 588–597.
- [49] B. T. Polyak, "Introduction to optimization," *Optim. Softw. Inc., New York*, vol. 1, 1987.
- [50] R. Bassily, A. Smith, and A. Thakurta, "Private empirical risk minimization: Efficient algorithms and tight error bounds," in *Proc. IEEE 55th Annu. Symp. Found. Comput. Sci.*, 2014, pp. 464–473.
- [51] S. B. Gelfand and S. K. Mitter, "Recursive stochastic algorithms for global optimization in R^d ," *SIAM J. Control Optim.*, vol. 29, no. 5, pp. 999–1018, 1991.
- [52] M. Raginsky, A. Rakhlin, and M. Telgarsky, "Non-convex learning via stochastic gradient Langevin dynamics: A nonasymptotic analysis," in *Proc. Conf. Learn. Theory*, 2017, pp. 1674–1703.
- [53] Y. Zhang, P. Liang, and M. Charikar, "A hitting time analysis of stochastic gradient Langevin dynamics," in *Proc. Conf. Learn. Theory*, 2017, pp. 1980–2022.
- [54] V. Ugrinovskii, I. R. Petersen, and I. Shames, "A robust control approach to asymptotic optimality of the heavy ball method for optimization of quadratic functions," *Automatica*, vol. 155, 2023, Art. no. 111129.
- [55] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. New York, NY, USA: Wiley-Interscience, 1972.
- [56] K. Ogata, *Discrete-Time Control Systems*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1994.
- [57] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Tradeoffs between convergence rate and noise amplification for momentum-based accelerated optimization algorithms," 2022, *arXiv:2209.11920*.
- [58] A. Fam and J. Meditch, "A canonical parameter space for linear systems design," *IEEE Trans. Autom. Control*, vol. 23, no. 3, pp. 454–458, Jun. 1978.
- [59] V. D. Blondel, M. Gurbuzbalaban, A. Megretski, and M. L. Overton, "Explicit solutions for root optimization of a polynomial family with one affine constraint," *IEEE Trans. Autom. Control*, vol. 57, no. 12, pp. 3078–3089, Dec. 2012.
- [60] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Variance amplification of accelerated first-order algorithms for strongly convex quadratic optimization problems," in *Proc. IEEE 57th Conf. Decis. Control*, Miami, FL, USA, 2018, pp. 5753–5758.
- [61] A. S. Nemirovsky and D. B. Yudin, *Problem Complexity and Method Efficiency in Optimization*. Wiley-Interscience, 1983.
- [62] R. Padmanabhan and P. Seiler, "Analysis of gradient descent with varying step sizes using integral quadratic constraints," *IEEE Trans. Automat. Control*, 2024, doi: [10.1109/TAC.2024.3438808](https://doi.org/10.1109/TAC.2024.3438808).
- [63] B. Hu, P. Seiler, and A. Rantzer, "A unified analysis of stochastic optimization methods using jump system theory and quadratic constraints," in *Proc. Conf. Learn. Theory*, 2017, pp. 1157–1189.
- [64] B. Hu, P. Seiler, and L. Lessard, "Analysis of biased stochastic gradient descent using sequential semidefinite programs," *Math. Program.*, vol. 187, no. 1, pp. 383–408, 2021.



Hesameddin Mohammadi (Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Southern California in Los Angeles, Los Angeles, CA, USA, in 2022.

His Ph.D. research includes robustness of gradient-based optimization algorithms and their performance and sample complexity for data-driven optimal control. He is currently working on the verification and validation of autonomous vehicle software with NVIDIA, Santa Clara, CA, USA.



Meisam Razaviyayn (Member, IEEE) received the Ph.D. degree in electrical engineering with minor in computer science from the University of Minnesota, Minneapolis, MN, USA, in 2014.

He is the Andrew and Erna Viterbi Early Career Chair and the Associate Professor in industrial and systems engineering, electrical engineering, and computer science with the University of Southern California (USC), Los Angeles, CA, USA. He is also the Associate Director and a Founding Member of the USC-Meta Center for

Research and Education in AI and Learning. Prior to joining the USC, he was a Postdoctoral Research Fellow with the Department of Electrical Engineering, Stanford University.



Mihailo R. Jovanović (Fellow, IEEE) received the Ph.D. degree in mechanical engineering from the University of California at Santa Barbara, Santa Barbara, CA, USA, in 2004.

He was a faculty member with the Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, from 2004 until 2017, and has held visiting positions with Stanford University, the Institute for Mathematics and its Applications, the Simons Institute for the Theory of Computing, and the University of

Belgrade. He is currently a Professor with the Ming Hsieh Department of Electrical and Computer Engineering and the Founding Director of the Center for Systems and Control, University of Southern California, Los Angeles, CA.

Prof. Jovanović serves as an Associate Editor for the *IEEE Control Systems Magazine*. He is a Fellow of the American Physical Society and was the recipient of a CAREER Award from the National Science Foundation in 2007, the George S. Axelby Outstanding Paper Award from the IEEE Control Systems Society in 2013, and the Distinguished Alumni Award from the Department of Mechanical Engineering, UC Santa Barbara in 2014.