2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Transient growth of accelerated first-order methods

Samantha Samuelson, Hesameddin Mohammadi, and Mihailo R. Jovanović

Abstract— **We examine transient responses of accelerated first-order optimization algorithms. By focusing on strongly convex quadratic problems, we identify the presence of modes whose algebraic growth induces large transient departure from the optimal solution. Leveraging the tools from linear systems theory, we explicitly quantify the transient growth caused by these resonant interactions. Our results demonstrate that both the time at which the transient response peaks and the largest value of the Euclidean distance between the optimization variable and the global minimizer are proportional to the square root of the condition number.**

Index Terms— **Acceleration, convex optimization, first-order optimization algorithms, gradient descent, Nesterov's accelerated algorithm, Polyak's heavy-ball method, transient growth.**

I. Introduction

First-order optimization algorithms are widely used in applications that arise in statistics, signal and image processing, control, and machine learning [1]–[6]. Accelerated first-order algorithms achieve faster convergence rate while preserving low per-iteration complexity of gradient descent. This has motivated the analysis of convergence under different stepsize selection rules [7]–[10]. There is also a growing body of literature that studies robustness of these algorithms for different types of uncertainties [11]–[17]. These studies demonstrate that acceleration increases sensitivity to uncertainty in gradient evaluation.

In addition to poor robustness in the face of uncertainty, even when the exact gradient is available, accelerated algorithms may exhibit undesirable transient behavior. This is in contrast to gradient descent which is a contraction mapping for strongly convex problems with suitable stepsize [18]. The existing results on convergence and robustness do not capture this transient growth and fail to provide insight into performance in non-asymptotic regimes with limited time budgets. For example, first-order algorithms are often used as a building block in multi-stage optimization including ADMM [19] and distributed optimization methods [20]. In these scenarios, variables are updated by performing only a few iterations at each stage. This motivates an in-depth study of the non-asymptotic behavior of first-order methods.

It is widely recognized that large transients of linear systems can arise from the presence of resonant modal interaction and non-normality of dynamical generators [21]. Even in the absence of exponentially growing normal modes these can induce large transient responses, significantly amplify background disturbances, and trigger departure from nominal operating conditions. For example, in fluid dynamics, such mechanisms may initiate departure from laminar flows and provide routes for transition to turbulence [22], [23].

In this paper, we examine non-asymptotic behavior of accelerated methods for strongly convex quadratic problems. In this case, the accelerated algorithms are linear dynamical systems and transient responses can be explicitly quantified. We characterize the transient behavior of two most commonly used accelerated algorithms: Polyak's heavy-ball and Nesterov's methods. We derive analytical expressions for the state-transition matrices and quantify the transient response in terms of convergence rate and iteration number. We also derive tight analytical bounds on magnitude and iteration number of the peak of the transient response in terms of the condition number κ and show that these quantities grow with the square root of κ. Finally, we demonstrate how a Lyapunov-based approach can be used to establish an upper bound on the transient response.

Recent reference [24] studied the transient growth of second-order systems and introduced a framework for establishing upper bounds, with a focus on real eigenvalues. The result was applied to the heavy-ball method but was not applicable to general quadratic problems in which the dynamical generator may have complex eigenvalues. We account for complex eigenvalues and generalize to Nesterov's accelerated algorithm. Furthermore, we provide tight upper and lower bounds on transient responses in terms of the condition number and identify the structure of the initial condition that induces largest transient growth.

The paper is structured as follows. In Section II, we provide background on the accelerated algorithms and discuss strongly convex quadratic problems. In Section III, we examine the transient growth of accelerated methods and provide analytical expressions and explicit bounds. In Section IV, we utilize a Lyapunov-based approach to upper bound the transient peak. In Section V, we offer concluding remarks and discuss future directions.

II. Motivation and background

The unconstrained optimization problem

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

where $f\colon \mathbb{R}^n \to \mathbb{R}$ is a convex function with an L-Lipschitz continuous gradient ∇f, can be solved using variety of techniques including the gradient descent method

$$x^{t+1} \;=\; x^t \;-\; \alpha^t \nabla f(x^t) \tag{2a}$$

Financial support from the National Science Foundation under Awards ECCS-1708906 and ECCS-1809833 is gratefully acknowledged.

The authors are with the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA 90089. E-mails: ({sasamuel, hesamedm, mihailo}@usc.edu).

978-1-5386-8266-1/$31.00 ©2020 AACC

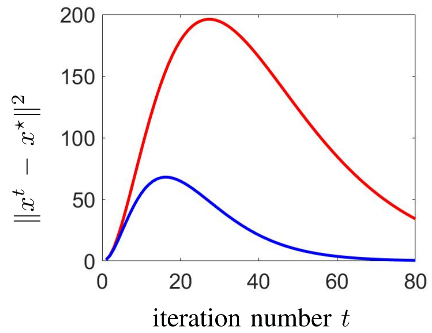

Fig. 1: Error in the optimization variable for Polyak's heavy-ball (black) and Nesterov's (red) methods with optimal parameters provided in Table I applied to a strongly convex quadratic problem with condition number $\kappa = 10^3$.

and its accelerated variants, namely Polyak's heavy-ball method

$$x^{t+2} \;=\; x^{t+1} \;+\; \beta^t(x^{t+1} \,-\, x^t) \;-\; \alpha^t \nabla f(x^{t+1}) \qquad \text{(2b)}$$

and Nesterov's accelerated algorithm

$$\begin{aligned} x^{t+2} \;=\; & x^{t+1} \;+\; \beta^t(x^{t+1} \,-\, x^t) \;- \\ & \alpha^t \nabla f\big(x^{t+1} \,+\, \beta^t(x^{t+1} \,-\, x^t)\big) \,. \end{aligned} \qquad \text{(2c)}$$

Here, t is the iteration index, $\alpha^t > 0$ is the stepsize, and $\beta^t \in (0, 1]$ is the acceleration parameter.

In the absence of strong convexity, Nesterov's accelerated method enjoys an optimal convergence rate among first-order algorithms. While the gradient descent with the stepsize $1/L$ provides $O(1/t)$ decay rate in the objective value,

$$f(x^t) \,-\, f(x^\star) \;\leq\; \frac{L}{2t}\,\|x^0 \,-\, x^\star\|^2$$

Nesterov's accelerated algorithm with $\alpha^t = 1/L$ and $\beta^t = (t-1)/(t+2)$ yields [9],

$$f(x^t) \,-\, f(x^\star) \;\leq\; \frac{4L}{(t\,+\,2)^2}\,\|x^0 \,-\, x^\star\|^2$$

where $x^\star$ is a global minimizer of f. Furthermore, when the objective function f is not only L-smooth but also m-strongly convex, i.e., when $f(x) - \frac{m}{2}||x||^2$ is convex, these algorithms can achieve a linear convergence rate $\rho < 1$,

$$\|x^t \,-\, x^\star\| \;\leq\; c\,\rho^t\,\|x^0 \,-\, x^\star\| \qquad (3)$$

where $c \geq 1$ is a constant. In this case, acceleration improves the rate of convergence from $O(1 - 1/\kappa)$ for gradient descent to $O(1 - 1/\sqrt{\kappa})$ for Nesterov's algorithm [25], where $\kappa := L/m$ is the condition number associated with the function f. In spite of significant improvement in the rate of convergence, acceleration may deteriorate performance on finite time intervals and lead to large transient responses. In particular, the constant c in (3) may become significantly larger than 1 for accelerated algorithms whereas $c = 1$ for gradient descent because of its contractive property for strongly convex problems. Figure 1 shows the transient growth of the error in the optimization variable for accelerated algorithms (2b) and (2c). A strongly convex quadratic problem with $\kappa = 10^3$ is considered and the parameters α and β in Table I that optimize the linear convergence rate are used.

Method	Parameter choice	Rate
Gradient	$\alpha = \frac{2}{L+m}$	$\frac{\kappa-1}{\kappa+1}$
Nesterov	$\alpha = \frac{4}{3L+m}$, $\beta = \frac{\sqrt{3\kappa+1}-2}{\sqrt{3\kappa+1}+2}$	$\frac{\sqrt{3\kappa+1}-2}{\sqrt{3\kappa+1}}$
Polyak	$\alpha = \frac{4}{(\sqrt{L}+\sqrt{m})^2}$, $\beta = \frac{(\sqrt{\kappa}-1)^2}{(\sqrt{\kappa}+1)^2}$	$\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}$

TABLE I: Optimal parameters and the linear convergence rate bounds for m-strongly convex quadratic objective functions with L-Lipschitz gradients and $\kappa := L/m$ [26].

In this paper, we examine the transient responses of accelerated algorithms for strongly convex quadratic optimization. For this class of problems, the function f in (1) is given by

$$f(x) \;=\; \frac{1}{2}\,x^T Q x \;+\; q^T x$$

where $Q = Q^T \succ 0$ is a positive definite matrix and $q \in \mathbb{R}^n$ is a vector. In this case, the condition number is determined by the ratio of the largest and smallest eigenvalues of the matrix Q, i.e., $\kappa = \lambda_{\max}(Q)/\lambda_{\min}(Q)$ and the constant values of parameters α and β provided in Table I yield the fastest rate of convergence for all three algorithms. Furthermore, since $\nabla f(x) = Qx + q$, algorithms (2) are linear time-invariant (LTI) systems and the dynamics of the error $x^t - x^\star$ are governed by the state-space model

$$\psi^{t+1} \;=\; A\,\psi^t \qquad \text{(4a)}$$

$$z^t \;=\; C\,\psi^t \qquad \text{(4b)}$$

where ψ^t is the state, $z^t = x^t - x^\star$ is the performance output, and A and C are constant matrices of appropriate dimensions. For gradient descent, we have

$$A \;=\; I \,-\, \alpha\,Q, \quad C \;=\; I, \quad \psi^t \;:=\; x^t \,-\, x^\star$$

and for the accelerated methods, we can let

$$\begin{aligned} \text{Polyak}: \quad A \;&=\; \left[\begin{array}{cc} 0 & I \\ -\beta I & (1+\beta)I - \alpha Q \end{array}\right] \\ \text{Nesterov}: \quad A \;&=\; \left[\begin{array}{cc} 0 & I \\ -\beta(I - \alpha Q) & (1+\beta)(I - \alpha Q) \end{array}\right]. \\ \psi^t \;&:=\; \left[\begin{array}{c} x^t \,-\, x^\star \\ x^{t+1} \,-\, x^\star \end{array}\right], \quad C \;=\; \left[\begin{array}{cc} I & 0 \end{array}\right] \end{aligned}$$

where I is the identity matrix.

For LTI system (4), the rate of convergence is determined by the spectral radius $\rho(A) := \max_i |\mu_i(A)|$, where μ_i are the eigenvalues of the matrix A. For α and β in Table I and $Q \succ 0$, $\rho(A) < 1$ and system (4) is stable in all three cases. While the convergence rate is a commonly used metric for evaluating performance of optimization algorithms, this quantity only determines the asymptotic behavior and it does not provide useful insight into transient responses. We demonstrate detrimental impact of acceleration on transient responses by quantifying performance of accelerated algorithms on finite-time intervals.

III. Transient responses of accelerated methods

The transient response is an important measure of performance of optimization algorithms, for the minimizer must be computed after a finite number of iterations. The response of LTI system (4) with an initial state ψ^0 is determined by

$$z^t \;=\; CA^t\psi^0 \;=\; \Phi(t)\psi^0.$$

Here, A^t is the state-transition matrix of system (4), i.e., the tth power of the matrix A, and $\Phi(t) := CA^t$ is the mapping from the initial condition ψ^0 to the performance output $z^t := x^t - x^\star$. At a fixed iteration t, we are interested in quantifying the transient growth, i.e., the worst case ratio of the energy of z^t to the energy of the initial condition ψ^0,

$$\sup_{\psi^0 \neq 0} \frac{\|z^t\|_2}{\|\psi^0\|_2} \;=\; \sup_{\|\psi^0\|_2 = 1} \|\Phi(t)\psi^0\|_2 \;=\; \sigma_{\max}(\Phi(t))$$

where $\sigma_{\max}(\cdot)$ is the largest singular value. Furthermore, the singular value decomposition of $\Phi(t)$ yields

$$z^t \;=\; U(t)\,\Sigma(t)\,V^T(t)\psi^0 \;=\; \sum_{j=0}^{r} \sigma_j(t)u_j(t)v_j^T(t)\psi^0.$$

The principal right singular vector $v_1(t)$ of the matrix $\Phi(t)$ determines the unit norm initial condition that yields the largest response at time t. The resulting response z^t is in the direction of the principal left singular vector $u_1(t)$ and the corresponding gain is determined by $\sigma_1(t) = \sigma_{\max}(\Phi(t))$. In what follows, we derive analytical bounds on the largest singular value of the matrix $\Phi(t)$.

A. Dimensionality reduction

We use the eigenvalue decomposition of $Q = V\Lambda V^T \succ 0$, where $\Lambda = \operatorname{diag}(\lambda_i)$ is a diagonal matrix of the eigenvalues

$$\lambda_1 \;\le\; \lambda_2 \;\le\; \cdots \;\le\; \lambda_n \tag{5}$$

and V is a unitary matrix of the eigenvectors of Q to bring A and C in (4) into a block diagonal form,

$$\hat{A} \,=\, \operatorname{diag}(\hat{A}_i),\; \hat{C} \,=\, \operatorname{diag}(\hat{C}_i),\; i \,=\, 1,\ldots,n.$$

In particular, the unitary coordinate transformation

$$\hat{x}^t \;:=\; V^T(x^t - x^\star) \tag{6}$$

brings the state-space model for the gradient descent into a diagonal form with

$$\hat{\psi}_i^t \,=\, \hat{x}_i^t,\; \hat{A}_i \,=\, 1 \,-\, \alpha\lambda_i,\; \hat{C}_i \,=\, 1.$$

Similarly, for accelerated algorithms, change of coordinates (6) in conjunction with a permutation of variables yield

$$\begin{aligned}
\text{Polyak}: \quad \hat{A}_i \;&=\; \begin{bmatrix} 0 & 1 \\ -\beta & (1+\beta) - \alpha\lambda_i \end{bmatrix} \\
\text{Nesterov}: \quad \hat{A}_i \;&=\; \begin{bmatrix} 0 & 1 \\ -\beta(1-\alpha\lambda_i) & (1+\beta)(1-\alpha\lambda_i) \end{bmatrix} \\
\hat{\psi}_i^t \;&=\; \begin{bmatrix} \hat{x}_i^t \\ \hat{x}_i^{t+1} \end{bmatrix},\; \hat{C}_i \;=\; \begin{bmatrix} 1 & 0 \end{bmatrix}.
\end{aligned} \tag{7}$$

Since both V and the permutation matrix are unitary, they preserve the 2-norm. It is thus easy to verify that

$$\sigma_{\max}(\Phi(t)) \;=\; \sigma_{\max}(\hat{C}\hat{A}^t) \;=\; \max_i \, \sigma_{\max}(\hat{\Phi}_i(t)) \tag{8}$$

where $\hat{\Phi}_i(t) := \hat{C}_i\hat{A}_i^t$ and the second equality follows from the block-diagonal structure of the matrices $\hat{A}$ and $\hat{C}$.

Each block $\hat{A}_i$ can be analyzed separately, reducing the dimensionality of the problem. It is worth noting that the matrices $\hat{A}_i$ in (7) are non-normal for accelerated methods. As we illustrate next, this non-normality can cause transient growth before eventual decay even in the absence of resonant modes.

B. An example

Let us consider a second-order system

$$\begin{bmatrix} \psi_1^{t+1} \\ \psi_2^{t+1} \end{bmatrix} \;=\; \begin{bmatrix} \mu_1 & \gamma \\ 0 & \mu_2 \end{bmatrix} \begin{bmatrix} \psi_1^t \\ \psi_2^t \end{bmatrix}$$

where μ_1, μ_2, and γ are the real numbers. For $\mu_1 \neq \mu_2$, the solution is given by

$$\begin{aligned}
\psi_1^t \;&=\; \mu_1^t\,\psi_1^0 \;+\; \frac{\gamma}{\mu_2 - \mu_1}\,(\mu_2^t - \mu_1^t)\,\psi_2^0 \\
\psi_2^t \;&=\; \mu_2^t\,\psi_2^0
\end{aligned}$$

and both states asymptotically decay to zero at a linear rate if $|\mu_i| < 1$. While ψ_2^t monotonically decays to zero, ψ_1^t experiences transient growth because of the interaction of two geometrically decaying modes μ_1 and μ_2. This growth is caused by the non-normality of the dynamical generator, i.e., the coupling from the second state to the first state and it increases linearly with the increase in parameter γ. We note that it occurs even in the absence of resonant interactions (i.e., $\mu_1 = \mu_2$) or near resonances (i.e., $\mu_1 \approx \mu_2$). For repeated eigenvalues, i.e., $\mu = \mu_1 = \mu_2$, we have

$$\begin{aligned}
\psi_1^t \;&=\; \mu^t\,\psi_1^0 \;+\; \gamma\,t\,\mu^{t-1}\,\psi_2^0 \\
\psi_2^t \;&=\; \mu^t\,\psi_2^0
\end{aligned}$$

and the transient growth arises from the increasing term t that initially dominates the decaying mode μ^{t-1}.

C. State-transition matrices

As discussed in the previous section, the state-space representation of accelerated algorithms can be brought into a block-diagonal form of n decoupled second-order systems,

$$\begin{aligned}
\hat{\psi}_i^{t+1} \;&=\; \hat{A}_i\,\hat{\psi}_i^t \\
\hat{z}_i^t \;&=\; \hat{C}_i\,\hat{\psi}_i^t.
\end{aligned}$$

The matrices $\hat{A}_i$ and $\hat{C}_i$ are given by

$$\hat{A}_i \;=\; \begin{bmatrix} 0 & 1 \\ a_i & b_i \end{bmatrix}, \; \hat{C}_i \;=\; \begin{bmatrix} 1 & 0 \end{bmatrix}$$

where ($a_i = -\beta$; $b_i = (1+\beta) - \alpha\lambda_i$) for the heavy-ball method and ($a_i = -\beta(1-\alpha\lambda_i)$; $b_i = (1+\beta)(1-\alpha\lambda_i)$) for Nesterov's method. The characteristic polynomial of $\hat{A}_i$ is

$$\det(sI - \hat{A}_i) \;=\; s^2 \,-\, b_i s \,-\, a_i \;=\; (s-\mu_{1i})(s-\mu_{2i})$$

and $\hat{A}_i$ can be represented in terms of its eigenvalues as

$$\hat{A}_i \;=\; \begin{bmatrix} 0 & 1 \\ -\mu_{1i}\mu_{2i} & \mu_{1i}+\mu_{2i} \end{bmatrix}.$$

Next lemma allows us to determine explicit expression for $\hat{A}_i^t$. This result can be obtained using eigenvalue/Jordan decomposition.

Lemma 1: Let μ_1 and μ_2 be the eigenvalues of the matrix

$$M \;=\; \begin{bmatrix} 0 & 1 \\ -\mu_1\mu_2 & \mu_1+\mu_2 \end{bmatrix}. \tag{9}$$

For any positive integer t, the matrix M^t is determined by

$$M^t \;=\; \frac{1}{\mu_2-\mu_1}\begin{bmatrix} \mu_1\mu_2(\mu_1^{t-1}-\mu_2^{t-1}) & \mu_2^t-\mu_1^t \\ \mu_1\mu_2(\mu_1^t-\mu_2^t) & \mu_2^{t+1}-\mu_1^{t+1} \end{bmatrix}$$

which for $\mu := \mu_1 = \mu_2$ simplifies to

$$M^t \;=\; \begin{bmatrix} (1-t)\,\mu^t & t\,\mu^{t-1} \\ -t\,\mu^{t+1} & (t+1)\,\mu^t \end{bmatrix}.$$

The explicit expressions in Lemma 1 allow us to obtain a tight upper bound on the norm of the first row of matrix M given by (9) in terms of the spectral radius of M.

Lemma 2: The matrix M in (9) satisfies

$$\|\begin{bmatrix} 1 & 0 \end{bmatrix} M^t\|_2^2 \;\le\; (t-1)^2\rho^{2t} \,+\, t^2\rho^{2t-2} \tag{10}$$

where ρ is the spectral radius of M. Moreover, (10) becomes equality if M has repeated eigenvalues.

Proof: Using Lemma 1, we can represent the first row of the matrix M^t in terms of the eigenvalues μ_1, μ_2 of M,

$$\begin{bmatrix} 1 & 0 \end{bmatrix} M^t \;=\; \begin{bmatrix} -\sum\limits_{i=0}^{t-2}\mu_1^{i+1}\mu_2^{t-1-i} & \sum\limits_{i=0}^{t-1}\mu_1^i\mu_2^{t-1-i} \end{bmatrix}.$$

Now, we can use the triangle inequality to write

$$|\sum_{i=0}^{t-2}\mu_1^{i+1}\mu_2^{t-1-i}| \le \sum_{i=0}^{t-2}|\mu_1^{i+1}\mu_2^{t-1-i}| \le \sum_{i=0}^{t-2}\rho^t \le (t-1)\rho^t$$

$$|\sum_{i=0}^{t-1}\mu_1^{i}\mu_2^{t-1-i}| \le \sum_{i=0}^{t-1}|\mu_1^{i}\mu_2^{t-1-i}| \le \sum_{i=0}^{t-1}\rho^{t-1} \le t\rho^{t-1}.$$

For repeated eigenvalues, $\rho = |\mu| = |\mu_1| = |\mu_2|$ and these inequalities become equalities. ■

D. Transient growth analysis

We next use Lemma 2 to establish analytical expression for the largest singular value of the matrix $\Phi(t) = CA^t$ associated with Polyak's and Nesterov's accelerated algorithms.

Theorem 1: For accelerated algorithms, the largest singular value of the matrix $\Phi(t) := CA^t$ satisfies

$$\sigma_{\max}^2(\Phi(t)) \;\le\; (t-1)^2\rho^{2t} \,+\, t^2\rho^{2(t-1)} \tag{11}$$

where ρ is the corresponding rate of convergence which only depends on the condition number κ. Moreover, (11) becomes equality for the optimal parameters provided in Table I.

Proof: Let μ_{1i} and μ_{2i} be the eigenvalues of the matrix $\hat{A}_i$ with the spectral radius $\hat{\rho}_i = \max\{|\mu_{1i}|, |\mu_{2i}|\}$. We can use Lemma 2 with $M := \hat{A}_i$ for each i to obtain

$$\begin{aligned} \sigma_{\max}^2(\Phi(t)) \;&=\; \max_i \; \sigma_{\max}^2(\begin{bmatrix} 1 & 0 \end{bmatrix}\hat{A}_i^t) \\ &\le\; \max_i \; (t-1)^2\hat{\rho}_i^{2t} \,+\, t^2\hat{\rho}_i^{2t-2} \\ &\le\; (t-1)^2\rho^{2t} \,+\, t^2\rho^{2t-2} \end{aligned} \tag{12}$$

where ρ is the spectral radius of the matrix A. Here, the equality follows from (8), the first inequality follows from Lemma 2, and the third inequality is a consequence of

$$\rho \;=\; \max_i \; \hat{\rho}_i.$$

For the parameters in Table I, it can be shown that $\hat{A}_1$ and $\hat{A}_n$, that correspond to the smallest and largest eigenvalues of Q, i.e., $\lambda_1 = m$ and $\lambda_n = L$, respectively, have the largest spectral radius [15, Eq. (64)],

$$\rho \;=\; \hat{\rho}_1 \;=\; \hat{\rho}_n \;\ge\; \hat{\rho}_i, \quad i \,=\, 2,\ldots,n-1. \tag{13}$$

Furthermore, it is straightforward to verify that the matrix $\hat{A}_1$ has repeated eigenvalues. Thus, we can write

$$\begin{aligned} \sigma_{\max}^2(\Phi(t)) \;&=\; \max_i \; \sigma_{\max}^2(\begin{bmatrix} 1 & 0 \end{bmatrix}\hat{A}_i^t) \\ &\ge\; \sigma_{\max}^2(\begin{bmatrix} 1 & 0 \end{bmatrix}\hat{A}_1^t) \;=\; (t-1)^2\hat{\rho}_1^{2t} \,+\, t^2\hat{\rho}_1^{2t-2} \\ &=\; (t-1)^2\rho^{2t} \,+\, t^2\rho^{2t-2}. \end{aligned} \tag{14}$$

Here, the second inequality follows from Lemma 2 applied to $M := \hat{A}_1$ and the last inequality follows from (13). Finally, combining (12) and (14) completes the proof. ■

Theorem 1 highlights the source of disparity between the long and short term behavior of the state-transition matrix. In particular, while the geometric decay of ρ^t drives $\Phi(t)$ to 0 as $t \to \infty$, early stages are dominated by the algebraic term which induces a transient growth. We next use this result to provide tight bounds on the time $t_{\max}$ at which the largest transient response takes place and the corresponding peak value $\sigma_{\max}(\Phi(t_{\max}))$. Even though the explicit expressions for these two quantities can be derived, our tight upper and lower bounds are more informative and easier to interpret.

Theorem 2: For accelerated algorithms with the parameters provided in Table I and $\rho \in [1/\mathrm{e}, 1)$, the time at which the largest transient response takes place $t_{\max} := \mathrm{argmax}_t\, \sigma_{\max}(\Phi(t))$ and the corresponding peak value

$\sigma_{\max}(\Phi(t_{\max})) := \max_t \sigma_{\max}(\Phi(t))$ satisfy

$$-\frac{1}{\log \rho} \;\le\; t_{\max} \;\le\; -\frac{1}{\log \rho} + 1$$

$$-\frac{\sqrt{2}\rho}{\mathrm{e}\log \rho} \;\le\; \sigma_{\max}(\Phi(t_{\max})) \;\le\; -\frac{\sqrt{2}}{\mathrm{e}\,\rho\log \rho}.$$

Proof: Let $a(t) := t\rho^t$. Then, Theorem 1 implies

$$\sigma^2_{\max}(\Phi(t)) \;=\; \rho^2 a^2(t-1) \;+\; \rho^{-2} a^2(t). \qquad (15)$$

For $\rho \in [1/\mathrm{e}, 1)$, we have $\max_t \; a(t) = -1/(\mathrm{e}\log\rho)$ and the maximum takes place at $t_{\max} = -1/\log\rho \ge 1$. Moreover, the terms on the right-hand side of (15) satisfy

$$\rho^2 a^2(t-1) \;\le\; \rho^{-2} a^2(t)$$

for all $t \ge 1$, which implies

$$2\rho^2 a^2(t-1) \;\le\; \sigma^2_{\max}(\Phi(t)) \;\le\; 2\rho^{-2} a^2(t). \qquad (16)$$

Substituting for $a(t)$ and $a(t-1)$ in (16) using their maximum values $-1/(\mathrm{e}\log\rho)$ establishes the bounds on $\max_t \sigma_{\max}(\Phi(t))$. Regarding the time at which maximum occurs, it is straightforward to verify that $\mathrm{d}\sigma_{\max}(\Phi(t))/\mathrm{d}t$ is positive at $t = -1/\log\rho$ and negative at $t = 1 - 1/\log\rho$. Moreover, it can be shown that the function $\sigma_{\max}(\Phi(t))$ has only one critical point for $t \ge 1$ and that this critical point is a maximizer. Based on these observations, we conclude that the maximizer must lie between $-1/\log\rho$ and $1 - 1/\log\rho$. This completes the proof. ∎

For accelerated algorithms with parameters provided in Table I, Theorem 2 can be used to determine the location of the transient peak in terms of condition number κ. Since for Polyak's heavy-ball method, $\rho = 1 - 2/(\sqrt{\kappa}+1)$ and for Nesterov's accelerated algorithm, $\rho = 1 - 2/(\sqrt{3\kappa+1})$, the Mercator series $\log(1+x) = \sum_{i=1}^{\infty} \frac{(-1)^{i-1}}{i} x^i, \quad |x| < 1$ in conjunction with Theorem 2 can be used to establish that both $t_{\max}$ and $\sigma_{\max}(\Phi(t_{\max}))$ scale as $\sqrt{\kappa}$ for $\kappa \gg 1$.

Proposition 1: For accelerated algorithms with the parameters provided in Table I and the condition number $\kappa \gg 1$, the time at which the largest transient response takes place $t_{\max} := \operatorname{argmax}_t \sigma_{\max}(\Phi(t))$ and the peak value $\sigma_{\max}(\Phi(t_{\max})) := \max_t \sigma_{\max}(\Phi(t))$ satisfy

(i) Polyak's heavy-ball method

$$(\sqrt{\kappa}+1)/2 \;\le\; t_{\max} \;\le\; (\sqrt{\kappa}+3)/2$$

$$\frac{\sqrt{\kappa}-1}{\sqrt{2}\,\mathrm{e}} \;\le\; \sigma_{\max}(\Phi(t_{\max})) \;\le\; \frac{\sqrt{\kappa}+1}{\sqrt{2}\,\mathrm{e}}$$

(ii) Nesterov's accelerated method

$$\frac{\sqrt{3\kappa+1}}{2} \;\le\; t_{\max} \;\le\; \frac{\sqrt{3\kappa+1}+2}{2}$$

$$\frac{\sqrt{3\kappa+1}-2}{\sqrt{2}\,\mathrm{e}} \;\le\; \sigma_{\max}(\Phi(t_{\max})) \;\le\; \frac{\sqrt{3\kappa+1}}{\sqrt{2}\,\mathrm{e}}.$$

Remark 1: The initial condition that leads to the largest transient growth on a fixed time interval is determined by

$$\hat{\psi}_i^0 \;=\; \frac{1}{\sigma_{\max}(\Phi(t))} \begin{bmatrix} (1-t)\,\rho^t \\ t\,\rho^{t-1} \end{bmatrix}.$$

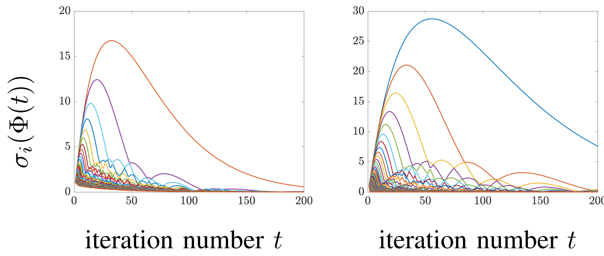

(a) Polyak's method (b) Nesterov's method

Fig. 2: Singular values of the matrix $\Phi(t) := CA^t$ for Polyak's and Nesterov's methods for a strongly convex quadratic problem with a Toeplitz matrix $Q \in \mathbb{R}^{100\times 100}$.

For accelerated algorithms with the parameters in Table I and $\kappa \gg 1$, the initial condition that induces the largest response is approximately given by

$$\hat{\psi}_i^0 \;\approx\; \frac{1}{\sigma_{\max}(\Phi(t_{\max}))} \begin{bmatrix} -t_{\max}\,\rho^{t_{\max}} \\ t_{\max}\,\rho^{t_{\max}} \end{bmatrix}$$

where bounds on $t_{\max}$ and $\sigma_{\max}(\Phi(t_{\max}))$ are established in Proposition 1. Thus, for large condition numbers, the worst case initial condition for the second part of the state points in the opposite direction relative to the first state component.

An example: For a strongly convex quadratic problem in which $Q \in \mathbb{R}^{100\times 100}$ is the Toeplitz matrix with the first row determined by $[2 \;\; -1 \;\; 0 \;\; \cdots \;\; 0]$, Fig. 2 shows the singular values of the matrix $\Phi(t) = CA^t$ for accelerated algorithms with parameters provided in Table I. In addition to the principal right singular vector which generates the largest transient response, many other singular vectors experience significant growth. This demonstrates that apart from resonant interactions that arise from repeated eigenvalues of the underlying dynamical generators, many other initial conditions can grow significantly on finite time intervals.

IV. Lyapunov approach for bounding $\sigma_{\max}(A^t)$

Quadratic Lyapunov functions can be used as an effective means to analyze the largest singular value of the state-transition matrix. If a positive definite matrix P satisfies

$$APA^T \;-\; P \;\preceq\; 0 \qquad (17)$$

then the following bound on the largest singular value of A^t

$$\sigma^2_{\max}(A^t) \;\le\; \lambda_{\max}(P)/\lambda_{\min}(P)$$

holds for all $t \ge 1$. To see this, let us define the LTI system $\psi^{t+1} = A^T \psi^t$ with an initial condition ψ^0 and consider the Lyapunov function candidate $V(\psi) \;=\; \psi^T P\,\psi$. It is easy to verify that under (17), the sequence $V(\psi^t)$ is non-increasing,

$$V(\psi^t) \;\le\; V(\psi^{t-1}) \;\le\; \cdots \;\le\; V(\psi^0). \qquad (18)$$

Moreover, since $P \succ 0$, for all ψ we have

$$\lambda_{\min}(P)\,\|\psi\|_2^2 \;\le\; V(\psi) \;\le\; \lambda_{\max}(P)\,\|\psi\|_2^2$$

which in conjunction with (18) yield

$$\|\psi^t\|_2^2 \;\leq\; \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)}\,\|\psi^0\|_2^2.$$

Thus, the state-transition matrix satisfies

$$\sigma_{\max}^2(A^t) \;=\; \sigma_{\max}^2((A^T)^t) \;=\; \sup_{\psi^0 \neq 0}\; \frac{\|\psi^t\|_2^2}{\|\psi^0\|_2^2} \;\leq\; \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)}.$$

Next, we specialize this result to the accelerated methods.

A. Application to accelerated methods

Herein, we employ the Lyapunov-based method to obtain bounds on the largest singular value of the state-transition matrix of the accelerated algorithms. The next lemma provides a solution to the Lyapunov inequality (17) associated with the matrices $\hat{A}_i$.

Lemma 3: Consider the matrices

$$M \;=\; \begin{bmatrix} 0 & 1 \\ a & b \end{bmatrix}, \quad P \;=\; \begin{bmatrix} 1 & b/(1-a) \\ b/(1-a) & 1 \end{bmatrix} \quad (19)$$

and let the spectral radius of M be strictly smaller than 1. Then, $P \succ 0$ and $MPM^T - P \preceq 0$.

The proof of Lemma 3 exploits the Schur complement. We omit it due to page limitations.

We can now use Lemma 3 to establish a bound on $\sigma_{\max}^2(M^t)$. In particular, the eigenvalues of P in (19) are given by $1 \pm b/(1-a)$. Thus, by Lemma 3 we have

$$\sigma_{\max}^2(M^t) \;\leq\; \frac{\lambda_{\max}(P)}{\lambda_{\min}(P)} \;=\; \max\left\{\gamma, \gamma^{-1}\right\}$$

where $\gamma \;:=\; (1-a+b)/(1-a-b)$ is the ratio between the eigenvalues of the matrix P.

The next proposition, builds on this result and derives bounds on the state-transition matrix of accelerated methods.

Proposition 2: For the parameters provided in Table I and $\kappa \gg 1$, the state-transition matrix is upper bounded by

$$\begin{aligned} \text{Polyak}: &\quad \sigma_{\max}(A^t) \leq \sqrt{\kappa} \\ \text{Nesterov}: &\quad \sigma_{\max}(A^t) \leq \sqrt{3\,\kappa - 1} \;-\; 1. \end{aligned}$$

While Proposition 2 provides only upper bounds on the norm of the state transition matrix, the lower bounds established in Proposition 1 imply that the upper bounds provided by this proposition are tight. In particular, this implies that, for both algorithms, the peak magnitude of $\sigma_{\max}^2(A^t)$ scales linearly with the condition number κ.

V. Concluding remarks

We have utilized a Lyapunov-based approach and spectral analysis to explicitly quantify and bound transient responses of accelerated methods for strongly convex quadratic problems. Our results demonstrate that the transient growth is caused by resonant interaction in the underlying dynamical generators, and that the largest transient peak scales as the square root of the condition number.

References

[1] L. Bottou and Y. Le Cun, "On-line learning for very large data sets," *Appl. Stoch. Models Bus. Ind.*, vol. 21, no. 2, pp. 137–151, 2005.

[2] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.

[3] Y. Nesterov, "Gradient methods for minimizing composite objective functions," *Math. Program.*, vol. 140, no. 1, pp. 125–161, 2013.

[4] 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, 2016.

[5] L. Bottou, F. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *SIAM Rev.*, vol. 60, no. 2, pp. 223–311, 2018.

[6] S. Hassan-Moghaddam and M. R. Jovanović, "Topology design for stochastically-forced consensus networks," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 3, pp. 1075–1086, September 2018.

[7] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proc. ICML*, 2013, pp. 1139–1147.

[8] 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.

[9] Y. Nesterov, "A method for solving the convex programming problem with convergence rate $O(1/k^2)$," in *Dokl. Akad. Nauk SSSR*, vol. 27, no. 2, 1983, pp. 543–547.

[10] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2013, vol. 87.

[11] O. Devolder, F. Glineur, and Y. Nesterov, "First-order methods of smooth convex optimization with inexact oracle," *Math. Program.*, vol. 146, no. 1-2, pp. 37–75, 2014.

[12] B. Hu and L. Lessard, "Dissipativity theory for Nesterov's accelerated method," in *Proc. ICML*, vol. 70, 2017, pp. 1549–1557.

[13] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Variance amplification of accelerated first-order algorithms for strongly convex quadratic optimization problems," in *Proceedings of the 57th IEEE Conference on Decision and Control*, Miami, FL, 2018, pp. 5753–5758.

[14] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Performance of noisy Nesterov's accelerated method for strongly convex optimization problems," in *Proceedings of the 2019 American Control Conference*, Philadelphia, PA, 2019, pp. 3426–3431.

[15] H. Mohammadi, M. Razaviyayn, and M. R. Jovanović, "Robustness of accelerated first-order algorithms for strongly convex optimization problems," *IEEE Trans. Automat. Control*, 2019, submitted; also arXiv:1905.11011.

[16] S. Michalowsky, C. Scherer, and C. Ebenbauer, "Robust and structure exploiting optimization algorithms: An integral quadratic constraint approach," 2019, arXiv:1905.00279.

[17] J. I. Poveda and N. Li, "Robust hybrid zero-order optimization algorithms with acceleration via averaging in continuous time," 2019, arXiv:1909.00265.

[18] D. P. Bertsekas, "Convex optimization algorithms," 2015.

[19] Y. Ouyang, Y. Chen, G. Lan, and E. Pasiliao, "An accelerated linearized alternating direction method of multipliers," *SIAM J. IMAGING. SCI.*, vol. 8, no. 1, pp. 644–681, 2015.

[20] W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin, "On the linear convergence of the admm in decentralized consensus optimization," *IEEE Trans. Signal Process.*, vol. 62, no. 7, pp. 1750–1761, 2014.

[21] L. N. Trefethen and M. Embree, *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton: Princeton University Press, 2005.

[22] M. R. Jovanović and B. Bamieh, "Componentwise energy amplification in channel flows," *J. Fluid Mech.*, vol. 534, pp. 145–183, July 2005.

[23] P. J. Schmid, "Nonmodal stability theory," *Annu. Rev. Fluid Mech.*, vol. 39, pp. 129–162, 2007.

[24] 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.

[25] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Kluwer Academic Publishers, 2004, vol. 87.

[26] L. Lessard, B. Recht, and A. Packard, "Analysis and design of optimization algorithms via integral quadratic constraints," *SIAM J. Opt.*, vol. 26, no. 1, pp. 57–95, 2016.