

Control Interpretations for First-Order Optimization Methods

Bin Hu¹

Laurent Lessard^{1,2}

American Control Conference, pp. 3114–3119, 2017

Abstract

First-order iterative optimization methods play a fundamental role in large scale optimization and machine learning. This paper presents control interpretations for such optimization methods. First, we give loop-shaping interpretations for several existing optimization methods and show that they are composed of basic control elements such as PID and lag compensators. Next, we apply the small gain theorem to draw a connection between the convergence rate analysis of optimization methods and the input-output gain computations of certain complementary sensitivity functions. These connections suggest that standard classical control synthesis tools may be brought to bear on the design of optimization algorithms.

1 Introduction

First-order iterative optimization methods have been widely applied in data science and machine learning [3, 16]. These methods only require access to first-order derivative information, and iterate on the data until satisfactory convergence is achieved. For example, the gradient method is

$$x^{k+1} = x^k - \alpha \nabla f(x^k). \quad (1)$$

Such simple methods are often favored over higher order methods such as Newton’s method when the dimension of the underlying space is large and computing Hessians is prohibitively expensive.

There has been significant recent interest in finding ways to accelerate the convergence of the gradient method while maintaining low iteration costs. For example, the *Heavy-ball* method includes an additional momentum term

$$x^{k+1} = x^k - \alpha \nabla f(x^k) + \beta(x^k - x^{k-1}). \quad (2)$$

This slight modification can yield a dramatic improvement in worst-case convergence rate if f is quadratic. A similar acceleration scheme, *Nesterov’s accelerated method*, can improve the convergence rate for strongly

convex f with smooth gradients [10, 13]. These convergence results are derived on a case-by-case basis, and the intuition behind the acceleration is still not fully understood.

Recent efforts have adopted a dynamical system (or differential equation) perspective in analyzing acceleration for convex objectives [17, 23], though a more general understanding of acceleration is still lacking (non-convex objectives, inexact computations, etc). This paper aims to bring new insights on how to accelerate first-order optimization methods for objective functions which are not convex in general. Our main contributions are as follows:

1. We pose the iterative optimization paradigm as an output regulation problem, which lends itself to a loop-shaping interpretation. In particular, we show that several popular optimization algorithms may be viewed as controllers which are composed of basic PID or lag compensation elements. We also demonstrate that existing parameter tuning guidelines for these optimization methods are consistent with the loop-shaping design guidelines in control theory.
2. Using the small gain theorem [24], we draw a connection between the convergence rate analysis of optimization methods (under sector-bounded assumptions) and the input-output gain computation for a particular complimentary sensitivity function. It follows that the *design* of optimization algorithms for sector-bounded functions can be interpreted as \mathcal{H}_∞ state feedback synthesis. This explains why acceleration typically requires stronger function assumptions (not necessarily convexity) beyond just sector-bounded gradients.

A related line of research has emerged in the distributed optimization literature [6, 8, 21, 22]. In [21, 22], a continuous-time differential equation was used to describe the dynamics of distributed optimization, leading to a natural iterative algorithm which may be interpreted as a PI controller. In [8], event-triggered control methods were tailored for distributed optimization over networks. In contrast with the work on distributed optimization, the present work is concerned with control-theoretic properties and interpretations of a big class of first-order optimization methods.

A second related line of research is the unified integral quadratic constraint framework in [10], which provides

¹B. Hu and L. Lessard are with the Optimization Group at the Wisconsin Institute for Discovery.

²L. Lessard is also with the Department of Electrical and Computer Engineering, University of Wisconsin–Madison, Madison, WI 53706, USA. Email: {bhu38, laurent.lessard}@wisc.edu

a numerical tool based on semidefinite programming for use in analyzing optimization algorithms. In contrast with this work, the present work uses a small gain approach with a simple interpretation that interfaces with existing results on complementary sensitivity integrals.

The paper is organized as follows. Section 2 explains notation and problem formulation, Section 3 describes our loop-shaping interpretation for first-order methods, and Section 4 presents our main results involving the small gain theorem and connections to complementary sensitivity.

2 Preliminaries

2.1 Spaces and operators

Let ℓ_{2e}^p denote the sequences $x := (x^0, x^1, \dots) \subseteq \mathbb{R}^p$, and let $\ell_2^p \subseteq \ell_{2e}^p$ be the set of square-summable sequences, so if $x \in \ell_2^p$, then $\sum_{k=0}^{\infty} \|x^k\|^2 < \infty$ where $\|x^k\|^2 := (x^k)^\top x^k$ denotes the standard Euclidean norm. We will omit the superscript p when it is implied by context. The gain of a causal operator $K : \ell_{2e} \rightarrow \ell_{2e}$ is defined as

$$\|K\| := \sup_{x \in \ell_2; x \neq 0} \frac{\|Kx\|}{\|x\|} \quad (3)$$

In addition, K is said to be bounded if it has a finite gain. Notice this gain is induced by ℓ_2 signals while the operator K itself is defined on ℓ_{2e} . This definition makes sense since any bounded operator on ℓ_2 to itself has a natural causal extension to the operator from ℓ_{2e} to ℓ_{2e} . Clearly, every bounded operator must map zero inputs to zero outputs.

2.2 Various classes of objective functions in optimization

Consider the unconstrained optimization problem

$$\min_{x \in \mathbb{R}^p} f(x) \quad (4)$$

where it is assumed that there exists a unique $x^* \in \mathbb{R}^p$ satisfying $\nabla f(x^*) = 0$. How to solve (4) and find x^* heavily depends on the assumptions about f . A simple assumption is that f is quadratic. Two other common assumptions are L -smoothness and strong convexity. A continuously differentiable function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is L -smooth if the following inequality holds for all $x, y \in \mathbb{R}^p$

$$\|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|. \quad (5)$$

We define $\mathcal{L}(L)$ to be the set of L -smooth functions. The continuously differentiable function f is m -strongly convex if the following inequality holds for all $x, y \in \mathbb{R}^p$

$$f(x) \geq f(y) + \nabla f(y)^\top (x - y) + \frac{m}{2}\|x - y\|^2. \quad (6)$$

Note that we recover ordinary convexity in (6) if $m = 0$. We define $\mathcal{F}(m, L)$ to be the set of functions that are both L -smooth and m -strongly convex. The class \mathcal{F} covers a large family of objective functions in machine learning,

including ℓ_2 -regularized logistic regression [20], smooth support vector machines [9], etc. Clearly, $\mathcal{F}(m, L) \subseteq \mathcal{L}(L)$. For all $f \in \mathcal{F}(m, L)$, the following inequality holds for all $x \in \mathbb{R}^p$

$$\begin{bmatrix} x - x^* \\ \nabla f(x) \end{bmatrix}^\top \begin{bmatrix} -2mLI_p & (L+m)I_p \\ (L+m)I_p & -2I_p \end{bmatrix} \begin{bmatrix} x - x^* \\ \nabla f(x) \end{bmatrix} \geq 0 \quad (7)$$

where I_p denotes the $p \times p$ identity matrix [10, Lemma 6]. On the other hand, a function satisfying the above inequality may not belong to $\mathcal{F}(m, L)$, and may not even be convex. The set of continuously differentiable functions satisfying (7) is denoted as $\mathcal{S}(m, L)$. This class of functions has sector-bounded gradients, and includes $\mathcal{F}(m, L)$ as a subset.

2.3 Review of first-order optimization methods

A classical way to solve (4) is the gradient descent method, which uses the iteration (1) to gradually converge to x^* . The intuition behind gradient descent method is as follows. At each step k , we find a quadratic approximation of f about x^k , which hopefully captures the local structure of f , and we solve the quadratic minimization problem

$$\min_{x \in \mathbb{R}^p} \left(f(x^k) + \nabla f(x^k)^\top (x - x^k) + \frac{1}{2\alpha} \|x^k - x\|^2 \right). \quad (8)$$

When $f \in \mathcal{S}(m, L)$, if α is chosen well, then there exists a constant $\rho \in (0, 1)$ and a constant $c \geq 1$ such that

$$\|x^k - x^*\| \leq c\rho^k \|x^0 - x^*\| \quad (9)$$

Thus the iterates $\{x^k\}$ converge exponentially to x^* . By convention, this is known as *linear convergence* in the optimization literature. For example, we can choose $\alpha = \frac{2}{L+m}$ and obtain $\rho = \frac{L-m}{L+m}$ and $c = 1$. Another popular choice is $\alpha = \frac{1}{L}$, which leads to $\rho = 1 - \frac{m}{L}$ and $c = 1$. These results are formally documented in [10, Section 4.4]. It is emphasized that the proofs of these results only require (7).

When $f \in \mathcal{F}(m, L)$, one can achieve a better convergence rate $\rho = \sqrt{1 - \frac{m}{L}}$ using Nesterov's accelerated method:

$$\begin{aligned} x^{k+1} &= y^k - \alpha \nabla f(y^k) \\ y^k &= (1 + \beta)x^k - \beta x^{k-1} \end{aligned} \quad (10)$$

where $\alpha = \frac{1}{L}$ and $\beta = \frac{\sqrt{L-m}}{\sqrt{L+m}}$. When L/m is large, Nesterov's accelerated method guarantees a much faster convergence rate compared to the gradient descent method. This fact was stated in [13, Theorem 2.2.3].

When f is a quadratic function, one can accelerate the gradient descent method by incorporating a momentum term into the iteration, such as the Heavy-ball method (2). Although the Heavy-ball method works extremely well for quadratic objective functions, it can fail to converge for other functions in $\mathcal{F}(m, L)$; see [10, Section 4.6].

The intuitions behind Nesterov’s accelerated method and the Heavy-ball method are still not fully understood. Hence, there is no intuitive way to modify these methods to accelerate the convergence when optimizing more general functions, i.e. $f \in \mathcal{S}(m, L)$ or $f \in \mathcal{S}(m, L) \cap \mathcal{L}(L)$. Gaining intuition for these methods can be beneficial for designing accelerated schemes for more general classes of objective functions.

Finally, it is worth mentioning that every optimization method mentioned in this section can be cast as the feedback interconnection $F_u(P, K)$ as shown in Fig 1. Here, $P := \nabla f$ is a static nonlinearity and K is a linear time-invariant (LTI) system (the algorithm).




Figure 1: Feedback representation for the optimization method $F_u(\nabla f, K)$.

Feedback representations for optimization methods are discussed in [10, Section 2]. In [10], P is a static nonlinear operator that maps u to $v = Pu$ as $v^k = \nabla f(u^k)$. However, this is not a bounded operator since it does not map zero inputs to zero outputs. For the convenience of our discussion, we will choose P to be the operator which maps u to $v = Pu$ as $v^k = \nabla f(u^k + x^*)$ where x^* is the unique point satisfying $\nabla f(x^*) = 0$. Then this choice of P leads to a bounded operator. One can perform a state shifting argument to the feedback representations in [10] and cast all the mentioned optimization methods as

$$\begin{aligned} \xi^{k+1} &= A\xi^k + Bv^k \\ u^k &= C\xi^k \\ v^k &= \nabla f(u^k + x^*) \end{aligned} \quad (11)$$

where (A, B, C) are the state matrices of K . For example, to rewrite the gradient descent method (1), one can set $\xi^k = u^k = x^k - x^*$ and $v^k = \nabla f(x^k) = \nabla f(u^k + x^*)$. Then (1) can be cast as (11) with $(A, B, C) = (I_p, -\alpha I_p, I_p)$. Notice here $\xi^k = x^k - x^*$ and hence the convergence rate of the optimization algorithm is equivalent to the rate at which ξ^k goes to 0. The optimization method converges to the optimum x^* at a rate ρ if and only if the model (11) drives ξ^k to 0 from any initial conditions at the same rate ρ . Using similar arguments, Nesterov’s accelerated method and the Heavy-ball method can be written as (11). In these two cases, the associated state matrices for K are the same as (2.5) and (2.7) in [10], although the states have been shifted by x^* . When $f \in \mathcal{S}(m, L)$, the inequality (7) imposes a sector bound on the input/output pair of P . Let $v = Pu$. Then the

following inequality holds for all k

$$\begin{bmatrix} u^k \\ v^k \end{bmatrix}^\top \begin{bmatrix} -2mLI_p & (L+m)I_p \\ (L+m)I_p & -2I_p \end{bmatrix} \begin{bmatrix} u^k \\ v^k \end{bmatrix} \geq 0. \quad (12)$$

The above inequality is important for further analysis of optimization methods.

2.4 Input-output stability and small gain theorem

The key analysis tool in this paper is the small gain theorem, which is now briefly reviewed. Suppose two causal operators $P : \ell_{2e} \rightarrow \ell_{2e}$ and $K : \ell_{2e} \rightarrow \ell_{2e}$ both map zero input to zero output. Let $[P, K]$ denote the feedback interconnection of P and K illustrated in Fig. 2:

$$\begin{cases} v = Pu + e \\ u = Kv + r. \end{cases} \quad (13)$$




Figure 2: Feedback interconnection with exogenous inputs

The interconnection $[P, K]$ is said to be *well-posed* if the map $(u, v) \mapsto (r, e)$ defined by (13) has a causal inverse on ℓ_{2e} . It is (input-output) *stable* if it is well-posed and this inverse causal map from (r, e) to (u, v) is bounded. Clearly, $u, v \in \ell_2$ for all $r, e \in \ell_2$ if $[P, \Delta]$ is stable. Well-posedness holds only if the solutions to (13) have no finite escape time. The small gain theorem states the following [7, 24].

Theorem 1 (small gain theorem). *Suppose P and K are bounded causal operators and $[P, K]$ is well-posed. If $\|P\|\|K\| < 1$, then $[P, K]$ is input-output stable.*

The small gain theorem can be used to check the input-output stability of $[P, K]$ when the gains of P and K are both known. Note that there are exogenous signals r, e in the setup of $[P, K]$ and zero initial conditions on K to ensure that K maps the zero input to a zero output. In contrast, $F_u(P, K)$ allows any initial condition for K , so the optimization method $F_u(P, K)$ can be initialized at any initial condition $\xi^0 \in \mathbb{R}^n$. For any optimization method $F_u(P, K)$ described by (11), one can form an associated interconnection $[P, K]$ by adding the signals (r, e) and fixing the initial condition of K to be zero, since the nonlinear static map P is set up in a way to map zero inputs to zero outputs. An important connection between the internal stability of $F_u(P, K)$ and the input-output stability of $[P, K]$ has been stated in [2, Proposition 5]. Consequently, one may apply the small gain theorem for the convergence rate analysis of optimization methods.

3 Loop-shaping interpretations for optimization methods

This section presents basic control interpretations for the gradient descent method, Nesterov’s accelerated method, and the Heavy-ball method with the hope of shedding light on the general principles underlying the design of first-order methods. The goal of the optimization method is to find x^* satisfying $\nabla f(x^*) = 0$. Hence the optimization method may be viewed as a controller that regulates the plant “ ∇f ” to zero. When viewing ∇f as the plant one wants to control, the unconstrained optimization problem (4) is an output regulation problem, and the LTI part K in the first-order optimization method can be viewed as a controller. The key issue for this output regulation problem is that the equilibrium point x^* is unknown.

Transfer functions for the controller K are listed in Table 1. We use the symbol \otimes to denote the Kronecker product. These products appear because the controllers corresponding to our algorithms of interest are repetitions of a single-input-single-output (SISO) system.

Table 1: Transfer function $K(z)$ for first-order methods

Optimization Method	Controller K
Gradient Descent	$I_p \otimes \frac{-\alpha}{z-1}$
Heavy-Ball	$I_p \otimes \frac{-\alpha z}{z^2 - (1+\beta)z + \beta}$
Nesterov’s Method	$I_p \otimes \frac{-\alpha(1+\beta)z + \alpha\beta}{z^2 - (1+\beta)z + \beta}$

For the gradient descent method, K is a pure integrator. Hence, the gradient descent regulates the nonlinear plant P via pure integral control. Integral action is necessary since the algorithm must converge to x^* , which amounts to having zero steady-state error when K tracks a step input.

The Heavy-ball method (2) differs from gradient descent in the inclusion of an additional momentum term. This momentum term may be viewed as a lag compensator. The Heavy-ball method corresponds to the following controller:

$$K = I_p \otimes \left(\frac{-\alpha}{z-1} \right) \left(\frac{z}{z-\beta} \right) \quad (14)$$

The first term provides integral action to ensure zero steady-state error as with the gradient method, while the second term is a discrete-time lag compensator. The lag compensation has the net effect of 1) boosting low-frequency response by a factor of roughly $\frac{1}{1-\beta}$, which improves the tracking speed of the controller and hence the convergence of the algorithm and 2) attenuating high-frequency response by a factor of roughly $\frac{1}{1+\beta}$. It intuitively makes sense that the Heavy-ball method can accelerate convergence for quadratic objectives, since the

plant P becomes a linear operator in this case. However, the lag compensator increases the slope of the loop gain near the crossover frequency, which may have a detrimental effect on the robustness of the closed loop. This qualitative observation is confirmed by the fact that the Heavy-ball method may fail to converge at all if the objective function is relaxed to include more general strongly convex functions [10].

Unlike the Heavy-ball method, Nesterov’s accelerated method performs well when applied to strongly-convex objective functions. A control interpretation is that Nesterov’s accelerated method includes derivative control to decrease the slope of the loop gain near the crossover frequency, which significantly improves the robustness of the algorithm for certain classes of nonlinearities. To see the derivative controller in Nesterov’s accelerated method, rewrite (10) as

$$y^{k+1} = y^k + \beta(y^k - y^{k-1}) - \alpha \nabla f(y^k) - \alpha\beta(\nabla f(y^k) - \nabla f(y^{k-1})) \quad (15)$$

The last term is a difference of the plant output ∇f , and can be viewed as a derivative control.

Nesterov’s method may also be interpreted as lag compensation together with integral action, as in the Heavy-ball case. The corresponding controller is

$$K = I_p \otimes \left(\frac{-\alpha}{z-1} \right) \left(\frac{(1+\beta)z - \beta}{z - \beta} \right), \quad (16)$$

which has a zero at $z = \frac{\beta}{1+\beta}$, and this helps increase the slope of the Bode plot near the crossover frequency. The control interpretations for different optimization methods are summarized in Table 2.

Table 2: Control interpretations for first-order methods

Optimization Methods	Control Structure
Gradient Descent	Integral Control
Heavy-Ball	Lag + Integral Control
Nesterov’s Method	Lag + PID Control

We now demonstrate that the design of state-of-the-art optimization methods is actually consistent with general loop-shaping principles from control theory. The loop-shaping principle states that the low-frequency loop gain should be sufficiently large to ensure good tracking performance while the high-frequency loop gain should be small enough for the purpose of noise rejection. In addition, the slope of the loop gain near the crossover frequency should be flat (typically around -20 dB/decade) to assure a proper phase margin and good robustness. A thorough discussion on loop-shaping can be found in standard references [4, 14, 15].

Given a function $f \in \mathcal{F}(m, L)$, the standard gradient descent stepsize is $\alpha = \frac{1}{L}$, and the standard parameter choice for Nesterov’s accelerated method is $\alpha = \frac{1}{L}$

and $\beta = \frac{\sqrt{L}-\sqrt{m}}{\sqrt{L}+\sqrt{m}}$. Other parameter choices, when f is quadratic for example, are documented in [10, Proposition 1]. Fig. 3 shows the Bode plots of the resultant controllers K (only the SISO part) for all these parameter choices under the assumption that $m = 0.01$ and $L = 1$.

The Bode plots are consistent with the properties of these optimization methods when a loop-shaping intuition is adopted. First, the gradient descent method with the standard stepsize $\alpha = \frac{1}{L}$ is known to be slower than other first-order methods when f is quadratic. This is reflected in the Bode plot, which shows that the gradient method has a relatively low gain particularly in the low-frequency region. Using the optimal tuning of $\alpha = \frac{2}{L+m}$ improves the gain slightly.

Second, the optimal quadratic tuning for all three methods leads to controllers whose crossover frequencies are roughly at 0.5 Hz. Intuitively, such tuning places excessive weight on tracking performance and is very fragile to noise at the output of the plant P . This is consistent with the known robustness properties of these methods as well. For example, the gradient method with $\alpha = \frac{1}{L}$ is known to be very robust to the noise in the gradient computation while the gradient method with $\alpha = \frac{2}{m+L}$ is known to be fragile to such noise [10, Section 5.2]. Comparing the high frequency responses of these two cases immediately leads to the same conclusion. Finally, the slope of Bode plot at the crossover frequency supports the fact that Nesterov's method works for a larger class of functions than the Heavy-ball method.

In summary, the intuition brought by the traditional loop-shaping theory is consistent with the known properties of the existing first-order methods. This suggests that loop-shaping intuition may be used as a general high-level guideline in the design of optimization methods. The control interpretations above also indicate that classical PID tuning [1] can be used for optimization algorithm design. For example, one could drop the lag compensation and simply use the PID controller:

$$K(z) = I_p \otimes \frac{-\alpha(1+\beta)z + \alpha\beta}{z(z-1)} \quad (17)$$

It remains an open question as to how to choose an appropriate $K(z)$ subject to different assumptions on the objective function. Since acceleration schemes for the optimization of quadratic functions or functions in $\mathcal{F}(m, L)$ already exist, we will focus on the case $f \in \mathcal{S}(m, L)$. We will derive one connection between such an optimization design problem and classical control synthesis theory.

4 Analysis and design of optimization methods using the small gain theorem

In this section, it is assumed that $f \in \mathcal{S}(m, L)$ and x^* is the unique point satisfying $\nabla f(x^*) = 0$.




Figure 3: Bode plots of K for various first-order methods with commonly-used parameter tunings.

4.1 New feedback representations for first-order methods

The feedback representation (11) for first-order methods involves a nonlinear operator P which belongs to the sector (m, L) (12). This does not coincide perfectly with a gain bound since upper and lower bounds don't match. We will therefore, use a loop-shifted $F_u(P', K')$ that ensures the small gain condition on P' captures the full sector. Choose P' to map u to $v = Pu$ as $v^k = u^k - \frac{2}{m+L} \nabla f(u^k + x^*)$. Then, substitute $\xi^k = u^k = x^k - x^*$ into the gradient descent method (1) to get an alternative feedback interconnection

$$\begin{aligned} \xi^{k+1} &= \left(1 - \frac{(m+L)\alpha}{2}\right) \xi^k + \frac{(m+L)\alpha}{2} v^k \\ u^k &= \xi^k \\ v^k &= u^k - \frac{2}{L+m} \nabla f(u^k + x^*) \end{aligned} \quad (18)$$

Direct manipulation of (7) shows that P' is in a sector $(-\frac{L-m}{L+m}, \frac{L-m}{L+m})$, which leads to gain bound $\|P'\| \leq \frac{L-m}{L+m}$.

Suppose $K = I_p \otimes \bar{K}$ where \bar{K} is a SISO LTI system. In general, a loop transformation argument can be used to show that any optimization method $F_u(P, K)$ can also be represented as $F_u(P', K')$ where $K' = I_p \otimes \bar{K}'$ and $\bar{K}' = \bar{K} / (\bar{K} - \frac{2}{m+L})$. Consequently, the feedback interconnection (18) provides another way to model first-order methods.

4.2 Main theorem

Our main approach is inspired by the loop transformation used in [2, 5]. For any $\rho \in (0, 1)$, the operators ρ^+ and ρ^- are defined as the time-domain, time-dependent multipliers ρ^k , ρ^{-k} , respectively. Here, superscripts indicate k -th power. Define $K'_\rho := \rho^- \circ K' \circ \rho^+$, and $P'_\rho := \rho^- \circ P' \circ \rho^+$. From [2, Section 3], one can conclude $F_u(P', K')$ converges at rate ρ if $[P'_\rho, K'_\rho]$ is input-output stable and $K'_\rho(z) = K'(\rho z)$. Similarly, define $\bar{K}'_\rho := \rho^- \circ \bar{K}' \circ \rho^+$, and one has $\bar{K}'_\rho(z) = \bar{K}'(\rho z)$. The main result of this paper is stated below.

Theorem 2. *Let $\bar{K}' = \bar{K}/(\bar{K} - \frac{2}{m+L})$, and $K' = I_p \otimes \bar{K}'$. If $\|\bar{K}'_\rho\| < \frac{L+m}{L-m}$, then the optimization method $F_u(P', K')$ has a linear convergence rate ρ .*

Proof. Notice P' is a pointwise nonlinearity, and hence one can use the small gain condition on P' to show $\|P'_\rho\| \leq \frac{L-m}{L+m}$ (see Section 5.1 in [2] or Section IV.C in [5] for detailed arguments). In addition, $\|K'_\rho\| = \|I_p \otimes \bar{K}'_\rho\| = \|\bar{K}'_\rho\| < \frac{L+m}{L-m}$, and $\|K'_\rho\| \|P'_\rho\| < 1$. By the small gain theorem, $[P'_\rho, K'_\rho]$ is input-output stable. By [2, Proposition 5], $F_u(P', K')$ converges at rate ρ . ■

The power of Theorem 2 is that it connects the convergence rate analysis of the optimization method to an input-output gain computation on a SISO system \bar{K}'_ρ . Note that the ℓ_2 -induced norm (3) of a stable LTI system is equal to its \mathcal{H}_∞ -norm [25]. In addition, we have $\bar{K}'_\rho(z) = \bar{K}'(\rho z)$. Hence we only need to verify that $\bar{K}'(\rho z)$ is stable and then compare the \mathcal{H}_∞ -norm of $\bar{K}'(\rho z)$ to $\frac{L+m}{L-m}$.

4.3 Recovery of rate results for gradient descent

As a sanity check, we apply Theorem 2 to recover convergence rate results for the gradient descent method applied to functions in $\mathcal{S}(m, L)$. Since $\bar{K}(z) = \frac{\alpha}{z-1}$, we have

$$\bar{K}'(z) = \frac{\alpha(m+L)}{2z-2+\alpha(m+L)} \quad (19)$$

If $\alpha = \frac{2}{m+L}$, it is straightforward to obtain

$$\bar{K}'(z) = \frac{1}{z}, \quad \bar{K}'(\rho z) = \frac{1}{\rho z} \quad (20)$$

Clearly, $\bar{K}'(\rho z)$ is stable for any $\rho > 0$. Moreover, $\|\bar{K}'(\rho z)\| = \rho^{-1}$. By Theorem 2, the gradient descent method converges for any $\rho > \frac{L-m}{L+m}$. This recovers the existing rate result for the gradient method with $\alpha = \frac{2}{L+m}$. Another popular choice for $\alpha = \frac{1}{L}$. In this case, the shifted controller is given by

$$\bar{K}'(z) = \frac{1+\kappa}{2z-\kappa+1} \quad (21)$$

where $\kappa := \frac{L}{m}$ is the condition number. Hence, one has

$$\bar{K}'(\rho z) = \frac{1+\kappa}{2\rho\kappa z - \kappa + 1} \quad (22)$$

When $\rho > \frac{1}{2}(1 - \frac{1}{\kappa})$, $\bar{K}'(\rho z)$ is stable. In addition, one can substitute $z = 1$ into (22) to obtain the peak frequency response (the \mathcal{H}_∞ norm) of $\bar{K}'(\rho z)$. To ensure this norm is smaller than $\frac{L+m}{L-m}$, one has the condition

$$\frac{1+\kappa}{2\rho\kappa - \kappa + 1} < \frac{\kappa+1}{\kappa-1}. \quad (23)$$

Upon simplifying (23) together with $\rho > \frac{1}{2}(1 - \frac{1}{\kappa})$, we finally obtain $\rho > 1 - \frac{1}{\kappa}$, which is the linear convergence rate for the gradient descent method when $\alpha = \frac{1}{L}$.

Remark 1. *A small technical issue in the above analysis is that the rate result proved by the small gain theorem is a strict inequality. This is due to the fact that for LTI systems, input-output stability is slightly stronger than the global uniform stability. See [5, Remark 1] for a detailed explanation. This issue is negligible from a practical standpoint.*

4.4 Connections to complementary sensitivity

Since we have $\bar{K}'(\rho z) = \bar{K}(\rho z)/(\bar{K}(\rho z) - \frac{2}{m+L})$, $\bar{K}'(\rho z)$ is the complementary sensitivity function of the closed-loop system $F_u(\bar{K}(\rho z), -\frac{m+L}{2})$. Accelerating optimization for $f \in \mathcal{S}(m, L)$ requires finding the smallest ρ such that there exists $\bar{K}(\rho z)$ that stabilizes $\bar{K}'(\rho z)$ while ensuring that $\|\bar{K}'(\rho z)\| < \frac{L+m}{L-m}$. One possible method would be to perform a bisection search on ρ . For each ρ , we can try to design \bar{K} to stabilize $\bar{K}'(\rho z)$ and minimize the \mathcal{H}_∞ -norm of $\bar{K}'(\rho z)$ at the same time. This subproblem may be reformulated as an \mathcal{H}_∞ state feedback synthesis problem since \bar{K} always contains a pure integrator and the state of the scaled dynamics $\frac{1}{\rho z-1}$ is accessible at every timestep. Consequently, the only design variable is the state feedback gain, which happens to be the stepsize of the gradient method. This explains why acceleration in this case is difficult even given memory of past iterates. It is worth noting that $\bar{K}(\rho z)$ always has an unstable pole at $z = \rho^{-1}$. There is a large body of discrete-time complementary sensitivity integral results [11, 12, 18, 19] which could potentially be used in studying the design limits of \bar{K} under the analytic constraints posed by the unstable pole at ρ^{-1} .

From the above connection, we can see that acceleration typically requires some function properties which can be decoded as constraints involving dynamics. The condition (7) is a static constraint, and it is hard to design accelerated schemes with this single constraint. However, it is still possible to accelerate non-convex optimization when other function properties are available, e.g. $f \in \mathcal{L}(L)$.

5 Conclusion

This paper discussed connections between the analysis of optimization algorithms and classical control-theoretic concepts. Specifically, the gradient method, the Heavy-ball method, and Nesterov's accelerated method were interpreted as combinations of PID and lag compensators.

A loop-shaping interpretation was also used to explain several well-known robustness properties of these algorithms.

We invoked the small gain theorem to show that finding worst-case convergence rates for algorithms amounts to computing the gain of a complementary sensitivity function. In addition, we demonstrated a connection between \mathcal{H}_∞ state feedback synthesis and stepsize selections of the gradient method. These observations are an encouraging first step toward leveraging tools from control theory for the analysis and eventual synthesis of robust optimization algorithms.

References

- [1] K. Ang, G. Chong, and Y. Li. PID control system analysis, design, and technology. *IEEE Transactions on Control Systems Technology*, 13(4):559–576, 2005.
- [2] R. Boczar, L. Lessard, and B. Recht. Exponential convergence bounds using integral quadratic constraints. In *IEEE Conference on Decision and Control*, pages 7516–7521, 2015.
- [3] S. Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- [4] G. Franklin, J. Powell, and A. Emami-Naeini. *Feedback control of dynamic systems*. Pearson Higher Education, Inc, 6th edition, 2010.
- [5] B. Hu and P. Seiler. Exponential decay rate conditions for uncertain linear systems using integral quadratic constraints. *IEEE Transactions on Automatic Control*, 61(11):3631–3637, 2016.
- [6] D. Jakovetić, J. Moura, and J. Xavier. Linear convergence rate of a class of distributed augmented Lagrangian algorithms. *IEEE Transactions on Automatic Control*, 60(4):922–936, 2015.
- [7] H.K. Khalil. *Nonlinear Systems*. Prentice Hall, third edition, 2001.
- [8] S. Kia, J. Cortés, and S. Martínez. Distributed convex optimization via continuous-time coordination algorithms with discrete-time communication. *Automatica*, 55:254–264, 2015.
- [9] Y. Lee and O. Mangasarian. SSVM: A smooth support vector machine for classification. *Computational Optimization and Applications*, 20(1):5–22, 2001.
- [10] L. Lessard, B. Recht, and A. Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- [11] R. Middleton. Trade-offs in linear control system design. *Automatica*, 27(2):281–292, 1991.
- [12] C. Mohtadi. Bode’s integral theorem for discrete-time systems. In *IEE Proceedings D-Control Theory and Applications*, volume 137, pages 57–66, 1990.
- [13] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer Academic Publishers, 2003.
- [14] N. Nise. *Control systems engineering*. John Wiley & Sons, 2007.
- [15] K. Ogata. *Modern control engineering*. Prentice-Hall, 5th edition, 2010.
- [16] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014.
- [17] W. Su, S. Boyd, and E. Candès. A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems 27*, pages 2510–2518, 2014.
- [18] H. Sung and S. Hara. Properties of sensitivity and complementary sensitivity functions in single-input single-output digital control systems. *International Journal of Control*, 48(6):2429–2439, 1988.
- [19] H. Sung and S. Hara. Properties of complementary sensitivity function in SISO digital control systems. *International Journal of Control*, 50(4):1283–1295, 1989.
- [20] C. Teo, A. Smola, S. Vishwanathan, and Q. Le. A scalable modular convex solver for regularized risk minimization. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 727–736, 2007.
- [21] J. Wang and N. Elia. Control approach to distributed optimization. In *Allerton Conference on Communication, Control, and Computing*, pages 557–561, 2010.
- [22] J. Wang and N. Elia. A control perspective for centralized and distributed convex optimization. In *IEEE Conference on Decision and Control*, pages 3800–3805, 2011.
- [23] A. Wibisono, A. Wilson, and M. I. Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.
- [24] G. Zames. On the input-output stability of time-varying nonlinear feedback systems part one: Conditions derived using concepts of loop gain, conicity, and positivity. *IEEE Transactions on Automatic Control*, 11(2):228–238, 1966.
- [25] K. Zhou, J.C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice-Hall, 1996.