



Heriot-Watt University
Research Gateway

On the connections between optimization algorithms, Lyapunov functions, and differential equations

Citation for published version:

Dobson, P, Sanz-Serna, JM & Zygalakis, K 2023 'On the connections between optimization algorithms, Lyapunov functions, and differential equations: theory and insights'.
<https://doi.org/10.48550/arXiv.2305.08658>

Digital Object Identifier (DOI):

[10.48550/arXiv.2305.08658](https://doi.org/10.48550/arXiv.2305.08658)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Early version, also known as pre-print

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

On the connections between optimization algorithms, Lyapunov functions, and differential equations: theory and insights

Paul Dobson
University of Edinburgh
pdobson@ed.ac.uk

Jesus Maria Sanz-Serna
Universidad Carlos III de Madrid
jmsanzserna@gmail.com

Konstantinos C. Zygalakis
University of Edinburgh
k.zygalakis@ed.ac.uk

Abstract

We study connections between differential equations and optimization algorithms for m -strongly and L -smooth convex functions through the use of Lyapunov functions by generalizing the Linear Matrix Inequality framework developed by Fazy-lab et al. in 2018. Using the new framework we derive analytically a new (discrete) Lyapunov function for a two-parameter family of Nesterov optimization methods and characterize their convergence rate. This allows us to prove a convergence rate that improves substantially on the previously proven rate of Nesterov's method for the standard choice of coefficients, as well as to characterize the choice of coefficients that yields the optimal rate. We obtain a new Lyapunov function for the Polyak ODE and revisit the connection between this ODE and the Nesterov's algorithms. In addition discuss a new interpretation of Nesterov method as an additive Runge-Kutta discretization and explain the structural conditions that discretizations of the Polyak equation should satisfy in order to lead to accelerated optimization algorithms.

1 Introduction

We are interested in solving the minimization problem

$$\min_{x \in \mathbb{R}^d} f(x), \quad (1)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is m -strongly convex and L -smooth (we denote the set of all such f as $\mathcal{F}_{m,L}$). Gradient descent is the simplest algorithm for obtaining the solution x^* of (1); it converges with a rate $1 - \mathcal{O}(1/\kappa)$, $\kappa = L/m$, for $f \in \mathcal{F}_{\mu,L}$ (Nesterov (2014)). This rate is unsatisfactory, because in many applications the condition number κ is $\gg 1$. It is possible to improve upon gradient descent by resorting to *accelerated* methods with rates $1 - \mathcal{O}(1/\sqrt{\kappa})$. An example of such a method is the well-known Nesterov's accelerated method Nesterov (1983) for $f \in \mathcal{F}_{m,L}$

$$x_{k+1} = y_k - \frac{1}{L} \nabla f(y_k), \quad (2a)$$

$$y_k = x_k + \frac{1 - \sqrt{1/\kappa}}{1 + \sqrt{1/\kappa}} (x_k - x_{k-1}). \quad (2b)$$

It is shown in (Nesterov, 2014, Theorem 2.2.3) that, if $y_0 = x_0$,

$$f(x_k) - f(x^*) \leq \left(1 - \sqrt{1/\kappa}\right)^k \left(f(x_0) - f(x^*) + \frac{m}{2} \|x_0 - x^*\|^2\right).$$

We can view the algorithm (2) as a discretization of the damped oscillator equation in Polyak (1964)

$$\ddot{x} + \bar{b}\sqrt{m}\dot{x} + \nabla f(x) = 0, \quad \bar{b} > 0. \quad (3)$$

A fixed point of (21) corresponds to the minimizer of f and moreover, the solutions $x(t)$ of (21) approach x^* as $t \rightarrow \infty$ if f is m -strongly convex (Wilson et al., 2021, Proposition 3). The connection between ordinary differential equations (ODE) and optimization algorithms has been a rewarding line of research. As was shown in Su et al. (2016); Scieur et al. (2017) finding an ODE which corresponds to the limit of an optimization algorithm can give insights about the algorithm. Conversely, there is a large number of research works (see e.g. Wibisono et al. (2016); Krichene et al. (2015)) which propose accelerated algorithms by suitable discretizations of second order dissipative ODEs both in Euclidean and non-Euclidean geometry. Furthermore, the links between such ODEs and Hamiltonian dynamics led to a number of research works that tried to construct or explain optimization algorithms using concepts such as shadowing Orvieto and Lucchi (2019), symplecticity Betancourt et al. (2018); Bravetti et al. (2019); Muehlebach and Jordan (2019, 2021); Shi et al. (2019), discrete gradients Ehrhardt et al. (2018), and backward error analysis Franca et al. (2021).

A common element of the analysis presented in many of the papers mentioned above is the construction of discrete Lyapunov functions used to identify the convergence rate of the underlying algorithm. Furthermore, Lessard et al. (2016) introduced a control theoretic view of optimization algorithms, that has been later connected with Lyapunov functions both in discrete and continuous time by Fazlyab et al. (2018) (see also Sanz-Serna and Zygalkakis (2021)).

In this work, we modify some of the conditions needed to obtain a Lyapunov function in the control theoretic framework presented in Fazlyab et al. (2018). This allows us to construct a new Lyapunov function for a two-parameter family of Nesterov optimization methods, which in turn is used to prove convergence rates that improve on those previously reported in the literature. It turns out that the coefficients in (2) are not optimal in terms of the convergence rate. We also derive a new Lyapunov function for the oscillator (21). The interpretation of Nesterov algorithms as discretizations of (21) is well known; however this discretization does not belong to the standard classes of integrators given by linear multistep or Runge-Kutta methods. We show that Nesterov algorithms are in fact instances of the class of additive Runge-Kutta integrators (Cooper and Sayfy (1980)). Finally we explain the structural conditions that discretizations of (21) have to satisfy in order to lead to accelerated optimization algorithms.

2 Linear Matrix Inequalities and Lyapunov functions

A useful viewpoint on optimization algorithms (Lessard et al. (2016)) is that they can often be represented as linear dynamical systems interacting with one or more static nonlinearities. We focus on the following state space representation

$$\xi_{k+1} = A\xi_k + Bu_k, \quad (4a)$$

$$u_k = \nabla f(y_k), \quad (4b)$$

$$y_k = C\xi_k, \quad (4c)$$

$$x_k = E\xi_k, \quad (4d)$$

where $\xi_k \in \mathbb{R}^n$ is the state, $u_k \in \mathbb{R}^d$ is the input ($d \leq n$), $y_k \in \mathbb{R}^d$ is the feedback output that is mapped to u_k by the nonlinear map ∇f . Here, A, B, C, E are constant matrices of appropriate sizes and x_k is the approximation to the minimizer x^* after k steps.

To study the convergence rate of optimization algorithms, Fazlyab et al. (2018) considers Lyapunov functions of the form

$$V_k(\xi) = \rho^{-2k} (a_0(f(x) - f(x^*)) + (\xi - \xi^*)^\top P(\xi - \xi^*)), \quad \rho \in (0, 1), \quad (5)$$

where $a_0 > 0$ and P is positive semi-definite (denoted by $P \succeq 0$). If along the trajectories of (18) $V_{k+1}(\xi_{k+1}) \leq V_k(\xi_k)$, we can conclude that $\rho^{-2k} a_0(f(x_k) - f(x^*)) \leq V_k(\xi_k) \leq V_0(\xi_0)$ so that we have the following convergence estimate:

$$f(x_k) - f(x^*) \leq \rho^{2k} \frac{V_0(\xi_0)}{a_0}.$$

Fazlyab et al. (2018) present a useful framework based on Linear Matrix Inequalities to check whether functions of the form (26) decay along trajectories, see, in particular, Theorem 3.2 there.

Our analysis here is based on the observation that, for $f \in \mathcal{F}_{\mu, L}$, we have that $f(x) - f(x_*) \geq (m/2)\|x - x_*\|^2$, a fact that, when constructing Lyapunov functions of the form (26), may be used

to modify the requirement $P \succeq 0$ demanded by Fazlyab et al. (2018) (see details in the supplementary material). More precisely, the following theorem is a useful alternative to Theorem 3.2 in Fazlyab et al. (2018) (where it is demanded that $P \succeq 0$). (The notations $\sigma(E^T E)$, $\sigma(\tilde{P})$ indicate eigenvalues.)

Theorem 1 *Suppose that, for (18), there exist $a_0 > 0$, $\rho \in (0, 1)$, $\ell > 0$, and a symmetric matrix P , with $\tilde{P} := P + (a_0 m/2)E^T E \succ 0$, such that*

$$T = M^{(0)} + a_0 \rho^2 M^{(1)} + a_0 (1 - \rho^2) M^{(2)} + \ell M^{(3)} \preceq 0, \quad (6)$$

where

$$M^{(0)} = \begin{bmatrix} A^T P A - \rho^2 P & A^T P B \\ B^T P A & B^T P B \end{bmatrix},$$

and

$$M^{(1)} = N^{(1)} + N^{(2)}, \quad M^{(2)} = N^{(1)} + N^{(3)}, \quad M^{(3)} = N^{(4)},$$

with

$$\begin{aligned} N^{(1)} &= \begin{bmatrix} EA - C & EB \\ 0 & I_d \end{bmatrix}^T \begin{bmatrix} \frac{\ell}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} EA - C & EB \\ 0 & I_d \end{bmatrix}, \\ N^{(2)} &= \begin{bmatrix} C - E & 0 \\ 0 & I_d \end{bmatrix}^T \begin{bmatrix} -\frac{m}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} C - E & 0 \\ 0 & I_d \end{bmatrix}, \\ N^{(3)} &= \begin{bmatrix} C^T & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{m}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & I_d \end{bmatrix}, \\ N^{(4)} &= \begin{bmatrix} C^T & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{mL}{m+L} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & -\frac{1}{m+L} I_d \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & I_d \end{bmatrix}. \end{aligned}$$

Then, for $f \in \mathcal{F}_{m,L}$, with V given by (26), the sequence $\{x_k\}$ satisfies

$$\|x_k - x_*\|^2 \leq \max \sigma(E^T E) \|\xi_k - \xi^*\|_{\tilde{P}} \leq \frac{\max \sigma(E^T E)}{\min \sigma(\tilde{P})} V(\xi_0, 0) \rho^{2k}. \quad (7)$$

2.1 Continuous-time systems

We will also consider continuous-time dynamical systems in state space form

$$\dot{\xi}(t) = \bar{A}\xi(t) + \bar{B}u(t), \quad y(t) = \bar{C}\xi(t), \quad u(t) = \nabla f(y(t)) \quad t \geq 0, \quad (8)$$

where $\xi(t) \in \mathbb{R}^n$ is the state, $y(t) \in \mathbb{R}^d$ ($d \leq n$) the output, and $u(t) = \nabla f(y(t))$ the continuous feedback input. Fixed points of (20) satisfy

$$0 = \bar{A}\xi^*, \quad y^* = \bar{C}\xi^*, \quad u^* = \nabla f(y^*);$$

in our context, $u^* = 0$ and $y^* = x^*$. Similarly, to the discrete case, Fazlyab et al. (2018) considers Lyapunov functions of the form

$$V(\xi, t) = e^{\lambda t} (f(y(t)) - f(y^*) + (\xi(t) - \xi^*)^T \bar{P} (\xi(t) - \xi^*)) \quad (9)$$

where \bar{P} is a positive definite matrix. Again, in a similar way as in the discrete case we can modify the positive-definiteness assumption to obtain the following Theorem.

Theorem 2 *Suppose that, for (20), there exist $\lambda > 0$, $\sigma \geq 0$ and a symmetric matrix \bar{P} with $\tilde{\bar{P}} := \bar{P} + (m/2)\bar{C}^T \bar{C} \succ 0$, that satisfy*

$$\bar{T} = \bar{M}^{(0)} + \bar{M}^{(1)} + \lambda \bar{M}^{(2)} + \sigma \bar{M}^{(3)} \preceq 0$$

where

$$\begin{aligned}\bar{M}^{(0)} &= \begin{bmatrix} \bar{P}\bar{A} + \bar{A}^T\bar{P} + \lambda\bar{P} & \bar{P}\bar{B} \\ \bar{B}^T\bar{P} & 0 \end{bmatrix}, \\ \bar{M}^{(1)} &= \frac{1}{2} \begin{bmatrix} 0 & (\bar{C}\bar{A})^T \\ \bar{C}\bar{A} & \bar{C}\bar{B} + \bar{B}^T\bar{C}^T \end{bmatrix}, \\ \bar{M}^{(2)} &= \begin{bmatrix} \bar{C}^T & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{m}{2}I_d & \frac{1}{2}I_d \\ \frac{1}{2}I_d & 0 \end{bmatrix} \begin{bmatrix} \bar{C} & 0 \\ 0 & I_d \end{bmatrix}, \\ \bar{M}^{(3)} &= \begin{bmatrix} \bar{C}^T & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{mL}{m+L}I_d & \frac{1}{2}I_d \\ \frac{1}{2}I_d & -\frac{1}{m+L}I_d \end{bmatrix} \begin{bmatrix} \bar{C} & 0 \\ 0 & I_d \end{bmatrix}.\end{aligned}$$

Then the following inequality holds for $f \in \mathcal{F}_{m,L}$, $t \geq 0$, with V given by (30)

$$\|y(t) - y_*\|^2 \leq \max \sigma(\bar{C}^T \bar{C}) \|\xi(t) - \xi^*\|_{\bar{P}} \leq \frac{\max \sigma(C^T C)}{\min \sigma(\tilde{P})} e^{-\lambda t} V(\xi(0), 0).$$

3 A Lyapunov function for a family of Nesterov's methods

We will now study optimization methods of the form

$$x_{k+1} = x_k + \beta(x_k - x_{k-1}) - \alpha \nabla f(y_k), \quad (10a)$$

$$y_k = x_k + \beta(x_k - x_{k-1}), \quad (10b)$$

$k = -1, 0, 1, \dots$, with parameters $\alpha, \beta > 0$; note that (2) arises from a particular choice of α, β . We write (10) in the state space form (18). In order to easily relate the material in this section to later developments, we first introduce as a new variable the divided difference, $k = 0, 1, \dots$,

$$d_k = \frac{1}{\delta}(x_k - x_{k-1}),$$

where the *steplength* $\delta = \sqrt{m\alpha}$ is *nondimensional*, in the sense that its numerical value does not change when f and x are scaled (β is also nondimensional). With the new variable, (10) becomes ($k = 0, 1, \dots$)

$$d_{k+1} = \beta d_k - \frac{\alpha}{\delta} \nabla f(y_k), \quad (11a)$$

$$x_{k+1} = x_k + \delta \beta d_k - \alpha \nabla f(y_k), \quad (11b)$$

$$y_k = x_k + \delta \beta d_k, \quad (11c)$$

and these equations are of the form (18) with $\xi_k = [d_k^T, x_k^T]^T \in \mathbb{R}^{2d}$ and

$$A = \begin{bmatrix} \beta I_d & 0 \\ \delta \beta I_d & I_d \end{bmatrix}, \quad B = \begin{bmatrix} -(\alpha/\delta)I_d \\ -\alpha I_d \end{bmatrix}, \quad C = [\delta \beta I_d \quad I_d], \quad E = [0 \quad I_d]. \quad (12)$$

According to Theorem 1, in order to identify a convergence rate for (19), it is sufficient to find numbers $a_0 > 0$, $\rho \in (0, 1)$, $\ell \geq 0$ and a matrix P with $P + (a_0 m/2)E^T E \succ 0$ such that T in (6) is $\preceq 0$. We set $\ell = 0$, as this does not have a significant impact on the value of ρ that results from the analysis. This, in turn, allows us to further simplify things, since T is homogeneous in P and a_0 and we may assume $a_0 = 1$. Then T is a function of P and ρ (and the method parameters β and δ).

The matrix A in (12) is a Kronecker product of a 2×2 matrix and I_d ,

$$A = \begin{bmatrix} \beta & 0 \\ \delta \beta & 1 \end{bmatrix} \otimes I_d;$$

the factor I_d originates from the dimensionality of the decision variable x and the 2×2 factor is independent of d and arises from the optimization algorithm. The matrices B , C and E have a similar Kronecker product structure. It is then natural to consider symmetric matrices P of the form

$$P = \hat{P} \otimes I_d, \quad \hat{P} = \begin{bmatrix} p_{11} & p_{12} \\ p_{12} & p_{22} \end{bmatrix}, \quad (13)$$

and then T will also have a Kronecker product structure

$$T = \widehat{T} \otimes I_d, \quad \widehat{T} = \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{12} & t_{22} & t_{23} \\ t_{13} & t_{23} & t_{33} \end{bmatrix}, \quad (14)$$

where the t_{ij} are explicitly given by the following complicated expressions obtained from (12) and the recipes for $M^{(0)}$, $M^{(1)}$ and $M^{(2)}$ in Theorem 1:

$$\begin{aligned} t_{11} &= \beta^2 p_{11} + 2\delta\beta^2 p_{12} + \delta^2 \beta^2 p_{22} - \rho^2 p_{11} - \delta^2 \beta^2 m/2, \\ t_{12} &= \beta p_{12} + \delta\beta p_{22} - \rho^2 p_{12} - \delta\beta m/2 + \rho^2 \delta\beta m/2, \\ t_{13} &= -\delta^{-1}\alpha\beta p_{11} - 2\alpha\beta p_{12} - \delta\alpha\beta p_{22} + \delta\beta/2, \\ t_{22} &= p_{22} - \rho^2 p_{22} - m/2 + \rho^2 m/2, \\ t_{23} &= -\delta^{-1}\alpha p_{12} - \alpha p_{22} + 1/2 - \rho^2/2, \\ t_{33} &= \delta^{-2}\alpha^2 p_{11} + 2\delta^{-1}\alpha^2 p_{12} + \alpha^2 p_{22} + \alpha^2 L/2 - \alpha. \end{aligned}$$

Our objective is to find $\rho \in (0, 1)$, p_{11} , p_{12} , and p_{22} that lead to $\widehat{T} \preceq 0$ and $\widehat{P} + (m/2)\widehat{E}^\top \widehat{E} \succ 0$ (which in turn imply $T \preceq 0$ and $P + (m/2)E^\top E \succ 0$). The algebra becomes simpler if we represent β and ρ^2 as $\beta = 1 - b\delta$, $\rho^2 = 1 - r\delta$. We choose the entries of \widehat{P} as (details may be found in the supplementary material)

$$\begin{aligned} p_{11} &= p_{22}\delta^2 - mr\delta + \frac{m}{2}, \\ p_{12} &= \frac{mr}{2} - \delta p_{22}, \\ p_{22} &= mr \frac{b^2\delta^3 - b^2\delta - 2rb\delta^3 + 2rb\delta + 3r\delta^2 - 2\delta - r}{(4\delta r - 4)}. \end{aligned}$$

With this choice, $t_{13} = t_{23} = 0$, $t_{33} = (\alpha^2 L - \alpha)/2$ and t_{11}, t_{12}, t_{22} depend only on δ, r, b . We next limit the steplength by imposing $\alpha \leq 1/L$, or equivalently $\delta \leq \delta_{max} = 1/\sqrt{\kappa}$, so as to have $t_{33} \leq 0$. Finally we demand that $t_{11}t_{22} - t_{12}^2 = 0$, which gives a relation $\varphi(r, b; \delta) = 0$ between the method parameter b and the rate r for each choice of δ . When this relation is satisfied, $t_{11} \leq 0$, which, in tandem with the other properties of the t_{ij} established before, guarantees $T \preceq 0$. For two values of δ_{max} , we plot in Figure 1a the curve $\varphi = 0$ for the most favourable steplength $\delta = \delta_{max}$ and, in addition, we compare with the analogous curve obtained in Sanz-Serna and Zygalkakis (2021) under the constraint $P \succeq 0$ required by the framework in Fazlyab et al. (2018). As we may see, by changing the constraint on P it is possible to prove a significantly better convergence rate. In particular, for the modified constraint in the present analysis, one can show that b may be chosen to get $r = \sqrt{2} - \mathcal{O}(\delta)$, which in turn implies that for $\delta = \delta_{max}$ in (7):

$$\rho^2 = 1 - \frac{\sqrt{2}}{\sqrt{\kappa}} + \mathcal{O}(\kappa^{-1}), \quad \kappa \rightarrow \infty.$$

Also note that the parameter values α and β in (2) lead to the best convergence rate that may be established with the approach in Sanz-Serna and Zygalkakis (2021); however the present analysis shows that higher convergence rates may be rigorously proved for alternative choices of β .

4 Connections with the differential equation

It is well known that (see e.g. Su et al. (2016)) that, if we set $h = \sqrt{\alpha}$ (h has the dimensions of t), and assume that in (10), the parameter $\beta = \beta_h$ changes smoothly with h in such a way that, for some constant $\bar{b} \in \mathbb{R}$, $\beta_h = 1 - \bar{b}\sqrt{m}h + o(h)$ we obtain (21) in the limit of $h \rightarrow 0$. Note that the friction coefficient \bar{b} is non-dimensional.

4.1 Lyapunov function

We will now proceed in a similar way as in the discrete setting to obtain a Lyapunov equation for (21). Our first step is to express (21) in the control theoretic framework (20). In particular, we define




Figure 1: The left panel shows the relationship between the rate r and the method parameter b in the discrete case when $\delta = \delta_{max} = 1/\sqrt{\kappa}$. The red curves correspond to the present analysis and the blue curves correspond to the hypothesis $P \succeq 0$. The solid curves are for $\kappa = 10^6$; the dashed curves are for $\kappa = 10^2$. The right panel shows the relationship between the rate \bar{r} and the parameter \bar{b} in the time-continuous case. The red and blue solid lines on the left are indistinguishable from the red and blue lines on the right. The green star denotes the choice of parameters used in Section 5.

$v = (1/\sqrt{m})\dot{x}$ and rewrite (21) as a first-order system

$$\dot{v} = -\bar{b}\sqrt{m}v - \frac{1}{\sqrt{m}}\nabla f(x), \quad (15a)$$

$$\dot{x} = \sqrt{m}v. \quad (15b)$$

The scaling factor \sqrt{m} has been introduced to ensure that the variables x and v have the same dimensions. If we set $\xi = [v^\top, x^\top]^\top$, then (15) is of the form (20) with

$$\bar{A} = \begin{bmatrix} -\bar{b}\sqrt{m}I_d & 0_d \\ \sqrt{m}I_d & 0_d \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} -(1/\sqrt{m})I_d \\ 0_d \end{bmatrix}, \quad \bar{C} = [0_d \quad I_d].$$

Similarly to the discrete case, the structure of the matrices \bar{A} , \bar{B} , and \bar{C} implies that we can look for a 2×2 $\hat{\bar{P}}$ and a 3×3 $\hat{\bar{T}}$ analogous to the equations (13) and (14), rather than for \bar{P} and \bar{T} .

We use Theorem 2 to find a Lyapunov function. Similarly to the discrete case, we will simplify the subsequent analysis by considering the case $\sigma = 0$. We find that the elements of $\hat{\bar{T}}$ are given by

$$\begin{aligned} \bar{t}_{11} &= -2\bar{b}\sqrt{m}\bar{p}_{11} + 2\sqrt{m}\bar{p}_{12} + \lambda\bar{p}_{11}, \\ \bar{t}_{12} &= -\bar{b}\sqrt{m}\bar{p}_{12} + \sqrt{m}\bar{p}_{22} + \lambda\bar{p}_{12}, \\ \bar{t}_{13} &= -(1/\sqrt{m})\bar{p}_{11} + \sqrt{m}/2, \\ \bar{t}_{22} &= \lambda\bar{p}_{22} - (m/2)\lambda, \\ \bar{t}_{23} &= -(1/\sqrt{m})\bar{p}_{12} + \lambda/2, \\ \bar{t}_{33} &= 0. \end{aligned}$$

We determine λ and $\hat{\bar{P}}$. The algebra is simplified if we set $\lambda = \sqrt{m}\bar{r}$. Since $\bar{t}_{33} = 0$, the requirement $\hat{\bar{T}} \preceq 0$ implies $\bar{t}_{13} = \bar{t}_{23} = 0$ which leads to

$$\bar{p}_{11} = m/2, \quad \bar{p}_{12} = (m/2)\bar{r}.$$

We now only have to deal with the leading 2×2 submatrix of $\hat{\bar{T}}$ and we need $\bar{t}_{11} < 0, \bar{t}_{22} < 0$ and $\Delta := \bar{t}_{11}\bar{t}_{22} - \bar{t}_{12}^2 \geq 0$. These conditions lead to $\bar{r} \leq 2\bar{b}/3, p_{22} \leq 1/2$ and

$$\Delta := -\sqrt{m}\bar{r} \left(\frac{3m^{3/2}\bar{r}}{2} - \bar{b}m^{3/2} \right) \left(\frac{m}{2} - p_{22} \right) - m \left(p_{22} + \frac{\bar{r}^2 m}{2} - \frac{\bar{b}\bar{r}m}{2} \right)^2 \geq 0.$$

We seek points in the (\bar{p}_{22}, \bar{r}) plane with $\Delta = 0$ and $(\partial/\partial\bar{p}_{22})\Delta = 0$. The second of these relations yields $\bar{p}_{22} = \bar{r}^2/4$. Then an analysis of the first relation shows that it is possible to get all rates \bar{r}

in the interval $(0, \sqrt{2})$ ($\bar{r} = \sqrt{2}$ has to be excluded because for this value \bar{P} does not satisfy the requirement on \tilde{P} in Theorem 2). Each value of $\bar{r} \in (0, \sqrt{2})$ may be achieved in two ways, the first by choosing $\bar{b} = 3\bar{r}/2 \in (0, 3\sqrt{2}/2)$ and the second by choosing $\bar{b} > 3\sqrt{2}/2$ such that $\bar{r} = \bar{b} - \sqrt{\bar{b}^2 - 4}$, see Figure 1b. As in the discrete case, the modification of the hypothesis on \bar{P} , allows to prove a significantly better convergence rate. Also note that the curve that relates \bar{r} and \bar{b} for the Polyak ODE is indistinguishable from the solid red curve in Figure 1a that relates b and r in the discrete case, a coincidence that will be explained in the next subsection.

If the objective function f is quadratic, it is of course possible to obtain a sharp bound for the convergence rate $\lambda = \sqrt{m\bar{r}}$ by solving (21) in terms of eigenvalues/vectors. (See (Lessard et al., 2016, Section 2.2) for the analysis in the discrete scenario.) For comparison, we have included in Figure 1b the rate for quadratic problems, which is maximized for $\bar{b} = 2$, where $\lambda = 2\sqrt{m}$. For non-quadratic targets, the rate that may be proved under the Fazlyab et al. (2018) hypothesis $\bar{P} \succeq 0$ is also maximized when $\bar{b} = 2$, where $\lambda = \sqrt{m}$. The present analysis establishes for non-quadratic targets rates arbitrarily close to $\lambda = \sqrt{2}\sqrt{m}$, by choosing \bar{b} close to $3\sqrt{2}/2$. Also note Figure 1b shows that the rate provided by our analysis cannot be improved when $\bar{b} \geq 3\sqrt{2}/2$.

4.2 Optimization algorithms as integrators

Systems of differential equations $(d/dt)z = g(z)$ in cases where it makes sense to decompose $g(z)$ as a sum $g(z) = \sum_{\nu=1}^N g^{[\nu]}(z)$ may be integrated by using Additive Runge-Kutta (ARK) algorithms, a generalization of the well-known Runge-Kutta (RK) integrators. In the RK case, the numerical solution is advanced over a time step $z_k \mapsto z_{k+1}$ by evaluating $g(z)$ at a sequence of so-called stage vectors $Z_{k,1}, \dots, Z_{k,s}$ and then setting $z_{k+1} = z_k + \sum_{i=1}^s b_i g(Z_{k,i})$, where the b_i are suitable weights. In turn, for the explicit algorithms we are interested in, the stages are computed successively, $i = 1, \dots, s$, as $Z_{k,i} = z_k + h \sum_{j=1}^{i-1} a_{i,j} g(Z_{k,j})$, with suitable coefficients $a_{i,j}$. ARK algorithms are entirely similar, but evaluate the individual pieces $g^{[\nu]}(z)$ rather than $g(z)$.

With $z = [v^\top, x^\top]^\top \in \mathbb{R}^{2d}$, the system (15) may be rewritten as

$$\frac{d}{dt}z = g^{[1]}(z) + g^{[2]}(z) + g^{[3]}(z) := \begin{bmatrix} -\bar{b}\sqrt{mv} \\ 0 \end{bmatrix} + \begin{bmatrix} -\frac{1}{\sqrt{m}}\nabla f(x) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \sqrt{mv} \end{bmatrix};$$

the three parts of $g(z)$ respectively represent the friction force, potential force and inertia in the oscillator. It is easily checked that, if we choose a steplength $h > 0$, and see d_k and x_k as approximations to $v(kh)$ and $x(kh)$ respectively, then a step $(d_k, x_k) \mapsto (d_{k+1}, x_{k+1})$ of the optimization algorithm (19) with parameters $\alpha = h^2$, $\beta = 1 - h\bar{b}\sqrt{m}$, $\delta = \sqrt{m}h$ is just one step $z_k \mapsto z_{k+1}$ of the ARK integrator for (15) given by:

$$\begin{aligned} Z_{k,1} &= z_k, \\ Z_{k,2} &= z_k + hg^{[1]}(Z_{k,1}), \\ Z_{k,3} &= z_k + hg^{[1]}(Z_{k,1}) + hg^{[3]}(Z_{k,2}), \\ Z_{k,4} &= z_k + hg^{[1]}(Z_{k,1}) + hg^{[3]}(Z_{k,2}) + hg^{[2]}(Z_{k,3}), \\ z_{k+1} &= z_k + hg^{[1]}(Z_{k,1}) + hg^{[2]}(Z_{k,3}) + hg^{[3]}(Z_{k,4}). \end{aligned}$$

The stage vectors have $Z_{k,1} = [d_k^\top, x_k^\top]^\top$, $Z_{k,2} = [\beta d_k^\top, x_k^\top]^\top$, $Z_{k,3} = [\beta d_k^\top, y_k^\top]^\top$, $Z_{k,4} = [d_{k+1}^\top, y_k^\top]^\top$, and therefore the computation of the second, third and fourth stages incorporate successively the contributions of friction, inertia and potential force.

If we now think that the value of $h > 0$ varies and consider the optimization algorithm (10) (or (19)) with $\alpha = h^2$, $\beta = 1 - h\bar{b}\sqrt{m}$, the standard theory of numerical integration of ODEs shows that, if the initial points x_{-1} and x_0 are chosen in such a way that, as $h \rightarrow 0$, x_0 and $(1/h)(x_0 - x_{-1})$ converge to limits A and B , then in the limit of $kh \rightarrow t$, x_k and $(1/h)(x_{k+1} - x_k)$ converge to $x(t)$ and $\dot{x}(t)$ respectively, where $x(t)$ is the solution of (21) with initial conditions $x(0) = A$ and $\dot{x}(0) = B$. In addition, the discrete Lyapunov function of the optimization algorithm in Section 3 may be shown to converge to the Lyapunov function of the ODE found in this section. Finally the discrete decay factor over k steps $(1 - \sqrt{m}rh)^k$ converges to the continuous decay factor $\exp(-\lambda t)$. These facts in particular imply that, in Figure 1, the graph of the relation between \bar{b} and \bar{r} that holds for the ODE is indistinguishable from the corresponding graph for the optimization algorithm when κ is large (κ being large corresponds to h being small).

4.3 Discretizations that do not succeed in getting acceleration

Many recent contributions have derived optimization algorithms by discretizing suitable chosen dissipative ODEs. It is well known that, unfortunately, many properties of ODEs are likely to be lost in the discretization process, even if high-order, sophisticated integrators are used. The archetypical example is provided by the discretization of the standard harmonic oscillator: most numerical methods, regardless of their accuracy, provide solutions that either decay to the origin or spiral out to infinity as the number of computed points grows unboundedly. Similarly, discretizations of (21) are likely not to share the favourable decay properties in Section 4.1.

Let us consider the following extension of the optimization algorithm (2):

$$x_{k+1} = x_k + \beta(x_k - x_{k-1}) - \alpha \nabla f(y_k), \quad (16a)$$

$$y_k = x_k + \gamma(x_k - x_{k-1}), \quad (16b)$$

with the additional parameter γ . The choice $\gamma = 0$ yields the *heavy ball* algorithm, which (see Sanz-Serna and Zygalkis (2021)) corresponds to a “natural” standard linear multistep discretization of the Polyak equation (21) but that does not provide acceleration. Sanz-Serna and Zygalkis (2021) shows that for $\gamma = 0$ (or more generally for $\gamma \neq \beta$), the optimization algorithm (16) does not inherit the Lyapunov functions of the Polyak ODE. The analysis in that paper hinges on a study of the nondimensional quantity $c := t_{11}/(m\delta)$, which for $\widehat{T} \leq 0$ has to be ≤ 0 and for a discretization of an ODE has a finite limit as $\delta \rightarrow 0$. When $\gamma = 0$, the expression for the quantity c includes a positive contribution $\delta(\kappa - 1)\beta^2/2$; for acceleration, δ has to be $\mathcal{O}(1/\sqrt{\kappa})$ which makes it impossible for c to be ≤ 0 .

The unwelcome presence of κ in t_{11} may be traced back to the appearance of L in the matrix $N^{(1)}$ in Theorem 1. Nesterov’s algorithms of the family (10) do not suffer from that appearance because for them the matrix $EA - C$ that multiplies $(L/2)I_d$ in the recipe for $N^{(1)}$ vanishes. The condition $EA - C = 0$ appears then to be of key importance in the success of Nesterov algorithms; we put it into words by saying that one has to impose that the point $y_k = C\xi_k$ where the gradient is evaluated has to coincide with the point $x_{k+1} = EA\xi_k$ that the algorithm would yield if $u_k = \nabla f(y_k)$ happened to vanish (see (18)). This suggests that the integrator has to treat the potential force and the friction force in the oscillator separately, something that may be achieved by ARK algorithms but not by more conventional linear multistep or RK methods that do not avail themselves of the separate pieces $g^{[1]}(z)$, $g^{[2]}(z)$, $g^{[3]}(z)$ but are rather formulated in terms of $g(z)$.

5 Numerical illustration

We illustrate Theorem 2 for the simple one dimensional function in $\mathcal{F}_{\mu,L}$ given by

$$f(x) = \frac{m}{2}x^2 + 4(L - m)\log(1 + e^{-x}). \quad (17)$$

We solve (15) using a very accurate RK algorithm and plot the results in Figure 2. In the left panel, where $\bar{b} = 3\sqrt{2}/2$, we plot, as a function of time t , the Lyapunov function $V(\xi(t), t)$ in (30) with the matrix \bar{P} found here and different values of $\lambda = \bar{r}\sqrt{m}$. In agreement with Section 4.1, V decays for $\bar{r} \leq \sqrt{2}$. In the central panel, we present the evolution of the function V when the matrix \bar{P} is chosen as in Sanz-Serna and Zygalkis (2021) to meet the more demanding requirement in Fazlyab et al. (2018); now V does not decay for $\bar{r} > 1$. Finally, the right panel corresponds to $\bar{b} = 3\bar{r}/2$, $\bar{r} = 1.4 < \sqrt{2}$, see Figure 1b, and plots $\|y(t) - y^*\|^2$ as a function of time as well as the bound provided by our analysis. The distance $\|y(t) - y^*\|$ decays non-monotonically, as expected in damped oscillators (and in Nesterov’ algorithms); however it is apparent that the decay rate λ provided by our analysis is fairly sharp. For this value of \bar{b} , the rate guaranteed by the analysis in Sanz-Serna and Zygalkis (2021) is significantly lower (Figure 1b).

6 Discussion

We have established improved non-asymptotic convergence guarantees for Nesterov optimization methods and the related Polyak ODE. For a range of parameter values b in the discrete case and \bar{b} in




Figure 2: Convergence of Polyak ODE for f given by (17) when $m = 10^{-3}$, $L = 1$. The first two panels display the evolution of the function V along trajectories.

the continuous case, we proved faster convergence rates than those established in previous studies. Our analysis indicates that the optimal choice of \bar{b} for the continuous setting is (close to) $3\sqrt{2}/2$ and this gives a convergence rate arbitrarily close to $\sqrt{2m}$. In the discrete setting, for a suitable choice of the algorithmic parameters (slightly different to the standard choices for Nesterov method), one may prove rates arbitrarily close to $1 - \sqrt{2}/\sqrt{\kappa}$ when $\kappa \gg 1$. The improved convergence rates obtained in this paper both in the continuous and discrete setting are a direct result of modifying the positive semidefiniteness condition for the matrices \bar{P} and P respectively when searching for appropriate Lyapunov functions.

The analysis in the discrete setting highlights an important point; namely that not all numerical discretizations of the Polyak ODE (15) lead to accelerated optimization algorithms. We have derived structural conditions that numerical integrators of (15) should satisfy in order to correspond to accelerated optimization schemes. Standard numerical methods for ODEs such as conventional linear multi-step or RK methods cannot satisfy those structural conditions. The success of Nesterov’s method may be easily understood when interpreting it as an Additive RK method.

Acknowledgments and Disclosure of Funding

P.D and K.C.Z acknowledges support from the EPSRC grant EP/V006177/1. J.M.S.-S. has been supported by Ministerio de Ciencia e Innovación (Spain) through project PID2019-104927GB-C21, MCIN/AEI/10.13039/501100011033, ERDF (“A way of making Europe”).

References

- M. Betancourt, M. I. Jordan, and A. C. Wilson. On symplectic optimization. *arXiv preprint*, 2018.
- A. Bravetti, M. L. Daza-Torres, H. Flores-Arguedas, and M. Betancourt. Optimization algorithms inspired by the geometry of dissipative systems. *arXiv preprint*, 2019.
- G.J. Cooper and A. Sayfy. Additive methods for the numerical solution of ordinary differential equations. *Mathematics of Computation*, 35(152):1159–1172, 1980.
- M. J. Ehrhardt, E. S. Riis, T. Ringholm, and C.-B. Schönlieb. A geometric integration approach to smooth optimisation: Foundations of the discrete gradient method. *arXiv preprint*, 2018.
- M. Fazlyab, A. Ribeiro, M. Morari, and V. M. Preciado. Analysis of optimization algorithms via integral quadratic constraints: nonstrongly convex problems. *SIAM J. Optim.*, 28(3):2654–2689, 2018.
- G. Franca, M. I. Jordan, and R. Vidal. On dissipative symplectic integration with applications to gradient-based optimization. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(4): 043402, 2021.
- W. Krichene, A. Bayen, and P. L. Bartlett. Accelerated mirror descent in continuous and discrete time. In *Advances in Neural Information Processing Systems 28*, pages 2845–2853. 2015.

- L. Lessard, B. Recht, and A. Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM J. Optim.*, 26(1):57–95, 2016.
- A. Megretski and A. Rantzer. System analysis via integral quadratic constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, 1997.
- M. Muehlebach and M. I. Jordan. A dynamical systems perspective on Nesterov acceleration. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4656–4662. PMLR, 2019.
- M. Muehlebach and M. I. Jordan. Optimization with momentum: Dynamical, control-theoretic, and symplectic perspectives. *J. Mach. Learn. Res.*, 22(1):1–50, 2021.
- Y. Nesterov. A method for solving the convex programming problem with convergence rate $O(k^{-2})$. *Dokl. Akad. Nauk SSSR*, 269:543–547, 1983.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 2014.
- A. Orvieto and A. Lucchi. Shadowing properties of optimization algorithms. In *Advances in Neural Information Processing Systems 32*, pages 12692–12703. 2019.
- B.T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, pages 1–17, 1964.
- J. M. Sanz-Serna and K. C. Zygalakis. The connections between lyapunov functions for some optimization algorithms and differential equations. *SIAM J. Numer. Anal.*, 59(3):1542–1565, 2021.
- D. Scieur, V. Roulet, F. R. Bach, and A. d’Aspremont. Integration methods and optimization algorithms. In *Advances in Neural Information Processing Systems*, volume 30, pages 1109–1118, 2017.
- B. Shi, S. S. Du, W. Su, and M. I. Jordan. Acceleration via symplectic discretization of high-resolution differential equations. In *Advances in Neural Information Processing Systems*, volume 32, pages 5744–5752, 2019.
- W. Su, S. Boyd, and E. J. Candès. A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights. *J. Mach. Learn. Res.*, 17(153):1–43, 2016.
- A. Wibisono, A. C. Wilson, and M. I. Jordan. A variational perspective on accelerated methods in optimization. *Proc. Natl. Acad. Sci. U.S.A.*, 113(47):E7351–E7358, 2016.
- A. C. Wilson, B. Recht, and M. I. Jordan. A lyapunov analysis of accelerated methods in optimization. *J. Mach. Learn. Res.*, 22(1):1–34, 2021.

A Control framework and linear matrix inequalities

Consider objective functions $f \in \mathcal{F}_{m,L}$, that is, f is continuously differentiable, m -strongly convex and ∇f is L -Lipschitz. The situation of interest is when the condition number $\kappa = L/m$ is large. A broad class of optimization algorithms which minimise f can be expressed as a linear dynamical system

$$\xi_{k+1} = A\xi_k + Bu_k, \quad (18a)$$

$$u_k = \nabla f(y_k), \quad (18b)$$

$$y_k = C\xi_k, \quad (18c)$$

$$x_k = E\xi_k, \quad (18d)$$

Here $\xi_k \in \mathbb{R}^n$ is the state, $u_k \in \mathbb{R}^d$ is the input, $y_k \in \mathbb{R}^d$ is the feedback output that is mapped to u_k by the function ∇f . The scheme is described by the matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times d}$, $C \in \mathbb{R}^{d \times n}$, and $E \in \mathbb{R}^{d \times n}$. All fixed points, ξ^* , of (18) satisfy

$$\xi^* = A\xi^* + Bu^*, \quad y^* = C\xi^*, \quad u^* = \nabla f(y^*), \quad x^* = E\xi^*.$$

In order for the fixed point to correspond to a minimiser we require that $x^* = y^*$ and $u^* = 0$.

The Gradient Descent algorithm with stepsize $\delta > 0$,

$$x_{k+1} = x_k - \delta \nabla f(x_k),$$

can be expressed in the form (18) by setting $n = d$, $A = I_d$, $B = -\delta I_d$, $C = I_d$, and $E = I_d$.

We can also express Nesterov's accelerated gradient method (NAG) in the form (18). NAG for the parameters $\alpha, \beta, \delta > 0$ is given by

$$d_{k+1} = \beta d_k - \frac{\alpha}{\delta} \nabla f(y_k), \quad (19a)$$

$$x_{k+1} = x_k + \delta \beta d_k - \alpha \nabla f(y_k), \quad (19b)$$

$$y_k = x_k + \delta \beta d_k. \quad (19c)$$

The standard choice of the parameters for strongly convex f is given by

$$\alpha = \frac{1}{L}, \quad \beta = \frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1}, \quad \delta = \frac{1}{\sqrt{\kappa}}.$$

If we write $\xi_k = [d_k^\top, x_k^\top]^\top$ then (19) is equivalent to (18) with

$$A = \begin{bmatrix} \beta I_d & 0_d \\ \delta \beta & I_d \end{bmatrix}, \quad B = \begin{bmatrix} -\frac{\alpha}{\delta} I_d \\ -\alpha I_d \end{bmatrix}, \quad C = [\delta \beta I_d, I_d], \quad E_k = [0_d, I_d].$$

There is an analogous framework for first order differential equations. Let $f \in \mathcal{F}_{m,L}$ and consider the system

$$\dot{\xi}(t) = \bar{A}\xi(t) + \bar{B}u(t), \quad y(t) = \bar{C}\xi(t), \quad u(t) = \nabla f(y(t)), \quad t \geq 0. \quad (20)$$

Here $\xi(t) \in \mathbb{R}^n$ is the state, $y(t) \in \mathbb{R}^d$ the output and $u(t) \in \mathbb{R}^d$ is the continuous feedback input. Similarly to the discrete setting all fixed points of (20) satisfy

$$0 = \bar{A}\xi^* + \bar{B}u^*, \quad y^* = \bar{C}\xi^*, \quad u^* = \nabla f(y^*).$$

We are interested in settings where $x^* = y^*$ is the minimiser of f and hence $u^* = 0$.

If $d = n$ and $A = 0_d, B = -I_d, C = I_d$ then (20) is the gradient flow of f .

An important ODE for our discussion of optimization algorithms is the Polyak damped oscillator ODE Polyak (1964)

$$\ddot{x} + \bar{b}\sqrt{m}\dot{x} + \nabla f(x) = 0, \quad \bar{b} > 0, \quad (21)$$

here \bar{b} is a nondimensional damping coefficient. The algorithm (19) can be seen as a discretization of this ODE. We can write (21) in the form (20) with $\xi(t) = [(1/\sqrt{m})\dot{x}(t)^\top, x(t)^\top]^\top$ and

$$\bar{A} = \begin{bmatrix} -\bar{b}\sqrt{m}I_d & 0_d \\ \sqrt{m}I_d & 0_d \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} -(1/\sqrt{m})I_d \\ 0_d \end{bmatrix}, \quad \bar{C} = [0_d, I_d].$$

B Expressing properties of $\mathcal{F}_{m,L}$ as matrix inequalities

In Megretski and Rantzer (1997) integral quadratic constraints (IQC) were proposed as a tool to describe classes of nonlinearities in control theory. This was adapted for optimization algorithms in Lessard et al. (2016).

The key idea here is to express concepts like m -convexity in terms of matrix inequalities. For example a function is m -strongly convex if and only if for all $x, y \in \mathbb{R}^d$

$$m\|x - y\|^2 \leq (x - y)^\top (\nabla f(x) - \nabla f(y)). \quad (22)$$

This is equivalent to the statement involving IQCs that f is m -strongly convex if and only if

$$\begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix}^\top \begin{bmatrix} -mI_d & \frac{1}{2}I_d \\ \frac{1}{2}I_d & 0_d \end{bmatrix} \begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix} \geq 0.$$

We will use two other inequalities for $f \in \mathcal{F}, x, y \in \mathbb{R}^d$, the first follows if ∇f is L -Lipschitz and the second for all $f \in \mathcal{F}_{m,L}$

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

$$\frac{mL}{m+L} \|x - y\|^2 + \frac{1}{m+L} \|\nabla f(x) - \nabla f(y)\|^2 \leq (\nabla f(x) - \nabla f(y))^\top (x - y)$$

we can express these two inequalities as the following IQCs

$$f(x) - f(y) \leq \begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix}^\top \begin{bmatrix} -\frac{L}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix},$$

$$\begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix}^\top \begin{bmatrix} -\frac{mL}{m+L} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & -\frac{1}{m+L} I_d \end{bmatrix} \begin{bmatrix} x - y \\ \nabla f(x) - \nabla f(y) \end{bmatrix} \geq 0.$$

From the above inequalities we obtain the following lemma in Fazlyab et al. (2018).

Lemma 1 (Fazlyab et al., 2018, Lemma 4.1) Fix $f \in \mathcal{F}_{m,L}$. Define $e_k = ((\xi_k - \xi^*)^\top, u_k^\top)^\top$ then the following inequalities hold for all k

$$e_k^\top M^{(1)} e_k \geq f(x_{k+1}) - f(x_k), \quad (23)$$

$$e_k^\top M^{(2)} e_k \geq f(x_{k+1}) - f(x^*), \quad (24)$$

$$e_k^\top M^{(3)} e_k \geq 0, \quad (25)$$

where

$$M^{(0)} = \begin{bmatrix} A^\top P A - \rho^2 P & A^\top P B \\ B^\top P A & B^\top P B \end{bmatrix},$$

and

$$M^{(1)} = N^{(1)} + N^{(2)}, \quad M^{(2)} = N^{(1)} + N^{(3)}, \quad M^{(3)} = N^{(4)},$$

with

$$N^{(1)} = \begin{bmatrix} EA - C & EB \\ 0 & I_d \end{bmatrix}^\top \begin{bmatrix} \frac{L}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} EA - C & EB \\ 0 & I_d \end{bmatrix},$$

$$N^{(2)} = \begin{bmatrix} C - E & 0 \\ 0 & I_d \end{bmatrix}^\top \begin{bmatrix} -\frac{m}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} C - E & 0 \\ 0 & I_d \end{bmatrix},$$

$$N^{(3)} = \begin{bmatrix} C^\top & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{m}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0 \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & I_d \end{bmatrix},$$

$$N^{(4)} = \begin{bmatrix} C^\top & 0 \\ 0 & I_d \end{bmatrix} \begin{bmatrix} -\frac{mL}{m+L} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & -\frac{1}{m+L} I_d \end{bmatrix} \begin{bmatrix} C & 0 \\ 0 & I_d \end{bmatrix}.$$

C Proof of the Theorems

C.1 Proof of Theorem 1

Fix $f \in \mathcal{F}_{m,L}$ and define

$$V_k(\xi) = \rho^{-2k} (a_0(f(E\xi) - f(x^*)) + (\xi - \xi^*)^\top P (\xi - \xi^*)), \quad \rho \in (0, 1). \quad (26)$$

Let P, ρ, σ be as in the statement of Theorem 1. In order to show that V is a Lyapunov function we need to establish that V is non-increasing along trajectories of the algorithm, i.e. $V_{k+1}(\xi_{k+1}) \leq V_k(\xi_k)$, and that $V_k(\xi_k)$ controls the convergence of $\|x_k - x^*\|$. Consider the second of these properties, which we establish by showing that V is a suitable upper bound for the distance between x_k and x^* .

Since f is m -strongly convex we have that

$$f(x) - f(x^*) \geq \frac{m}{2} \|x - x^*\|^2.$$

Therefore we can bound $V_k(\xi)$ from below with

$$V_k(\xi_k) \geq \rho^{-2k} \left(a_0 \frac{m}{2} \|x_k - x^*\|^2 + (\xi_k - \xi^*)^\top P (\xi_k - \xi^*) \right).$$

Writing $\tilde{P} = \frac{1}{2} a_0 m E^\top E + P$ and using that $x_k - x^* = E(\xi_k - \xi^*)$ we have

$$V_k(\xi_k) \geq \rho^{-2k} (\xi_k - \xi^*)^\top \tilde{P} (\xi_k - \xi^*).$$

By the assumptions of Theorem 1 we have that $\tilde{P} \succ 0$ so the quadratic form $(\xi - \xi^*)^\top \tilde{P} (\xi - \xi^*)$ is non-degenerate and is bounded from below by the minimum eigenvalue of \tilde{P} . Hence

$$V_k(\xi_k) \geq \min \sigma(\tilde{P}) \rho^{-2k} \|\xi_k - \xi^*\|^2.$$

Using that $x_k - x^* = E(\xi_k - \xi^*)$ we obtain the following bound

$$\|x_k - x^*\|^2 \leq \frac{\max \sigma(E^\top E)}{\min \sigma(\tilde{P})} \rho^{2k} V_k(\xi_k). \quad (27)$$

It remains to establish that $V_k(\xi_k)$ is non-increasing. Indeed, if $V_{k+1}(\xi_{k+1}) \leq V_k(\xi_k)$ then we have the desired bound

$$\|x_k - x^*\|^2 \leq \frac{\max \sigma(E^\top E)}{\min \sigma(\tilde{P})} \rho^{2k} V_0(\xi_0). \quad (28)$$

Let $M^{(0)}$ be given in the statement of Theorem 1 then we have

$$e_k^\top M^{(0)} e_k = (\xi_{k+1} - \xi_{k+1}^*)^\top P (\xi_{k+1} - \xi^*) - \rho^2 (\xi_{k+1} - \xi_k^*)^\top P (\xi_k - \xi^*), \quad (29)$$

where $e_k = ((\xi_k - \xi^*)^\top, u_k^\top)^\top$. Indeed, by using (18) we have

$$\begin{aligned} (\xi_{k+1} - \xi_{k+1}^*)^\top P (\xi_{k+1} - \xi^*) &= (\xi_k - \xi_k^*)^\top A^\top P A (\xi_k - \xi^*) + (u_k)^\top B^\top P A (\xi_k - \xi^*) \\ &\quad + (\xi_k - \xi_k^*)^\top A^\top P B u_k + u_k^\top B^\top P B u_k \\ &= e_k^\top M^{(0)} e_k + \rho^2 (\xi_{k+1} - \xi_k^*)^\top P (\xi_k - \xi^*). \end{aligned}$$

Using (23)-(24), and (29) we have

$$\begin{aligned} V_{k+1}(\xi_{k+1}) - V_k(\xi_k) &= a_0 \rho^{-2(k+1)} \left(\rho^2 ((f(E\xi_{k+1}) - f(E\xi_k)) + (1 - \rho^2)(f(E\xi_{k+1}) - f(x^*))) \right. \\ &\quad \left. + \rho^{-2(k+1)} ((\xi_{k+1} - \xi_{k+1}^*)^\top P (\xi_{k+1} - \xi^*) - \rho^2 (\xi_k - \xi_k^*)^\top P (\xi_k - \xi^*)) \right) \\ &\leq \rho^{-2(k+1)} e_k^\top \left(M^{(0)} + a_0 \rho^2 M^{(1)} + a_0 (1 - \rho^2) M^{(2)} \right) e_k. \end{aligned}$$

Since $e_k^\top M^{(3)} e_k$ is positive by (25) we can add $\ell e_k^\top M^{(3)} e_k$ to obtain

$$V_{k+1}(\xi_{k+1}) - V_k(\xi_k) \leq \rho^{-2(k+1)} e_k^\top \left(M^{(0)} + a_0 \rho^2 M^{(1)} + a_0 (1 - \rho^2) M^{(2)} + \ell M^{(3)} \right) e_k = \rho^{-2(k+1)} e_k^\top T e_k.$$

By the assumptions of Theorem 1 we have that $T \preceq 0$ so we have that $V_{k+1}(\xi_{k+1}) \leq V_k(\xi_k)$.

C.2 Proof of Theorem 2

Fix $f \in \mathcal{F}_{m,L}$, and define

$$V(\xi, t) = e^{\lambda t} \left(f(y(t)) - f(y^*) + (\xi(t) - \xi^*)^\top \tilde{P} (\xi(t) - \xi^*) \right). \quad (30)$$

As in the discrete case we need to establish two properties:

1. The function $V(\xi(t), t)$ is non-increasing in t ;
2. We can bound $V(\xi, t)$ from below with $\|y - y^*\|^2$, i.e.

$$\|y - y^*\| \leq e^{-\lambda t} \frac{\max \sigma(\tilde{C}^\top \tilde{C})}{\min \sigma(\hat{P})} V(\xi, t)$$

From these two properties the desired non-asymptotic bound in Theorem 2 follows immediately, see the proof of Theorem 1 for details. The proof of the second property follows from an analogous argument to the proof of Theorem 1.

It remains to show that $t \mapsto V(\xi(t), t)$ is non-increasing, which follows if $\dot{V}(\xi(t), t) \leq 0$. Using that $y = \bar{C}\xi$ we can differentiate V to find

$$\dot{V}(\xi, t) = e^{\lambda t} \left(\lambda(f(y) - f(y^*)) + \lambda(\xi - \xi^*)^\top \bar{P}(\xi - \xi^*) + \nabla f(y)^\top \bar{C}\dot{\xi} + \dot{\xi}^\top \bar{P}(\xi - \xi^*) + (\xi - \xi^*)^\top \bar{P}\dot{\xi} \right).$$

Replacing $\nabla f(y)$ with u and $\dot{\xi}$ with the equation (20)

$$\begin{aligned} e^{-\lambda t} \dot{V}(\xi, t) &= \lambda(f(y) - f(y^*)) + \lambda(\xi - \xi^*)^\top \bar{P}(\xi - \xi^*) + u^\top \bar{C}(\bar{A}\xi + \bar{B}u) \\ &\quad + (\bar{A}\xi + \bar{B}u)^\top \bar{P}(\xi - \xi^*) + (\xi - \xi^*)^\top \bar{P}(\bar{A}\xi + \bar{B}u). \end{aligned}$$

Define $e(t) = [(\xi(t) - \xi^*)^\top, u(t)^\top]^\top$ then we can write

$$e^{-\lambda t} \dot{V}(\xi, t) = \lambda(f(y) - f(y^*)) + e(t)^\top \begin{bmatrix} \lambda \bar{P} + \bar{A}^\top \bar{P} + \bar{P} \bar{A} & \frac{1}{2}(\bar{C} \bar{A})^\top + \bar{P} \bar{B} \\ \frac{1}{2} \bar{C} \bar{A} + \bar{B}^\top \bar{P} & \frac{1}{2}(\bar{C} \bar{B} + (\bar{C} \bar{B})^\top) \end{bmatrix} e(t).$$

Recall the matrix $M^{(0)}$ and $M^{(1)}$ defined in Theorem 2, with this notation we have

$$e^{-\lambda t} \dot{V}(\xi, t) = \lambda(f(y) - f(y^*)) + e(t)^\top (M^{(0)} + M^{(1)})e(t).$$

In order to control the f terms we use the convex inequality

$$f(y_1) - f(y_2) \leq \nabla f(y_1)^\top (y_1 - y_2) - \frac{m}{2} \|y_1 - y_2\|^2. \quad (31)$$

Note that setting $y_1 = x_{k+1}$ and $y_2 = x^*$ in (31) corresponds to (24) in the discrete setting, similarly we can rewrite (31) with $y_1 = y$ and $y_2 = y^*$ as

$$f(y) - f(y^*) \leq \begin{bmatrix} \xi - \xi^* \\ u \end{bmatrix}^\top \begin{bmatrix} \bar{C} & 0_d \\ 0_d & I_d \end{bmatrix}^\top \begin{bmatrix} -\frac{m}{2} I_d & \frac{1}{2} I_d \\ \frac{1}{2} I_d & 0_d \end{bmatrix} \begin{bmatrix} \bar{C} & 0_d \\ 0_d & I_d \end{bmatrix} \begin{bmatrix} \xi - \xi^* \\ u \end{bmatrix} = e(t)^\top M^{(2)} e(t).$$

Therefore we have

$$e^{-\lambda t} \dot{V}(\xi, t) \leq e(t)^\top (M^{(0)} + M^{(1)} + \lambda M^{(2)})e(t).$$

By (25) we have that $M^{(3)} \succeq 0$ hence for any $\sigma > 0$ we have

$$e^{-\lambda t} \dot{V}(\xi, t) \leq e(t)^\top (M^{(0)} + M^{(1)} + \lambda M^{(2)} + \sigma M^{(3)})e(t) = e(t)^\top T e(t).$$

By the assumptions of Theorem 2 we have that $T \preceq 0$ so $\dot{V}(\xi, t) \leq 0$ and is non-increasing.

D More details for the discrete case

We are required to find $p_{11}, p_{12}, p_{22} \in \mathbb{R}, \rho \in (0, 1)$ such that $\hat{T} \preceq 0$ where

$$\hat{P} = \begin{bmatrix} p_{11} & p_{12} \\ p_{12} & p_{22} \end{bmatrix}, \quad \hat{T} = \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{12} & t_{22} & t_{23} \\ t_{13} & t_{23} & t_{33} \end{bmatrix}.$$

Setting $\beta = 1 - b\delta$ and $\rho^2 = 1 - r\delta$ we find

$$\begin{aligned} t_{11} &= (1 - b\delta)^2 p_{11} + 2\delta(1 - b\delta)^2 p_{12} + \delta^2(1 - b\delta)^2 p_{22} - (1 - r\delta)^2 p_{11} - \delta^2(1 - b\delta)^2 m/2, \\ t_{12} &= (1 - b\delta) p_{12} + \delta(1 - b\delta) p_{22} - (1 - r\delta)^2 p_{12} - \delta(1 - b\delta) m/2 + (1 - r\delta)^2 \delta \beta m/2, \\ t_{13} &= -\delta^{-1} \alpha(1 - b\delta) p_{11} - 2\alpha(1 - b\delta) p_{12} - \delta \alpha(1 - b\delta) p_{22} + \delta(1 - b\delta)/2, \\ t_{22} &= r\delta p_{22} - r\delta m/2, \\ t_{23} &= -\delta^{-1} \alpha p_{12} - \alpha p_{22} + r\delta/2, \\ t_{33} &= \delta^{-2} \alpha^2 p_{11} + 2\delta^{-1} \alpha^2 p_{12} + \alpha^2 p_{22} + \alpha^2 L/2 - \alpha. \end{aligned}$$

Observe that in the limit as $\alpha \rightarrow 0$ we have that T converges pointwise to \bar{T} the continuous Lyapunov function for the Polyak ODE. In the limit as $\alpha \rightarrow 0$ we have that $t_{33} = 0$ and moreover $t_{33} = \mathcal{O}(\alpha)$ for α small. For this reason in order to have $T \preceq 0$ it is necessary that t_{13} and t_{23} vanish in the limit as $\alpha \rightarrow 0$. We ensure this holds by setting $t_{13}, t_{23} = 0$ which leads to the following conditions:

$$\begin{aligned} p_{11} &= p_{22}\delta^2 - mr\delta + \frac{m}{2}, \\ p_{12} &= \frac{mr}{2} - \delta p_{22} \end{aligned}$$

Next we need to ensure that $t_{33} \leq 0$, with this choice of p_{11} and p_{12} we have

$$t_{33} = \frac{1}{2}\alpha(L\alpha - 1)$$

so we require $\alpha \leq 1/L$. Note this justifies the earlier arguments that we are interested in α small.

With these choices the matrix \hat{T} has the form

$$\hat{T} = \begin{bmatrix} \hat{T}_{1:2} & 0 \\ 0 & t_{33} \end{bmatrix}$$

with $t_{33} \leq 0$ so it remains to show that the submatrix $\hat{T}_{1:2}$ is negative semi-definite.

In order to have that $\hat{T} \preceq 0$ it is necessary and sufficient to find $p_{22} \in \mathbb{R}$ and $r > 0$ such that $t_{22} \leq 0$ and $\Delta = \det(\hat{T}_{1:2}) \geq 0$.

We choose p_{22} which maximises Δ by solving $\partial_{p_{22}} \Delta = 0$ for p_{22} . Indeed Δ is a quadratic in p_{22} which tends to $-\infty$ for p_{22} large so there is a unique extremum which is a maxima of Δ . This gives the following expression for p_{22}

$$p_{22} = mr \frac{b^2\delta^3 - b^2\delta - 2rb\delta^3 + 2rb\delta + 3r\delta^2 - 2\delta - r}{(4\delta r - 4)}.$$

Note that t_{22} is non-positive provided $p_{22} \leq m/2$. It remains to find the solutions of $\Delta = 0$ and verify that $\hat{P} \succ 0$ and $p_{22} \leq m/2$. This can be done with a numerical solver, the results are given in Figure 1 of the paper.