Near-optimal Closed-loop Method via Lyapunov Damping for Convex Optimization

Severin Maier   **Camille Castera**[*]
Department of Mathematics
University of Tübingen
Germany

Peter Ochs
Department of Mathematics and Computer Science
Saarland University
Germany

ABSTRACT

We introduce an autonomous system with closed-loop damping for first-order convex optimization. While, to this day, optimal rates of convergence are only achieved by non-autonomous methods via open-loop damping (*e.g.*, Nesterov's algorithm), we show that our system is the first one featuring a closed-loop damping while exhibiting a rate arbitrarily close to the optimal one. We do so by coupling the damping and the speed of convergence of the system via a well-chosen Lyapunov function. We then derive a practical first-order algorithm called LYDIA by discretizing our system, and present numerical experiments supporting our theoretical findings.

1   Introduction

We consider the unconstrained minimization of a smooth convex real-valued and lower-bounded function f over a Hilbert space $\mathcal{H}$:

$$\min_{x \in \mathcal{H}} f(x).$$

One of the most famous algorithms for such optimization problems is Nesterov's Accelerated Gradient method (NAG) [25], which is known to achieve the optimal rate of convergence for first-order methods on convex functions with Lipschitz-continuous gradient. Among several ways to explain the efficiency of NAG, Su et al. [31] studied the algorithm through the lens of Ordinary Differential Equations (ODEs) and proposed the following model:

$$\ddot{x}(t) + \frac{a}{t}\dot{x}(t) + \nabla f(x(t)) = 0, \qquad \forall t \geq t_0, \tag{AVD$_a$}$$

where $a > 0$, $t_0 \geq 0$ and ∇f denotes the gradient of the smooth real-valued convex function f. Here, $\dot{x} \stackrel{\text{def}}{=} \frac{\mathrm{d}x}{\mathrm{d}t}$, respectively $\ddot{x} \stackrel{\text{def}}{=} \frac{\mathrm{d}^2 x}{\mathrm{d}t^2}$, denotes the first, resp. second, time-derivative (or velocity, resp. acceleration) of the solution x of the ODE. AVD stands for Asymptotically Vanishing Damping

---

[*]Corresponding author: `camille.castera@protonmail.com`

and relates to the coefficient a/t. NAG is obtained by (non-straightforward) discretization of (AVD$_a$). Conversely (AVD$_a$) can be seen as NAG with infinitesimal step-sizes. Following [31], many works studied the system (AVD$_a$), and notably proved that when $a > 3$, the function values along a solution (or trajectory) of (AVD$_a$) converge with the asymptotic rate $o\left(\frac{1}{t^2}\right)$ to the optimal value as $t \to +\infty$ [31, 6, 24]. This matches the rate of NAG in the case of discrete (or iterative) algorithms [25]. We further discuss additional results for (AVD$_a$) and other possible choices of a later in Section 2.

The interest in the connection between ODEs (or dynamical systems) and algorithms comes from the abundance of theory and tools for analyzing ODEs and the insights that they provide for understanding algorithms. In particular, continuous-time analyses relying on Lyapunov functions can often be adapted to the discrete setting, see *e.g.*, [6]. In this work, we are interested in a special case of the Inertial Damped Gradient (IDG) system:

$$\ddot{x}(t) + \gamma(t)\dot{x}(t) + \nabla f(x(t)) = 0, \qquad \forall t \geq t_0, \qquad \text{(IDG}_\gamma\text{)}$$

where γ is a positive function and is called "damping" by analogy with mechanics. There are two inherently different ways of designing the damping coefficient γ: the so-called open- and closed-loop manners. While obtaining fast convergence rates $o\left(\frac{1}{t^2}\right)$ with open-loop damping is known, for example, when $\gamma(t) = \frac{a}{t}$ with $a > 3$, which exactly yields (AVD$_a$), we propose the first closed-loop damping that provides (IDG$_\gamma$) with near-optimal rate of convergence. Our special instance of (IDG$_\gamma$) is called (LD) and is introduced hereafter. We first discuss the problem setup in more details.

1.1 Problem setting

The damping $\frac{a}{t}$ in (AVD$_a$) depends explicitly on the time variable t, making the system non-autonomous. We say that a damping with such explicit dependence on the time t is *open-loop*. In contrast, a damping γ that does not explicitly depend on t is called *closed-loop* and makes (IDG$_\gamma$) an autonomous ODE (since the ODE is then independent of t_0): the feedback for the system, in terms of γ, may only depend on the state x (and its derivatives) but *not explicitly* on the time t. For optimization, autonomous ODEs are often preferable over non-autonomous ones, since the dependence on the time t, and hence on the initial time t_0, is removed. For example, even though the asymptotic rate of convergence of (AVD$_a$) does not depend on t_0, the trajectory does depend on the choice of t_0: when choosing a large t_0, the damping a/t is very small at all time and the "damping effect" is almost completely lost. We illustrate this in Figure 1 where we show how the choice of t_0 heavily influences the solution of (AVD$_a$). Additionally some important tools for analyzing ODEs only hold (or are simpler to use) for autonomous ODEs, see *e.g.*, the Hartman–Grobman Theorem [20, 19], which relates the behavior of autonomous ODEs around their equilibrium points to that of a linear system. Another such example is the so-called ODE method [23, 13], which formally states that algorithms with vanishing step-sizes asymptotically behave like solutions of corresponding ODEs.

In this work we therefore tackle the following open question:

> *Can one design the damping γ in (IDG$_\gamma$) in a closed-loop manner (so as to make the ODE autonomous) while still achieving the same optimal convergence rate as (AVD$_a$)?*

This question hides a circular argument that makes it hard to solve. Indeed, according to Cabot et al. [16], a sufficient condition on γ for the convergence of the function values is that $\int_{t_0}^{+\infty} \gamma(t)\mathrm{d}t = +\infty$. The straightforward choice is then the constant damping $\gamma(t) = a > 0$, which in (IDG$_\gamma$) yields
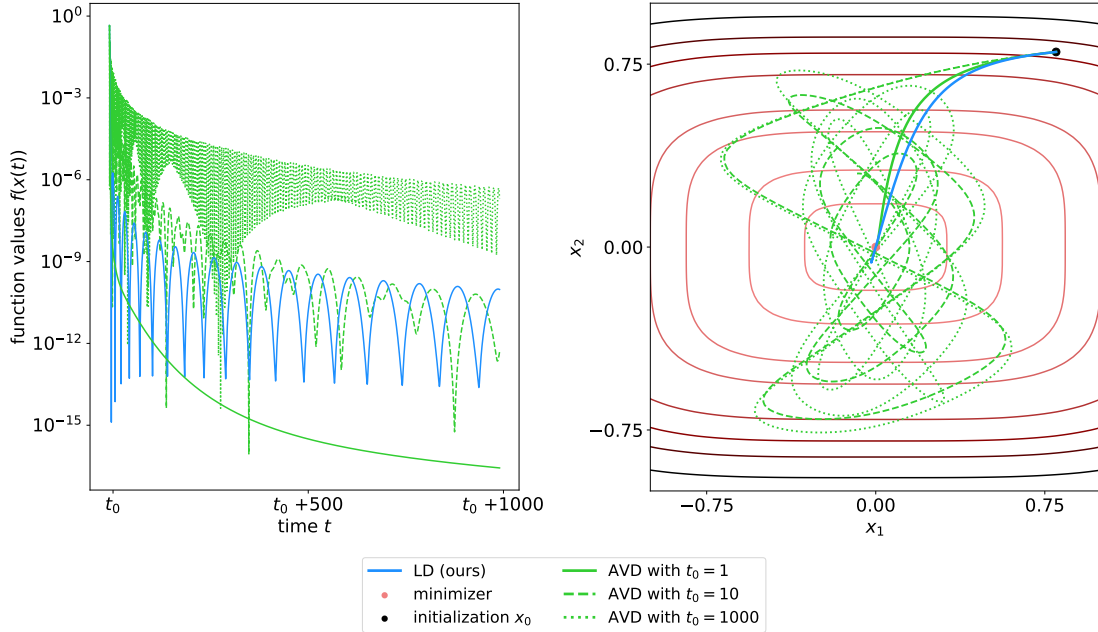
Figure 1: Comparison of our autonomous system (LD) and the non-autonomous one (AVD_a) for different initial times t_0 on the 2D function $f(x_1, x_2) = x_1^4 + 0.1x_2^4$. The left plot shows the evolution of the function values over time. The right plot shows the trajectories in the space (x_1, x_2). Different initial times heavily affect the solution of (AVD_a), but not (LD). We approximated the solutions using NAG for (AVD_a) and LYDIA (see Algorithm 1) for (LD), both with very small step-sizes.

the Heavy Ball with Friction (HBF) algorithm [28], known to be sub-optimal for convex functions. We therefore seek a closed-loop damping that converges to zero, but not "too fast". In the open-loop setting, the natural choice is $\gamma(t) = \frac{a}{t}$, $a > 0$ because this is the fastest polynomial decay that converges to zero while being non-integrable. Further [4] showed that any damping of the form $\frac{a}{t^\beta}$ with $\beta \in (0, 1)$ yields a sub-optimal rate. Therefore we formally seek to design a closed-loop damping γ that behaves like $\frac{1}{t}$. Because it cannot depend directly on the variable t, our closed-loop damping γ must be built from other quantities of the system that may converge to zero, for example $\|\dot{x}(t)\|$. *However, the rate at which such quantities converge to zero depend themselves on the choice of γ.* To escape this loop of thought, we design γ using quantities that are known to behave asymptotically like $\frac{1}{t}$ in the case of (AVD_a) and investigate whether this choice of γ still gives fast convergence rates. This idea takes inspiration from [9] which used $\|\dot{x}(t)\|$ as damping, with a key difference that we now explain.

1.2 Our contribution

In the case of (AVD_a) with $a > 3$, it is known that under standard assumptions the quantity $E(t) \stackrel{\text{def}}{=} f(x(t)) - f^\star + \frac{1}{2} \|\dot{x}(t)\|^2$ (where $f^\star \stackrel{\text{def}}{=} \min_{\mathcal{H}} f$) is such that $E(t) = o\left(\frac{1}{t^2}\right)$ [24]. Therefore, following the aforementioned intuition, we propose a closed-loop damping defined for all $t \geq t_0$ by $\gamma(t) \stackrel{\text{def}}{=} \sqrt{E(t)}$. Our version of ($\text{IDG}_\gamma$) then reads:

$$\ddot{x}(t) + \sqrt{E(t)}\dot{x}(t) + \nabla f(x(t)) = 0, \qquad \forall t \geq t_0, \tag{LD}$$

3

and is called (LD) for *Lyapunov Damping*, since it will turn out later that E is non-increasing, making it a so-called "Lyapunov function" for (LD). Note that (LD) assumes the availability of the optimal value $f^\star$, we further discuss this later in Remark 4.10. Our main result is then that our system (LD) which features a closed-loop damping yields the following rate of convergence.

Theorem 1.1. Assume that f is a continuously differentiable convex function and that $\operatorname{argmin}_{\mathcal{H}} f \neq \emptyset$. Then, for any bounded solution x of (LD) and for any $\delta > 0$,

$$f(x(t)) - f^\star = o\left(\frac{1}{t^{2-\delta}}\right).$$

This means that the rate of our method is arbitrarily close to $o\left(\frac{1}{t^2}\right)$, the optimal one achieved by (AVD$_a$) with $a > 3$. To the best of our knowledge, this is the first system with closed-loop damping featuring such rate. Additionally, note that (LD) does not make use of the hyper-parameter a, whose choice is crucial in (AVD$_a$).

1.3 Organization

In Section 2 we review related work regarding ODEs for optimization, closed-loop damping, and known results in the open-loop case. In Section 3 we make the setting precise and show the existence of solutions to (LD). Section 4 is devoted to showing our main result. In Section 5 we derive a practical algorithm from (LD) and use it to perform numerical experiments. We finish by drawing conclusions and further discussing our results.

2 Related Work

ODEs for optimization. There is a long line of work in exploiting the interplay between ODEs and optimization algorithms, going back, at least, to the work of Polyak's [28] Heavy Ball with Friction (HBF) method for acceleration. As previously stated, Su et al. [31] linked NAG [25] to the differential equation (AVD$_a$), hence providing a new view on the heavily used, yet not perfectly understood algorithm. NAG is however obtained via a non-straightforward discretization of (AVD$_a$) since the gradient of f is evaluated at an "extrapolated point". Recently, Alecsa et al. [2] proposed a model with "implicit Hessian", whose Euler explicit discretization yields NAG. Higher-order ODEs have also been introduced to better understand NAG: Attouch et al. [7] proposed a model with "Hessian damping" based on [3], while Shi et al. [30] similarly considered higher-resolution ODEs allowing to better distinguish NAG from other (IDG$_\gamma$) systems.

Closed-loop damping. The closest work to ours is [9], which proposed a closed-loop damping of the form $\gamma = \|\dot{x}\|^p$ for several values of $p \in \mathbb{R}$. Our work builds on [9] since our damping $\sqrt{E}$ also involves $\|\dot{x}\|$. Yet, Attouch et al. [9] could provide counterexamples for which their systems do not achieve near-optimal rates, unlike ours. A key difference is that our damping $\sqrt{E}$ is non-increasing, while $\|\dot{x}\|$ may oscillate heavily. They nonetheless derived rates under additional assumptions (*e.g.* strong convexity or the Kurdyka-Łojasiewicz (KL) property).

Lin and Jordan [22] introduced more complex systems of closed-loop dampings relying on the gradient ∇f and the speed $\dot{x}$. This was followed by Attouch et al. [10] who considered a system that

is shown to generalize the systems in [22]. Further Attouch et al. [10] showed that their system can be interpreted as a time-rescaled gradient flow combined with so-called "time averaging" techniques. Their time-scaling is expressed in terms of the inverse of the norms of the speed $\dot{x}$ and the gradient ∇f. However, the improvement achieved by time rescaling cannot carry through discretization since it amounts to following the same trajectory but faster, which for discrete algorithms would boil down to increasing the step-size, yielding numerical instabilities. Nonetheless, by combining time rescaling and time averaging, they managed to build a practical method achieving acceleration compared to the $\mathcal{O}\left(\frac{1}{t}\right)$ rate of gradient flow, but still sub-optimal compared to (AVD$_a$).

Adly et al. [1] proposed to replace the term $\gamma(t)\dot{x}(t)$ in (IDG$_\gamma$) by a "non-smooth" potential. Their non-smooth potential provides the solutions of their differential inclusion with finite length and convergence to a point very close to a minimizer of f (if such a point exists). However, since the solution does not converge to the minimal value, it does not have the optimal rate that we seek.

Our damping makes use of the optimality gap $f(x(t)) - f^\star$ (see Remark 4.10 for further discussion). This idea is not new and is used, for example, in the Polyak step-size $\frac{f(x_k)-f^\star}{\|\nabla f(x_k)\|^2}$ [27]. In our system (LD) the optimality gap is rather used for designing the damping γ rather than the step-size.

Open-loop damping and proof techniques. The convergence properties of (AVD$_a$) (and more generally (IDG$_\gamma$) in the open-loop setting) have been intensively studied. First, for (AVD$_a$) with $a > 3$, following [31], the convergence rate $o\left(\frac{1}{t^2}\right)$ for function values and the convergence of the trajectories were proved in [24, 6, 17]. In the setting $a < 3$, [11, 8] derived the sub-optimal rate $\mathcal{O}\left(\frac{1}{t^{\frac{2a}{3}}}\right)$. For the critical value $a = 3$ only the rate $\mathcal{O}\left(\frac{1}{t^2}\right)$ is known and whether this rate can be improved to $o\left(\frac{1}{t^2}\right)$ and whether the trajectories converge remains open. Rates have also been derived under additional assumptions such as the KL property [12].

Regarding (IDG$_\gamma$), Attouch and Cabot [4] developed general conditions for the convergence of the values and of the trajectories, which unify several results mentioned above. All these results have in common that they rely on the analysis of Lyapunov functions, like E previously introduced. We refer to [32, 33] for more details on Lyapunov analyses. The proof of our main result takes inspiration from those in [15] where sufficient conditions to derive optimal rates are provided. They must nonetheless be significantly adapted since some of the conditions in [15] do not hold generally for closed-loop dampings like ours. We replace them by using a specific property of our system (see Lemma 4.1 hereafter). A more detailed comparison follows our main analysis in Remark 4.9.

3 Preliminaries and Existence of Solutions

Throughout the paper we fix a real Hilbert space $\mathcal{H}$ with inner product $\langle \cdot, \cdot \rangle$ and induced norm $\|\cdot\|$. We make the following assumptions on the function f.

Assumption 1. The function $f : \mathcal{H} \to \mathbb{R}$ is

(i) convex and continuously differentiable with locally Lipschitz-continuous gradient ∇f;

(ii) bounded from below by $f^\star \overset{\text{def}}{=} \inf_{\mathcal{H}} f$.

We also fix an initial time $t_0 \geq 0$ and initial conditions $x(t_0) = x_0 \in \mathcal{H}$, and $\dot{x}(t_0) = \dot{x}_0 \in \mathcal{H}$.

Definition 1. A function $x : [t_0, +\infty[\to \mathcal{H}$, which is twice continuously differentiable on $]t_0, +\infty[$ and continuously differentiable on $[t_0, +\infty[$, is called a (global) *solution* or *trajectory* to (IDG$_\gamma$), resp. (LD), if it satisfies (IDG$_\gamma$), resp. (LD), for all $t > t_0$ and satisfies the initial conditions previously mentioned.

Given the setting above, we can ensure the existence and uniqueness of the solutions of (LD).

Theorem 3.1. Under Assumption 1, there exists a unique solution x to (LD) with initial conditions $(x_0, \dot{x}_0) \in \mathcal{H} \times \mathcal{H}$ and initial time $t_0 \geq 0$.

The proof relies on the Picard–Lindelöf Theorem and is postponed to Appendix A.

Remark 3.2. The local Lipschitz-continuity of ∇f is only required to guarantee the existence and uniqueness of the solutions on $[t_0, +\infty[$ (see below), but is not used elsewhere in our analysis.

We finally recall a special case of the Landau notation for asymptotic comparison that we use heavily in the sequel.

Definition 2. For any non-negative function $g \colon \mathbb{R} \to \mathbb{R}$ and $\forall \alpha \geq 0$,

$$g(t) = o\left(\frac{1}{t^\alpha}\right) \iff \lim_{t \to +\infty} t^\alpha g(t) = 0.$$

4 Main Result: Convergence Rates for (LD)

This section is devoted to proving our main result Theorem 1.1, which states that the solution of (LD) achieves a convergence rate that is arbitrarily close to the optimal one for convex functions. First we show that E is indeed a Lyapunov function for the system (LD).

4.1 The Lyapunov function of (LD)

Throughout what follows, let x be the solution of (LD) with the initial conditions stated in Section 3. Recall that the damping coefficient in (LD) is $\sqrt{E(t)}$ where for all $t \geq t_0$, $E(t) = f(x(t)) - f^\star + \frac{1}{2} \|\dot{x}(t)\|^2$. We first show identities that are specific to (LD) and that play a crucial role in what follows.

Lemma 4.1. Under Assumption 1 the solution x of (LD) is such that E is continuously differentiable for all $t \geq t_0$ and

$$\frac{\mathrm{d}E(t)}{\mathrm{d}t} = -\sqrt{E(t)} \, \|\dot{x}(t)\|^2, \tag{1}$$

so in particular E is non-increasing. Furthermore, for all $t \geq t_0$ it holds that

$$\frac{\mathrm{d}\sqrt{E(t)}}{\mathrm{d}t} = -\frac{1}{2} \, \|\dot{x}(t)\|^2, \quad \text{or equivalently,} \quad \int_{t_0}^t \|\dot{x}(s)\|^2 \, \mathrm{d}s = 2\sqrt{E(t_0)} - 2\sqrt{E(t)}. \tag{2}$$

Proof. We apply the chain rule and use the fact that x solves (LD) to obtain:

$$\frac{\mathrm{d}E(t)}{\mathrm{d}t} = \langle \nabla f(x(t)), \dot{x}(t) \rangle + \langle \ddot{x}(t), \dot{x}(t) \rangle$$

$$= \left\langle \nabla f(x(t)) - \sqrt{E(t)}\dot{x}(t) - \nabla f(x(t)), \dot{x}(t) \right\rangle = -\sqrt{E(t)} \, \|\dot{x}(t)\|^2, \tag{3}$$

which proves (1). As for the second part, the chain rule and (1) imply

$$\frac{\mathrm{d}\sqrt{E(t)}}{\mathrm{d}t} = \frac{1}{2\sqrt{E(t)}}\frac{\mathrm{d}E(t)}{\mathrm{d}t} = -\frac{1}{2\sqrt{E(t)}}\sqrt{E(t)}\,\|\dot{x}(t)\|^2 = -\frac{1}{2}\,\|\dot{x}(t)\|^2\,.$$

Finally, by the Fundamental Theorem of Calculus and the continuity of $\dot{x}$ this is equivalent to

$$\int_{t_0}^{t}\|\dot{x}(s)\|^2\,\mathrm{d}s = 2\sqrt{E(t_0)} - 2\sqrt{E(t)}. \qquad \square$$

The function E describes a quantity which is non-increasing along the trajectory x, and note that $E(t) = 0$ if, and only if, $f(x(t)) = f^\star$ and $\dot{x}(t) = 0$. Such a function is called *a Lyapunov function* for the system (LD).

4.2 Preliminary convergence results

We make the following assumptions, which are consistent with those in Theorem 1.1.

Assumption 2. We assume that

    (i) $\operatorname{argmin}_{\mathcal{H}} f \neq \emptyset$;

    (ii) the solution x of (LD) is uniformly bounded on $[t_0, +\infty[$.

Remark 4.2. Assumption 2-(ii) holds, for example, when f is coercive.

We begin our analysis by showing that the trajectory x minimizes the function f. We do so by showing that $E(t)$ tends to zero as $t \to +\infty$.

Theorem 4.3. Under Assumptions 1 and 2, $E(t)$ converges to zero as $t \to +\infty$. This implies in particular that $f(x(t)) \xrightarrow[t \to +\infty]{} f^\star$ and $\|\dot{x}(t)\| \xrightarrow[t \to +\infty]{} 0$.

We make use of the following classical result to prove Theorem 4.3.

Lemma 4.4. If g is a non-negative continuous function on $[t_0, +\infty[$ and $\int_{t_0}^{+\infty} g(t)\mathrm{d}t$ is finite, then either $\lim\limits_{t \to +\infty} g(t)$ does not exist or $\lim\limits_{t \to +\infty} g(t) = 0$.

The proof of Lemma 4.4 is postponed to Appendix B. We now prove Theorem 4.3.

Proof of Theorem 4.3. Let $z \in \operatorname{argmin}_{\mathcal{H}} f$, and for all $t \geq t_0$ define the so-called "anchor function" $h_z(t) \stackrel{\text{def}}{=} \frac{1}{2}\,\|x(t) - z\|^2$. Then h_z is twice differentiable for all $t \geq t_0$ and we have:

$$\dot{h}_z(t) = \langle x(t) - z, \dot{x}(t) \rangle. \tag{4}$$

$$\begin{aligned}
\ddot{h}_z(t) &= \|\dot{x}(t)\|^2 + \langle x(t) - z, \ddot{x}(t) \rangle \\
&\stackrel{\text{(LD)}}{=} \|\dot{x}(t)\|^2 + \left\langle x(t) - z, -\sqrt{E(t)}\dot{x}(t) - \nabla f(x(t)) \right\rangle \\
&= \|\dot{x}(t)\|^2 + \left\langle x(t) - z, -\sqrt{E(t)}\dot{x}(t) \right\rangle - \langle x(t) - z, \nabla f(x(t)) \rangle \\
&\stackrel{(4)}{\leq} \|\dot{x}(t)\|^2 - \sqrt{E(t)}\dot{h}_z(t) + f^\star - f(x(t)),
\end{aligned} \tag{5}$$

7

where we used the first-order characterization of the convexity of f in the last step. So $h_z(t)$ fulfills the following differential inequality:

$$\ddot{h}_z(t) + \underbrace{f(x(t)) - f^\star + \frac{1}{2}\|\dot{x}(t)\|^2}_{=E(t)} \leq \|\dot{x}(t)\|^2 + \frac{1}{2}\|\dot{x}(t)\|^2 - \sqrt{E(t)}\dot{h}_z(t),$$

or equivalently,

$$E(t) \leq \frac{3}{2}\|\dot{x}(t)\|^2 - \ddot{h}_z(t) - \sqrt{E(t)}\dot{h}_z(t). \tag{6}$$

We integrate (6) from t_0 to $T > t_0$:

$$\int_{t_0}^{T} E(t)\mathrm{d}t \leq \frac{3}{2}\int_{t_0}^{T}\|\dot{x}(t)\|^2\,\mathrm{d}t - \int_{t_0}^{T}\ddot{h}_z(t)\mathrm{d}t - \int_{t_0}^{T}\sqrt{E(t)}\dot{h}_z(t)\mathrm{d}t$$

$$= 3\left(\sqrt{E(t_0)} - \sqrt{E(T)}\right) - \dot{h}_z(T) + \dot{h}_z(t_0) - \sqrt{E(t)}h_z(T) + \sqrt{E(0)}h_z(t_0) \tag{7}$$

$$- \frac{1}{2}\int_{t_0}^{T}\|\dot{x}(t)\|^2 h_z(t)\mathrm{d}t,$$

where we performed integration by parts on the last integral and used Lemma 4.1. Note that by the boundedness of E and the continuity of f, $f(x) - f^\star$ and $\|\dot{x}\|$ are uniformly bounded on $[t_0, +\infty[$. This implies together with Assumption 2 that h_z and $\dot{h}_z$ are uniformly bounded as well. Therefore, $\dot{h}_z(T)$, $\sqrt{E(T)}$ and $\sqrt{E(T)}h_z(T)$ are uniformly bounded from above for all $T \in [t_0, +\infty[$. Further $-\int_{t_0}^{T}\|\dot{x}(s)\|^2 h_z(s)\mathrm{d}s \leq 0$ and $\dot{h}_z(t_0)$ and $\sqrt{E(t_0)}h_z(t_0)$ are constants. So the right-hand side in (7) is uniformly bounded from above for all $T \in [t_0, +\infty[$. Therefore we deduce that

$$\int_{t_0}^{+\infty} E(t)\mathrm{d}t < +\infty. \tag{8}$$

Finally, E is non-negative and non-increasing by (1), so it *converges* to some value $E_\infty \in [0, E(t_0)]$ as $t \to \infty$. Further since E is continuous we can conclude by Lemma 4.4 that $\lim_{t \to +\infty} E(t) = 0$. $\qquad\square$

4.3 Rates of convergence for E

The proof of our main result Theorem 1.1 relies on the following lemma.

Lemma 4.5. Let $\alpha \geq 0$: If $\int_{t_0}^{+\infty} t^\alpha E(t)\mathrm{d}t < +\infty$, then $\lim_{t \to +\infty} t^{\alpha+1}E(t) = 0$.

This result follows from standard calculus arguments and is proved later in Appendix B. Observe that (8) already provides us with a first rate by applying Lemma 4.5 with $\alpha = 0$.

Proposition 4.6. Under Assumption 1 and 2, it holds that

$$E(t) = o\left(\frac{1}{t}\right).$$

Since E is a sum of non-negative quantities, any convergence rate of E translates into a rate for f, hence in particular Proposition 4.6 implies $f(x(t)) - f^\star = o\left(\frac{1}{t}\right)$. The proof of Theorem 1.1 follows similar steps as that of Theorem 4.3. The main idea is that by combining (8) with (2) we can improve the rate from $o\left(\frac{1}{t}\right)$ to $o\left(\frac{1}{t^{3/2}}\right)$ and iteratively repeat this process. This is stated in the following theorem.

Theorem 4.7. Under Assumptions 1 and 2 for all $0 < \varepsilon < \frac{1}{2}$ and all $n \in \mathbb{N}_{\geq 1}$ it holds that

$$E(t) = o\left(\frac{1}{t^{\alpha_n - \alpha_{n-1}\varepsilon}}\right), \text{ where } \alpha_n \stackrel{\text{def}}{=} 2 - \left(\frac{1}{2}\right)^{n-1}. \tag{9}$$

Before proving Theorem 4.7, we show that Theorem 1.1 is a direct consequence of this result.

Proof of Theorem 1.1. For any $0 < \delta < 1$, choose $0 < \varepsilon < \frac{\delta}{2}$ and observe that

$$\lim_{n\to+\infty} \alpha_n = 2, \text{ hence } \lim_{n\to+\infty} (\alpha_n - \alpha_{n-1}\varepsilon) = 2 - 2\varepsilon > 2 - \delta,$$

which shows that there exists $N \in \mathbb{N}_{\geq 1}$, such that $\alpha_N - \alpha_{N-1}\varepsilon > 2 - \delta$ and therefore by Theorem 4.7

$$0 = \lim_{t\to+\infty} t^{\alpha_N - \alpha_{N-1}\varepsilon} E(t) \geq \lim_{t\to+\infty} t^{2-\delta} E(t) \geq 0,$$

hence

$$\lim_{t\to+\infty} t^{2-\delta} E(t) = 0.$$

Finally the case $\delta \geq 1$ is covered by Proposition 4.6 and, as previously discussed, the following implication holds for all $\delta > 0$:

$$E(t) = o\left(\frac{1}{t^{2-\delta}}\right) \implies f(x(t)) - f^\star = o\left(\frac{1}{t^{2-\delta}}\right). \qquad \square$$

It now only remains to prove Theorem 4.7. We make use of the following lemma.

Lemma 4.8. Let g be continuous and non-negative on $[t_0, +\infty[$. Then for any $\beta > 1$,

$$\lim_{t\to+\infty} t^\beta g(t) \text{ exists and is finite } \implies \int_{t_0}^{+\infty} g(t)\mathrm{d}t < +\infty.$$

The proof of Lemma 4.8 is postponed to Appendix B

Proof of Theorem 4.7. Fix $0 < \varepsilon < \frac{1}{2}$. We show (9) by induction over $n \in \mathbb{N}_{>1}$. The case $n = 1$ holds from Proposition 4.6. Observe that for all $n \geq 1$, we have the following:

$$\alpha_n - 1 = 1 - \left(\frac{1}{2}\right)^{n-1} = \frac{1}{2}\left(2 - \left(\frac{1}{2}\right)^{n-2}\right) = \frac{\alpha_{n-1}}{2}. \tag{10}$$

Let us now assume that there exists $n \geq 1$ such that:

$$E(t) = o\left(\frac{1}{t^{\alpha_n - \alpha_{n-1}\varepsilon}}\right) \iff \sqrt{E(t)} = o\left(\frac{1}{t^{\frac{\alpha_n}{2} - \frac{\alpha_{n-1}}{2}\varepsilon}}\right). \tag{I.H.}$$

To proceed by induction we need to show that (I.H.) implies

$$E(t) = o\left(\frac{1}{t^{\alpha_{n+1} - \alpha_n\varepsilon}}\right). \tag{11}$$

9

To this aim, it is actually sufficient to prove that

$$\int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}E(t)\mathrm{d}t < +\infty. \tag{12}$$

Indeed, according to Lemma 4.5 if (12) holds, then

$$0 = \lim_{t\to+\infty} t^{\frac{\alpha_n}{2}+1-\alpha_n\varepsilon}E(t) \overset{(10)}{=} \lim_{t\to+\infty} t^{\alpha_n+1-\alpha_n\varepsilon}E(t),$$

which is equivalent to (11).

To show (12) we multiply both sides of (6) by $t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}$ and integrate from t_0 to $T \geq t_0$:

$$\int_{t_0}^{T} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}E(t)\mathrm{d}t \leq \frac{3}{2}\int_{t_0}^{T} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\left\|\dot{x}(t)\right\|^2\mathrm{d}t - \int_{t_0}^{T} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\ddot{h}_z(t)\mathrm{d}t$$

$$-\int_{t_0}^{T} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\sqrt{E(t)}\dot{h}_z(t)\mathrm{d}t. \tag{13}$$

Now it only remains to show that the limit as $T \to +\infty$ of each term on the right-hand side of (13) is uniformly bounded from above for $T \in [t_0, +\infty[$, which will imply that the left-hand side remains uniformly bounded as $T \to +\infty$. This approach is similar to the proof of Theorem 4.3, but the key equation (2) allows deducing further integrability results for $\left\|\dot{x}(t)\right\|^2$ by using the induction hypothesis (I.H.), which is new and specific to the system (LD).
We first make the following observation using (10):

$$\frac{t^{1+\varepsilon}}{t^{1-\frac{\alpha_n}{2}+\alpha_n\varepsilon}}\sqrt{E(t)} = t^{\frac{\alpha_n}{2}-\frac{\alpha_{n-1}}{2}\varepsilon}\sqrt{E(t)} \xrightarrow[t\to+\infty]{(\text{I.H.})} 0.$$

Therefore according to Definition 2, we have

$$\frac{1}{t^{1-\frac{\alpha_n}{2}+\alpha_n\varepsilon}}\sqrt{E(t)} = o\left(\frac{1}{t^{1+\varepsilon}}\right), \tag{14}$$

and using Lemma 4.8 with $g(t) = \frac{1}{t^{1-\frac{\alpha_n}{2}+\alpha_n\varepsilon}}\sqrt{E(t)}$ and $\beta = 1 + \varepsilon$,

$$\int_{t_0}^{+\infty} \frac{1}{t^{1-\frac{\alpha_n}{2}+\alpha_n\varepsilon}}\sqrt{E(t)}\mathrm{d}t < +\infty. \tag{15}$$

We now use (15) to show the integrability of $t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\left\|\dot{x}(t)\right\|^2$, which shows the uniform boundedness of the first term on the right-hand side of (13) as $T \to +\infty$. For any $T \geq t_0$:

$$\int_{t_0}^{T} t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\left\|\dot{x}(t)\right\|^2\mathrm{d}t \overset{I.P.,(2)}{=} \left[-2t^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\sqrt{E(t)}\right]_{t_0}^{T} - \int_{t_0}^{T}\left(\frac{\alpha_n}{2}-\alpha_n\varepsilon\right)t^{\frac{\alpha_n}{2}-1-\alpha_n\varepsilon}(-2\sqrt{E(t)})\mathrm{d}t$$

$$\overset{(10)}{=} -2T^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\sqrt{E(T)} + 2t_0^{\frac{\alpha_n}{2}-\alpha_n\varepsilon}\sqrt{E(t_0)} + 2\left(\frac{\alpha_n}{2}-\alpha_n\varepsilon\right)\int_{t_0}^{T}\frac{1}{t^{1-\frac{\alpha_n}{2}+\alpha_n\varepsilon}}\sqrt{E(t)}\mathrm{d}t,$$

10

where the second term is constant, the last integral is finite as $T \to +\infty$ by (15) and

$$-2 \lim_{T \to +\infty} T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \sqrt{E(T)} \leq 0.$$

Therefore the first term in (13) is bounded:

$$\int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \|\dot{x}(t)\|^2 \, \mathrm{d}t < +\infty. \tag{16}$$

Looking at the second term in (13) we have:

$$-\int_{t_0}^{T} t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \ddot{h}_z(t) \mathrm{d}t \overset{I.P.}{=} -\left[t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \dot{h}_z(t) \right]_{t_0}^{T} + \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \int_{t_0}^{T} t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} \dot{h}_z(t) \mathrm{d}t$$

$$\overset{I.P.}{=} -T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \dot{h}_z(T) + t_0^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \dot{h}_z(t_0) + \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \left[t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} h_z(t) \right]_{t_0}^{T}$$

$$- \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \left(\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon \right) \int_{t_0}^{T} t^{\frac{\alpha_n}{2} - 2 - \alpha_n \varepsilon} h_z(t) \mathrm{d}t. \tag{17}$$

Recall that there exists a $0 \leq M < +\infty$ such that $\forall t \geq t_0$ $h_z(t) \in [0, M]$, since by Assumption 2, x is bounded. Due to $-\underbrace{\left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right)}_{>0} \underbrace{\left(\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon \right)}_{<0} > 0$ and $\frac{\alpha_n}{2} - 2 - \alpha_n \varepsilon < -1$, we can bound the limit of the last term in (17) as $T \to +\infty$:

$$- \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \left(\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon \right) \int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2} - 2 - \alpha_n \varepsilon} h_z(t) \mathrm{d}t$$

$$\leq -M \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \left(\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon \right) \int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2} - 2 - \alpha_n \varepsilon} \mathrm{d}t < +\infty.$$

Again by boundedness of x and by the definition of $\dot{h}_z$ we get:

$$- \lim_{T \to +\infty} T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \dot{h}_z(T) \overset{C.S.}{\leq} \lim_{T \to +\infty} T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \|x(T) - z\| \, \|\dot{x}(T)\|$$

$$\leq \sqrt{M} \lim_{T \to +\infty} T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \|\dot{x}(T)\| = 0.$$

Indeed,

$$0 \overset{(I.H.)}{=} \lim_{t \to +\infty} t^{\alpha_n - \alpha_{n-1} \varepsilon} E(t) = \lim_{t \to +\infty} t^{\alpha_n - \alpha_{n-1} \varepsilon} \left(f(x(t)) - f^\star + \frac{1}{2} \|\dot{x}(t)\|^2 \right)$$

$$\geq \lim_{t \to +\infty} t^{\alpha_n - \alpha_{n-1} \varepsilon} \|\dot{x}(t)\|^2 \geq 0,$$

hence:

$$\lim_{t \to +\infty} t^{\alpha_n - \alpha_{n-1} \varepsilon} \|\dot{x}(t)\|^2 = 0 \iff \lim_{t \to +\infty} t^{\frac{\alpha_n}{2} - \frac{\alpha_{n-1}}{2} \varepsilon} \|\dot{x}(t)\| = 0$$

$$\Rightarrow \lim_{t \to +\infty} t^{\frac{\alpha_n}{2} - \frac{\alpha_{n-1}}{2} \varepsilon - \varepsilon} \|\dot{x}(t)\| = 0 \overset{(10)}{\iff} \lim_{t \to +\infty} t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \|\dot{x}(t)\| = 0.$$

11

To conclude that the limit $T \to +\infty$ of the right-hand side of (17) is bounded from above, we observe that two of the three remaining terms are constant and

$$\lim_{T \to +\infty} \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) T^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} h_z(T) = 0,$$

since $\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon < 0$ and h_z is bounded. We finish the proof by bounding the third term of (13). For $T \geq t_0$ we have:

$$-\int_{t_0}^T t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \sqrt{E(t)} \dot{h}_z(t) \mathrm{dt} \overset{I.P.,(2)}{=} - \left[t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \sqrt{E(t)} h_z(t) \right]_{t_0}^T$$

$$+ \int_{t_0}^T \left(\left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} \sqrt{E(t)} - \frac{1}{2} t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \left\| \dot{x}(t) \right\|^2 \right) h_z(t) \mathrm{dt}$$

$$= -T^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \sqrt{E(T)} h_z(T) + t_0^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \sqrt{E(t_0)} h_z(t_0) + \left(\frac{\alpha_n}{2} - \alpha_n \varepsilon \right) \int_{t_0}^T t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} \sqrt{E(t)} h_z(t) \mathrm{dt}$$

$$- \frac{1}{2} \int_{t_0}^T t^{\frac{\alpha_n}{2} - \alpha_n \varepsilon} \left\| \dot{x}(t) \right\|^2 h_z(t) dt. \quad (18)$$

Only the third term on the right-hand side of (18) is neither constant nor non-positive. Using again the boundedness of h_z, it holds that:

$$\int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} \sqrt{E(t)} h_z(t) \mathrm{dt} \leq M \int_{t_0}^{+\infty} t^{\frac{\alpha_n}{2} - 1 - \alpha_n \varepsilon} \sqrt{E(t)} \mathrm{dt} \overset{(15)}{<} +\infty.$$

Hence the limit as $T \to +\infty$ of the third term of (13) is bounded from above as well. Therefore all limits as $T \to +\infty$ of the terms on the right-hand side of (13) are bounded from above, hence (12) holds. By induction, we conclude that (9) holds for all $n \in \mathbb{N}_{\geq 1}$. $\qquad \square$

Remark 4.9. The induction step of Theorem 4.7 resembles the proof of [15, Lemma 3.10]. Yet, their result is not directly applicable to (LD) because it requires the knowledge of an asymptotic lower-bound on the damping, which we do not know a priori. Furthermore, Cabot and Frankel [15] showed how to iteratively repeat the application of this result but their proof relies on a strong integrability statement on $\| \dot{x}(t) \|$ that may not hold in general. The specific nature of our system (LD) allows us to deduce this integrability statement instead of assuming it.

4.4 Drawbacks of the approach

Now that we presented the main benefits of using a Lyapunov function as damping in (LD), we discuss some drawbacks of our approach, starting with the following remark.

Remark 4.10. Our system (LD) makes use of the optimal value $f^\star$. The latter might be unknown in general but is known in several practical cases such as over-determined regression with squared loss, or empirical risk minimization of some over-parameterized machine learning problems [18]. We refer to [14] for more examples and further discussion. All the results above do not require the knowledge of $f^\star$, only its existence. However Algorithm 1 presented in Section 5 is only practical when $f^\star$ is known.

The second drawback of our approach is that we can lower bound the convergence rate of our damping $\sqrt{E(t)}$. This not surprising as the bound actually matches the sufficient condition from [16] mentioned in the introduction.

Proposition 4.11. The convergence rate of the damping $\sqrt{E}$ of the dynamical system (LD) *cannot* be asymptotically faster than $\frac{1}{t}$, *i.e.*, there exists no $\alpha > 1$ such that $\sqrt{E(t)} = o\left(\frac{1}{t^\alpha}\right)$ for a non-constant bounded solution of (LD).

Remark 4.12. Proposition 4.11 expresses that E can never vanish faster than $o(\frac{1}{t^2})$. However, this does not mean that the convergence rate of $f(x(t)) - f^\star$ can never be faster than the worst-case $o(\frac{1}{t^2})$; one can sometimes obtain faster convergence of function values, as we will observe in the numerical experiments.

Proof of Proposition 4.11. Assume that $\sqrt{E(t)} = o\left(\frac{1}{t^\alpha}\right)$ for some $\alpha > 1$, then by Lemma 4.8,

$$\int_{t_0}^{\infty} \sqrt{E(t)} \mathrm{d}t < +\infty. \tag{19}$$

Let us make a similar computation as in [16, Prop. 3.5]:

$$\frac{\mathrm{d}E(t)}{\mathrm{d}t} + 2\sqrt{E(t)}E(t) \stackrel{(1)}{=} -\sqrt{E(t)}\,\|\dot{x}(t)\|^2 + 2\sqrt{E(t)}\left(f(x(t)) - f^\star + \frac{1}{2}\,\|\dot{x}(t)\|^2\right)$$

$$= 2\sqrt{E(t)}(f(x(t)) - f^\star) \geq 0. \tag{20}$$

Multiply (20) by $e^{2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s} > 0$ to obtain

$$0 \leq e^{2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s}\frac{\mathrm{d}E(t)}{\mathrm{d}t} + 2\sqrt{E(t)}e^{2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s}E(t) = \frac{\mathrm{d}}{\mathrm{d}t}\left[e^{2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s}E(t)\right]. \tag{21}$$

Now integrate 21 from t_0 to t:

$$0 \leq e^{2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s}E(t) - E(t_0)$$

$$\iff E(t) \geq e^{-2\int_{t_0}^{t}\sqrt{E(s)}\mathrm{d}s}E(t_0).$$

This implies that

$$\lim_{t\to\infty} E(t) \geq e^{-2\int_{t_0}^{\infty}\sqrt{E(s)}\mathrm{d}s}E(t_0) \stackrel{(19)}{>} 0,$$

which contradicts Proposition 4.3, since $\lim_{t\to\infty} E(t) = 0$. Therefore, we conclude that

$$\int_{t_0}^{\infty} \sqrt{E(t)}\mathrm{d}t = \infty, \tag{22}$$

and by Lemma 4.8 the limit of $t^\alpha\sqrt{E(t)}$ for $t \to +\infty$ does not exist or is not finite for any $\alpha > 1$, hence $E(t) \neq o\left(\frac{1}{t^\alpha}\right)$ for $\alpha > 1$. □

5 Algorithms and Numerical Experiments

5.1 Practical algorithms from (LD)

We first detail how we discretize (LD). We use an explicit discretization with fixed step-size $\sqrt{s} > 0$: for $k \in \mathbb{N}$ we approximate the solution x of (LD) at times $t_k = k\sqrt{s}$ and define $x_k \stackrel{\text{def}}{=} x(t_k)$. We use the approximations $\dot{x}(t) \approx \frac{x_k - x_{k-1}}{\sqrt{s}}$ and $\ddot{x}(t) \approx \frac{x_{k+1} - 2x_k + x_{k-1}}{s}$. We also define the discrete version of the damping accordingly by first defining

$$E_k \stackrel{\text{def}}{=} f(x_k) - f^\star + \frac{1}{2} \left\| \frac{x_k - x_{k-1}}{\sqrt{s}} \right\|^2. \tag{23}$$

and our damping then reads $\gamma(x_k, x_{k-1}) \stackrel{\text{def}}{=} \sqrt{E_k}$. Using this in (LD) we then propose the following discretization scheme for all $k \in \mathbb{N}$:

$$\frac{x_{k+1} - 2x_k + x_{k-1}}{s} + \gamma(x_k, x_{k-1})\frac{x_k - x_{k-1}}{\sqrt{s}} + \nabla f(y_k) = 0$$

$$\iff x_{k+1} = x_k + \left(1 - \sqrt{s}\gamma(x_k, x_{k-1})\right)\left[x_k - x_{k-1}\right] - s\nabla f(y_k), \tag{24}$$

where the gradient ∇f is evaluated at $y_k = x_k + (1 - \sqrt{s}\gamma(x_k, x_{k-1}))\left[x_k - x_{k-1}\right]$, in the same fashion as it is done in NAG. One can optionally rather evaluate the gradient at x_k.

We call the resulting algorithm LYDIA, for **LY**apunov **D**amped **I**nertial **A**lgorithm, which is summarized in Algorithm 1.

Algorithm 1: LYDIA

---
input: $x_0, x_{-1} \in \mathbb{R}^n$, step-size $s > 0$, $k_{\max} \in \mathbb{N}$

1 **for** $k = 1$ *to* k_{max} **do**

2 $\quad y_k \leftarrow x_k + \left(1 - \frac{\sqrt{E_k}}{\sqrt{E_0}}\right)\left[x_k - x_{k-1}\right]$

3 $\quad x_{k+1} \leftarrow y_k - s\nabla f(y_k)$

4 **end**

Note that compared to (24), we actually scaled the damping term $\sqrt{s}\sqrt{E_k}$ by $\sqrt{s}\sqrt{E_0}$ in Algorithm 1. This scaling is optional but improves numerical stability[1] as it ensures that the coefficient in front of $[x_k - x_{k-1}]$ remains non-negative. More specifically, we have the following lemma that shows that E_k is non-increasing under standard assumptions.

Lemma 5.1. Let f be a continuously differentiable convex functions whose gradient is Lipschitz continuous with constant $L > 0$. Let $(x_k, y_k)_{k\in\mathbb{N}}$ be the sequence generated by Algorithm 1, where $s < \frac{1}{L}$. Then, for all $k \in \mathbb{N}$,

$$E_{k+1} \le E_k - \frac{s}{2} \left\| \nabla f(x_k) - \nabla f(y_k) \right\|^2. \tag{25}$$

In particular $(E_k)_{k\in\mathbb{N}}$ is non-increasing, hence it is a "discrete" Lyapunov function for Algorithm 1.

Proof. Consider for any $x, y \in \mathcal{H}$ and $s \in (0, 1/L]$ the following refined version of the Descent Lemma [5, Remark 5.2.]:

$$f(y - s\nabla f(y)) \le f(x) + \langle \nabla f(y), y - x \rangle - \frac{s}{2} \left\| \nabla f(y) \right\|^2 - \frac{s}{2} \left\| \nabla f(x) - \nabla f(y) \right\|^2. \tag{26}$$

---
[1]An alternative to scaling is to choose s so that $\sqrt{s}\sqrt{E_0} \le 1$.

Take $y = y_k$ and $x = x_k$ in (26) and subtract $f^\star$ on both sides:

$$f(x_{k+1}) - f^\star \leq f(x_k) - f^\star + \langle \nabla f(y_k), y_k - x_k \rangle - \frac{s}{2} \|\nabla f(y_k)\|^2 - \frac{s}{2} \|\nabla f(x_k) - \nabla f(y_k)\|^2. \quad (27)$$

Observe that we have the following identities:

$$\begin{aligned}
\|x_{k+1} - x_k\|^2 &= \|s\nabla f(y_k) - (y_k - x_k)\|^2 \\
&= s^2 \|\nabla f(y_k)\|^2 - 2s \langle \nabla f(y_k), y_k - x_k \rangle + \|y_k - x_k\|^2, \quad \text{and,} \quad (28)
\end{aligned}$$

$$\|y_k - x_k\|^2 \overset{(1)}{=} \left\| \left(1 - \frac{\sqrt{E_k}}{\sqrt{E_0}}\right)(x_k - x_{k-1}) \right\|^2. \quad (29)$$

Therefore,

$$\begin{aligned}
\langle \nabla f(y_k), y_k - x_k \rangle &= \frac{s}{2} \|\nabla f(y_k)\|^2 - \frac{1}{2s} \|x_{k+1} - x_k\|^2 + \frac{1}{2s} \|y_k - x_k\|^2 \\
&= \frac{s}{2} \|\nabla f(y_k)\|^2 - \frac{1}{2s} \|x_{k+1} - x_k\|^2 + \frac{1}{2s} \left\| \left(1 - \frac{\sqrt{E_k}}{\sqrt{E_0}}\right)(x_k - x_{k-1}) \right\|^2.
\end{aligned}$$

Substitute this in (27) to get

$$\begin{aligned}
f(x_{k+1}) - f^\star + \frac{1}{2s} \|x_{k+1} - x_k\|^2 &\leq f(x_k) - f^\star + \frac{1}{2s} \left\| \left(1 - \frac{\sqrt{E_k}}{\sqrt{E_0}}\right)(x_k - x_{k-1}) \right\|^2 \\
&\quad - \frac{s}{2} \|\nabla f(x_k) - \nabla f(y_k)\|^2. \quad (30)
\end{aligned}$$

Observe that for the first iteration $k = 0$, (30) reads

$$\begin{aligned}
f(x_1) - f^\star &+ \frac{1}{2s} \|x_1 - x_0\|^2 \\
&\leq f(x_0) - f^\star + \frac{1}{2s} \left\| \left(1 - \frac{\sqrt{E_0}}{\sqrt{E_0}}\right)(x_0 - x_{-1}) \right\|^2 - \frac{s}{2} \|\nabla f(x_0) - \nabla f(y_0)\|^2
\end{aligned}$$

which is equivalent to $E_1 \leq f(x_0) - f^\star - \frac{s}{2} \|\nabla f(x_0) - \nabla f(y_0)\|^2$ and therefore

$$\begin{aligned}
E_1 &\leq f(x_0) - f^\star + \frac{1}{2s} \|x_0 - x_{-1}\|^2 - \frac{s}{2} \|\nabla f(x_0) - \nabla f(y_0)\|^2 \\
&= E_0 - \frac{s}{2} \|\nabla f(x_0) - \nabla f(y_0)\|^2.
\end{aligned}$$

So $E_1 \leq E_0$. Let $k \in \mathbb{N}_{>0}$ such that $\forall\, 1 \leq i \leq k$, $E_i \leq E_{i-1}$, then in particular $E_k \leq E_0$ and we deduce from (29) that:

$$\left\| \left(1 - \frac{\sqrt{E_k}}{\sqrt{E_0}}\right)(x_k - x_{k-1}) \right\|^2 \leq \|x_k - x_{k-1}\|^2, \quad (31)$$

which simplifies (30) into

$$f(x_{k+1}) - f^\star + \frac{1}{2s} \|x_{k+1} - x_k\|^2 \leq f(x_k) - f^\star + \frac{1}{2s} \|x_k - x_{k-1}\|^2 - \frac{s}{2} \|\nabla f(x_k) - \nabla f(y_k)\|^2.$$

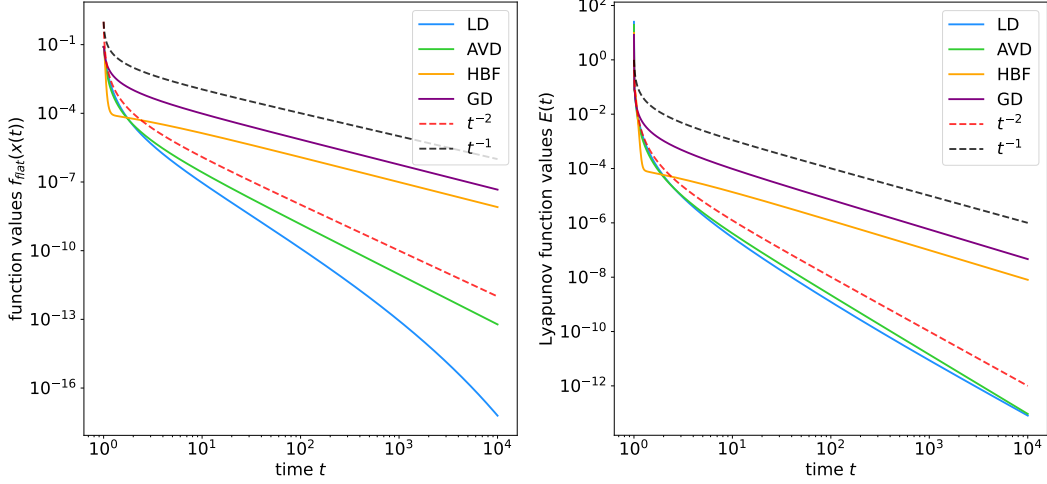This is exactly (25) after using (23), hence by induction, (25) holds for all $k \in \mathbb{N}$. $\qquad\square$

Figure 2: Rate of convergence of LD, AVD, HBF and GD on the test-function f_{flat}. Left: Evolution of the function values. Right: Evolution of the Lyapunov function.
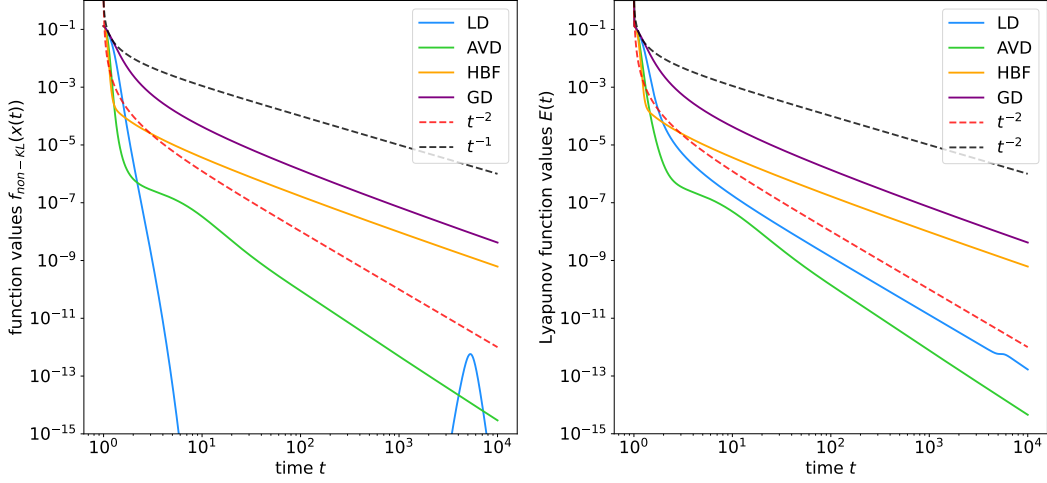

Figure 3: Rate of convergence of LD, AVD, HBF and GD on the test-function $f_{\text{non-KL}}$. Left: Evolution of the function values. Right: Evolution of the Lyapunov function.

5.2 Numerical experiments

Methodology. We evaluate the performance of our system and empirically corroborate Theorem 1.1. We consider the following 1-dimensional test functions, defined for all $x \in \mathbb{R}$ by:

$$f_{\text{flat}}(x) \stackrel{\text{def}}{=} x^{24}, \qquad f_{\text{non-KL}}(x) \stackrel{\text{def}}{=} \begin{cases} 0, & \text{if } x = 0; \\ \exp\left(-\frac{1}{x^2}\right), & \text{else}, \end{cases}$$

$$f_{\text{uneven}}(x) \stackrel{\text{def}}{=} \begin{cases} x^3, & \text{if } x > 0; \\ x^2, & \text{if } x \leq 0, \end{cases} \qquad f_{\text{contmin}}(x) \stackrel{\text{def}}{=} \begin{cases} (x - \varepsilon)^2, & \text{if } x > \varepsilon; \\ (x + \varepsilon)^2, & \text{if } x < \varepsilon; \\ 0, & \text{else}. \end{cases}$$
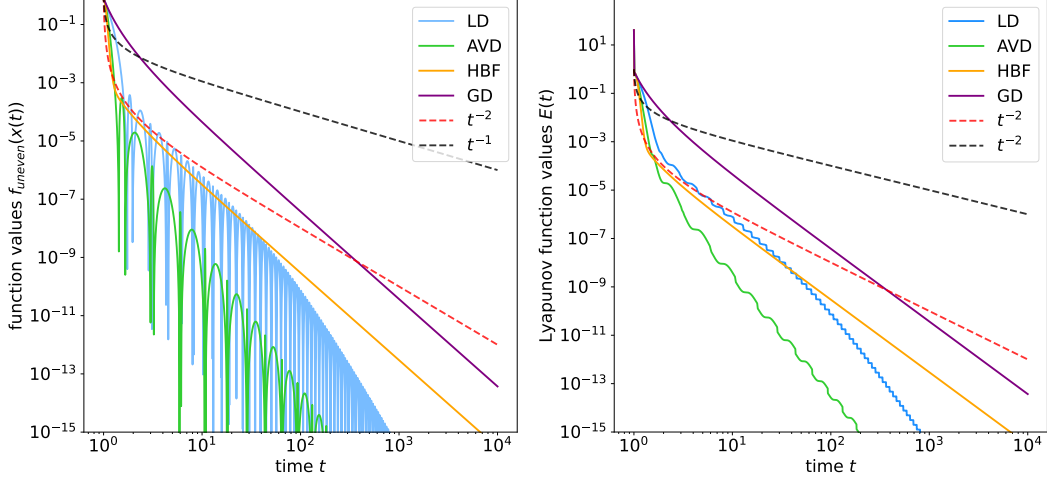
16

Figure 4: Rate of convergence of LD, AVD, HBF and GD on the test-function f_{uneven}. Left: Evolution of the function values. Right: Evolution of the Lyapunov function.
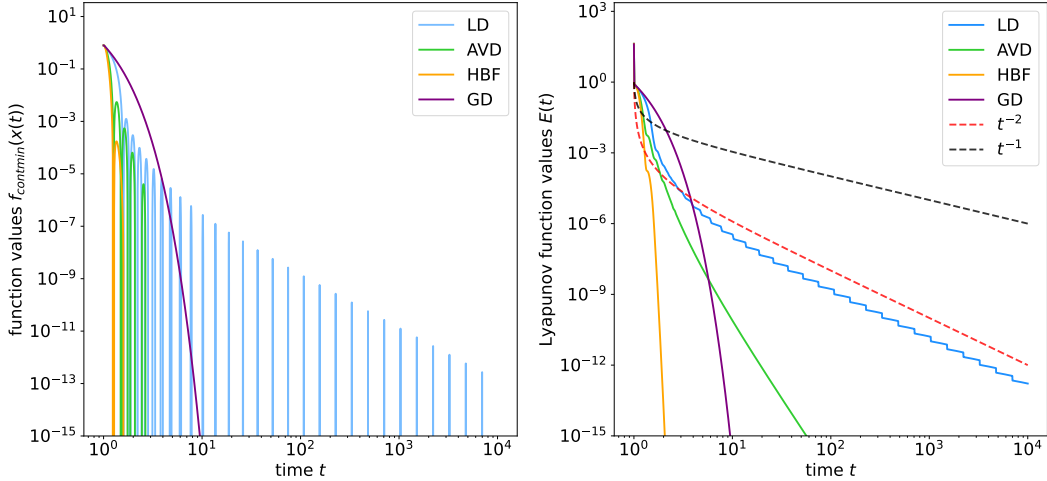

Figure 5: Rate of convergence of LD, AVD, HBF and GD on the test-function f_{contmin}. Left: Evolution of the function values. Right: Evolution of the Lyapunov function.

Since (LD) and (AVD$_a$) are second-order ODEs that do not have closed-form solutions in general, we approximate them by using the LYDIA algorithm and NAG respectively, both with small step-sizes. We take $a = 3.1$ in NAG. We also consider Gradient Descent (GD) and the HBF algorithms mentioned in the introduction. For each test function[2], we present the evolution of the optimality gap $f(x(t)) - f^\star$ as a function of the time t. Since this value is not monotonous and may oscillate heavily, we also display the evolution of the (discretized) Lyapunov function E.

Results. The function f_{flat} is a convex polynomial that is very flat around its minimizer. As discussed in [6, Example 2.12] such polynomials of large degree allow emphasizing the worst-case bounds on the rates of convergence. Indeed, we observe on the right of Figure 2 that the rate for GD and HBF

---

[2]Some of these functions do not have a globally Lipschitz continuous gradient which could cause numerical instabilities. We overcome this by using small-enough step-sizes to ensure boundedness of the iterates of all the algorithms considered.

is asymptotically close to $\mathcal{O}\left(\frac{1}{t}\right)$ while that of LD and (AVD$_a$) gets close to $\mathcal{O}\left(\frac{1}{t^2}\right)$ which empirically validates our main result Theorem 1.1.

Some convex functions are not strongly convex but possess the KL property which allows deriving faster rates of convergence (as done in the closed-loop setting by [9]). Since Theorem 1.1 does not assume such a property, we evaluate the algorithms on a function that is known for not having the KL property on Figure 3. The results again match the theoretical ones. Similarly, properties such as f being even (*i.e.*, $f(-x) = f(x),\ \forall x$) or $\operatorname{argmin}_{\mathcal{H}} f$ being a singleton can allow one to derive faster convergence results than for general convex functions. Our last-two test functions in Figures 4 and 5 illustrate that the rate of (LD) also holds without such properties.

Finally, while Figures 2 and 3 evidence cases where the solution of (LD) yields very fast minimization of function values, Figure 4 shows a case where (LD) is faster than GD and HBF but not as fast as (AVD$_a$) and Figure 5 shows a case where (LD) is behind all other methods. Nonetheless, in all cases observe that E is non-increasing and that it exhibits the rate predicted by Theorem 1.1.

6 Conclusions and Perspectives

We introduced a new system (LD) which, as initially intended, is independent of the choice of the initial time t_0, unlike (AVD$_a$), and which does require the hyper-parameter $a > 3$ compared to the latter, but however assumes knowledge of the optimal value $f^\star$. We showed that it is possible to design closed-loop dampings that achieve near-optimal rates of convergence for convex optimization. The key ingredient is the coupling between the damping and the speed of convergence of the system. This yielded in particular the identity (1), specific to (LD) and essential to adapt the proofs from the open-loop to the closed-loop setting. This coupling provides a new understanding on how to choose the damping term in second-order ODEs for optimization which may prove to be useful beyond the specific cases of (LD) and (AVD$_a$). We provided numerical experiment corroborating our theoretical results as well as a new practical first-order algorithm LYDIA, derived from (LD), and for which we showed that E a still a (discrete) Lyapunov function.

As for future work, it remains open to know whether we can improve the rate from "arbitrarily close to optimal" to optimal (*i.e.*, can we take $\delta = 0$ in Theorem 1.1?). We also suspect that one can drop Assumption 2-(ii) as it is not required for (AVD$_a$) [24, 21], but we could not yet do it in the closed-loop setting. On the algorithmic side, our experiments suggest that the rate of (LD) is transferred to the LYDIA algorithm, which remains to be shown. One might also consider other discretization schemes. Finally, it is worth investigating whether one can find systems and algorithms with the same properties that do not require the optimal value $f^\star$.

Acknowledgment

Appendices

A Proof of Existence and Uniqueness

This section is devoted to proving Theorem 3.1. We recall some result from the theory of ODEs.

Theorem A.1. [29, Theorem 4.2.6] (Classical Version of Picard–Lindelöf)
Let $\mathcal{H}$ be a Hilbert space, let $\Omega \subseteq \mathbb{R} \times \mathcal{H}$ be open. Let $G : \Omega \to \mathcal{H}$ be continuous and let $(t_0, x_0) \in \Omega$. Assume that there exists $L \geq 0$ such that for all $(t, x), (t, y) \in \Omega$ we have

$$\|G(t, x) - G(t, y)\| \leq L \, \|x - y\| \, .$$

Then, there exists $\delta > 0$ such that the initial value problem

$$\begin{cases} u'(t) = G(t, u(t)) & (t \in (t_0, t_0 + \delta)), \\ u(t_0) = x_0, \end{cases}$$

admits a unique continuously differentiable solution $u : [t_0, t_0 + \delta[\to \mathcal{H}$, which satisfies $(t, u(t)) \in \Omega$ for all $t \in [t_0, t_0 + \delta[$.

Definition 3. [26, Definition 3.20] Let $\phi : I \to \mathbb{R}^n$, $n \in \mathbb{N}$, be a solution to an ODE defined on some open interval $I \subset \mathbb{R}$.

(a) We say ϕ admits an extension or continuation to the right (resp. to the left) if, and only if, there exists a solution $\psi : J \to \mathbb{R}^n$ to the same ODE, defined on some open interval $J \supset I$ such that $\psi|_I = \phi$ and

$$\sup J > \sup I \quad (\text{resp. } \inf J < \inf I)$$

An extension or continuation of ϕ is a function that is an extension to the right or an extension to the left (or both).

(b) The solution ϕ is called a *maximal* solution if, and only if, it does not admit any extension in the sense defined in (a).

Theorem A.2. [26, Theorem 3.22] Every solution $\phi_0 : I_0 \to \mathbb{R}^n$ to some ODE, defined on an open interval $I_0 \subset \mathbb{R}$, can be extended to a maximal solution of this ODE.

Remark A.3. Definition 3 is easily generalized to Hilbert spaces and Theorem A.2 is valid for an ODE in an infinite dimensional space as well, since the proof relies on Zorn's Lemma.

Proof of Theorem 3.1. We first reduce the system (LD) to an equivalent first-order system:

$$\frac{\mathrm{d}}{\mathrm{d}t} \begin{pmatrix} x(t) \\ \dot{x}(t) \end{pmatrix} = \begin{pmatrix} \dot{x}(t) \\ -\sqrt{f(x(t)) - f^\star + \frac{1}{2} \|\dot{x}(t)\|^2} \dot{x}(t) - \nabla f(x(t)) \end{pmatrix} \overset{\mathrm{def}}{=} G(t, (x(t), \dot{x}(t))). \qquad (32)$$

We first show the claim in the case $\mathrm{argmin}_{\mathcal{H}} f \neq \emptyset$ and later for the case $\mathrm{argmin}_{\mathcal{H}} f = \emptyset$.

First, if $(x_0, \dot{x}_0) \in \mathrm{argmin}_{\mathcal{H}} f \times \{0\}$, then $(x(t), \dot{x}(t)) \equiv (x_0, 0)$ is the unique global solution, since $\ddot{x}(t) = 0$ for all $t \geq t_0$. Now assume that $(x_0, \dot{x}_0) \notin \mathrm{argmin}_{\mathcal{H}} f \times \{0\}$ and let us check that for all $t \geq t_0$, the function

$$(u, v) \mapsto G(t, (u, v)) = \left(\begin{array}{c} v \\ -\sqrt{f(u) - f^\star + \frac{1}{2} \|v\|^2} v(t) - \nabla f(u) \end{array} \right)$$

is Lipschitz continuous on the open set

$$\mathcal{A} \overset{\mathrm{def}}{=} (\mathcal{H} \times \mathcal{H}) \setminus (\mathrm{argmin}_{\mathcal{H}} f \times \{0\}).$$

The first coordinate of G is clearly locally Lipschitz continuous w.r.t. (u, v) and therefore it suffices to check that the second coordinate of G is also locally Lipschitz continuous w.r.t. (u, v) on $\mathcal{A}$. The square root function is locally Lipschitz continuous everywhere except at 0 and $f(u) - f^\star + \frac{1}{2} \|v\|^2 \neq 0$ for all $(u, v) \in \mathcal{A}$. Therefore $(u, v) \in \mathcal{A} \mapsto \sqrt{f(u) - f^\star + \frac{1}{2} \|v\|^2} v$ is locally Lipschitz continuous as a product of locally Lipschitz-continuous functions. Finally ∇f is locally Lipschitz continuous by Assumption 1-(i) , hence G is locally Lipschitz continuous on $\mathcal{A}$. We can now apply the Picard–Lindelöf Theorem A.1 to any open and bounded neighborhood of the initial conditions $(t_0, (x_0, \dot{x}_0))$ to find some $\delta > 0$ and a unique solution $(x, \dot{x})$ to (32) on $[t_0, t_0 + \delta[$.

It now remains to show that the solution exists on $[t_0, +\infty[$. Let $T \in]t_0, +\infty]$ be the maximal time of existence (as stated in Theorem A.2) of the solution $(x, \dot{x})$ of (32). Assume that $T < +\infty$, then the limits of x and $\dot{x}(t)$ exist as $t \to T$. Indeed, first note that the uniform boundedness of E implies that $\dot{x}$ is bounded on $[t_0, T[$ and by (LD) $\ddot{x}$ is bounded on $[t_0, T[$ as well. Now consider a Cauchy sequence $(t_n)_{n \in \mathbb{N}}$ converging to T and let $m, n \in \mathbb{N}$. Then by the mean value Theorem we have:

$$\|x(t_n) - x(t_m)\| \leq \sup_{s \in]t_0, T[} \|\dot{x}(s)\| \, |t_n - t_m| , \text{ and}$$

$$\|\dot{x}(t_n) - \dot{x}(t_m)\| \leq \sup_{s \in]t_0, T[} \|\ddot{x}(s)\| \, |t_n - t_m| .$$

Therefore $(x(t_n))_{n \in \mathbb{N}}$ and $(\dot{x}(t_n))_{n \in \mathbb{N}}$ are Cauchy sequences which implies that $\lim_{t \to T} x(t)$ and $\lim_{t \to T} \dot{x}(t)$ exist.

Then, if $(x(T), \dot{x}(T)) \in \mathcal{A}$ then we can extend the solution by applying Picard–Lindelöf to an open and bounded neighborhood of $((T, (x(T), \dot{x}(T)))$ which contradicts the fact that T was the maximal time of existence. If $(x(T), \dot{x}(T)) \in \partial \mathcal{A} \subset \mathrm{argmin}_{\mathcal{H}} f \times \{0\}$, where $\partial \mathcal{A}$ denotes the boundary of $\mathcal{A}$, then $\ddot{x}(T) = 0$ and we can extend the solution on $[T, T + \delta[$, for some $\delta > 0$, by taking $(x(t), \dot{x}(t)) = (x(T), 0)$. This is again a contradiction, therefore we necessarily have $T = +\infty$.

Finally, in the case where $\mathrm{argmin}_{\mathcal{H}} f = \emptyset$, we see similarly to before that $G(t, (u, v))$ is locally Lipschitz continuous w.r.t. (u, v) everywhere on $\mathcal{H} \times \mathcal{H}$. So the Picard–Lindelöf Theorem implies the existence of a unique local solution for any initial value, which we can again extend to a global solution on $[t_0, +\infty[$. $\qquad \square$

B Missing Proofs of Lemmas

Proof of Lemma 4.4. Assume that there exists $c > 0$ such that $\lim_{t \to +\infty} g(t) = c$, which implies the existence of $t_1 \geq t_0$ such that $\forall t \geq t_1$, $g(t) \geq \frac{c}{2}$. Then

$$\int_{t_0}^{+\infty} g(t)\mathrm{d}t \geq \int_{t_1}^{+\infty} g(t)\mathrm{d}t \geq \int_{t_1}^{+\infty} \frac{c}{2}\mathrm{d}t = +\infty,$$

which contradicts the assumption that $\int_{t_0}^{+\infty} g(t)\mathrm{d}t < +\infty$. □

Proof of Lemma 4.5. Denote the positive and negative parts of a function g by

$$[g(\cdot)]_+ \overset{\text{def}}{=} \max(g(\cdot), 0) \text{ and } [g(\cdot)]_- \overset{\text{def}}{=} \max(-g(\cdot), 0) \geq 0,$$

so that $\forall x \in \mathcal{H}$,

$$g(x) = [g(x)]_+ - [g(x)]_- .$$

Then note that

$$\left[\left(t^{\alpha+1} E(t) \right)' \right]_+ \overset{(1)}{=} \left[(\alpha + 1) t^\alpha E(t) - t^{\alpha+1} \sqrt{E(t)} \, \|\dot{x}(t)\|^2 \right]_+ = (\alpha + 1) t^\alpha E(t), \tag{33}$$

and that,

$$\begin{aligned}
\int_{t_0}^{+\infty} \left(t^{\alpha+1} E(t) \right)' \mathrm{d}t &= \int_{t_0}^{+\infty} \left[\left(t^{\alpha+1} E(t) \right)' \right]_+ - \left[\left(t^{\alpha+1} E(t) \right)' \right]_- \mathrm{d}t \\
&\leq \int_{t_0}^{+\infty} \left[\left(t^{\alpha+1} E(t) \right)' \right]_+ \mathrm{d}t \overset{(33)}{=} (\alpha + 1) \int_{t_0}^{+\infty} t^\alpha E(t)\mathrm{d}t \overset{Ass.}{<} +\infty.
\end{aligned} \tag{34}$$

Therefore by definition of the improper integral, $\lim_{t \to +\infty} t^{\alpha+1} E(t)$ admits a finite limit $l \in \mathbb{R}_{\geq 0}$.

Now, assume that $l > 0$, then for any $\varepsilon \in]0, l[$ there exists $t_1 \geq t_0$ such for all $t \geq t_1$ we have $E(t) > \frac{l - \varepsilon}{t^{\alpha+1}}$. Then,

$$\int_{t_0}^{+\infty} t^\alpha E(t)\mathrm{d}t \geq \int_{t_1}^{+\infty} t^\alpha E(t)\mathrm{d}t > \int_{t_1}^{+\infty} t^\alpha \frac{l - \varepsilon}{t^{\alpha+1}}\mathrm{d}t = (l - \varepsilon) \int_{t_1}^{+\infty} \frac{1}{t}\mathrm{d}t = +\infty,$$

which contradicts the integrability assumption on $t^\alpha E(t)$. □

Proof of Lemma 4.8. Let $l \overset{\text{def}}{=} \lim_{t \to \infty} t^\beta g(t) \geq 0$. Then for any $\varepsilon > 0$ we can find $t_1 \geq t_0$ such that $\forall t \geq t_1$,

$$g(t) \leq \frac{(l + \varepsilon)}{t^\beta} .$$

Therefore,

$$\int_{t_0}^{+\infty} g(t)\mathrm{d}t = \underbrace{\int_{t_0}^{t_1} g(t)\mathrm{d}t}_{\overset{\text{def}}{=} I < +\infty} + \int_{t_1}^{+\infty} g(t)\mathrm{d}t \leq I + (l + \varepsilon) \int_{t_1}^{+\infty} \frac{1}{t^\beta}\mathrm{d}t < +\infty,$$

since $\beta > 1$, which proves the result. □

References

[1] Samir Adly, Hedy Attouch, and Alexandre Cabot. Finite time stabilization of nonlinear oscillators subject to dry friction. In *Nonsmooth Mechanics and Analysis*, pages 289–304, 2006.

[2] Cristian Daniel Alecsa, Szilárd Csaba László, and Titus Pinţa. An extension of the second order dynamical system that models Nesterov's convex gradient method. *Applied Mathematics & Optimization*, 84:1687–1716, 2021.

[3] Felipe Alvarez, Hedy Attouch, Jérôme Bolte, and Patrick Redont. A second-order gradient-like dissipative dynamical system with Hessian-driven damping: Application to optimization and mechanics. *Journal de mathématiques pures et appliquées*, 81(8):747–779, 2002.

[4] Hedy Attouch and Alexandre Cabot. Asymptotic stabilization of inertial gradient dynamics with time-dependent viscosity. *Journal of Differential Equations*, 263(9):5412–5458, 2017.

[5] Hedy Attouch and Jalal Fadili. From the Ravine method to the Nesterov method and vice versa: a dynamical system perspective. *SIAM Journal on Optimization*, 32(3):2074–2101, 2022.

[6] Hedy Attouch, Zaki Chbani, Juan Peypouquet, and Patrick Redont. Fast convergence of inertial dynamics and algorithms with asymptotic vanishing damping. *Mathematical Programming*, 168: 123–175, 2015.

[7] Hedy Attouch, Juan Peypouquet, and Patrick Redont. Fast convex optimization via inertial dynamics with Hessian driven damping. *Journal of Differential Equations*, 261(10):5734–5783, 2016.

[8] Hedy Attouch, Zaki Chbani, and Hassan Riahi. Rate of convergence of the Nesterov accelerated gradient method in the subcritical case $\alpha \leq 3$. *ESAIM: Control, Optimisation and Calculus of Variations*, 25:2, 2019.

[9] Hedy Attouch, Radu Ioan Boţ, and Ernö Robert Csetnek. Fast optimization via inertial dynamics with closed-loop damping. *Journal of the European Mathematical Society*, 25(5):1985–2056, 2022.

[10] Hedy Attouch, Radu Ioan Boţ, and Dang-Khoa Nguyen. Fast convex optimization via closed-loop time scaling of gradient dynamics. *arXiv preprint arXiv:2301.00701*, 2023.

[11] Jean-François Aujol and Charles Dossal. Optimal rate of convergence of an ODE associated to the fast gradient descent schemes for $b > 0$. *HAL Preprint hal-01547251*, 2017.

[12] Jean-Francois Aujol, Charles Dossal, and Aude Rondepierre. Optimal convergence rates for Nesterov acceleration. *SIAM Journal on Optimization*, 29(4):3131–3153, 2019.

[13] Michel Benaïm, Josef Hofbauer, and Sylvain Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 2005.

[14] Stephen Boyd, Lin Xiao, and Almir Mutapcic. Subgradient methods. *Lecture notes of EE392o, Stanford University*, 2003.

[15] Alexandre Cabot and Pierre Frankel. Asymptotics for some semilinear hyperbolic equations with non-autonomous damping. *Journal of Differential Equations*, 252(1):294–322, 2012.

[16] Alexandre Cabot, Hans Engler, and Sébastien Gadat. On the long time behavior of second order differential equations with asymptotically small dissipation. *Transactions of the American Mathematical Society*, 361(11):5983–6017, 2009.

[17] Antonin Chambolle and Charles Dossal. On the convergence of the iterates of the "fast iterative shrinkage/thresholding algorithm". *Journal of Optimization theory and Applications*, 166:968–982, 2015.

[18] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[19] David M Grobman. Homeomorphism of systems of differential equations. *Doklady Akademii Nauk SSSR*, 128(5):880–881, 1959.

[20] Philip Hartman. A lemma in the theory of structural stability of differential equations. *Proceedings of the American Mathematical Society*, 11(4):610–620, 1960.

[21] Mohamed Ali Jendoubi and Ramzi May. Asymptotics for a second-order differential equation with nonautonomous damping and an integrable source term. *Applicable Analysis*, 94(2):435–443, 2015.

[22] Tianyi Lin and Michael I Jordan. A control-theoretic perspective on optimal high-order optimization. *Mathematical Programming*, 195(1):929–975, 2022.

[23] Lennart Ljung. Analysis of recursive stochastic algorithms. *IEEE transactions on automatic control*, 22(4):551–575, 1977.

[24] Ramzi May. Asymptotic for a second-order evolution equation with convex potential and vanishing damping term. *Turkish Journal of Mathematics*, 41(3):681–685, 2017.

[25] Yurii Nesterov. A method of solving a convex programming problem with convergence rate $O\left(\frac{1}{k^2}\right)$. In *Doklady Akademii Nauk*, volume 269(3), pages 543–547, 1983.

[26] Peter Philip. Ordinary differential equations, lecture notes, 2023. URL `https://www.math.lmu.de/~philip/publications/lectureNotes/philipPeter_ODE.pdf`. [Online; accessed 17.07.2023].

[27] Boris Polyak. *Introduction to optimization*. New York, Optimization Software, 1987.

[28] B.T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.

[29] Christian Seifert, Sascha Trostorff, and Marcus Waurick. *Evolutionary Equations: Picard's Theorem for Partial Differential Equations, and Applications*. Springer Nature, 2022.

[30] Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *Mathematical Programming*, 195(1):79–148, 2021.

[31] Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling Nesterov's accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27, pages 1–9. Curran Associates, Inc., 2014.

[32] Andre Wibisono, Ashia C Wilson, and Michael I Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences*, 113(47):E7351–E7358, 2016.

[33] Ashia C Wilson, Ben Recht, and Michael I Jordan. A Lyapunov analysis of accelerated methods in optimization. *Journal of Machine Learning Research*, 22(1):5040–5073, 2021.