On the stability of optimization algorithms given by discretizations of the Euler-Lagrange ODE

Rachel Walker
Central Washington University
Ellensburg, WA

Emily Zhang
Massachusetts Institute of Technology
Cambridge, MA

***Abstract*—We study discretizations of an Euler-Lagrange equation which generate a large class of accelerated methods whose convergence rate is $O(\frac{1}{t^p})$ in continuous-time, where parameter p is the order of the optimization method. Specifically, we address the question asking why a naive explicit-implicit Euler discretization of this solution produces an unstable algorithm, even for a strongly convex objective function. We prove that for a strongly convex L-smooth quadratic objective function and step size $\delta < \frac{1}{L}$, the naive discretization will exhibit stable behavior when the number of iterations k satisfies the inequality $k < (\frac{4}{Lp^2\delta^p})^{\frac{1}{p-2}}$.**

I. Introduction

The phenomenon of acceleration is currently a heavily researched topic in convex optimization. Su et. al. first explored the concept of taking continuous time limits of optimization methods in an attempt to better understand acceleration [5]. More recently, high resolution continuous-time ODEs have been derived, and they shine light on how gradient correction leads to a faster convergence rate [3]. This new perspective has also motivated the use of various discretization schemes on continuous-time problems to generate new families of optimization algorithms [4].

Wibisono et. al. derived a second order Euler-Lagrange ODE whose solution minimizes an objective function f at an exponential rate with order p, for any distance generating function. When attempting to discretize this ODE, the authors found that the system of two update equations given by an explicit-implicit Euler discretization of the ODE initially converges to the minimizer of f, oscillates around the minimizer, and eventually diverges as shown in Figure (1). This occurs even for strongly convex quadratic functions and a Euclidian distance generating function. The reason for divergence here is unclear. In their work, the authors solve the problem of instability by introducing a rate matching discretization, which utilizes a third update sequence. However, as evident for the $p = 3$ case, the implementation of this third sequence is difficult [2].

Recently, there have been explorations of the convergence rates and stability of various discretization methods. We adopt ideas from numerical analysis and recent work on stabilizing gradient descent to analyze the end behavior of the explicit and implicit Euler method applied to the Euler-Lagrange ODE [1]. We hope that this work will provide better insight into the behavior of the explicit-implicit discretization scheme and also provide easier-to-implement alternatives to the rate matching discretization.

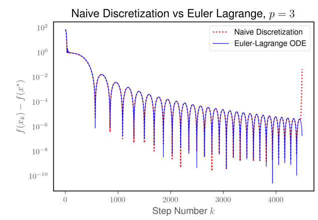

Fig. 1: The explicit-implicit Euler method (naive discretization) eventually diverges

A. Problem Setting

Throughout this paper, we consider the optimization problem

$$x^* = \arg\min_{x \in \mathbb{R}^d} f(x), \tag{I.1}$$

where $f(x) = \frac{1}{2}(x - x^*)^T A(x - x^*)$ is a convex function $f : \mathbb{R}^d \to \mathbb{R}$ with some unique minimizer x^* that satisfies the optimality condition $\nabla f(x^*) = Ax^* = 0$, and A is a symmetric $d \times d$ matrix. We mainly focus on the quadratic objective function as the linear gradient allows for easier analysis. The importance of minimizing quadratic objective functions has many applications in machine learning, such as a least squares loss function for a neural network.

B. A Continuous Time Solution

Wibisono et. al. derived the following Euler-Lagrange ODE, whose solution minimizes f at an exponential rate [6]. When in the Euclidean setting, this ODE is

$$\ddot{X}_t + \frac{p+1}{t}\dot{X}_t + Cp^2 t^{p-2} \nabla f(X_t) = 0 \tag{I.2}$$

where $X_t := X(t)$, $C > 0$ is a constant, and $p \geq 2$ is the parameter which describes the order of the optimization method. Let X^* be the minimizer of the objective function f. The authors show that in continuous time, (I.2) has the convergence rate of

$$f(X_t) - f(X^*) < O\left(\frac{1}{t^p}\right). \tag{I.3}$$

If $p = 2$, equation (I.2) is the continuous time limit of Nesterov's method derived by Su et. al. [5]. When $p = 3$, equation (I.2) is the Euclidean case of the continuous time limit of cubic-regularized Newtons method [2]. Due to the order p exponential convergence rate of this Euler-Lagrange ODE, it is of interest to derive a discretization of this ODE with a convergence rate that matches the one given in (I.3). However, as mentioned previously, the explicit-implicit discretization eventually diverges even for a strongly convex quadratic objective function. Note that (I.2) can also be written as a system of two first order ODEs

$$\begin{aligned} \dot{X}_t &= \frac{p}{t}(Z_t - X_t) \\ \dot{Z}_t &= -Cpt^{p-1}\nabla f(X_t). \end{aligned} \tag{I.4}$$

II. Our Approach

In this section, we describe the approach that we take to analyze the behavior of various discretizations of the Euler-Lagrange ODE. We are primarily interested in determining whether a system of update equations given by a certain discretization scheme has converging, diverging, or stable long-term behavior. To be precise, we give the following definitions. In this paper, a system of update equations given by a discretization method is

(a) **converging** to the minimizer if the *upper bound* on $|x_k - x^*|$ is decreasing as k increases, [1]
(b) **diverging** from the minimizer if the *upper bound* on $|x_k - x^*|$ is increasing as k increases, and
(c) **stable** if, for sufficiently large N, $|x_k - x^*| = |x_{k+1} - x^*|$ for all $k > N$.

Oftentimes, an optimization problem has very large dimensions; that is x, the value that we are updating, is multi-dimensional. For our purposes, however, the analysis of a system of update equations on one-dimensional x is sufficient to study the behavior of a certain discretization scheme applied to the Euler-Lagrange ODE. This is stated more formally and proved in the following proposition.

Proposition II.1. *In order to study the stability of update equations derived from various discretization methods, we can focus on cases where x is one-dimensional without loss of generality.*

Proof. We rewrite $f(x)$, a general objective function where x is d-dimensional and A is symmetric, as follows:

$$\begin{aligned} f(x) &= \frac{1}{2}(x - x^*)^T A (x - x^*) \\ &= \frac{1}{2}(x - x^*)^T PDP^T (x - x^*) \\ &= \frac{1}{2}(P^T(x - x^*))^T DP^T (x - x^*) \\ &= \frac{1}{2}\tilde{x}^T D \tilde{x} \end{aligned}$$

where $\tilde{x} := P^T(x - x^*)$, P is the matrix of eigenvectors of A, and D is the diagonal matrix of eigenvalues of A.

Since all dimensions of $\tilde{x}$ update independently of each other, the case where $\tilde{x}$ and x are one-dimensional is without loss of generality. □

[1]Note that $|x_k - x^*|$ does not have to be a monotonically decreasing sequence in order to be converging.

We now make several definitions which help us set up the framework that we will use to analyze various discretizations of the Euler-Lagrange ODE. We let $u_i := \begin{bmatrix} \tilde{x_i} \\ z_i \end{bmatrix}$, where $\tilde{x}$ is defined as in the proof of Proposition 2.1, and consider discretizations of the Euler-Lagrange ODE of the form

$$u_{k+1} = M_k u_k. \tag{II.1}$$

Additionally, we define $M_\infty := \lim_{k\to\infty} M_k$ and $u_\infty := \lim_{k\to\infty} u_k$. Finally, we define the stability function, which tells us the end behavior of systems of update equations in the form given by equation (II.1). Proposition (II.2) shows how the stability function determines end behavior.

Definition II.1. *We define $R(M_k) := |\lambda_{k,\max}|$ where $\lambda_{k,\max}$ is the eigenvalue of M_k with the largest magnitude, and the **stability function** is given by*

$$R(M_\infty) = \lim_{k\to\infty} R(M_k)$$

Proposition II.2. *A discretization method will be*

(a) converging to the minimizer when $R(M_\infty) < 1$.
(b) stable when $R(M_\infty) = 1$.

Proof. Computing u_k from u_0, we have

$$u_k = M_{k-1} M_{k-2} \dots M_1 M_0 u_0. \tag{II.2}$$

When all the eigenvalues of M_i have magnitude less than 1, then $\|u_i\| < \|u_{i-1}\|$. Since $\|\tilde{x_i}\| \leq \|u_i\|$, the upper bound on $\|\tilde{x_i}\|$ is also strictly decreasing when all eigenvalues' magnitudes are less than 1. Letting k go to ∞ proves part (a) of the proposition.

When $R(M_\infty) = 1$, the part of x_∞ that lies along the eigenvector of M_∞ associated with the eigenvalue(s) equal to 1 will always remain the same size. Parts of x_∞ that lie along other eigenvector(s) will go to 0. Thus, the value of x_k for sufficently large k will not change, and the iterations are stable. □

III. Explicit-Implicit Euler Method

The update equations given by the discretization of (I.4) using a explicit-implicit Euler method and the identification $t = \delta k$ are

$$\begin{aligned} \frac{x_{k+1} - x_k}{\delta} &= \frac{p}{t}(z_k - x_k) \\ \frac{z_k - z_{k-1}}{\delta} &= -Cpt^{p-1}\nabla f(x_k). \end{aligned} \tag{III.1}$$

This set of update equations eventually diverges after approaching and oscillating around the minimizer, yet it is unknown why this occurs [6]. We present the following theorem, which describes the behavior of the explicit-implicit discretization, and an outline of our proof.

Theorem III.1. *Let $f(x) : \mathbb{R}^d \to \mathbb{R}$ be an L-smooth function defined as*

$$f(x) = \frac{1}{2}(x - x^*)^T A (x - x^*) \tag{III.2}$$

where $x^* \in \mathbb{R}^d$ is the unique minimizer with $\nabla f(x^*) = \vec{0}$ and A is a positive definite, symmetric $d \times d$ matrix. Let $\delta < \frac{1}{L}$ and $\epsilon = \delta^p$. Then, after we go out enough iterations in the system of update equations given by equation $(III.1)$ such that $k > p$ and take $C < \frac{1}{\epsilon L}$, we have the following properties:

(a) If $p = 2$, the naive method exhibits stable end behavior.

(b) If $p > 2$, the naive method will exhibit stable behavior when

$$k < \left(\frac{4}{CLp^2\epsilon}\right)^{\frac{1}{p-2}}.$$

Proof. In Proposition II.1, we showed that the case where x and $\tilde{x}$ are one-dimensional case is without loss of generality. Thus, we begin by considering the problem $f(x) = \frac{1}{2}A(x - x^*)^2 = \frac{1}{2}A(\tilde{x})^2$ where x and $\tilde{x}$ are one-dimensional, and later generalize our results to d-dimensional x. We now define the update equations as done in II.1 where

$$M_k = \begin{bmatrix} 1 - \frac{p}{k} & \frac{p}{k} \\ -Cp\epsilon(k+1)^{p-1}(\frac{k-p}{k})A & 1 - Cp\epsilon(k+1)^{p-1}(\frac{p}{k})A \end{bmatrix}.$$

Next, we analyze the end behavior of the this algorithm for various p by looking at the eigenvalues of $M_\infty = \lim_{k\to\infty} M_k$ for one-dimensional x and determining the stability function.

Case 1, $p = 2$. We solve for the eigenvalues of M_∞ by setting the characteristic polynomial of this matrix equal to 0. In the characteristic equation, we omit terms that go to 0 as $k \to \infty$. We have

$$\begin{aligned} 0 &= \det(M_\infty - \lambda I) \\ &= \lambda^2 + \lambda(4C\epsilon A - 2) + 1. \end{aligned}$$

Now, let $c = 4C\epsilon A > 0$. Since we make the assumption that $C < \frac{1}{\epsilon L}$, we have that $C < \frac{1}{\epsilon A}$ for a one dimensional problem. Thus, we have $c < 4$. This gives the following eigenvalues:

$$\begin{aligned} \lambda_1 &= \frac{-c + 2 + \sqrt{c^2 - 4c}}{2} = \frac{2-c}{2} + \frac{\sqrt{4c - c^2}i}{2} \\ \lambda_2 &= \frac{-c + 2 - \sqrt{c^2 - 4c}}{2} = \frac{2-c}{2} - \frac{\sqrt{4c - c^2}i}{2}. \end{aligned}$$

Because $|\lambda_1| = |\lambda_2| = 1$, the stability function $R(M_\infty) = 1$. Thus, $\begin{bmatrix}\tilde{x}_\infty \\ z_\infty\end{bmatrix}$ will be stable when $p = 2$.

Case 2, $p > 2$. We solve for the eigenvalues of M_∞, once again omitting terms in the characteristic equation that go to 0 as $k \to \infty$. We have

$$\begin{aligned} 0 &= \det(M_\infty - \lambda I) \\ &= \lim_{k\to\infty} \left(\lambda^2 + \lambda(Cp^2\epsilon k^{p-2}A - 2) + 1\right) \\ &= \lambda^2 + \lambda(c-2) + 1, \end{aligned}$$

where $c = \lim_{k\to\infty} Cp^2\epsilon k^{p-2}A$. Thus we have the eigenvalues

$$\lambda_1 = \frac{-c + 2 + \sqrt{c^2 - 4c}}{2},\ \lambda_2 = \frac{-c + 2 - \sqrt{c^2 - 4c}}{2}.$$

Note that $R(M_\infty) = |\lambda_2| \gg 1$. Thus, the solution for $\begin{bmatrix}\tilde{x}_\infty \\ z_\infty\end{bmatrix}$ is unstable. In order to determine on which iteration the explicit-implicit method starts to diverge for each $p > 2$, we find the eigenvalues of M_k. Let $a_k = \frac{p}{k}$ and let $b_k = Cp\epsilon(k+1)^{p-1}A$. This gives us

$$M_k = \begin{bmatrix} 1 - a_k & a_k \\ -b_k + a_k b_k & 1 - a_k b_k \end{bmatrix}.$$

The characteristic equation for M_k is

$$\lambda^2 + \lambda(a_k b_k + a_k - 2) + (1 - a_k) = 0$$

which gives the eigenvalues

$$\begin{aligned} \lambda_1 &= \frac{-a_k b_k - a_k + 2 + \sqrt{(a_k b_k + a_k - 2)^2 - 4(1 - a_k)}}{2} \\ \lambda_2 &= \frac{-a_k b_k - a_k + 2 - \sqrt{(a_k b_k + a_k - 2)^2 - 4(1 - a_k)}}{2}. \end{aligned}$$

In order to determine which of the eigenvalues describe the divergence, we make the following claims.

Claim 1. After going out enough iterations such that $k > p$, we never have divergence when $|\lambda_1| = |\lambda_2|$.

Proof of Claim 1. We only have $|\lambda_1| = |\lambda_2|$ when the eigenvalues are complex or when the eigenvalues are the same real value. That is $|\lambda_1| = |\lambda_2|$ implies

$$(a_k b_k + a_k - 2)^2 - 4(1 - a_k) \le 0. \qquad \text{(III.3)}$$

When (III.3) is true, the eigenvalues can be written as

$$\begin{aligned} \lambda_1 &= \frac{-a_k b_k - a_k + 2}{2} + \frac{\sqrt{4(1 - a_k) - (a_k b_k + a_k - 2)^2}}{2}i \\ \lambda_2 &= \frac{-a_k b_k - a_k + 2}{2} - \frac{\sqrt{4(1 - a_k) - (a_k b_k + a_k - 2)^2}}{2}i. \end{aligned}$$

Thus, the magnitudes of these eigenvalues are equivalent. Simplifying the magnitude, we get

$$\begin{aligned} |\lambda_1| = |\lambda_2| &= \sqrt{1 - a_k} \\ &= \sqrt{1 - \frac{p}{k}} \\ &< 1. \end{aligned}$$

□

Claim 2. We never have divergence when $|\lambda_1| > |\lambda_2|$.

Proof of Claim 2. In order to have $|\lambda_1| > |\lambda_2|$, we must have $-a_k b_k - a_k + 2 > 0$. Suppose for the sake of contradiction that we have $|\lambda_1| > 1$ which implies $-a_k b_k - a_k + 2 > 0$ and $\lambda_1 > 0$. Then we have

$$\begin{aligned} |\lambda_1| &> 1 \\ (a_k b_k + a_k - 2)^2 - 4(1 - a_k) &> (a_k b_k + a_k)^2 \\ -4a_k b_k &> 0. \qquad \Rightarrow\Leftarrow \end{aligned}$$

□

By Claim 1 and Claim 2, we can only have divergence when $|\lambda_2| > |\lambda_1| \implies -a_k b_k - a_k + 2 < 0$. Thus, it is enough to look at the magnitude of λ_2, when it is real, to determine when divergence happens. Therefore we have divergence when $|\lambda_2| > 1$. Through algebraic simplification, we can equivalently say that divergence occurs when

$$a_k b_k + a_k + \sqrt{(a_k b_k + a_k - 2)^2 - 4(1 - a_k)} > 4. \quad \text{(III.4)}$$

Our goal is to get an expression that determines the number of iterations k allowed for a given p, ϵ, and A. To do so, we begin with the inequality (III.4). Note that we can rewrite this as

$$a_k b_k + a_k + \sqrt{(a_k b_k + a_k)^2 - 4 a_k b_k} > 4.$$

Now let $x = a_k b_k$ and $y = a_k b_k + a_k = x + a_k$. We have that the iterations of the update equation will be converging or stable when

$$y + \sqrt{y^2 - 4x} \leq 4.$$

To simplify this inequality, we consider a right triangle with hypotenuse y and sidelengths $s_1 = 2\sqrt{x}$ and $s_2 = \sqrt{y^2 - s_1^2} = \sqrt{y^2 - 4x}$. A visual representation of this triangle is shown below. Our inequality for when convergence or stability is achieved becomes

$$y + s_2 \leq 4.$$

By the Triangle Inequality, $y + s_2 \leq 4$ implies $s_1 < 4$. Thus, we have that stability or convergence is achieved when

$$\begin{aligned} s_1 &< 4 \\ a_k b_k &< 4 \\ \frac{p}{k} C p \epsilon (k+1)^{p-1} A &< 4 \\ k &< \left(\frac{4}{C A p^2 \epsilon} \right)^{\frac{1}{p-2}}. \end{aligned}$$

Finally, we generalize the one-dimensional result to a d-dimensional problem. Consider the following problem where x and $\tilde{x}$ are d-dimensional vectors. Written as in Proposition II.1, we have

$$f(x) = \frac{1}{2}(x - x^*)^T A (x - x^*) = \frac{1}{2} \tilde{x}^T D \tilde{x}.$$

This d-dimensional problem will be converging to the minimizer or stable when each of its dimensions are doing so. Thus, iterations are converging or stable when k satisfies

$$k < \left(\frac{4}{C D_i p^2 \epsilon} \right)^{\frac{1}{p-2}}$$

for all integer i in the range 1 to d, inclusive. From this, it is easy to see that the largest eigenvalue of A dictates when the iterations become unstable. Thus, if $f(x)$ is L-smooth, the explicit-implicit method will exhibit stable behavior when

$$k < \left(\frac{4}{C L p^2 \epsilon} \right)^{\frac{1}{p-2}}.$$

□

IV. Numerical Results and Discussion

For the case where $\delta = .01, L = 10, p = 4$, Theorem III.1 says that the explicit-implicit method should be stable when $k < 1,582$. Figure (2) below visualizes the effect of this bound and compares the convergence rate to Nesterov's accelerated gradient descent for convex functions.

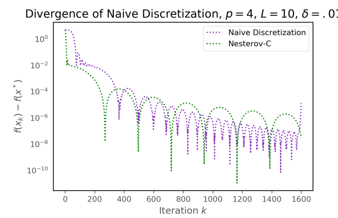

Fig. 2: Naive discretization diverging at predicted bound

In addition, we have identified several possible future directions to take. Empirically, we see that a fourth-order Runge-Kutta discretization of the Euler-Lagrange ODE is able to run for more iterations than the explicit-implicit method before it begins to diverge. For this reason, it would be of interest to use some of the approaches discussed in this paper to bound the number of iterations of guaranteed stable behavior. Furthermore, while we showed where the explicit-implicit method is converging, we have not showed that this convergence rate matches that of the Euler-Lagrange ODE. Showing that each of the discretization methods discussed in this paper achieves the $O\left(\frac{1}{t^p}\right)$ convergence rate before they diverge could result in a more useful algorithm. We also note that our current analysis restricts the objective function to be quadratic and is only analyzed in the Euclidean setting. It would be of interest to expand our analysis to a more general context.

V. Acknowledgements

This research was conducted as part of the 2019 REU program at Georgia Institute of Technology and was supported by NSF grant DMS1851843. We would like to thank our adviser Professor Rachel A. Kuske for her guidance and Dr. Andre Wibisono for his help during the research.

References

[1] Armin Eftekhari, Bart Vandereycken, Gilles Vilmart, and Konstantinos C Zygalakis. Explicit stabilised gradient descent for faster strongly convex optimisation. *arXiv preprint arXiv:1805.07199*, 2018.

[2] Yu Nesterov. Accelerating the cubic regularization of newtons method on convex problems. *Mathematical Programming*, 112(1):159–181, 2008.

[3] Bin Shi, Simon S Du, Michael I Jordan, and Weijie J Su. Understanding the acceleration phenomenon via high-resolution differential equations. *arXiv preprint arXiv:1810.08907*, 2018.

[4] Bin Shi, Simon S Du, Weijie J Su, and Michael I Jordan. Acceleration via symplectic discretization of high-resolution differential equations. *arXiv preprint arXiv:1902.03694*, 2019.

[5] Weijie Su, Stephen Boyd, and Emmanuel Candes. A differential equation for modeling nesterovs accelerated gradient method: Theory and insights. In *Advances in Neural Information Processing Systems*, pages 2510–2518, 2014.

[6] 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.