

Fast Monte Carlo Analysis for 6-DoF Powered-Descent Guidance via GPU-Accelerated Sequential Convex Programming

Govind M. Chari*, Abhinav G. Kamath*, Purnanand Elango*, and Behçet Açıkmeşe†
University of Washington, Seattle, WA 98195, USA

We introduce a GPU-accelerated Monte Carlo framework for nonconvex, free-final-time trajectory optimization problems. This framework makes use of the prox-linear method, which belongs to the larger family of sequential convex programming (SCP) algorithms, in conjunction with a constraint reformulation that guarantees inter-sample constraint satisfaction. Key features of this framework are: (1) continuous-time constraint satisfaction; (2) a matrix-inverse-free solution method; (3) the use of the proportional-integral projected gradient (PIPG) method, a first-order convex optimization solver, customized to the convex subproblem at hand; and, (4) an end-to-end, library-free implementation of the algorithm. We demonstrate this GPU-based framework on the 6-DoF powered-descent guidance problem, and show that it is faster than an equivalent serial CPU implementation for Monte Carlo simulations with over 1000 runs. To the best of our knowledge, this is the first GPU-based implementation of a general-purpose nonconvex trajectory optimization solver.

Notation

$\mathbb{R}^n = \mathbb{R}^{n \times 1}$	Vectors are matrices with one column
$1_n, 0_n$	Vector of ones and zeros, respectively, in \mathbb{R}^n (subscript inferred whenever omitted)
$0_{n \times m}$	Matrix of zeros in $\mathbb{R}^{n \times m}$
I_n	Identity matrix in $\mathbb{R}^{n \times n}$
$V_{[1:n]}$	Collection of vectors or matrices V_k , for $k = 1, \dots, n$
$ v _+$	$\square \mapsto \max\{0, \square\}$ applied element-wise for vector v
v^2	$\square \mapsto \square^2$ applied element-wise for vector v
(u, v)	Concatenation of vectors $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^m$ to form a vector in \mathbb{R}^{n+m}
$[A \ B]$	Concatenation of matrices A and B with same number of rows
$\begin{bmatrix} A \\ B \end{bmatrix}$	Concatenation of matrices A and B with same number of columns

I. Introduction

Monte Carlo analysis is critical for certifying the safety and performance of guidance, navigation, and control (GNC) systems [1]. For a 6-DoF powered-descent guidance algorithm, Monte Carlo analysis would involve generating hundreds or even thousands of dispersed initial conditions around the nominal initial condition, computing trajectories with these initial conditions, and assessing statistics such as solve-times, percentage of converged trajectories, and propellant consumption. However, generating thousands of trajectories serially can be time consuming, and can lead to increased development times. Since Monte Carlo simulation is *embarrassingly parallel* [2], we propose performing Monte Carlo analyses on a graphics processing unit (GPU) to exploit its parallel computing capability.

Due to their parallel computing capability and high memory bandwidth, GPUs have become increasingly popular in fields that require solving large-scale problems (in which parallelism can be exploited). Most notably, GPUs have become a staple in machine learning for training deep learning models. For example, OpenAI has a supercomputer with 10,000 GPUs to train its large language models [3], and Tesla, reportedly, also has a 10,000-GPU cluster to perform training for its machine learning-based full self-driving technology [4]. Additionally, many modern machine learning tools such as JAX, TensorFlow, and PyTorch, provide programmers with user-friendly, high-level interfaces to leverage the compute power of GPUs [5–7].

*Ph.D. Student, William E. Boeing Department of Aeronautics & Astronautics; {gchari, agkamath, pelango}@uw.edu

†Professor, William E. Boeing Department of Aeronautics & Astronautics; behcet@uw.edu

GPUs have also been used outside of machine learning for a variety of tasks. For example, [8] uses a GPU for running projectile simulations, i.e., propagating the equations of motion for a projectile, in parallel. In [9], the authors develop a framework to perform trajectory optimization on GPUs; however, they consider a restrictive trajectory optimization problem template that does not model dynamics or generic nonconvex path constraints. The only constraints considered are maximum speed, maximum acceleration, and keepout zone constraints. As a result, the generated trajectories are not guaranteed to obey the equations of motion of the vehicle.

In this paper, we demonstrate a GPU-based parallelized implementation for Monte Carlo analysis of trajectory optimization problems with general nonlinear dynamics and nonconvex path constraints. The solutions obtained are guaranteed to be feasible with respect to the original nonlinear dynamics, and all the imposed constraints are guaranteed to be satisfied at all points in time, i.e., in continuous-time (subject to successful convergence and up to numerical integration and optimization tolerances). This final guarantee on continuous-time constraint satisfaction is a result of the problem formulation proposed in [10]. To demonstrate this framework, we consider the 6-DoF powered-descent guidance algorithm [11].

The propellant-optimal powered-descent guidance (PDG) problem is concerned with finding a feasible sequence of control inputs that allows a vehicle to decelerate itself from some initial position and velocity, to perform a soft landing at some specified target on the ground, while satisfying a variety of state and control constraints. This problem, in general, is of great importance to companies like SpaceX that land their boosters propulsively, allowing them to make launches cheaper [12], and for NASA JPL, to land their rovers near important scientific targets on Mars [13], for example. Solving the 3-DoF PDG problem, where only translational dynamics are considered, is one approach to solving this problem. However, since this method does not explicitly consider attitude dynamics, attitude constraints, such as line-of-sight, cannot be enforced directly. For applications such as NASA’s Safe and Precise Landing — Integrated Capabilities Evolution (SPLICE) project, the terrain relative navigation (TRN) cameras and hazard detection LiDAR (HDL) require constraining the attitude of the rocket, necessitating the use of 6-DoF PDG algorithms [14].

General trajectory optimization problems, such as the 6-DoF PDG problem, can be posed as nonconvex optimization problems in continuous-time [15]. Given the nonconvexity of these problems, it is typically prohibitively expensive, if not outright impossible, to obtain globally optimal trajectories. As a result, we must turn to solution procedures, such as sequential convex programming (SCP) [16], that yield locally optimal solutions. At a high level, SCP involves (1) convexification (often via linearization) of all nonconvex constraints about some reference trajectory to yield an infinite-dimensional convex optimization problem; (2) discretization, to convert the problem to a numerically tractable finite-dimensional convex optimization problem; and finally (3) finding a solution to this finite-dimensional convex subproblem to update the reference trajectory. This process is applied iteratively until convergence. As the name suggests, this process is inherently sequential in nature. However, we can leverage parallelism to simultaneously solve for many trajectories in Monte Carlo analysis.

There are a number of factors that make developing a GPU-based parallelized implementation of SCP challenging. SCP requires the use of many linear algebra operations and the solution to a sequence of convex optimization subproblems. For a CPU-based implementation, we can use linear algebra packages such as OpenBLAS [17] and open-source convex optimization solvers such as OSQP and ECOS [18, 19]. However, OpenBLAS binaries are compiled for CPU architectures and NVIDIA’s linear algebra libraries, cuBLAS and cuSPARSE, are written such that functions from these libraries must be called by the CPU, even though the computations are performed on the GPU. This is undesirable for our implementation of SCP, since we require all computations for a single trajectory to be performed by a single GPU thread. Similarly, although a GPU-accelerated implementation of OSQP exists, the solver itself must be called by the CPU [20]. Although off-the-shelf convex optimization solvers could be modified to work on the GPU, custom parser frameworks are typically required—to put the convex subproblem in the canonical form prior to being passed to the solver—as well [21, 22]. Thus, for our GPU-based implementation of SCP, we implement Level 1, 2, and 3 BLAS operations for discretization and other such operations, and the first-order proportional-integral projected gradient (PIPG) solver [23, 24], customized to the trajectory optimization problem template [25, 26], allowing us to bypass parsing entirely. Further, as in the sequential conic optimization (SeCO) framework described in [26], this framework completely eliminates the need for sparse linear algebra operations, is entirely matrix-factorization/inversion-free, and is implemented using a library-free codebase.

II. Problem Formulation

We consider a continuous-time dynamical system of the form:

$$\frac{d\xi(t)}{dt} = \dot{\xi}(t) = F(\xi(t), \zeta(t)) \quad (1)$$

for $t \in [0, t_f]$, with state $\xi \in \mathbb{R}^{n_\xi}$ and control input $\zeta \in \mathbb{R}^{n_\zeta}$, which are subject to path constraints:

$$g(\xi(t), \zeta(t)) \leq 0_{n_g} \quad (2a)$$

$$h(\xi(t), \zeta(t)) = 0_{n_h} \quad (2b)$$

where g and h are vector-valued functions. The operators “ \leq ” and “ $=$ ” are interpreted element-wise. Boundary conditions at the initial and final times could also be specified.

The goal of trajectory optimization is to generate a solution $\xi(t), \zeta(t)$ for (1) that satisfies the path constraints, boundary conditions, and optimizes a linear function of the terminal state $\xi(t_f)$, i.e., solve the following Mayer form optimal control problem [27, Sec. 3.3.2]:

$$\underset{t_f, \xi, \zeta}{\text{minimize}} \quad \xi(t_f)^\top e_{\text{cost}} \quad (3a)$$

$$\text{subject to} \quad \dot{\xi}(t) = F(\xi(t), \zeta(t)) \quad t \in [0, t_f] \quad (3b)$$

$$g(\xi(t), \zeta(t)) \leq 0 \quad t \in [0, t_f] \quad (3c)$$

$$h(\xi(t), \zeta(t)) = 0 \quad t \in [0, t_f] \quad (3d)$$

$$E_i \xi(0) = z_i, \quad E_f \xi(t_f) = z_f \quad (3e)$$

where $E_i \in \mathbb{R}^{n_i \times n_\xi}$ and $E_f \in \mathbb{R}^{n_f \times n_\xi}$ select components of the initial and final state, respectively.

A. Time-Dilation and Constraint Reformulation

When the final time, t_f , is an unknown, we adopt a well-known technique called time-dilation [10, 24, 28] to convert a free-final-time problem to an equivalent fixed-final-time problem:

$$\frac{d\xi(t(\tau))}{d\tau} = \overset{\circ}{\xi}(\tau) = \frac{d\xi(\tau)}{dt} s(\tau) = s(\tau) F(\xi(\tau), \zeta(\tau)) \quad (4)$$

for $\tau \in [0, 1]$, where $t(\tau)$ is a strictly increasing function with domain $[0, 1]$. Note that ξ and ζ are re-defined to be functions of $\tau \in [0, 1]$ instead of t . The dilation factor:

$$s(\tau) = \frac{dt(\tau)}{d\tau}$$

is treated as an additional control input and is required to be positive. The control input is then re-defined to be $u = (\zeta, s)$, and the final time is given by:

$$t_f = \int_0^1 s(\tau) d\tau$$

Next, to address the long-standing issue of inter-sample constraint violation in direct methods [29], we augment the path constraints to the time-dilated system dynamics (4) as follows:

$$\overset{\circ}{x}(\tau) = \begin{bmatrix} \overset{\circ}{\xi}(\tau) \\ \overset{\circ}{y}(\tau) \end{bmatrix} = s(\tau) \begin{bmatrix} F(\xi(\tau), \zeta(\tau)) \\ 1^\top |g(\xi(\tau), \zeta(\tau))|_+^2 + 1^\top h(\xi(\tau), \zeta(\tau))^2 \end{bmatrix} = f(x(\tau), u(\tau)) \quad (5)$$

for $\tau \in [0, 1]$, where $x = (\xi, y)$ is the re-defined state. Periodic boundary condition on the constraint violation integrator:

$$y(0) = y(1) \quad (6)$$

ensure that the path constraints are satisfied for all $\tau \in [0, 1]$. We refer the reader to [10, 30] for further details.

B. Discretization

Discretization is a crucial step in direct numerical optimal control for computational tractability, i.e., for obtaining an approximate finite-dimensional optimization problem from the original infinite-dimensional optimal control problem. The first step is to generate a grid with N nodes within the interval $[0, 1]$:

$$0 = \tau_1 < \dots < \tau_N = 1 \quad (7)$$

The time duration over sub-interval $[\tau_k, \tau_{k+1}]$ is given by:

$$\Delta t_k = \int_{\tau_k}^{\tau_{k+1}} s(\tau) d\tau$$

for $k = 1, \dots, N - 1$, and the following constraints are specified on the dilation factor and time duration:

$$\begin{aligned} s_{\min} &\leq s(\tau) \leq s_{\max} \\ t_f &= \sum_{k=1}^{N-1} \Delta t_k \leq t_{f,\max} \\ \Delta t_{\min} &\leq \Delta t_k \leq \Delta t_{\max} \end{aligned}$$

Next, we parameterize the control input with a first-order-hold (FOH):

$$u(\tau) = \left(\frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \right) u_k + \left(\frac{\tau - \tau_k}{\tau_{k+1} - \tau_k} \right) u_{k+1}$$

for $\tau \in [\tau_k, \tau_{k+1}]$ and $k = 1, \dots, N - 1$. In practice, the combination of time-dilation and FOH parameterization for control input is sufficient for accurately modeling a large class of optimal control problems. See [1] for a detailed discussion on the relative merits of popular discretization and parameterization techniques in numerical optimal control.

We treat the node-point values of the state and control input, x_k and u_k , respectively, as decision variables, and re-write the system dynamics (5) equivalently as:

$$x_{k+1} = x_k + \int_{\tau_k}^{\tau_{k+1}} f(x(\tau), u(\tau)) d\tau \quad (8)$$

for $k = 1, \dots, N - 1$, where $x(\tau)$ is the solution to (5) over $[\tau_k, \tau_{k+1}]$ with initial condition x_k , and $u(\tau)$ is FOH-parameterized using u_k and u_{k+1} . Then, (8) can be compactly expressed through function $f_k : \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ as follows:

$$x_{k+1} = f_k(x_k, u_k, u_{k+1}) \quad (9)$$

The boundary condition (6) is relaxed to:

$$y_{k+1} - y_k \leq \epsilon$$

for $k = 1, \dots, N - 1$ and a sufficiently small $\epsilon > 0$, so that all feasible solutions do not violate LICQ [10]. Note that y_k is the value of the state y in (5) at node τ_k .

The discretized optimal control problem with parameterized control input is given by:

$$\underset{x_k, u_k}{\text{minimize}} \quad x_N^\top \tilde{e}_{\text{cost}} \quad (10a)$$

$$\text{subject to} \quad x_{k+1} = f_k(x_k, u_k, u_{k+1}) \quad k = 1, \dots, N - 1 \quad (10b)$$

$$E_y(x_{k+1} - x_k) \leq \epsilon \quad k = 1, \dots, N - 1 \quad (10c)$$

$$u_{\min} \leq u_k \leq u_{\max} \quad k = 1, \dots, N \quad (10d)$$

$$\tilde{E}_i x_1 = (z_i, 0), \quad \tilde{E}_f x_N = z_f \quad (10e)$$

where

$$\tilde{e}_{\text{cost}} = (e_{\text{cost}}, 0) \quad (11)$$

$$E_y = [0_{1 \times n_\xi} \quad 1] \quad (12)$$

$$\tilde{E}_i = \begin{bmatrix} E_i & 0_{n_\xi} \\ 0_{1 \times n_\xi} & 1 \end{bmatrix} \quad (13)$$

$$\tilde{E}_f = [E_f \quad 0_{n_\xi}] \quad (14)$$

C. Sequential Convex Programming

We adopt the prox-linear method [31], an SCP algorithm, to solve:

$$\underset{x_k, u_k}{\text{minimize}} \quad x_N^\top \tilde{e}_{\text{cost}} + w_{\text{ep}} \sum_{k=1}^{N-1} \|x_{k+1} - f_k(x_k, u_k, u_{k+1})\|_1 \quad (15a)$$

$$\text{subject to} \quad E_y(x_{k+1} - x_k) \leq \epsilon \quad (15b)$$

$$u_{\min} \leq u_k \leq u_{\max} \quad (15c)$$

$$\tilde{E}_i x_1 = (z_i, 0), \quad \tilde{E}_f x_N = z_f \quad (15d)$$

where the (only) nonconvex constraint of (10) is exactly penalized with ℓ_1 -norm. For a large enough, finite w_{ep} , we can show that a solution of (15) that is feasible with respect to (10b) is also a solution of (10). See [32, Chap. 17] for a detailed discussion.

The penalized problem (15) can be compactly stated as:

$$\underset{Z \in \mathcal{Z}}{\text{minimize}} \quad L(Z) + w_{\text{ep}} \|G(Z)\|_1 \quad (16)$$

where $Z = (x_1, \dots, x_N, u_1, \dots, u_N)$ and \mathcal{Z} is a convex set corresponding to constraints (15b), (15c), and (15d).

The prox-linear method determines a stationary point of (16) by solving a sequence of convex subproblems. At iteration $j + 1$, it solves the following convex problem:

$$\underset{Z \in \mathcal{Z}}{\text{minimize}} \quad L(Z) + w_{\text{ep}} \|G(Z^j) + \nabla G(Z^j)(Z - Z^j)\|_1 + \frac{w_{\text{prox}}}{2} \|Z - Z^j\|_2^2 \quad (17)$$

where Z^j is the solution from iteration j . The sequence Z^j can be shown to converge for an appropriate choice of the proximal term weight w_{prox} . Further, if the converged point is feasible for (10), then it is also a KKT point.

There are two important steps in the construction of each convex subproblem—(1) computation of ∇G : the gradient of f_k with respect to the previous iterate, and (2) prescaling and preconditioning.

1. Gradient of Discretized Dynamics

At each iteration of SCP, the discretized dynamics constraint (10b) is linearized with respect to the previous iterate, denoted by \bar{x}_k, \bar{u}_k , for $k = 1, \dots, N$.

For each $k = 1, \dots, N - 1$, let $\bar{x}^k(\tau)$ denote the solution to (5) over $[\tau_k, \tau_{k+1}]$ generated with initial condition \bar{x}_k and control input $\bar{u}(\tau)$, which is parameterized using \bar{u}_k , for $k = 1, \dots, N$. Note that, in general, $\bar{x}^k(\tau_{k+1}) \neq \bar{x}_{k+1}$ before SCP converges [24].

The Jacobians of f in (5), evaluated with respect to $\bar{x}^k(\tau), \bar{u}(\tau)$, are denoted by:

$$A^k(\tau) = \frac{\partial f(\bar{x}^k(\tau), \bar{u}(\tau))}{\partial x}$$

$$B^k(\tau) = \frac{\partial f(\bar{x}^k(\tau), \bar{u}(\tau))}{\partial u}$$

for $\tau \in [\tau_k, \tau_{k+1}]$.

The discretized dynamics constraint (8) is linearized as follows:

$$x_{k+1} = A_k x_k + B_k^- u_k + B_k^+ u_{k+1} + w_k \quad (18)$$

for $k = 1, \dots, N - 1$, where A_k, B_k^-, B_k^+, w_k result from the solution to the following initial value problem over $\tau \in [\tau_k, \tau_{k+1}]$:

$$\begin{aligned} \dot{\Phi}_x(\tau, \tau_k) &= A^k(\tau) \Phi_x(\tau, \tau_k) \\ \dot{\Phi}_u^-(\tau, \tau_k) &= A^k(\tau) \Phi_u^-(\tau, \tau_k) + B^k(\tau) \left(\frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \right) \\ \dot{\Phi}_u^+(\tau, \tau_k) &= A^k(\tau) \Phi_u^+(\tau, \tau_k) + B^k(\tau) \left(\frac{\tau - \tau_k}{\tau_{k+1} - \tau_k} \right) \\ \Phi_x(\tau_k, \tau_k) &= I_{n_x} \\ \Phi_u^-(\tau_k, \tau_k) &= 0_{n_x \times n_u} \\ \Phi_u^+(\tau_k, \tau_k) &= 0_{n_x \times n_u} \\ A_k &= \Phi_x(\tau_{k+1}, \tau_k) \\ B_k^- &= \Phi_u^-(\tau_{k+1}, \tau_k) \\ B_k^+ &= \Phi_u^+(\tau_{k+1}, \tau_k) \\ w_k &= \bar{x}^k(\tau_{k+1}) - A_k \bar{x}_k - B_k^- \bar{u}_k - B_k^+ \bar{u}_{k+1} \end{aligned}$$

2. Convex Subproblem with Prescaling and Preconditioning

The final step in the construction of the convex subproblem involves scaling the decision variables and preconditioning the constraint parameters. In particular, we treat \hat{x}_k and \hat{u}_k as the decision variables, which are linearly

related to x_k and u_k as follows:

$$x_k = P_x \hat{x}_k + \bar{x}_k \quad (19a)$$

$$u_k = P_u \hat{u}_k + \bar{u}_k \quad (19b)$$

where P_x and P_u are selected to ensure that \hat{x}_k and \hat{u}_k , for $k = 1, \dots, N$, are roughly on the same order of magnitude, which greatly aids in convergence. Next, we define:

$$\hat{A}_k^- = P_x^{-1} A_k P_x \quad (20a)$$

$$\hat{A}_k^+ = -I_{n_x} \quad (20b)$$

$$\hat{B}_k^- = P_x^{-1} B_k^- P_u \quad (20c)$$

$$\hat{B}_k^+ = P_x^{-1} B_k^+ P_u \quad (20d)$$

$$\hat{w}_k = P_x^{-1} (\bar{x}^k(\tau_{k+1}) - \bar{x}_{k+1}) \quad (20e)$$

Finally, the convex subproblem (17), which is a quadratic program (QP), can be stated as:

$$\underset{\hat{x}_k, \hat{u}_k}{\text{minimize}} \quad w_{\text{cost}} \hat{x}_N^\top e_{\text{cost}} + \frac{w_{\text{prox}}}{2} \sum_{k=1}^N \|\hat{x}_k\|_2^2 + \|\hat{u}_k\|_2^2 + w_{\text{ep}} \sum_{k=1}^{N-1} 1_{n_x}^\top (\mu_k^+ + \mu_k^-) \quad (21a)$$

$$\text{subject to} \quad \hat{A}_k^- \hat{x}_k + \hat{A}_k^+ \hat{x}_{k+1} + \hat{B}_k^- \hat{u}_k + \hat{B}_k^+ \hat{u}_{k+1} + \mu_k^+ - \mu_k^- + \hat{w}_k = 0 \quad 1 \leq k \leq N-1 \quad (21b)$$

$$E_y (\hat{x}_{k+1} - \hat{x}_k) \leq \hat{\epsilon}_k \quad 1 \leq k \leq N-1 \quad (21c)$$

$$\mu_k^+ \geq 0, \mu_k^- \geq 0 \quad 1 \leq k \leq N-1 \quad (21d)$$

$$\hat{u}_{k,\min} \leq \hat{u}_k \leq \hat{u}_{k,\max} \quad 1 \leq k \leq N \quad (21e)$$

$$\tilde{E}_i \hat{x}_1 = \hat{z}_i, \tilde{E}_f \hat{x}_N = \hat{z}_f \quad (21f)$$

where

$$\hat{\epsilon}_k = \epsilon E_y P_x^{-1} E_y^\top - E_y P_x^{-1} (\bar{x}_{k+1} - \bar{x}_k)$$

$$\hat{u}_{k,\min/\max} = P_u^{-1} (u_{\min/\max} - \bar{u}_k)$$

$$\hat{z}_i = \tilde{E}_i P_x^{-1} \tilde{E}_i^\top ((z_i, 0) - \tilde{E}_i \bar{x}_1)$$

$$\hat{z}_f = \tilde{E}_f P_x^{-1} \tilde{E}_f^\top (z_f - \tilde{E}_f \bar{x}_N)$$

III. Convex Optimization Solver

The convex subproblem (21) is a quadratic program (QP), which can be represented in the canonical form as:

$$\text{minimize} \quad \frac{1}{2} z^\top \hat{P} z + \hat{p}^\top z \quad (22a)$$

$$\text{subject to} \quad \hat{G} z = \hat{g} \quad (22b)$$

$$\hat{H} z \leq \hat{h} \quad (22c)$$

where

$$z = (\hat{x}_1, \dots, \hat{x}_N, \hat{u}_1, \dots, \hat{u}_N, \mu_1^-, \dots, \mu_{N-1}^-, \mu_1^+, \dots, \mu_{N-1}^+) \in \mathbb{R}^{(n_x+n_u)N+2n_x(N-1)} \quad (23a)$$

$$\hat{P} = \text{blkdiag}(w_{\text{tr}} I_{(n_x+n_u)N}, 0_{2n_x(N-1) \times 2n_x(N-1)}) \quad (23b)$$

$$\hat{p} = w_{\text{cost}}(0_{n_x(N-1)}, e_x, 0_{n_u N + 2n_x(N-1)}) \quad (23c)$$

$$+ w_{\text{vc}}(0_{(n_x+n_u)N}, 1_{2n_x(N-1)}) \quad (23d)$$

$$\hat{G} = \begin{bmatrix} \hat{G}_x & \hat{G}_u & \hat{G}_\mu \end{bmatrix} \quad (23e)$$

$$\hat{G}_x = \begin{bmatrix} \hat{A}^- & 0_{n_x(N-1) \times n_x} \end{bmatrix} + \begin{bmatrix} 0_{n_x(N-1) \times n_x} & \hat{A}^+ \end{bmatrix} \quad (23f)$$

$$\hat{G}_u = \begin{bmatrix} \hat{B}^- & 0_{n_x(N-1) \times n_u} \end{bmatrix} + \begin{bmatrix} 0_{n_x(N-1) \times n_u} & \hat{B}^+ \end{bmatrix} \quad (23g)$$

$$\hat{G}_\mu = \begin{bmatrix} I_{n_x(N-1)} & -I_{n_x(N-1)} \end{bmatrix} \quad (23h)$$

$$\hat{A}^- = \text{blkdiag}(\hat{A}_1^-, \dots, \hat{A}_{N-1}^-) \quad (23i)$$

$$\hat{A}^+ = \text{blkdiag}(\hat{A}_1^+, \dots, \hat{A}_{N-1}^+) \quad (23j)$$

$$\hat{B}^- = \text{blkdiag}(\hat{B}_1^-, \dots, \hat{B}_{N-1}^-) \quad (23k)$$

$$\hat{B}^+ = \text{blkdiag}(\hat{B}_1^+, \dots, \hat{B}_{N-1}^+) \quad (23l)$$

$$\hat{g} = -(\hat{w}_1, \dots, \hat{w}_{N-1}) \quad (23m)$$

$$\hat{H}_y = -[I_{N-1} \otimes E_y \ 0_{n_y(N-1) \times n_x}] + [0_{n_y(N-1) \times n_x} \ I_{N-1} \otimes E_y] \quad (23n)$$

$$\hat{H} = \begin{bmatrix} \hat{H}_y & 0_{n_y(N-1) \times n_u N + 2n_x(N-1)} \end{bmatrix} \quad (23o)$$

$$\hat{h} = (\hat{u}_{1,\text{max}}, \dots, \hat{u}_{N,\text{max}}, -\hat{u}_{1,\text{min}}, \dots, -\hat{u}_{N,\text{min}}, 0_{2n_x(N-1)}, \hat{\varepsilon}_1, \dots, \hat{\varepsilon}_{N-1}) \quad (23p)$$

We adopt the proportional-integral projected gradient (PIPG) algorithm [23] with extrapolation (shown in Algorithm 1) to solve (22). The input to the algorithm consists of maximum iterations j_{max} , the extrapolation parameter ρ , the step-size α , which depends on the singular values of \hat{P} and \hat{H} , and the step-size ratio ω , which is chosen heuristically to balance the convergence rates of the primal and dual iterates.

Algorithm 1 PIPG

Inputs: $j_{\text{max}}, \rho, \alpha, \omega$

Initialize: ζ^1, η^1, χ^1

- 1: $\beta \leftarrow \omega\alpha$
- 2: **for** $j = 1, \dots, j_{\text{max}}$ **do**
- 3: $z^{j+1} \leftarrow \zeta^j - \alpha(\hat{P}\zeta^j + \hat{p} + \hat{G}^\top \eta^j + \hat{H}^\top \chi^j)$
- 4: $\tilde{E}_i \hat{x}_1^{j+1} \leftarrow \hat{z}_i$
- 5: $\tilde{E}_f \hat{x}_K^{j+1} \leftarrow \hat{z}_f$
- 6: $\hat{u}_k^{j+1} \leftarrow \max\{\hat{u}_{k,\text{min}}, \min\{\hat{u}_{k,\text{max}}, \hat{u}_k^{j+1}\}\} \quad k = 1, \dots, N$
- 7: $(\mu_k^+)^{j+1} \leftarrow \max\{0, (\mu_k^+)^{j+1}\} \quad k = 1, \dots, N-1$
- 8: $(\mu_k^-)^{j+1} \leftarrow \max\{0, (\mu_k^-)^{j+1}\} \quad k = 1, \dots, N-1$
- 9: $w^{j+1} \leftarrow \eta^j + \beta(\hat{G}(2z^{j+1} - \zeta^j) - \hat{g})$
- 10: $v^{j+1} \leftarrow \max\{0, \chi^j + \beta(\hat{H}(2z^{j+1} - \zeta^j) - \hat{h})\}$
- 11: $\zeta^{j+1} \leftarrow (1-\rho)\zeta^j + \rho z^{j+1}$
- 12: $\eta^{j+1} \leftarrow (1-\rho)\eta^j + \rho w^{j+1}$
- 13: $\chi^{j+1} \leftarrow (1-\rho)\chi^j + \rho v^{j+1}$
- 14: **end for**

Return: z^j, w^j, v^j

Algorithm 1 can be customized to Algorithm 3 (referred to as PIPG_{custom}), which exploits the sparsity structure of the parameters (23) of (22). The recursive calls to PIPG_{custom} within SCP are warm-started with the primal-dual solution from the previous subproblem to accelerate its convergence.

The customized power iteration method shown in Algorithm 2, called before PIPG_{custom}, estimates the square of the maximum singular value of \hat{H} , denoted by σ . The maximum eigenvalue of \hat{P} is trivially known, given its construction as a diagonal matrix. Furthermore, PIPG_{custom} employs a customized termination criterion, as shown in Algorithm 4.

The primary design consideration in Algorithms 2, 3, and 4 is to eliminate matrix factorizations, inversions, and sparse linear algebra; only simple operations on small-sized dense vectors and matrices are used. See [23, 25, 26] for further details on custom, real-time implementations of PIPG.

Algorithm 2 Customized power iteration method

Inputs: $\hat{A}_{[1:N-1]}^-$, $\hat{A}_{[1:N-1]}^+$, $\hat{B}_{[1:N-1]}^-$, $\hat{B}_{[1:N-1]}^+$, $\hat{x}_{[1:N]}$, $\hat{u}_{[1:N]}$, $\mu_{[1:N-1]}^+$, $\mu_{[1:N-1]}^-$, ϵ_{abs} , ϵ_{rel} , ϵ_{buff} , j_{max}

Require: \hat{x}_k , \hat{u}_k , μ_k^+ , μ_k^- are not all 0

```

1:  $\sigma \leftarrow 0$ 
2: for  $k = 1, \dots, N - 1$  do
3:    $\sigma \leftarrow \sigma + \|\hat{x}_k\|_2^2 + \|\hat{u}_k\|_2^2 + \|\mu_k^+\|_2^2 + \|\mu_k^-\|_2^2$ 
4: end for
5:  $\sigma \leftarrow \sigma + \|\hat{x}_N\|_2^2 + \|\hat{u}_N\|_2^2$ 
6:  $\sigma \leftarrow \sqrt{\sigma}$ 
7: for  $j = 1, \dots, j_{\text{max}}$  do
8:   for  $k = 1, \dots, N - 1$  do
9:      $\phi_k \leftarrow \frac{1}{\sigma} (\hat{A}_k^- \hat{x}_k + \hat{A}_k^+ \hat{x}_{k+1} + \hat{B}_k^- \hat{u}_k + \hat{B}_k^+ \hat{u}_{k+1} + \mu_k^+ - \mu_k^-)$ 
10:     $\theta_k \leftarrow \frac{1}{\sigma} E_y (\hat{x}_{k+1} - \hat{x}_k)$ 
11:   end for
12:    $\hat{x}_1 \leftarrow \hat{A}_1^{-\top} \phi_1 - E_y^\top \theta_1$ 
13:    $\hat{u}_1 \leftarrow \hat{B}_1^{-\top} \phi_1$ 
14:    $\mu_1^+ \leftarrow \phi_1$ 
15:    $\mu_1^- \leftarrow -\phi_1$ 
16:   for  $k = 2, \dots, N - 1$  do
17:      $\hat{x}_k \leftarrow \hat{A}_k^{-\top} \phi_k + \hat{A}_{k-1}^{+\top} \phi_{k-1} - E_y^\top v_k + E_y^\top v_{k-1}$ 
18:      $\hat{u}_k \leftarrow \hat{B}_k^{-\top} \phi_k + \hat{B}_{k-1}^{+\top} \phi_{k-1}$ 
19:      $\mu_k^+ \leftarrow \phi_k$ 
20:      $\mu_k^- \leftarrow -\phi_k$ 
21:   end for
22:    $\hat{x}_N \leftarrow \hat{A}_{N-1}^{+\top} \phi_{N-1} + E_y^\top v_{N-1}$ 
23:    $\hat{u}_N \leftarrow \hat{B}_{N-1}^{+\top} \phi_{N-1}$ 
24:    $\sigma^* \leftarrow 0$ 
25:   for  $k = 1, \dots, N - 1$  do
26:      $\sigma^* \leftarrow \sigma^* + \|\hat{x}_k\|_2^2 + \|\hat{u}_k\|_2^2 + \|\mu_k^+\|_2^2 + \|\mu_k^-\|_2^2$ 
27:   end for
28:    $\sigma^* \leftarrow \sigma + \|\hat{x}_N\|_2^2 + \|\hat{u}_N\|_2^2$ 
29:    $\sigma^* \leftarrow \sqrt{\sigma^*}$ 
30:   if  $|\sigma^* - \sigma| \leq \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \max\{\sigma^*, \sigma\}$  then ▷ stopping criterion
31:     break
32:   else if  $j < j_{\text{max}}$  then
33:      $\sigma \leftarrow \sigma^*$ 
34:   end if
35: end for
36:  $\sigma \leftarrow (1 + \epsilon_{\text{buff}}) \sigma^*$  ▷ buffer the estimated maximum singular value

```

Return: σ

Inputs: $w_{\text{cost}}, e_{\text{cost}}, w_{\text{ep}}, w_{\text{prox}}$

$$\hat{A}_{[1:N-1]}^-, \hat{A}_{[1:N-1]}^+, \hat{B}_{[1:N-1]}^-, \hat{B}_{[1:N-1]}^+, \hat{w}_{[1:N-1]}$$

$$\hat{u}_{[1:N], \min}, \hat{u}_{[1:N], \max}, E_y, \hat{\epsilon}_{[1:N-1]}, \hat{E}_i, \hat{E}_f, \hat{z}_i, \hat{z}_f$$

$$\sigma, \omega, \rho, \epsilon_{\text{abs}}, \epsilon_{\text{rel}}, j_{\text{check}}, j_{\text{max}},$$

$$\hat{x}_{[1:N]}^*, \hat{u}_{[1:N]}^*, \mu_{[1:N-1]}^{+,*}, \mu_{[1:N-1]}^{-,*}, \phi_{[1:N-1]}^*, \theta_{[1:N-1]}^* \quad \triangleright \text{warm start}$$

- 1: $\tilde{x}_{[1:N]}^1 \leftarrow \hat{x}_{[1:N]}^*$ \triangleright initialize primal variables
- 2: $\tilde{u}_{[1:N]}^1 \leftarrow \hat{u}_{[1:N]}^*$
- 3: $\nu_{[1:N-1]}^{+,1} \leftarrow \mu_{[1:N-1]}^{+,*}$
- 4: $\nu_{[1:N-1]}^{-,1} \leftarrow \mu_{[1:N-1]}^{-,*}$
- 5: $\varphi_{[1:N-1]}^1 \leftarrow \phi_{[1:N-1]}^*$ \triangleright initialize dual variables
- 6: $\vartheta_{[1:N-1]}^1 \leftarrow \theta_{[1:N-1]}^*$
- 7: $\alpha \leftarrow \frac{2}{w_{\text{prox}} + \sqrt{w_{\text{prox}}^2 + 4\omega\sigma}}$ \triangleright step-sizes
- 8: $\beta \leftarrow \omega\alpha$
- 9: **for** $j = 1, \dots, j_{\text{max}}$ **do**
- 10: $\tilde{x}_1^j \leftarrow \tilde{x}_1^j - \alpha (w_{\text{prox}} \tilde{x}_1^j + \hat{A}_1^{-\top} \varphi_1^j - E_y^\top \vartheta_1^j)$
- 11: $\tilde{E}_i \hat{x}_1^j \leftarrow \hat{z}_i$
- 12: $\tilde{u}_1^j \leftarrow \max\{\hat{u}_{1, \min}, \min\{\hat{u}_{1, \max}, \tilde{u}_1^j - \alpha (w_{\text{prox}} \tilde{u}_1^j + \hat{B}_1^{-\top} \varphi_1^j)\}\}$
- 13: **for** $k = 2, \dots, N-1$ **do** \triangleright projected gradient step
- 14: $\tilde{x}_k^{j+1} \leftarrow \tilde{x}_k^j - \alpha (w_{\text{prox}} \tilde{x}_k^j + \hat{A}_k^{-\top} \varphi_k^j + \hat{A}_{k-1}^{+\top} \varphi_{k-1}^j - E_y^\top \vartheta_k^j + E_y^\top \vartheta_{k-1}^j)$
- 15: $\tilde{u}_k^{j+1} \leftarrow \max\{\hat{u}_{k, \min}, \min\{\hat{u}_{k, \max}, \tilde{u}_k^j - \alpha (w_{\text{prox}} \tilde{u}_k^j + \hat{B}_k^{-\top} \varphi_k^j + \hat{B}_{k-1}^{+\top} \varphi_{k-1}^j)\}\}$
- 16: **end for**
- 17: $\hat{x}_N^{j+1} \leftarrow \tilde{x}_N^j - \alpha (w_{\text{prox}} \tilde{x}_N^j + w_{\text{cost}} e_{\text{cost}} + \hat{A}_{N-1}^{+\top} \varphi_{N-1}^j + E_y^\top \vartheta_{N-1}^j)$
- 18: $\hat{E}_f \hat{x}_N^{j+1} \leftarrow \hat{z}_f$
- 19: $\hat{u}_N^{j+1} \leftarrow \max\{\hat{u}_{N, \min}, \min\{\hat{u}_{N, \max}, \tilde{u}_N^j - \alpha (w_{\text{prox}} \tilde{u}_N^j + \hat{B}_{N-1}^{+\top} \varphi_{N-1}^j)\}\}$
- 20: $\mu_{[1:N-1]}^{\pm, j} \leftarrow \max\{0, \nu_{[1:N-1]}^{\pm, j} - \alpha (w_{\text{ep}} \mathbf{1}_{n_x} + \varphi_{[1:N-1]}^j)\}$
- 21: **for** $k = 1, \dots, N-1$ **do** \triangleright PI feedback of affine equality constraint violation
- 22: $\phi_k^{j+1} \leftarrow \varphi_k^j + \beta (\hat{A}_k^- (2\hat{x}_k^{j+1} - \hat{x}_k^j) + \hat{A}_k^+ (2\hat{x}_{k+1}^{j+1} - \hat{x}_{k+1}^j) + \hat{B}_k^- (2\hat{u}_k^{j+1} - \hat{u}_k^j) + \hat{B}_k^+ (2\hat{u}_{k+1}^{j+1} - \hat{u}_{k+1}^j)$
- 23: $\quad + (2\mu_k^{+, j+1} - \nu_k^{+, j}) - (2\mu_k^{-, j+1} - \nu_k^{-, j}) + \hat{w}_k)$
- 24: $\theta_k^{j+1} \leftarrow \max\{0, \vartheta_k^j + \beta ((2E_y \hat{x}_{k+1}^{j+1} - E_y \hat{x}_{k+1}^j) - (2E_y \hat{x}_k^{j+1} - E_y \hat{x}_k^j) - \hat{\epsilon}_k)\}$
- 25: **end for**
- 26: $\hat{x}_{[1:N]}^{j+1} \leftarrow (1-\rho) \tilde{x}_{[1:N]}^j + \rho \hat{x}_{[1:N]}^{j+1}$ \triangleright extrapolate primal variables
- 27: $\hat{u}_{[1:N]}^{j+1} \leftarrow (1-\rho) \tilde{u}_{[1:N]}^j + \rho \hat{u}_{[1:N]}^{j+1}$
- 28: $\nu_{[1:N-1]}^{\pm, j+1} \leftarrow (1-\rho) \nu_{[1:N-1]}^{\pm, j} + \rho \mu_{[1:N-1]}^{\pm, j+1}$
- 29: $\varphi_{[1:N-1]}^{j+1} \leftarrow (1-\rho) \varphi_{[1:N-1]}^j + \rho \phi_{[1:N-1]}^{j+1}$ \triangleright extrapolate dual variables
- 30: $\vartheta_{[1:N-1]}^{j+1} \leftarrow (1-\rho) \vartheta_{[1:N-1]}^j + \rho \theta_{[1:N-1]}^{j+1}$
- 31: **if** $j \bmod j_{\text{check}} = 0$ **then** \triangleright check stopping criterion every j_{check} iterations
- 32: $\text{TERMINATE} \leftarrow \text{STOPPING}_{\text{custom}}(\hat{x}_{[1:N]}^{j+1}, \hat{u}_{[1:N]}^{j+1}, \mu_{[1:N-1]}^{\pm, j+1}, \phi_{[1:N-1]}^{j+1}, \theta_{[1:N-1]}^{j+1})$
- 33: $\hat{x}_{[1:N]}^j, \hat{u}_{[1:N]}^j, \mu_{[1:N-1]}^{\pm, j}, \phi_{[1:N-1]}^j, \theta_{[1:N-1]}^j, \epsilon_{\text{abs}}, \epsilon_{\text{rel}}$
- 34: **if** $\text{TERMINATE} = \text{TRUE}$ **then** \triangleright stopping criterion
- 35: **break**
- 36: **end if**
- 37: **end if**
- 38: **end for**
- 39: $\hat{x}_{[1:N]}^* \leftarrow \hat{x}_{[1:N]}^{j+1}$ \triangleright update primal variables
- 40: $\hat{u}_{[1:N]}^* \leftarrow \hat{u}_{[1:N]}^{j+1}$
- 41: $\mu_{[1:N-1]}^{\pm, *} \leftarrow \mu_{[1:N-1]}^{\pm, j+1}$
- 42: $\phi_{[1:N-1]}^* \leftarrow \phi_{[1:N-1]}^{j+1}$ \triangleright update dual variables
- 43: $\theta_{[1:N-1]}^* \leftarrow \theta_{[1:N-1]}^{j+1}$

Return: $\hat{x}_{[1:N]}^*, \hat{u}_{[1:N]}^*, \mu_{[1:N-1]}^{\pm, *}, \phi_{[1:N-1]}^*, \theta_{[1:N-1]}^*$

Algorithm 4 Customized stopping criterion evaluation:

$$\text{STOPPING}_{\text{custom}}(\hat{x}_{[1:N]}^{j+1}, \hat{u}_{[1:N]}^{j+1}, \mu_{[1:N-1]}^{+,j+1}, \mu_{[1:N-1]}^{-,j+1}, \phi_{[1:N-1]}^{j+1}, \theta_{[1:N-1]}^{j+1}, \hat{x}_{[1:N]}^j, \hat{u}_{[1:N]}^j, \mu_{[1:N-1]}^{+,j}, \mu_{[1:N-1]}^{-,j}, \phi_{[1:N-1]}^j, \theta_{[1:N-1]}^j, \epsilon_{\text{abs}}, \epsilon_{\text{rel}})$$

Inputs: $\hat{x}_{[1:N]}^{j+1}, \hat{u}_{[1:N]}^{j+1}, \mu_{[1:N-1]}^{+,j+1}, \mu_{[1:N-1]}^{-,j+1}, \phi_{[1:N-1]}^{j+1}, \theta_{[1:N-1]}^{j+1}, \hat{x}_{[1:N]}^j, \hat{u}_{[1:N]}^j, \mu_{[1:N-1]}^{+,j}, \mu_{[1:N-1]}^{-,j}, \phi_{[1:N-1]}^j, \theta_{[1:N-1]}^j, \epsilon_{\text{abs}}, \epsilon_{\text{rel}}$

```

1:  $z_{\infty}^{j+1} \leftarrow \max \left\{ \|\hat{x}_{[1:N]}^{j+1}\|_{\infty}, \|\hat{u}_{[1:N]}^{j+1}\|_{\infty}, \|\mu_{[1:N-1]}^{+,j+1}\|_{\infty}, \|\mu_{[1:N-1]}^{-,j+1}\|_{\infty} \right\}$ 
2:  $z_{\infty}^j \leftarrow \max \left\{ \|\hat{x}_{[1:N]}^j\|_{\infty}, \|\hat{u}_{[1:N]}^j\|_{\infty}, \|\mu_{[1:N-1]}^{+,j}\|_{\infty}, \|\mu_{[1:N-1]}^{-,j}\|_{\infty} \right\}$ 
3:  $z_{\infty}^{\Delta j} \leftarrow \max \left\{ \|\hat{x}_{[1:N]}^{j+1} - \hat{x}_{[1:N]}^j\|_{\infty}, \|\hat{u}_{[1:N]}^{j+1} - \hat{u}_{[1:N]}^j\|_{\infty}, \|\mu_{[1:N-1]}^{+,j+1} - \mu_{[1:N-1]}^{+,j}\|_{\infty}, \|\mu_{[1:N-1]}^{-,j+1} - \mu_{[1:N-1]}^{-,j}\|_{\infty} \right\}$ 
4:  $r_{\infty}^{j+1} \leftarrow \max \left\{ \|\phi_{[1:N-1]}^{j+1}\|_{\infty}, \|\theta_{[1:N-1]}^{j+1}\|_{\infty} \right\}$ 
5:  $r_{\infty}^j \leftarrow \max \left\{ \|\phi_{[1:N-1]}^j\|_{\infty}, \|\theta_{[1:N-1]}^j\|_{\infty} \right\}$ 
6:  $r_{\infty}^{\Delta j} \leftarrow \max \left\{ \|\phi_{[1:N-1]}^{j+1} - \phi_{[1:N-1]}^j\|_{\infty}, \|\theta_{[1:N-1]}^{j+1} - \theta_{[1:N-1]}^j\|_{\infty} \right\}$ 
7: if  $z_{\infty}^{\Delta j} \leq \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \max \{z_{\infty}^{j+1}, z_{\infty}^j\}$  and  $r_{\infty}^{\Delta j} \leq \epsilon_{\text{abs}} + \epsilon_{\text{rel}} \max \{r_{\infty}^{j+1}, r_{\infty}^j\}$  then
8:   TERMINATE  $\leftarrow$  TRUE
9: else
10:  TERMINATE  $\leftarrow$  FALSE
11: end if

```

Return: TERMINATE

IV. 6-DoF Powered Descent Guidance

The 6-DoF rocket model is based on [28], with the inclusion of an independent torque control input by means of reaction control thrusters (as shown in [26]), and with the aerodynamics forces and moments ignored.

The system dynamics (1) and constraint function (2) are given by:

$$\xi = (m, \mathbf{r}_{\mathcal{I}}, \mathbf{v}_{\mathcal{I}}, \mathbf{q}_{\mathcal{I} \leftarrow \mathcal{B}}, \boldsymbol{\omega}_{\mathcal{B}}) \quad (24a)$$

$$\zeta = (\mathbf{T}_{\mathcal{B}}, \boldsymbol{\gamma}_{\mathcal{B}}) \quad (24b)$$

$$F(\xi, \zeta) = \begin{bmatrix} -\alpha_{\dot{m}} \|\mathbf{T}_{\mathcal{B}}\|_2 \\ \mathbf{v}_{\mathcal{I}} \\ \frac{1}{m} C_{\mathcal{I} \leftarrow \mathcal{B}} \mathbf{T}_{\mathcal{B}} + \mathbf{g}_{\mathcal{I}} \\ \frac{1}{2} \Omega(\boldsymbol{\omega}_{\mathcal{B}}) \mathbf{q}_{\mathcal{I} \leftarrow \mathcal{B}} \\ J_{\mathcal{B}}^{-1} (\mathbf{r}_{\mathcal{I}, \mathcal{B}} \times \mathbf{T}_{\mathcal{B}} - \boldsymbol{\omega}_{\mathcal{B}} \times J_{\mathcal{B}} \boldsymbol{\omega}_{\mathcal{B}} + \boldsymbol{\gamma}_{\mathcal{B}}) \end{bmatrix} \quad (24c)$$

$$g(\xi, \zeta) = \begin{bmatrix} m_{\text{dry}} - m \\ -\mathbf{e}_1^{\top} \mathbf{r}_{\mathcal{I}} \\ \|\mathbf{v}_{\mathcal{I}}\|_2^2 - v_{\text{max}}^2 \\ 4 \|H_{\theta} \mathbf{q}_{\mathcal{I} \leftarrow \mathcal{B}}\|_2^2 + (1 - \cos \theta_{\text{max}})^2 \\ \|\boldsymbol{\omega}_{\mathcal{B}}\|_2^2 - \omega_{\text{max}}^2 \\ \|\mathbf{T}_{\mathcal{B}}\|_2 - \mathbf{e}_1^{\top} \mathbf{T}_{\mathcal{B}} \sec \delta_{\text{max}} \\ \|\mathbf{T}_{\mathcal{B}}\|_2 - T_{\text{max}} \\ -\|\mathbf{T}_{\mathcal{B}}\|_2 + T_{\text{min}} \\ \|\boldsymbol{\gamma}_{\mathcal{B}}\|_2^2 - \gamma_{\text{max}}^2 \end{bmatrix} \quad (24d)$$

The numerical values for the parameters within the system dynamics ($\alpha_{\dot{m}}, \mathbf{g}_{\mathcal{I}}, J_{\mathcal{B}}, \mathbf{r}_{\mathcal{I}, \mathcal{B}}, H_{\theta}$), boundary conditions (z_i, z_f), and constraint bounds ($m_{\text{dry}}, \gamma_{\text{gs}}, v_{\text{max}}, \theta_{\text{max}}, T_{\text{max}}, T_{\text{min}}$) correspond to the baseline problem in [11]. Note that the non-differentiability of the constraint functions for thrust magnitude bounds at $\mathbf{T}_{\mathcal{B}} = \mathbf{0}_3$ does not affect the linearization operations within SCP in practice with an appropriate initial guess for the thrust vector. Furthermore, the box constraint on the scaled control input in Lines 12, 15 and 19 of Algorithm 3 are only imposed on the dilation factor.

V. GPU Architecture and Code Implementation

A. Overview

Loosely speaking, CPUs are designed to execute code serially. Consider the problem of adding two vectors. A CPU is designed to add these two vectors by looping through the elements of each vector and adding the corresponding elements. In contrast, GPUs have many specialized hardware units that allow them to perform many floating point operations in parallel. A GPU would add two vectors together by adding the corresponding elements in parallel by assigning each scalar addition to a different compute thread.

In GPU programming, a single sequence of instructions is called a *thread*. Back to the example of adding two vectors on a GPU, a single thread would execute a single scalar addition. Threads in turn are organized into a *grid* of *thread-blocks*. In the vector addition example, we would want to have the same number of threads as the number of elements in each vector. To write programs for the GPU, NVIDIA provides the compute unified device architecture (CUDA) platform, which is an extension of the C programming language [33].

CUDA enables the programmer to write special GPU functions called *kernels*, which, when called, are executed a specified number of times in parallel. When the programmer writes a kernel, they are specifying how a single thread should execute the given code; when calling a kernel, the programmer specifies the number of threads per thread-block and the number of thread-blocks in the grid. The product of these two numbers is the number of threads being called. Each thread of execution has a unique identifier based on its thread-block index within the grid and its thread index within the thread-block. The thread ID in a kernel is an integer in $[0, N)$, where N is the number of threads called. However, before performing computations on the GPU, we must first ensure that the data the kernel is operating on is in GPU memory rather than in CPU memory. This typically requires the programmer to allocate memory on the GPU to store this data, using the `cudaMalloc()` function, and to then copy the data from the CPU to the GPU using the `cudaMemcpy()` function.

When the kernel is called, the thread-blocks are then distributed to *streaming multiprocessors* (SM), which contain a collection of CUDA cores. These CUDA cores execute the threads within a thread-block. When a SM finishes executing all threads on a thread-block, the GPU allocates it another thread-block.

In all of this, the programmer only has to worry about creating a grid of thread-blocks when they call the kernel. The GPU is responsible for allocating thread-blocks to SMs, which are in turn responsible for allocating threads to CUDA cores. The thread hierarchy is summarized by Figure 1.

As an example, consider the code in Listing 1. The first function is how one would write a function to add two vectors on a CPU. The second function is how one would write a CUDA kernel to add two vectors in parallel. When launching this kernel, each GPU thread executes the code within the kernel in parallel. The first line in this kernel computes the unique thread ID for the thread; subsequently this thread will perform one scalar addition for the index that matches its thread ID. The `main()` function shows how to call the serial version of this vector addition and how to call the CUDA kernel. We must first allocate the memory for arrays a , b , and c on the GPU and then copy the data in the arrays from the CPU to the GPU. Finally, the kernel can be called with the desired grid size and thread-block size.

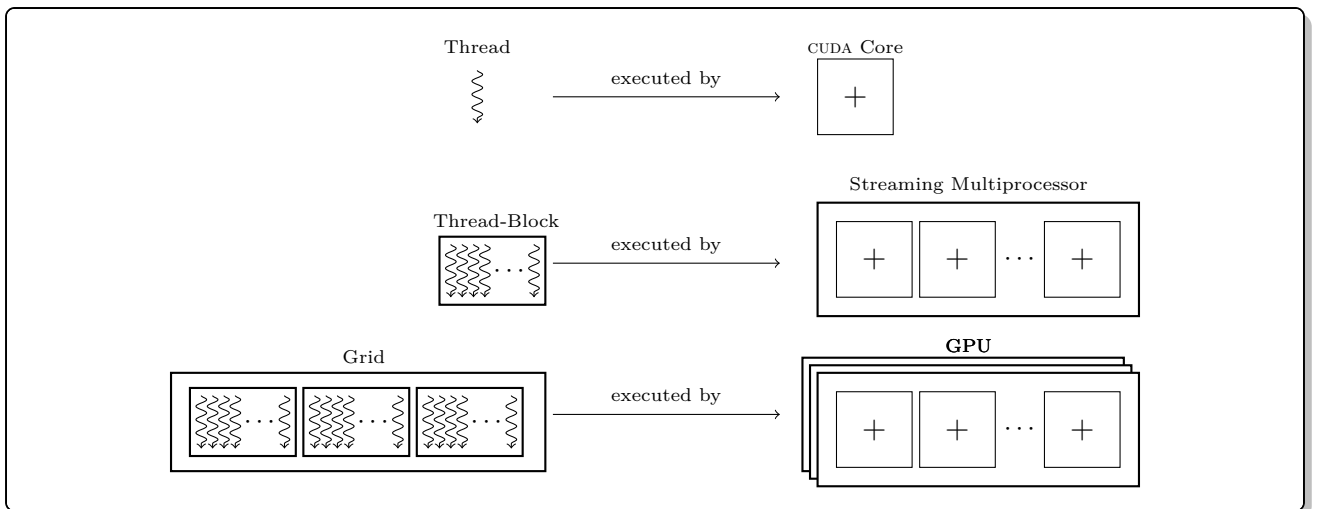


Fig. 1 Thread hierarchy in the GPU

```

1  /* Performs the vector addition  $c = a + b$  serially, where all vectors have  $N$  elements. */
2  void add_vector_serial(int* a, int* b, int* c, int N)
3  {
4      for(int i = 0; i < N; i++)
5      {
6          c[i] = a[i] + b[i];
7      }
8  }
9
10 /* CUDA kernel to perform  $c = a + b$ , where all vectors have  $N$  elements. */
11 __global__ void add_vector_parallel(int* a, int* b, int* c, int N)
12 {
13     /* Get the ID of the thread executing this function. */
14     int thread_id = blockIdx.x*blockDim.x + threadIdx.x;
15
16     /* Safeguard if the kernel is launched with more threads than the number of elements. */
17     if (i < N)
18     {
19         c[i] = a[i] + b[i];
20     }
21 }
22
23 int main()
24 {
25     /* Declare pointers for arrays in CPU memory. */
26     int* a[N];
27     int* b[N];
28     int* c[N];
29
30     /* Helper function to fill a and b with data. */
31     initialize(a, b);
32
33     /* Declare pointers for arrays in GPU memory. */
34     int* d_a, d_b, d_c;
35
36     /* Allocate memory for arrays in GPU memory. */
37     cudaMalloc((void**) &d_a, N * sizeof(int));
38     cudaMalloc((void**) &d_b, N * sizeof(int));
39     cudaMalloc((void**) &d_c, N * sizeof(int));
40
41     /* Copy data in array from CPU to GPU memory. */
42     cudaMemcpy(d_a, a, N * sizeof(int), cudaMemcpyHostToDevice);
43     cudaMemcpy(d_b, b, N * sizeof(int), cudaMemcpyHostToDevice);
44     cudaMemcpy(d_c, c, N * sizeof(int), cudaMemcpyHostToDevice);
45
46     /* Call serial addition function on CPU. */
47     add_vector_serial(a, b, c, N);
48
49     /* Call parallel addition kernel on GPU. */
50     add_vector_parallel<<<grid_size, block_size>>>(d_a, d_b, d_c, N);
51 }

```

Listing 1 Code for serial (CPU) and parallelized (GPU) vector addition

B. Code Structure

The data transfer bandwidth between the CPU and the GPU is much slower than that between the CUDA cores and the GPU memory. Thus to achieve peak performance, it is important to minimize data transfer between the CPU and GPU [33]. Consequently, we chose the following execution model: create an array of problem instances with varying initial conditions on the CPU, allocate the necessary memory on the GPU to store these problem instances, copy the problem instance data to the GPU, launch a kernel that solves all of these trajectory optimization problem instances on the GPU, and finally copy back the trajectories from the GPU to the CPU. In this paradigm, we are assigning each thread one trajectory to compute. Each thread is able to access the problem instance it needs to solve by indexing into the array of problem data using its unique thread ID. Figure 2 shows the flow of data in our implementation.

By taking this approach we only have to incur the costly overhead of CPU-GPU communication once when sending the problem instances to the GPU and once when copying the computed trajectories back from the GPU. As mentioned in Section I, a key difficulty with implementing this procedure on the GPU is the need to implement all operations from scratch in CUDA, since we are unable to use off-the-shelf linear algebra routines and convex optimization solvers. Therefore, all code necessary to solve for a trajectory was implemented in CUDA, including




Fig. 2 Data flow between the CPU and the GPU

Level 1, 2, and 3 BLAS operations and the $\text{PIPG}_{\text{custom}}$ algorithm.

VI. Numerical Results

This section presents and analyzes Monte Carlo results for the 6-DoF powered-descent guidance problem outlined in Section IV using the solution procedure in Section II. We then compare the computation time between a serial CPU implementation and parallel GPU implementation for Monte Carlo simulations of various sizes. All results were obtained on a machine with an AMD Ryzen 9 7950X3D CPU and an NVIDIA RTX 3090 GPU.

When conducting the Monte Carlo, we disperse the initial position according to the following uniform distribution:

$$\mathbf{r}_{\mathcal{I}}(0) \sim \begin{bmatrix} \mathcal{U}(6, 9) \\ \mathcal{U}(3, 6) \\ \mathcal{U}(1, 2) \end{bmatrix} \quad (25)$$

For all simulations, we use the same weights, w_{cost} , w_{tr} , and w_{vc} , defined in Equation 21a, we set the maximum number of SCP iterations to 25, and we use 2500 PIPG iterations with the same hyperparameters ω and ρ for each subproblem solve. We find that SCP converges within 25 iterations for all 256 of these trajectories, which is the first demonstration of robustness to variations in initial conditions for SCP with PIPG as a subproblem solver. Figures 3 and 4 show the trajectories obtained. An extensive Monte Carlo analysis of the 6-DoF PDG problem is provided in [1].






Fig. 3 Trajectories generated from 256 Monte Carlo runs

Figure 5 depicts timing results for serial CPU and parallel GPU Monte Carlos for varying Monte Carlo batch sizes and for various thread-block sizes. For the GPU simulations, only the solve-time is reported since execution times for `malloc` and `memcpy` were an order of magnitude smaller than the solve-time.

As mentioned in Section V, calling a kernel requires two parameters: thread-block size and grid size, which is the number of thread-blocks in the grid. The product of these two parameters dictates the number of threads




(a) Mass Depletion




(b) Distribution of SCP Iterations

Fig. 4 Statistics from 256 Monte Carlo runs



(a) Log-log timing plot



(b) Semi-log timing plot

Fig. 5 CPU and GPU Monte Carlo solve-times. GPU results are plotted for the indicated thread-block sizes.

launched. This gives users a choice in choosing the thread-block and grid size when launching a kernel. For example, if we wanted to perform a Monte Carlo with a batch size of 512, we could choose a thread-block size of 128 and grid size of four or a thread-block size of 256 and grid size of two. Here, we perform Monte Carlo runs on the GPU with maximum thread-block sizes of 64, 128, and 256. For batch sizes below the maximum thread-block size, we set the thread-block size to be the batch size, and the grid size to one. For example, if we had a maximum thread-block size of 64 and were conducting a Monte Carlo run with a batch size of four, we would set the thread-block size to 4 and grid size to one. On the other hand, if we were conducting a Monte Carlo run with a batch size of 128, we would set the thread-block size to 64 and the grid size to two.

From Figure 5, it is evident that the Monte Carlo run on the GPU is slower than that on the CPU for small batch sizes, and faster for larger batch sizes, with a crossover point roughly between 1024 and 2048 Monte Carlo runs. Since CPUs, unlike GPUs, are optimized for serial tasks such as SCP, it is unsurprising that for smaller runs, the CPU is faster. However, as the batch size increases, the GPU is able to leverage parallelism to provide acceleration. In other words, although the GPU takes longer than the CPU to solve a single SCP problem, it is able to accelerate the Monte Carlo run when given a large number of problems to solve.

The log-log timing plot for the CPU implementation, shown in Figure 5a, is linear with a constant slope, as expected, since doubling the number of runs will double the solve-time for a serial implementation. However, the timing plots for the GPU exhibit three phases as the batch size increases. Initially, the solve-time increases until the maximum block size is reached, then the solve-time stays relatively constant, and finally, it increases again. In the first phase, since all threads are in a single thread-block, all the computations for the Monte Carlo are being performed by a single streaming multiprocessor. After the number of runs exceeds the maximum size of the

thread-block, a second thread-block is populated, which allows a second streaming multiprocessor to compute these trajectories in parallel with the first thread-block. This results in the second phase, wherein the solve-time curve is flat. The number of runs after which the curve flattens out is dictated by the maximum thread-block size. The third phase, where increasing the number of runs increases the solve-time again, is likely due to the GPU not being able to allocate any more streaming multiprocessors to the thread-blocks, forcing them to execute thread-blocks serially.

VII. Conclusion

We introduce a GPU-accelerated Monte Carlo framework for general nonconvex trajectory optimization problems using sequential convex programming. By comparing a parallel GPU-based implementation of a Monte Carlo framework for the 6-DoF powered-descent guidance problem to a serial CPU-based implementation, we show that for simulations over ~ 1000 runs, the GPU-based implementation is faster. Future work will investigate parallelized multi-start for nonconvex trajectory optimization, wherein a number of initial guesses are provided, the resulting problems are solved in parallel, and the converged trajectory with the lowest cost is chosen.

Acknowledgments

This research was supported by NASA grant NNX17AH02A, AFOSR grant FA9550-20-1-0053, ONR grant N00014-20-1-2288, and Blue Origin, LLC; Government sponsorship is acknowledged.

References

- [1] Malyuta, D., Reynolds, T., Szmuk, M., Mesbahi, M., Açıkmese, Behçet., and Carson III, J. M., “Discretization performance and accuracy analysis for the rocket powered descent guidance problem,” *AIAA SciTech 2019 Forum*, Reston, Virginia, 2019, pp. 1–20. <https://doi.org/10.2514/6.2019-0925>.
- [2] Herlihy, M., Shavit, N., Luchangco, V., and Spear, M., *The Art of Multiprocessor Programming*, Newnes, 2020.
- [3] Langston, J., “Microsoft announces new supercomputer, lays out vision for future AI work,” <https://news.microsoft.com/source/features/innovation/openai-azure-supercomputer/>, Accessed: 2023-12-02.
- [4] Shilov, A., “Tesla’s \$300 million AI cluster is going live today,” , Aug 2023. URL <https://www.tomshardware.com/news/teslas-dollar300-million-ai-cluster-is-going-live-today>.
- [5] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X., “TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems,” , 2015. URL <https://www.tensorflow.org/>, software available from tensorflow.org.
- [6] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S., “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [7] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q., “JAX: composable transformations of Python+NumPy programs,” , 2018. URL <http://github.com/google/jax>.
- [8] Ilg, M., Rogers, J., and Costello, M., “Projectile Monte-Carlo Trajectory Analysis Using a Graphics Processing Unit,” *AIAA Atmospheric Flight Mechanics Conference, Guidance, Navigation, and Control and Co-located Conferences*, American Institute of Aeronautics and Astronautics, 2011, pp. 1–18.
- [9] Rastgar, F., Masnavi, H., Kruusamäe, K., Aabloo, A., and Singh, A. K., “GPU Accelerated Batch Trajectory Optimization for Autonomous Navigation,” *2023 American Control Conference (ACC)*, IEEE, 2023, pp. 718–725.
- [10] Elango, P., Luo, D., Kamath, A. G., Uzun, S., Kim, T., and Açıkmese, B., “Successive Convexification for Trajectory Optimization with Continuous-Time Constraint Satisfaction,” , 2024. <https://doi.org/10.48550/ARXIV.2404.16826>, URL <https://arxiv.org/abs/2404.16826>.
- [11] Szmuk, M., Reynolds, T. P., and Açıkmese, B., “Successive Convexification for Real-Time Six-Degree-of-Freedom Powered Descent Guidance with State-Triggered Constraints,” *Journal of Guidance, Control, and Dynamics*, Vol. 43, No. 8, 2020, pp. 1399–1413.

- [12] Blackmore, L., “Autonomous precision landing of space rockets,” in *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2016 Symposium*, Vol. 46, 2016, pp. 15–20.
- [13] Martin, M. S., Mendeck, G. F., Brugarolas, P. B., Singh, G., Serricchio, F., Lee, S. W., Wong, E. C., and Essmiller, J. C., “In-flight experience of the Mars Science Laboratory guidance, navigation, and control system for entry, descent, and landing,” *CEAS Space Journal*, Vol. 7, 2015, pp. 119–142.
- [14] Carson, J. M., Munk, M. M., Sostaric, R. R., Estes, J. N., Amzajerjian, F., Blair, J. B., Rutishauser, D. K., Restrepo, C. I., Dwyer-Cianciolo, A. M., Chen, G., et al., “The SPLICE project: Continuing NASA development of GN&C technologies for safe and precise landing,” *AIAA SciTech 2019 Forum*, 2019, p. 0660.
- [15] Malyuta, D., Yu, Y., Elango, P., and Açıkmeşe, B., “Advances in trajectory optimization for space vehicle control,” *Annual Reviews in Control*, Vol. 52, 2021, pp. 282–315.
- [16] Malyuta, D., Reynolds, T. P., Szmuk, M., Lew, T., Bonalli, R., Pavone, M., and Açıkmeşe, B., “Convex Optimization for Trajectory Generation: A Tutorial on Generating Dynamically Feasible Trajectories Reliably and Efficiently,” *IEEE Control Systems*, Vol. 42, No. 5, 2022, pp. 40–113. <https://doi.org/10.1109/mcs.2022.3187542>.
- [17] Zhang, X., “OpenBLAS,” <https://github.com/OpenMathLib/OpenBLAS>, 2023.
- [18] Stellato, B., Banjac, G., Goulart, P., Bemporad, A., and Boyd, S., “OSQP: an operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, Vol. 12, No. 4, 2020, pp. 637–672. <https://doi.org/10.1007/s12532-020-00179-2>.
- [19] Domahidi, A., Chu, E., and Boyd, S., “ECOS: An SOCP solver for embedded systems,” *2013 European Control Conference (ECC)*, IEEE, 2013, pp. 3071–3076.
- [20] Schubiger, M., Banjac, G., and Lygeros, J., “GPU acceleration of ADMM for large-scale quadratic programming,” *Journal of Parallel and Distributed Computing*, Vol. 144, 2020, pp. 55–67. <https://doi.org/10.1016/j.jpdc.2020.05.021>.
- [21] Lofberg, J., “YALMIP: A toolbox for modeling and optimization in MATLAB,” *2004 IEEE International Conference on Robotics and Automation*, IEEE, 2004, pp. 284–289.
- [22] Diamond, S., and Boyd, S., “CVXPY: A Python-embedded modeling language for convex optimization,” *The Journal of Machine Learning Research*, Vol. 17, No. 1, 2016, pp. 2909–2913.
- [23] Yu, Y., Elango, P., Açıkmeşe, B., and Topcu, U., “Extrapolated Proportional-Integral Projected Gradient Method for Conic Optimization,” *IEEE Control Systems Letters*, Vol. 7, 2023, pp. 73–78. <https://doi.org/10.1109/LCSYS.2022.3186647>.
- [24] Kamath, A. G., Elango, P., Yu, Y., Mceowen, S., Chari, G. M., Carson III, J. M., and Açıkmeşe, Behçet., “Real-Time Sequential Conic Optimization for Multi-Phase Rocket Landing Guidance,” *IFAC-PapersOnLine*, Vol. 56, No. 2, 2023, pp. 3118–3125. <https://doi.org/10.1016/j.ifacol.2023.10.1444>, 22nd IFAC World Congress.
- [25] Elango, P., Kamath, A. G., Yu, Y., Carson III, J. M., Mesbahi, M., and Açıkmeşe, B., “A Customized First-Order Solver for Real-Time Powered-Descent Guidance,” *AIAA SciTech 2022 Forum*, 2022, p. 0951.
- [26] Kamath, A. G., Elango, P., Kim, T., Mceowen, S., Yu, Y., Carson III, J. M., Mesbahi, M., and Açıkmeşe, B., “Customized real-time first-order methods for onboard dual quaternion-based 6-DoF powered-descent guidance,” *AIAA SciTech 2023 Forum*, 2023, p. 2003.
- [27] Liberzon, D., *Calculus of Variations and Optimal Control Theory: A Concise Introduction*, Princeton University Press, 2011.
- [28] Szmuk, M., Reynolds, T. P., and Açıkmeşe, Behçet., “Successive Convexification for Real-Time Six-Degree-of-Freedom Powered Descent Guidance with State-Triggered Constraints,” *Journal of Guidance, Control, and Dynamics*, Vol. 43, No. 8, 2020, pp. 1399–1413. <https://doi.org/10.2514/1.G004549>.
- [29] Betts, J. T., *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd ed., Advances in Design and Control, SIAM, 2010. <https://doi.org/10.1137/1.9780898718577>.
- [30] Sargent, R. W. H., and Sullivan, G. R., “The development of an efficient optimal control package,” *Proceedings of the 8th IFIP Conference on Optimization Techniques, Würzburg, September 5–9, 1977*, Vol. 7, Springer-Verlag, Berlin/Heidelberg, 1977, pp. 158–168. <https://doi.org/10.1007/bfb0006520>.
- [31] Drusvyatskiy, D., and Lewis, A., “Error bounds, quadratic growth, and linear convergence of proximal methods,” *Mathematics of Operations Research*, Vol. 43, No. 3, 2018, pp. 919–948.
- [32] Nocedal, J., and Wright, S., *Numerical Optimization*, 2nd ed., Springer, 2006.
- [33] “CUDA C Programming Guide,” <https://docs.nvidia.com/cuda/cuda-c-programming-guide/contents.html>, 2023. Accessed: 2023-11-28.