

TOP: Trajectory Optimization via Parallel Optimization towards Constant Time Complexity

Jiajun Yu^{†1,2}, Nanhe Chen^{†1,2}, Guodong Liu^{2,3}, Chao Xu^{1,2}, Fei Gao^{1,2}, and Yanjun Cao^{1,2}

Abstract—Optimization has been widely used to generate smooth trajectories for motion planning. However, existing trajectory optimization methods show weakness when dealing with large-scale long trajectories. Recent advances in parallel computing have accelerated optimization in some fields, but how to efficiently solve trajectory optimization via parallelism remains an open question. In this paper, we propose a novel trajectory optimization framework based on the Consensus Alternating Direction Method of Multipliers (CADMM) algorithm, which decomposes the trajectory into multiple segments and solves the subproblems in parallel. The proposed framework reduces the time complexity to $O(1)$ per iteration with respect to the number of segments, compared to $O(N)$ of the state-of-the-art (SOTA) approaches. Furthermore, we introduce a closed-form solution that integrates convex linear and quadratic constraints to speed up the optimization, and we also present a numerical solution for general convex inequality constraints. A series of simulations and experiments demonstrate that our approach outperforms the SOTA approach in terms of efficiency and smoothness. Especially for a large-scale trajectory, with one hundred segments, achieving over a tenfold speedup. To fully explore the potential of our algorithm on modern parallel computing architectures, we deploy our framework on a GPU and show high performance with thousands of segments.

Index Terms—Motion and Path Planning, Aerial Systems: Applications, Parallel Trajectory optimization

I. INTRODUCTION

TRAJECTORY optimization is an essential part of the motion planning for the mobile robots. It has been extensively studied in the control community for general robotics systems [1]. Its primary purpose is to refine a path generated by path planning into a time-parameterized trajectory that satisfies various constraints and performance goals. Thanks to the great advantage of differential flatness [2], we can recover all the states and inputs of a flat system from its flat output with finite derivatives. This property is widely applied in the trajectory optimization for mobile robots [3] [4], making the problem tractable. However, the traditional trajectory optimization methods still suffer from high time complexity, especially for large-scale tasks. Such as large-scale inspection, environmental monitoring [5], and precision agriculture [6], which inherently demand long trajectories. TrajOpt [7], solves a non-linear programming (NLP) problem using non-linear optimization algorithms like sequential

[†] Equal contribution

This work was supported by National Nature Science Foundation of China under Grant 62103368. (Corresponding authors: Yanjun Cao, Chao Xu. {yanjunhi, cxu}@zju.edu.cn)

¹ State Key Laboratory of Industrial Control Technology, Institute of Cyber-Systems and Control, Zhejiang University, Hangzhou, 310027, China.

² Huzhou Institute of Zhejiang University and Huzhou Key Laboratory of Autonomous Systems, Huzhou, 313000, China.

³ Tsinghua University, Beijing, 100084, China.




Fig. 1. A real-world experiment of a quadrotor executing a large-scale trajectory (top) generated by the proposed method and the visualized parallel optimization process (bottom).

quadratic programming (SQP). Due to its problem structure, the computational complexity of this NLP is at least $O(N^2)$. Even the state-of-the-art (SOTA) algorithm, GCOPTER [8], which achieves $O(N)$ time complexity operations on this representation under various planning requirements, still underperforms on large-scale problems. Here, N denotes the number of pieces, referred to as segments in this paper.

Inspired by the development of modern parallel computing architectures, many researchers have been working on the parallel optimization algorithms in the robotics field. Wang et al. [9] achieve a parallel structure by splitting the manipulator's joint angle trajectory optimization into a sequence of convex quadratic programming problems (QPs) using the Consensus Alternating Direction Method of Multipliers (CADMM) algorithm [10]. However, the safety constraints lack convex guarantees, which may result in the ADMM algorithm failing to converge. Furthermore, its low-order continuity makes it unsuitable for high-performance mobile robots requiring smooth motion, such as agile motion control for multicopters.

To enhance the performance of parallel trajectory optimization, we propose a novel framework named **TOP**. The foundation of **TOP** lies in the optimality conditions [8] for minimizing control effort and the CADMM algorithm framework for high-order continuity. Specifically, we reformulate the trajectory optimization problem in a piecewise manner, where the optimization of each segment is formulated as a

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

convex quadratically constrained quadratic program (QCQP) problem. The high-order continuity between segments is satisfied by a set of corresponding consensus variables. In this way, our approach reduces the time complexity of trajectory optimization from $O(N)$ to $O(1)$ per iteration. Then we address the convex linear and quadratic inequality constraints using the indicator functions and projection methods for a closed-form solution. Furthermore, we also provide a numerical solution for readers to handle the general convex nonlinear constraints. Finally, we conduct a series of simulations and experiments to demonstrate that our approach outperforms the SOTA approach in terms of efficiency and smoothness. Particularly for a large-scale trajectory with one hundred segments, achieving a speedup of over ten times. To fully explore the potential of our algorithms on modern parallel computing architectures, we deploy our framework on a GPU and show high performance with thousands of segments. As shown in Fig.1, a real-world experiment of a quadrotor flying a single large-scale trajectory of over 500 meters with 1000 segments, where the optimization time is less than one second.

The contributions are summarized as follows:

- We reformulate the trajectory optimization in a parallel optimization framework, which reduces the time complexity from $O(N)$ to $O(1)$ while maintaining high-order continuity.
- We propose two approaches to address the inequality constraints: one for convex linear and quadratic inequality constraints with a closed-form solution to speed up the optimization, and the other for general convex inequality constraints with a numerical solution.
- A series of simulations and experiments demonstrate that our approach outperforms the SOTA method in efficiency and smoothness. Especially for a large-scale trajectory, with one hundred segments, achieving over a tenfold speedup. The framework is also deployed on GPU to explore the potential of our algorithm on modern parallel computing architectures.

II. RELATED WORK

A. Motion Planning

Motion planning algorithms are generally categorized into graph search-based methods, sampling-based methods, and trajectory optimization methods. Graph search-based methods, such as Dijkstra [11], A* [11], and JPS, discretize the configuration space into a graph and use search algorithms to find the optimal path. They are effective in low-dimensional space but suffer from the curse of dimensionality. Sampling-based methods, such as PRM, RRT, and RRT*, construct trees or graphs via random sampling in continuous space, enabling rapid exploration of feasible paths. They are scalable to high-dimensional problems but may take unaffordable time cost to get optimal solutions in real-time. Trajectory optimization methods, such as TrajOpt [7], GPOPS-II [12], and GCOPTER [8], numerically optimize trajectories in continuous space to balance various constraints and control goals (e.g., minimal time or minimal energy). They are designed for high-quality applications but suffer from high time complexity. For example, TrajOpt, solves an NLP using non-linear optimization

algorithms like SQP. The underlying NLP structure exhibits at least $O(N^2)$ time complexity. The collocation method used in GPOPS-II has a time complexity of $O(N^3)$ with respect to the number of collocation points, due to solving large systems of equations for each collocation point. The GCOPTER, reduces the complexity of trajectory optimization to $O(N)$ per iteration by utilizing a novel trajectory representation based on optimality conditions [8]. However, the complexity still grows linearly with the number of segments, making it unsuitable for large-scale tasks.

B. Parallel Optimization

Recently, parallel optimization algorithms have attracted great attention in the robotics field. Existing methods mainly focus on multi-agent scenarios [13], with limited research applying to single-agent trajectory optimization. Singh et al. [14] propose a distributed framework by decomposing the mobile-manipulator trajectory optimization into a sequence of convex QPs. However, the concept of splitting the trajectory into multiple segments is not addressed in this work. Wang et al. [9] achieve a parallel structure by splitting the manipulator's joint angle trajectory into multiple segments, and using CADMM to guarantee the continuity between the segments. However, the running speed of the algorithm framework has not been fully explored. Furthermore, its low-order continuity makes it unsuitable for high-quality control for multicopters. Leu et al. [15] consider the scenarios requiring large-scale motion planning but still provide a nonsmooth trajectory. In this paper, we reformulate the trajectory optimization problem into a parallel framework while enforcing the high-order continuity between the segmented trajectory pieces. Our algorithm reduces the time complexity of trajectory optimization from $O(N)$ to $O(1)$ in a single iteration.

III. PARALLEL TRAJECTORY OPTIMIZATION ALGORITHM

In this section, we present our parallel trajectory optimization algorithm **TOP** based on the CADMM. First, we reconstruct the constrained minimum control effort problem in a piecewise manner (Sec. III-A). Then, we reformulate the origin problem into a parallel consensus problem by adding consensus variables to ensure the continuity of the trajectory (Sec. III-B). Next, we propose two approaches to address both convex linear and nonlinear constraints: one is a closed-form solution, and the other is a numerical solution (Sec. III-C). Finally, we present the implementation details of our method (Sec. III-D).

A. Preliminaries

By leveraging the property of differential flatness, we can generate an optimal trajectory for the dynamic's flat output $\sigma(t) : \mathbb{R} \mapsto \mathbb{R}^m, t \in [0, T]$, which enables us to know the full states and inputs of the dynamic system.

The trajectory optimization problem can be formulated as a minimum control effort problem, i.e.,

$$\min_{\sigma(t)} \int_0^T \sigma^{(p)}(t)^\top \mathbf{W} \sigma^{(p)}(t) dt, \quad (1a)$$

$$s.t. \sigma^{[p-1]}(0) = \bar{\sigma}_0, \sigma^{[p-1]}(T) = \bar{\sigma}_N, \quad (1b)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 2. Consensus optimization for multiple segments from an initialized path at Splitting Points, which allows the adjacent segments to converge towards shared Consensus Points.

$$\mathcal{G}(\sigma(t), \dot{\sigma}(t), \dots, \sigma^{(p)}(t)) \preceq \mathbf{0}, \quad \forall t \in [0, T], \quad (1c)$$

where the chain of n -integrators of the flat output is denoted as $\sigma^{[n]} = (\sigma^\top, \dot{\sigma}^\top, \dots, \sigma^{(n)\top})^\top$. Here, $\mathbf{W} \in \mathbb{R}^{m \times m}$ is a diagonal matrix that scales the cost across different dimensions, $\bar{\sigma}_0 \in \mathbb{R}^{mp}$ represents the initial condition, $\bar{\sigma}_N \in \mathbb{R}^{mp}$ represents the terminal condition, $\sigma^{(p)}$ denotes the control input and \mathcal{G} consists of equivalent constraints.

Given a set of time durations $\mathcal{T} = \{T_1, T_2, \dots, T_N\}$, where $T_i \in \mathbb{R}_{>0}$ and N is the number of segments, we reformulate the problem in a piecewise manner:

$$\min_{\sigma(t)} \sum_{i=1}^N \left(\int_0^{T_i} \sigma_i^{(p)}(t)^\top \mathbf{W} \sigma_i^{(p)}(t) dt \right), \quad (2a)$$

$$s.t. \sigma_i^{[d-1]}(T_i) = \sigma_{i+1}^{[d-1]}(0), i \in \{1, \dots, N-1\}, \quad (2b)$$

$$\sigma_1^{[p-1]}(0) = \bar{\sigma}_0, \sigma_N^{[p-1]}(T_N) = \bar{\sigma}_N, \quad (2c)$$

$$\mathcal{G}_i(\sigma_i(t), \dot{\sigma}_i(t), \dots, \sigma_i^{(p)}(t)) \preceq \mathbf{0}, \quad \forall t \in [0, T_i], \quad (2d)$$

where d is the number of derivatives, and Eq. (2b) ensures the corresponding order of continuity.

B. Problem Formulation

1) *Reformulate as a consensus problem:* Here we consider a general consensus problem formulation, with local variables $x_i \in \mathbb{R}^{n_i}$, $i \in \{1, \dots, N\}$ and a global consensus variable $z \in \mathbb{R}^n$:

$$\min \sum_{i=1}^N f_i(x_i), \quad (3a)$$

$$s.t. h(x_i) = z, \quad (3b)$$

where $h : \mathbb{R}^{n_i} \mapsto \mathbb{R}^n$ is the mapping from local variables to the global consensus variable. The consensus constraints in Eq. (3b) ensure that all x_i reach an agreement.

Motivated by this approach, we reformulate the original problem into a consensus problem by considering the following aspects:

- **Independent Problem:** The objective function of each segment depends solely on its local variables, allowing it to be divided into independent local objectives. Meanwhile, the absence of coupling constraints in Eq. (2d) between segments enables each segment to be solved independently.

- **Consensus Constraints:** By introducing consensus variables, continuity constraints in Eq. (2b) are transformed into consensus constraints as described in Eq.(3b).

Based on the previous analysis, the original problem can be reformulated into the consensus problem, i.e.,

$$\min_{\sigma(t)} \sum_{i=1}^N \left(\int_0^{T_i} \sigma_i^{(p)}(t)^\top \mathbf{W} \sigma_i^{(p)}(t) dt \right), \quad (4)$$

$$s.t. \tilde{\sigma}_i^{[d-1]} = \tilde{z}_i,$$

where $\tilde{\sigma}_i^{[d-1]} = (\sigma_i^{[d-1]}(0)^\top, \sigma_i^{[d-1]}(T_i)^\top)^\top$ represents the derivatives of the flat output at both boundaries. The consensus variable $z_i \in \mathbb{R}^{md}$ includes the consensus point along with its derivatives up to the $(d-1)$ -th order, while $\tilde{z}_i = (z_{i-1}^\top, z_i^\top)^\top$ represents the boundaries of the i -th segment. For simplicity, only continuous constraints are considered in the Eq. (4).

2) *CADMM Parallel Optimization:* We solve the consensus problem in the context of the CADMM algorithm [10]. To apply the algorithm, we first give the augmented Lagrangian function for the subproblem of each segment i :

$$\mathcal{L}_i = \int_0^{T_i} \sigma_i^{(p)}(t)^\top \mathbf{W} \sigma_i^{(p)}(t) dt + \frac{\rho}{2} \|\tilde{\sigma}_i^{[d-1]} - \tilde{z}_i + u_i\|_2^2. \quad (5)$$

Let $u_i \in \mathbb{R}^{2md}$ be the scaled dual variable for continuity constraints, defined as $u_i = y_i/\rho$, where ρ is the penalty parameter. So the overall function is given by $\mathcal{L}_\rho = \sum_{i=1}^N \mathcal{L}_i$.

The CADMM method to solve the consensus problem iteratively can be given as:

$$\sigma_i^{k+1} = \arg \min_{\sigma_i} \mathcal{L}_i(\sigma_i, z_i^k, u_i^k), \quad (6a)$$

$$z_i^{k+1} = \frac{1}{2} (\sigma_i^{[d-1]^{k+1}}(T_i) + \sigma_{i+1}^{[d-1]^{k+1}}(0)), \quad (6b)$$

$$u_i^{k+1} = u_i^k + \tilde{\sigma}_i^{[d-1]^{k+1}} - \tilde{z}_i^{k+1}, \quad (6c)$$

where z_i , $i \in \{1, \dots, N-1\}$ are updated by Eq. (6b), and z_0 and z_N , initialized by $\bar{\sigma}_0$ and $\bar{\sigma}_N$ respectively, are kept fixed during the iterations. Fig. 3 shows how local variables σ , consensus variables z , and scaled dual variables u participate in the algorithm during iterations following the circular flow (the slack variables s will be introduced in the next section). All variables are tensorized to make the process parallel.

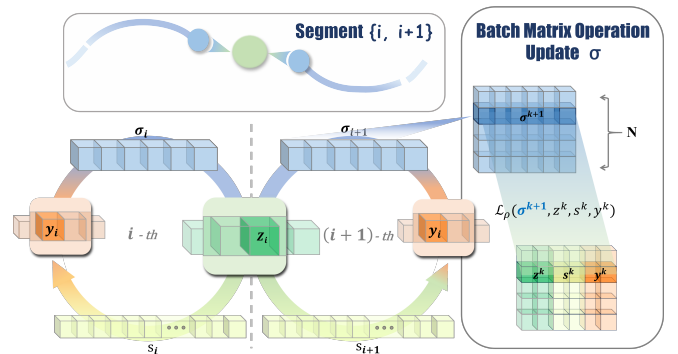


Fig. 3. The parallel optimization process at $(k+1)$ -th iteration to pull two segments $\{i, i+1\}$ together. The circular flow illustrates an iterative process for updating related variables, supported by a batch matrix update of the local trajectory (σ).

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

C. Convex Inequality Constraints

1) *Linear Inequality Constraints*: To address the linear inequality constraints, we introduce slack variables, converting the inequality constraints into equality constraints, making the procedure similar to Eq. (6). In trajectory optimization, collision-free space is often represented as a sequence of convex overlapping polytopes, defined by linear inequality constraints. Hereafter, we take a simple linear inequality constraint as an example:

$$\begin{aligned} & \arg \min_x f(x), \\ & \text{s.t. } \mathbf{A}x \preceq b, \end{aligned} \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and optimization variable $x \in \mathbb{R}^n$. After introducing slack variable $s \in \mathbb{R}^m$, the inequality constraint can be transformed into equality constraint:

$$\begin{aligned} & \arg \min_{x,s} f(x) + I_{s \geq 0}(s), \\ & \text{s.t. } \mathbf{A}x + s = b, \end{aligned} \quad (8)$$

where $I_{s \geq 0}(s)$ is the indicator function, formulated as:

$$I_{s \geq 0}(s) = \begin{cases} +\infty & s < 0, \\ 0 & s \geq 0. \end{cases} \quad (9)$$

Finally, the update rule is the same as that for linear equality constraint in Eq. (6), with the update rule for the slack variable s in the $(k+1)$ -th iteration as:

$$s^{k+1} = \max(0, -(\mathbf{A}x^{k+1} - b + v^k)), \quad (10)$$

where v is the corresponding scaled dual variable.

2) *Quadratic Inequality Constraints*: Besides the linear inequality constraints, convex quadratic inequality constraints are also widely used for trajectory optimization, such as dynamic feasibility, target tracking, and minimal trajectory length. In this work, we address these constraints using a projection method that provides a closed-form solution. An example of a convex quadratic inequality constraint is provided below:

$$\begin{aligned} & \arg \min_x f(x), \\ & \text{s.t. } \|\mathbf{B}x\|^2 \leq c, \end{aligned} \quad (11)$$

where $\mathbf{B} \in \mathbb{R}^{m \times n}$, $c \in \mathbb{R}_{>0}$ and optimization variable $x \in \mathbb{R}^n$. We introduce the $\phi = \mathbf{B}x$, $\phi \in \mathbb{R}^m$, and replace the quadratic inequality constraint with the indicator function: $I_{\|\phi\|^2 \leq c}(\phi)$. Then the original problem Eq. (11) can be transformed into:

$$\begin{aligned} & \arg \min_x f(x) + I_{\|\phi\|^2 \leq c}(\phi), \\ & \text{s.t. } \phi = \mathbf{B}x. \end{aligned} \quad (12)$$

Therefore, the update rule is identical to that for linear equality constraint in Eq. (6). Notably, the update rule for ϕ is:

$$\phi^{k+1} = \arg \min_{\phi} (I_{\|\phi\|^2 \leq c}(\phi) + \frac{\rho}{2} \|\mathbf{B}x^{k+1} - \phi + w^k\|_2^2), \quad (13)$$

where w is the scaled dual variable. The process is equivalent to projecting the vector $\mathbf{B}x + w$ onto the constraint set $\{\phi \in \mathbb{R}^m \mid \|\phi\|^2 \leq c\}$, which can be solved in a closed-form way:

$$\phi = \begin{cases} \sqrt{\frac{c}{\|\mathbf{B}x + w\|^2}} (\mathbf{B}x + w) & \text{if } \|\mathbf{B}x + w\|^2 > c, \\ \mathbf{B}x + w & \text{otherwise.} \end{cases} \quad (14)$$

3) *General Inequality Constraints*: Although the closed-form solution holds for most constraints in multicopter trajectory optimization, certain specialized ones may not be satisfied. Giesen et al. [16] provide a universal approach capable of directly handling arbitrary non-linear convex inequality constraints. However, its mathematical form limits the possibility of obtaining a closed-form solution. In this work, we employ the L-BFGS [17] algorithm to solve the subproblems when updating σ_i during the iterations. Specifically, it replaces the inequality constraints with an equality form by

$$g(x) = \max\{0, g_0(x)\}^2, \quad (15)$$

where $g_0(x)$ is original inequality constraint. By using Eq.(15), we transform the convex inequality constraints into convex equality constraints, differing from the approach of introducing slack variables. The $g(x)$ inherits convexity from $g_0(x)$, but the equality constraint $g(x) = 0$ is no longer linear. In summary, we consider the equality-constrained problem form:

$$\begin{aligned} & \min_x f(x), \\ & \text{s.t. } g(x) = 0. \end{aligned} \quad (16)$$

The augmented Lagrangian function for Eq. (16) is:

$$L_{\rho}(x, y) = f(x) + \frac{\rho}{2} \|g(x)\|^2 + y^{\top} g(x), \quad (17)$$

where $y \in \mathbb{R}^p$ is the Lagrange multiplier. Then the dual variable y^k will be updated similarly to the dual variable of the equality constraint Eq. (6c) in the $(k+1)$ -th iteration.

This method benefits from solving general convex inequality constraints while avoiding the introduction of additional variables and iteration substeps to the original problem. However, numerical solvers like L-BFGS slow down the running speed, which is detailed in the experiments part Sec. IV-A2. Above all, it's an alternative option for users to choose the one that best matches their requirements.

D. Implementation Details

1) *Problem Definition of Minimum-jerk Trajectory*: In this section, we construct a minimum-jerk trajectory ($p = 3$) through the proposed closed-form solution as an example. This trajectory simultaneously enforces collision-free and dynamical feasibility constraints. Meanwhile, we parameterize the trajectory with a set of polynomials. The flat output of the i -th segment $\sigma_i(t)$ is a linear combination of the basis function $\beta(t) = (1, t, \dots, t^d)^{\top}$, i.e.,

$$\sigma_i(t) = \mathbf{c}_i^{\top} \cdot \beta(t - t_{i-1}), \quad t \in [t_{i-1}, t_i], \quad (18)$$

where $\mathbf{c}_i \in \mathbb{R}^{(d+1) \times m}$ represents the coefficients of the i -th polynomial trajectory, and d (set to 5) is the degree of the basis function. As for the collision-free constraints, we use the FIRI [18] to generate the flight corridor, which can be formulated as multiple linear inequality constraints, addressed by Eq. (8). We then apply the projection method in Eq. (12) to enforce dynamic feasibility constraints, using maximum velocity as an example. For the convenience of closed-form operation, we squeeze the \mathbf{c}_i matrix into a vector $c_i \in \mathbb{R}^{m(d+1)}$. Considering

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

the i -th segment ($m = 3$), the consensus subproblem is formulated as:

$$\begin{aligned} \arg \min_{c_i} c_i^\top \widetilde{\mathbf{Q}}(T_i) c_i, \\ \text{s.t. } \widetilde{\mathbf{M}}(T_i) c_i = \widetilde{z}_i, \mathbf{A}_i^o c_i \preceq b_i, \\ \|\mathbf{A}^v(T_{ij}) c_i\|^2 \leq v_{max}^2, \end{aligned} \quad (19)$$

where $\widetilde{\mathbf{Q}}(t) = \mathbf{I}_{m \times m} \otimes \mathbf{Q}(t)$. $\mathbf{Q}(t) = \int_0^t \beta^{(p)}(\tau) \beta^{(p)}(\tau)^\top d\tau \in \mathbb{R}^{(d+1) \times (d+1)}$ is a positive semidefinite matrix. The matrix $\widetilde{\mathbf{M}}(T_i) = \mathbf{I}_{m \times m} \otimes \mathbf{M}(T_i)$ maps the coefficients to the high-order derivatives at both boundaries of the i -th segment, where $\mathbf{M}(t) = (\beta^0(0), \dots, \beta^{(d-1)}(0), \beta^0(t), \dots, \beta^{(d-1)}(t))^\top \in \mathbb{R}^{2d \times (d+1)}$. The free space is encoded as a convex polytope $\mathcal{P}_i^{\mathcal{H}} = \{c_i \in \mathbb{R}^{m(d+1)} \mid \mathbf{A}_i^o c_i \preceq b_i\}$, which is part of the safety flight corridor. The matrix $\mathbf{A}_i^o \in \mathbb{R}^{MK \times m(d+1)}$ and $b_i \in \mathbb{R}^{MK}$, where M represents the number of sample points per segment and K represents the number of hyperplanes per polytope. As for the dynamic constraints, T_{ij} is the timestamp of the j -th intra-sample point of the i -th segment, $j \in \{1, 2, \dots, M\}$, $\mathbf{A}^v(T_{ij}) \in \mathbb{R}^{m \times m(d+1)}$, and v_{max} is the maximal velocity. Finally, the update rule of the $(k+1)$ -th iteration is:

$$\begin{aligned} c_i^{k+1} = \left[2\widetilde{\mathbf{Q}}(T_i) + \rho \widetilde{\mathbf{M}}(T_i)^\top \widetilde{\mathbf{M}}(T_i) + \rho (\mathbf{A}_i^o)^\top \mathbf{A}_i^o + \right. \\ \left. \rho \sum_{j=1}^M \mathbf{A}^v(T_{ij})^\top \mathbf{A}^v(T_{ij}) \right]^{-1} \cdot \left[\widetilde{\mathbf{M}}(T_i)^\top (\widetilde{z}_i - u_i) + \right. \\ \left. (\mathbf{A}_i^o)^\top (b_i - s_i^k - v_i^k) + \sum_{j=1}^M \mathbf{A}^v(T_{ij})^\top (\phi_{ij}^k - w_{ij}^k) \right], \end{aligned} \quad (20a)$$

$$z_i^{k+1} = \frac{1}{2} \left(\widetilde{\beta}(T_i) \cdot c_i^{k+1} + \widetilde{\beta}(0) \cdot c_{i+1}^{k+1} \right), \quad (20b)$$

$$u_i^{k+1} = u_i^k + \widetilde{\mathbf{M}}(T_i) c_i^{k+1} - \widetilde{z}_i^{k+1}, \quad (20c)$$

$$s_i^{k+1} = \max(0, -(\mathbf{A}_i^o c_i^{k+1} - b_i + v_i^k)), \quad (20d)$$

$$v_i^{k+1} = v_i^k + \mathbf{A}_i^o c_i^{k+1} + s_i^{k+1} - b_i, \quad (20e)$$

$$\begin{aligned} \phi_{ij}^{k+1} = \sqrt{\max\left(\|\mathbf{A}^v(T_{ij}) c_i^{k+1} + w_{ij}^k\|^2 / v_{max}^2, 1\right)^{-1}} \\ \cdot (\mathbf{A}^v(T_{ij}) c_i^{k+1} + w_{ij}^k), \end{aligned} \quad (20f)$$

$$w_{ij}^{k+1} = w_{ij}^k + \mathbf{A}^v(T_{ij}) c_i^{k+1} - \phi_{ij}^{k+1}, \quad (20g)$$

where the matrix $\widetilde{\beta}(t) = \mathbf{I}_{m \times m} \otimes (\beta^0(t), \dots, \beta^{(d-1)}(t))^\top \in \mathbb{R}^{md \times m(d+1)}$.

Note that each subproblem updating c_i at each iteration is an overdetermined system, with $m \times 2d$ continuity constraints imposed on an objective with $m \times (d+1)$ degrees of freedom. However, a solution exists due to the relaxation of the constraints in the augmented Lagrangian term.

2) *Solution Existence and Convergence Analysis*: We guarantee the existence of the solution referring to the *Optimality Conditions* [8], which provides the necessary and sufficient conditions for the optimal solution of the minimum control effort problem in Eq. (1) without Eq. (1c), as follows:

- The degree of piecewise polynomial $d = 2p - 1$.
- The boundary conditions are constrained to $(p - 1)$ -th order derivative in Eq. (2c).
- The continuously differentiable constraints in Eq. (2b) are

Algorithm 1: TOP Algorithm Framework

Input: path \mathbf{P} , flight corridor $\mathcal{P}^{\mathcal{H}}$
Output: optimized polynomial coefficients \mathbb{C}

```

1 begin
2   Init( $\mathbf{P}$ ,  $\mathcal{P}^{\mathcal{H}}$ )
3   while  $\|r_p^k\|^2 < \mathcal{E}_r^2$  and  $\|r_d^k\|^2 < \mathcal{E}_d^2$  do
4     Parallel for  $i = 1$  to  $N$  do
5       if closed-form then
6          $c_i^{k+1}, \phi_i^{k+1}, w_i^{k+1} \leftarrow$ 
           ClosedFormSolve( $c_i^k, z_i^k, \phi_i^k, u_i^k, s_i^k, v_i^k, w_i^k$ )
7       else
8          $c_i^{k+1}, w_i^{k+1} \leftarrow$ 
           NumericSolve( $c_i^k, z_i^k, u_i^k, s_i^k, v_i^k, w_i^k$ )
9       end
10       $z_i^{k+1} \leftarrow$  UpdateConsensus( $c_i^{k+1}$ )
11       $r_p^{k+1}, r_d^{k+1} \leftarrow$  UpdateResidual( $c_i^{k+1}, z_i^{k+1}$ )
12       $s_i^{k+1} \leftarrow$  ProjectNonNeg( $c_i^{k+1}, v_i^k$ )
13       $u_i^{k+1}, v_i^{k+1} \leftarrow$ 
        UpdateDual( $c_i^{k+1}, z_i^{k+1}, s_i^{k+1}, u_i^k, v_i^k$ )
14       $\rho^{k+1} \leftarrow$  UpdateRho( $\rho^k$ )
15    end
16  end
17  return  $\mathbb{C}$ 
18 end
```

constrained to $(2p - 1)$ -th order derivative.

The only difference in our method is that we do not constrain the intermediate conditions of the *Optimality Conditions*. The optimal solution lies within our solution space, which is larger, and ensures the solution can be found. Moreover, the convexity of the formulated problem and the foundation of the ADMM [10] framework ensure that the objective converges to an optimal solution as the number of iterations k approaches infinity. Meanwhile, result from [19] guarantees a global linear convergence of our formulation, achieving an independent number of iterations with the segment number.

3) *Stopping Criterion*: The stopping criterion for **TOP** is determined by the primal and dual residuals. The primal residual is defined as the gaps between state variables at the splitting points:

$$\begin{aligned} r_p^{k+1} = \left(\widetilde{\beta}(T_1) c_1^{k+1} - \widetilde{\beta}(0) c_2^{k+1}; \dots; \right. \\ \left. \widetilde{\beta}(T_{N-1}) c_{N-1}^{k+1} - \widetilde{\beta}(0) c_N^{k+1} \right). \end{aligned} \quad (21)$$

The dual residual reflects the convergence degree of the dual variables update, defined as:

$$\begin{aligned} r_d^{k+1} = -\rho \widetilde{\mathbf{M}}(T_i)^\top (\widetilde{z}_1^{k+1} - \widetilde{z}_1^k; \dots; \\ \widetilde{z}_{N-1}^{k+1} - \widetilde{z}_{N-1}^k; \widetilde{z}_N^{k+1} - \widetilde{z}_N^k). \end{aligned} \quad (22)$$

The algorithm will stop when the primal residual and dual residual are both smaller than the predefined tolerance.

$$\|r_p^k\|^2 < \mathcal{E}_r^2, \quad \|r_d^k\|^2 < \mathcal{E}_d^2, \quad (23)$$

where $\mathcal{E}_r = \mathcal{E}_d = N\mathcal{E}$. In practice, we recommend $\mathcal{E} = 0.05$.

4) *Adaptive Rho*: To speed up the convergence of the algorithm, we update the penalty parameter ρ^k adaptively,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 4. Benchmark on trajectory optimization, comparing our method (yellow) with GCOPTER (blue). Our approach requires less computation time (55 ms for TOP compared to 797 ms for GCOPTER) and achieves a lower objective cost (38.6 for TOP compared to 235 for GCOPTER), which corresponds to a smoother trajectory.

making the optimization less sensitive to the initial value ρ^0 . The penalty parameter ρ is updated in each iteration according to the following rule:

$$\rho^{k+1} = \begin{cases} \tau^{\text{incr}} \rho^k & \text{if } \|r_p^k\|_2 > \mu \|r_d^k\|_2, \\ \rho^k / \tau^{\text{decr}} & \text{if } \|r_d^k\|_2 > \mu \|r_p^k\|_2, \\ \rho^k & \text{otherwise,} \end{cases} \quad (24)$$

where $\mu > 1$, $\tau^{\text{incr}} > 1$, and $\tau^{\text{decr}} > 1$. In our implementation, we set $\mu = 10$, $\tau^{\text{incr}} = \tau^{\text{decr}} = 1.1$, and $\rho^0 = 1$. The strategy aims to maintain the primal and dual residuals in the same order of magnitude to accelerate the optimization process.

5) *Algorithm Summary*: The framework of **TOP** is outlined in Alg. 1. It takes the path $\mathbf{P} \in \mathbb{R}^{m \times (N+1)}$ found by the RRT* and the flight corridor $\mathcal{P}^{\mathcal{H}}$ generated by the FIRI as inputs. Line 2 (**Init**) sets up all variables for each segment. After initialization, all trajectory segments are optimized in parallel.

If the closed-form solution is chosen, line 6 (**ClosedFormSolve**) applies Eqs. (20a, 20f, 20g) to update c_i, ϕ_i, w_i in the closed-form. Constraints that cannot be addressed by the aforementioned method can be solved numerically (**NumericSolve**) using Eq. (17). Afterward, the consensus variable z_i is updated in line 10 using Eq. (20b). Line 11 (**UpdateResidual**) updates the primal and dual residuals using Eqs. (21, 22). Line 12 (**ProjectNonNeg**) projects the slack variable s_i into the non-negative space using Eq. (20d). The dual variables u_i and v_i are subsequently updated (**UpdateDual**) using Eqs. (20c, 20e) in line 13. Finally, line 14 updates the penalty parameter ρ via Eq. (24). The algorithm terminates once both the primal and dual residuals drop below the predefined tolerance, returning the optimized polynomial coefficients $\mathbb{C} = \{c_1, c_2, \dots, c_N\}$.

IV. EXPERIMENTS

In this section, we conduct a series of experiments with a quadrotor as the target platform to evaluate the performance of the proposed method. The flat output of the quadrotor is defined as $\sigma(t) = (p_x(t), p_y(t), p_z(t), \psi(t))^{\top}$, where $(p_x(t), p_y(t), p_z(t))^{\top}$ is the position of the quadrotor in the world frame, and $\psi(t)$ is yaw angle of the quadrotor. Similar to GCOPTER [8], we optimize the position of the quadrotor in experiments while fixing the yaw angle during optimization.

A. Simulations

TABLE I: BENCHMARK ON THE TRAJECTORY OPTIMIZATION

Piece Num	Method	Obj. Cost		Iter. Count		Time Cost (ms)	
		Avg	Std	Avg	Std	Avg	Std
30~39	TOP (Parallel)	34.13	26.02	297.32	132.66	78.29	17.62
	TOP (Sequential)	34.13	26.02	297.32	132.66	235.10	60.28
	GCOPTER	179.61	42.05	855.90	300.77	557.30	181.91
40~49	TOP (Parallel)	30.96	12.37	305.31	120.06	88.98	33.95
	TOP (Sequential)	30.96	12.37	305.31	120.06	315.62	139.26
	GCOPTER	232.04	20.22	946.55	216.45	745.70	242.41
50~59	TOP (Parallel)	60.93	26.82	302.50	76.89	108.86	25.81
	TOP (Sequential)	60.93	26.82	302.50	76.89	485.66	140.62
	GCOPTER	296.76	28.18	1075.00	289.21	877.00	527.07
60~69	TOP (Parallel)	64.40	27.66	305.40	59.58	118.40	22.07
	TOP (Sequential)	64.40	27.66	305.40	59.58	564.63	120.18
	GCOPTER	297.15	58.28	1101.30	186.62	1288.80	407.83

1) *Benchmark*: In this experiment, we compare the performance of our method with GCOPTER in three dimensions: optimized objective cost, optimization iteration count and optimization time cost, as shown Tab. I. TOP (Parallel) is a standard implementation of our method by using Thread Building Blocks (TBB) to parallelize the computation on a CPU. Considering GCOPTER’s architecture does not effectively support parallelization, we also implement a sequential version of our method TOP (Sequential) for comparison. All methods are implemented in C++ and tested on a computer with an AMD Ryzen Threadripper Pro 5995wx CPU. The distinction between TOP (Parallel) and TOP (Sequential) lies solely in how the per-segment updates are executed within each ADMM iteration. TOP (Parallel) updates the variables for each index in parallel, such as solving the Eq. (20a) for $i \in \{1, \dots, N\}$ using multiple threads. In contrast, TOP (Sequential) executes all the updates for the N subproblems in one thread. In Tab. I, “Obj. Cost” denotes the integral of the squared jerk of the trajectory. The smaller the value, the smoother the trajectory. “Avg” and “Std” refer to the average and standard deviation over multiple trials. The result shows that our proposed method outperforms GCOPTER in both smoothness and time cost. The optimized trajectories are visualized in Fig. 4.

Note that the iteration count of our method remains nearly constant as the number of segments increases, while the time cost in Tab. I slightly increases. It seems that the results do

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

not reflect the ideal constant-time complexity per iteration. This is because a CPU is a general-purpose processor with a limited number of cores compared to a GPU. Furthermore, its performance is constrained by the scheduling overhead from the operating system kernel, which usually performs complex, high-overhead context switches to mix our computation with numerous other system-level processes. However, due to the superior single-thread performance of CPU cores, it is often the better choice for small-scale problems with a limited number of segments. In the next section, we focus on the GPU implementation to demonstrate the advantages of our parallel optimization framework for large-scale problems.

2) *Ablation Study*: In this experiment, we evaluate the effectiveness of the **UpdateRho** function in Alg. 1 (line 14) by comparing the performance of the proposed method with several variations. Specifically, we remove the **UpdateRho** function and fix ρ to different values (1, 10, and 100) as control groups. The **UpdateRho** dynamically adjusts ρ during optimization to balance the convergence of the primal and dual residuals, allowing our proposed method to achieve faster convergence compared to the fixed ρ method. For these comparisons, the number of segments is fixed at 54. The details are presented in Tab. II. The hardware setup remains consistent with the experiment described in Sec. IV-A1.

TABLE II: ABLATION STUDY ON UPDATING ρ

Method	Iter. Num.	Pri. Res.	Dual Resi.	Obj. Cost	Time Cost (ms)
Adaptive ρ	416	0.7721	0.1817	66.8705	149.1764
Fixed ρ (1)	607	$1.9738e^{-2}$	0.7696	228.6242	195.4008
Fixed ρ (10)	5120	$1.3566e^{-3}$	0.7741	282.6921	1598.8727
Fixed ρ (100)	23778	$9.0375e^{-13}$	0.7741	273.4577	7440.5748

TABLE III: COMPARISON BETWEEN CLOSED-FORM METHOD AND NUMERICAL METHOD

Method	Mean Iter. Num.	Mean Time Cost (ms)
Closed-form	230.9	88.6334
Numerical	157.2	1169.7684

Next, we compare the performance of the proposed method using both the closed-form and numerical methods, as outlined in Alg. 1 (line 5-9). We also reuse the same hardware setting as the experiment in Sec. IV-A1. The results are listed in Tab. III. We find that the general constraints formulation allows the numerical method to converge after fewer iterations. However, the closed-form method is significantly more time-efficient than the numerical method, making it better suited for real-time applications.

Finally, we compare the performance of our approach with different segment numbers. We uniformly sample $N+1$ initial points from a parameterized curve and enforce the optimized trajectories to pass through the sampled points, forming N segments. In this experiment, we implement our code in Python and leverage GPU for parallelization. The computation is conducted on a computer with a NVIDIA GeForce RTX 4060 Laptop GPU. In this experiment, we compare not only the computation time of optimization under different segment numbers but also evaluate the approximation error of the

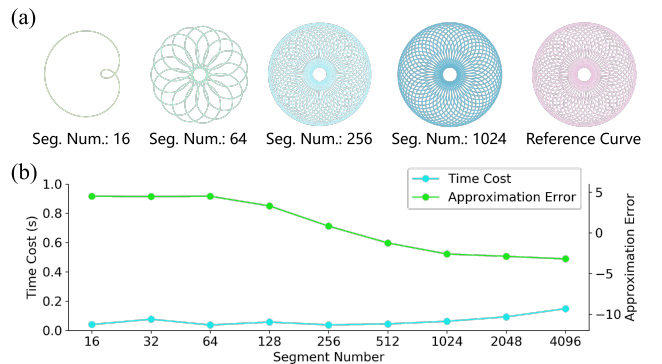


Fig. 5. Performance evaluation with increasing segment numbers on a GPU. (a) The result trajectories for different segment numbers (N) and the reference curve (rightmost). (b) The trends of computation time and approximation error with N .

optimized trajectories to the reference curve. The metric is defined as $e_{ae} = \lg(\sum_{i=1}^M |p_i - p_i^{ref}|)$, where M denotes the number of uniformly sampled points used for evaluation, set to 8192 in practice, and p_i and p_i^{ref} represent the i -th point of the optimized trajectory and the reference curve, respectively. Fig. 5 shows an almost constant runtime as segments increase, demonstrating effective parallelism. Meanwhile, the Approximation Error shows a clear downward trend as the number of segments (N) increases, achieving higher-fidelity trajectories.

B. Real-world Tests

To validate the performance of the proposed method in real-world applications, we conduct a series of experiments with a customized quadrotor platform. The quadrotor is equipped with a NVIDIA Orin NX for on-board computation, and a Livox MID360 for perception. We adopt a modified version of FAST-LIO2 [20] with a relocalization module for the robot's state estimation.

1) *Traverse in Forest*: In this experiment, we test the performance of the proposed method in a forest environment, where the quadrotor is required to traverse through trees and obstacles. The point cloud map is built offline, then we relocalize the quadrotor's pose in the map for planning. We use RRT* to generate a path and then apply the proposed method to optimize the trajectory, producing a smooth and agile trajectory with a maximum velocity of 5 m/s. The trajectory length is 118 m and the computation time is about 50 ms. The result is shown in Fig. 6, where the quadrotor successfully traverses through the trees without collisions. To demonstrate that the position gap between trajectory segments (gap_{TOP}) is acceptable, we compare it with the position gap of sampled discretized state gap_{samp} and the tracking error err_{track} (Fig. 6(c)). First, we observe the gap_{TOP} is consistently smaller than the gap_{samp} along the trajectory, indicating that it does not impact the controller's ability to track the trajectory. Second, the fluctuations in the err_{track} show no evidence of being correlated with gap_{TOP} , as observed in Fig. 6(c). Furthermore, around 15 centimeters position's tracking errors err_{track} are acceptable considering 5 m/s speed flying. Therefore, we prove the gap_{TOP} between segments on the trajectory is acceptable for quadrotor control in real-world tests.

2) *Paint Complex Geometry in the Sky*: In this experiment, we generate a 538.14m dynamical feasible trajectory with

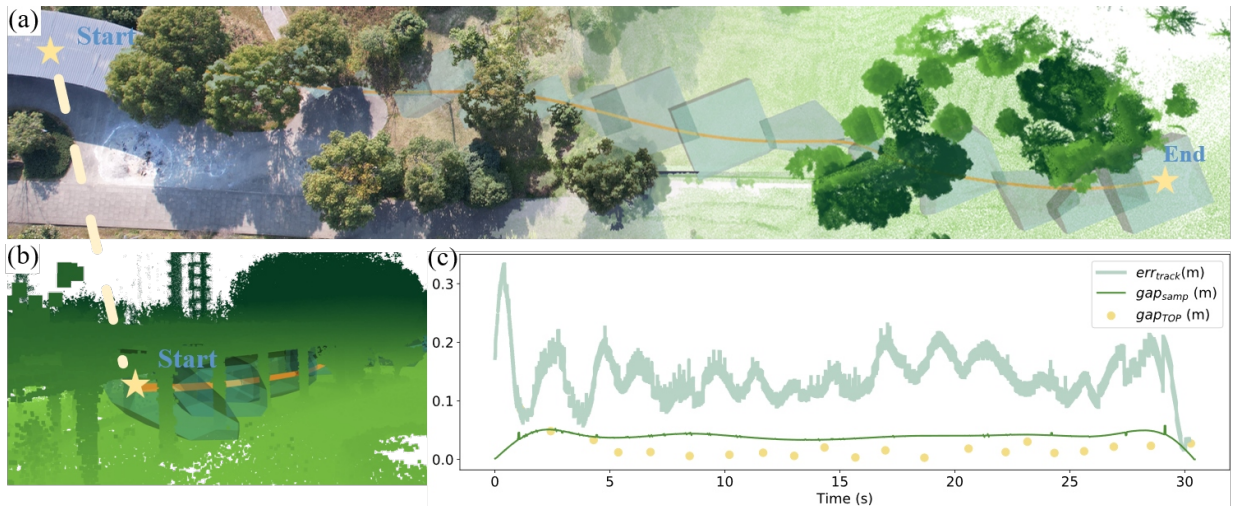


Fig. 6. The top-down (a) and side (b) views of the optimized trajectory and flight corridor in a real-world environment and point cloud map. In (c), as the gap_{TOP} is lower than the gap_{samp} in most cases and there are no sudden changes in err_{track} when encountering gap_{TOP} , we argue that the gap gap_{TOP} would not influence the performance of the quadrotor substantially.

complex geometry via the proposed method. The trajectory consists of 1000 segments and requires only 0.037 seconds for computation. The Fig. 1 demonstrates that our method produces high-quality trajectories that are friendly for the quadrotors. Note that the complex geometry of the trajectory requires numerous segments to recover the path, highlighting the advantages of parallelism in our method.

V. CONCLUSION

In this paper, we propose **TOP**, a trajectory optimization method via parallel optimization towards constant time complexity. The proposed method is based on the CADMM framework, enabling parallel optimization of multiple trajectory segments and reducing the time complexity to a constant level. Experimental results show that the proposed method is tenfold times faster than the SOTA approach when optimizing a large-scale trajectory with one hundred segments, while maintaining superior trajectory quality. Additionally, we introduce two approaches to solve the subproblem: a closed-form solution and a numerical solution. The former is more time-efficient, while the latter is more general and can handle various constraints. In the future, we will explore adaptive resplitting and remerging of trajectories within the environment to further enhance the efficiency of the proposed method.

REFERENCES

- [1] J. T. Betts, *Practical methods for optimal control and estimation using nonlinear programming*. SIAM, 2010.
- [2] M. Fliess, J. Lévine, P. Martin, and P. Rouchon, “Flatness and defect of non-linear systems: introductory theory and examples,” *International journal of control*, vol. 61, no. 6, pp. 1327–1361, 1995.
- [3] D. Mellinger and V. Kumar, “Minimum snap trajectory generation and control for quadrotors,” in *IEEE international conference on robotics and automation*. IEEE, 2011, pp. 2520–2525.
- [4] B. Mu and P. Chirarattananon, “Trajectory generation for underactuated multirotor vehicles with tilted propellers via a flatness-based method,” in *2019 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE, 2019, pp. 1365–1370.
- [5] A. Fascista, “Toward integrated large-scale environmental monitoring using wsn/uav/crowdsensing: A review of applications, signal processing, and future perspectives,” *Sensors*, vol. 22, no. 5, p. 1824, 2022.
- [6] H. Dong, X. Ma, and S. Zhang, “Multi-flight path planning for a single agricultural drone in a regular farmland area,” *Sustainability*, vol. 17, no. 6, p. 2433, 2025.
- [7] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, “Finding locally optimal, collision-free trajectories with sequential convex optimization,” in *Robotics: science and systems*, vol. 9, no. 1. Berlin, Germany, 2013, pp. 1–10.
- [8] Z. Wang, X. Zhou, C. Xu, and F. Gao, “Geometrically constrained trajectory optimization for multicopters,” *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3259–3278, 2022.
- [9] C. Wang, J. Bingham, and M. Tomizuka, “Trajectory splitting: A distributed formulation for collision avoiding trajectory optimization,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8113–8120.
- [10] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein *et al.*, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [11] L. Quan, L. Han, B. Zhou, S. Shen, and F. Gao, “Survey of uav motion planning,” *IET Cyber-systems and Robotics*, vol. 2, no. 1, pp. 14–21, 2020.
- [12] M. A. Patterson and A. V. Rao, “Gpops-ii: A matlab software for solving multiple-phase optimal control problems using hp-adaptive gaussian quadrature collocation methods and sparse nonlinear programming,” *ACM Transactions on Mathematical Software (TOMS)*, vol. 41, no. 1, pp. 1–37, 2014.
- [13] H. Xu, P. Liu, X. Chen, and S. Shen, “d{2} slam: Decentralized and distributed collaborative visual-inertial slam system for aerial swarm,” *IEEE Transactions on Robotics*, 2024.
- [14] A. K. Singh, A. Ahonen, R. Ghabcheloo, and A. Mueller, “Introducing multi-convexity in path constrained trajectory optimization for mobile manipulators,” in *2020 European Control Conference (ECC)*. IEEE, 2020, pp. 1178–1185.
- [15] J. Leu, M. Wang, and M. Tomizuka, “Long-horizon motion planning via sampling and segmented trajectory optimization,” in *2022 European Control Conference (ECC)*. IEEE, 2022, pp. 538–545.
- [16] J. Giesen and S. Laue, “Distributed convex optimization with many convex constraints,” *arXiv preprint arXiv:1610.02967*, 2016.
- [17] D. C. Liu and J. Nocedal, “On the limited memory bfgs method for large scale optimization,” *Mathematical programming*, vol. 45, no. 1, pp. 503–528, 1989.
- [18] Q. Wang, Z. Wang, M. Wang, J. Ji, Z. Han, T. Wu, R. Jin, Y. Gao, C. Xu, and F. Gao, “Fast iterative region inflation for computing large 2-d/3-d convex regions of obstacle-free space,” *arXiv preprint arXiv:2403.02977*, 2024.
- [19] D. Han, D. Sun, and L. Zhang, “Linear rate convergence of the alternating direction method of multipliers for convex composite programming,” *Mathematics of Operations Research*, vol. 43, no. 2, pp. 622–637, 2018.
- [20] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, “Fast-lio2: Fast direct lidar-inertial odometry,” *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.