

# Model Predictive Planning: Trajectory Planning in Obstruction-Dense Environments for Low-Agility Aircraft

Matthew T. Wallace<sup>1</sup>, Brett Streetman<sup>2</sup>, and Laurent Lessard<sup>3</sup>

**Abstract**—We present Model Predictive Planning (MPP), a trajectory planner for low-agility vehicles such as a fixed-wing aircraft to navigate obstacle-laden environments. MPP consists of (1) a multi-path planning procedure that identifies candidate paths, (2) a raytracing procedure that generates linear constraints around these paths to enforce obstacle avoidance, and (3) a convex quadratic program that finds a feasible trajectory within these constraints if one exists. Low-agility aircraft cannot track arbitrary paths, so refining a given path into a trajectory that respects the vehicle’s limited maneuverability and avoids obstacles often leads to an infeasible optimization problem. The critical feature of MPP is that it efficiently considers multiple candidate paths during the refinement process, thereby greatly increasing the chance of finding a feasible and trackable trajectory. We demonstrate the effectiveness of MPP on a longitudinal aircraft model.

## I. INTRODUCTION

We consider the problem of autonomous navigation and obstacle avoidance for a low-agility vehicle, such as a fixed-wing aircraft. In this context, it is important to distinguish between a *path* and a *trajectory*.

A *path* describes *where* the vehicle moves in space as it travels from its starting position to its final destination. In this paper, paths are parameterized as a sequence of spatial coordinates and include the straight lines connecting them. Paths are only concerned with the geometric course that the vehicle follows and do not specify the timing, speed, or orientation of the vehicle as it moves along. A *feasible path* is one that connects the starting position and the final destination and does not intersect any obstacles. Methods like Rapidly-exploring Random Trees\* (RRT\*) iteratively construct paths that are guaranteed to converge to the global shortest feasible path and can handle many geometric constraints [1].

A *trajectory* describes *how* the vehicle moves in space. It specifies how the full state of the vehicle evolves as a function of time. A *feasible trajectory* is one that connects the starting position and the final destination, satisfies the equations of motion of the vehicle, and does not collide with any obstacles. For high-agility vehicles such as quadrotors, we can navigate around obstacles by generating a feasible path and using a reference-tracking controller to track the path. However, this approach may fail with low-agility

vehicles since they cannot change directions abruptly and may collide with nearby obstacles as they attempt to follow the prescribed path.

In general, trajectory optimization amounts to solving a nonlinear optimal control problem [2] and is far more difficult than finding a feasible path. Examples of tools that can perform offline trajectory optimization include PROPT [3] and GPOPS-II [4]. As with most nonlinear optimization solvers, these methods find locally optimal solutions and are highly dependent on having good initial guesses. Including obstacle avoidance can be problematic because it can create more local minima and further increase the reliance on initial guesses. For example, if an obstacle were to split the space in half so that a vehicle could go either left or right, and only one of these two choices leads to a feasible trajectory, the wrong initial guess would cause the solver to fail. See Fig. 1 for an illustration of some difficult scenarios for trajectory optimization. The strong dependence on initial guesses persists when performing trajectory optimization in an in-the-loop fashion using a traditional Model Predictive Control (MPC) architecture.

To address the problem of dependence on initial guess, we propose a new approach for in-the-loop trajectory optimization and obstacle avoidance called Model Predictive Planning (MPP). The main components of MPP are as follows.

- 1) A multi-path planner based on RRT\*-AR [5] that generates several candidate paths.
- 2) A raytracing procedure that, for a given candidate path, generates linear constraints that enforce obstacle avoidance.
- 3) A convex quadratic program that uses the linear constraints to refine each candidate path into a feasible trajectory if one exists. We then select the best feasible trajectory among those that were found.

MPP works in an in-the-loop fashion similar to MPC. However, since MPP does not require solving nonlinear optimization problems, it can more efficiently evaluate whether a given candidate path can lead to a feasible trajectory. Therefore, the available computational effort can be spent evaluating numerous candidate paths. MPP is analogous to a breadth-first search (search many paths) rather than the depth-first approach employed in trajectory optimization (use a single path to initialize a nonlinear solver).

A number of motion planning approaches employ multiple trajectories. *Motion primitive planning* is one such method that has been demonstrated on quadrotors [6], robotic manipulators [7], and fixed wing drones [8]. This method reliably avoids obstacles, but can produce inefficient trajec-

This work was supported by a Draper Scholarship

<sup>1</sup>M.T. Wallace is a Draper Scholar and is with the dept. of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115, USA

<sup>2</sup>B. Streetman is with The Charles Stark Draper Laboratory, Cambridge, MA 02139, USA

<sup>3</sup>L. Lessard is with the dept. of Mechanical and Industrial Engineering, Northeastern University, Boston, MA 02115, USA

l.lessard@northeastern.edu




Fig. 1. A visualization of some challenges of trajectory optimization around obstacles. **Left:** the trajectory is caught in a flytrap; the vector pointing to the goal directly opposes the obstacle constraint vector. **Middle:** the constraint vector for the obstacle directly opposes the dynamic constraint. These two cases result in optimizers returning no solution. **Right:** optimizer chooses the long way around the obstacle and gets stuck in a local minimum.

tories since it involves selecting among a limited subset of possible control actions. Many trajectories are also used in *stochastic optimization* methods such as Stochastic Trajectory Optimization for Motion Planning (STOMP) [9] and Model Predictive Path Integral control (MPPI) [10]. These methods refine a baseline trajectory by sampling nearby trajectories, which differs from MPP in that they cannot explore topologically distinct trajectories. Another relevant work is the Fast and Safe Trajectory Planner (FASTER) [11], which plans two trajectories (aggressive and safe) for a quadrotor navigating in an unknown environment.

A few works explicitly address using multiple trajectories to avoid problems with local minima. Zhou et al. [12] use a topological search to generate multiple paths for a quadrotor. These paths are refined into B-spline trajectories. Our work differs in our choice of path planner, an RRT\* variant rather than a topological planner, and that we use full trajectory optimization rather than path smoothing. Using the RRT\* variant could potentially generate more diverse paths in a field with many obstacles as the topological search was limited to only 5 paths due to computational limitations. Using full trajectory optimization allows our method to apply to low-agility systems like fixed-wing aircraft.

*Differential Dynamic Programming* (DDP) [13] is a popular numerical approach for solving the Hamilton–Jacobi–Bellman equation via alternating shooting and back-propagation. Maximum Entropy DDP [14], a recent adaptation of DDP, is able to find multiple paths through an obstacle field. The strategy has the benefit of combining the path planning and optimization steps used in our algorithm, but the optimization problem proposed lacks the guarantees of convex optimization and additionally cannot handle constraints explicitly, but instead uses soft constraints.

Topological-driven Model Predictive Control (T-MPC) [15] uses a *probabilistic roadmap* and a depth-first search to generate paths through a field of moving obstacles and then solves local nonconvex optimization problems to refine each path. Our approach uses convex optimization, which is generally more computationally efficient and reliable.

None of the aforementioned works explicitly treat the problem of obstacle avoidance for fixed-wing aircraft. Some works on fixed-wing aircraft include [16], [17], but these focus on modifying an existing flight path to avoid a detected obstacle rather than reoptimizing the full trajectory.

Section II presents our system architecture, Section III develops the path planner and trajectory optimizer algorithms, Section IV demonstrates MPP on a longitudinal flight model, and Section V states conclusions and future work.

## II. SYSTEM MODEL

Our system architecture, shown in Fig. 2, is inspired by Ryll et al. [6]. We discuss our plant and dynamics in Section II-A. The high-frequency (HF) controller operates on information provided by the airspeed and gyro sensors to regulate the vehicle, which is discussed in Section II-B. The lidar sensor model provides information about nearby obstacles, which is refined into an occupancy grid in the perception model, discussed in Section II-C.

### A. Aircraft Dynamics Model

We use the following longitudinal flight dynamics model. The states are downrange distance, vertical height, airspeed, pitch, pitch rate, and flight path angle  $(x, z, v, \theta, \dot{\theta}, \gamma)$ . The inputs are thrust and elevator deflection  $(T, \delta_e)$  [18]. The equations of motion are

$$\begin{aligned} \dot{x} &= v \cos \gamma, & \dot{v} &= -\frac{D}{m} - g \sin \gamma + \frac{T}{m} \cos \alpha, \\ \dot{z} &= v \sin \gamma, & \dot{\gamma} &= \frac{L}{mv} - \frac{1}{v} \cos \gamma - \frac{T}{mv} \sin \alpha, \\ \ddot{\theta} &= I_{yy}^{-1} M(\delta_e). \end{aligned}$$

The angle of attack  $(\alpha := \theta - \gamma)$  is defined as the angle between the wings and oncoming airflow. Drag  $(D)$ , lift  $(L)$ , and pitching moment  $(M)$  are aerodynamic terms that depend on the states and inputs (see Appendix). Mass  $(m)$ , moment of inertia  $(I_{yy})$ , and gravity  $(g)$  are fixed parameters. We modeled our system as having airspeed and gyro sensors, which provide measurements of airspeed and pitch, respectively, each with Gaussian errors. See Fig. 3 for a diagram of the modeled forces and angles.

### B. High-Frequency Controller

The high-frequency (HF) controller consists of a Linear Quadratic Regulator (LQR) and an Extended Kalman Filter (EKF). The LQR regulates the flight path angle response of the system using elevator action. We linearize around straight and level flight conditions. The HF controller is responsible for tracking the desired pitch angle determined by MPP. The HF controller states are  $(v, \theta, \dot{\theta}, \gamma)$ , input is  $\delta_e$ , and the associated  $Q$  and  $R$  matrices for the LQR controller were chosen as  $Q = \text{diag}(1, 1, 0, 1000)$  and  $R = \frac{1}{2}$ . We used  $Q_{33} = 0$  because  $\dot{\theta}$  is not directly measured and its estimate is noisy. We used  $Q_{44} = 1000$  because the regulation of flight path angle is critical for obstacle avoidance. This HF control loop does not attempt to correct for positional error, which is only considered in the trajectory planner. We used the continuous time LQR solver from the Python Control Systems Library [19] to calculate our LQR gains.

The EKF estimates the true states based on pitch and airspeed indicator measurements. The estimated state is used both in the HF controller and trajectory planner.

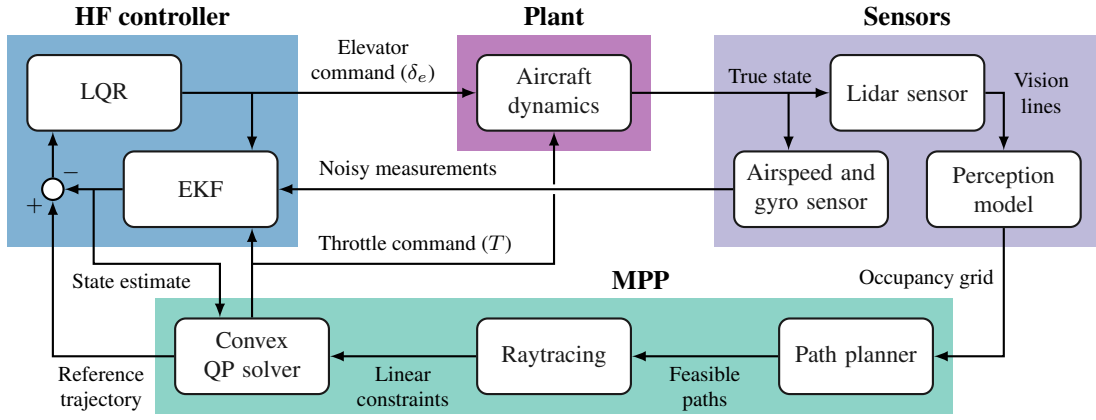


Fig. 2. The system architecture, which consists of four components: the plant, sensor models, high-frequency (HF) controller, and model predictive planner (MPP). The HF controller regulates the system and also provides a high-frequency (100 Hz) estimate of the airspeed and angles for trajectory optimization. MPP begins with planning a path from the occupancy grid, which is then used to determine constraints on an optimization problem through raytracing, and finally a convex quadratic program (QP) is solved to find the trajectory.




Fig. 3. Free body diagram of forces acting on the system.

### C. Perception model

We modeled our perception system using a generated obstacle field and simulated lidar. The model identified objects at a particular angle and range with an error proportional to the range. These ray measurements are converted into an occupancy grid: the  $(x, z)$  coordinates where lidar rays intersect obstacles in the field of view. The first pane of Fig. 4 shows the output occupancy grid from the perception model.

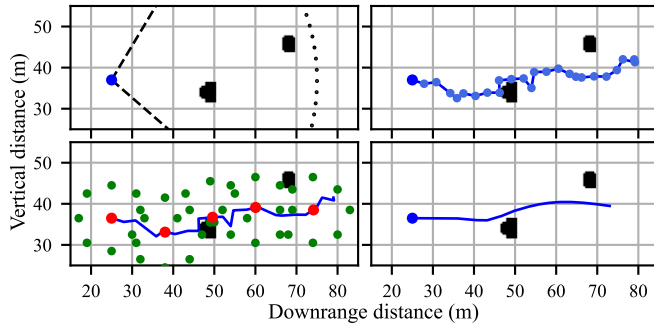


Fig. 4. The MPP trajectory creation process. **Top-left:** Lidar sweep used to create an occupancy grid. **Top-right:** Feasible path generated from the occupancy grid and RRT\*-AR. **Bottom-left:** Raytracing around the path. For clarity, only 1/100 of the collocation points are shown (in red); see Fig. 5 for more details. **Bottom-right:** Trajectory produced from the raytracing by solving a convex QP.

### III. MPP ALGORITHM

Pseudocode for the MPP algorithm is given in Algorithm 1. The inputs to this method are an estimated dynamic state and occupancy grid, supplied by the EKF and perception model respectively. The path planner, a variation on RRT discussed in Section III-A, takes in the occupancy grid and identifies paths through this space towards the goal state of 54 meters directly ahead. The raytracing procedure converts these paths and the occupancy grid to convex optimization constraints and is discussed in Section III-B. The resulting optimization problem is solved in CVXOPT, which is discussed in Section III-C. Fig. 4 shows the algorithm run with a single path. Because a path determined by a path planning may not correspond to a feasible trajectory, we optimize several paths rather than a single one. This approach is shown in Fig. 6.

#### Algorithm 1 Model Predictive Planning

```

 $x_0 \leftarrow$  EKF
OccupancyGrid  $\leftarrow$  PerceptionModel
Paths  $\leftarrow$  RRT*-AR(OccupancyGrid)
Trajectory*
for Path in Paths do
    PositionConstraints  $\leftarrow$  Raytrace(Path, OG)
    Trajectory  $\leftarrow$  CVX( $x_0$ , PositionConstraints)
    Trajectory*  $\leftarrow$  mincost(Trajectory, Trajectory*)
end for

```

#### A. Path Planning

Path Planning is carried out with the RRT variant RRT\*-AR, which differs from RRT\* in that it adds a cost penalty to close nodes sharing the same parent node [5]. Like standard RRT, the RRT\*-AR outputs obstacle-free paths. This approach produces several suboptimal alternate routes to the goal. We found running RRT\*-AR a single time provided insufficient dispersion of paths, but that sufficient dispersion was achieved by rerunning RRT\*-AR with a low number of

samples each time. The results of a single run are shown in the second pane of Fig. 4 and for 50 sampled paths in the first pane of Fig. 6.




Fig. 5. Visualization of raytracing (see bottom-left pane in Fig. 4). **Left:** Collocation points are placed along the path (equally spaced in time). **Middle:** Rays (in green) are drawn in different directions from a collocation point (in red) until they reach an obstacle or travel a given distance. **Right:** Each ray is the normal vector to a linear constraint and these constraints define a convex region. This process is repeated with all collocation points. Each point is allowed to move within its prescribed convex region subject to satisfying vehicle dynamics in order to produce an improved path.

### B. Raytracing

The raytracing procedure determines a feasible region around a path in which to search for a trajectory. While the optimization problem over the whole space is not convex, this local search problem is convex and thus much more practical to solve online. Our approach is inspired by the algorithm from Meerpohl et al. [20].

Beginning from points on the path, we check the occupancy grid (stored as a binary matrix) along 8 directions, returning when an occupied point is encountered. This process is visualized in Fig. 5. Hit points are shown in the third pane of Fig. 4. Hit points are translated to their position in a linearized straight and level vehicle frame and then an inner product is taken with the search direction, resulting in a linear constraint. Constraints are padded slightly to account for rounding error and provide margin for tracking.




Fig. 6. Visualization of the multi-trajectory process. The process is similar to Fig. 4, but many paths are planned instead of just one. Each candidate path is refined into a trajectory. Infeasible trajectories are rejected and the best feasible trajectory is chosen.

### C. Trajectory Optimization

The trajectory optimization problem we solve is a convex quadratic program (QP). For an introduction on convex optimization for trajectory generation, see Malyuta et al. [21]. This method is a form of direct optimization, solving for a perturbation off an initial guess rather than exploiting the Pontryagin principle. Since our dynamics are near-linear, we were able to obtain good results with a single perturbation step using straight and level flight as an initial guess without the need for iterative linearization as is often necessary in highly nonlinear cases. We used CVXOPT [22] to solve the QP. The optimization variable is  $\bar{x}$ , the perturbation of states and controls from straight and level flight over the forward time horizon. The problem parameters are the: quadratic cost  $\bar{Q}$ , equality constraint from the dynamics  $\bar{A}$ , vector of zeros in the equality constraint  $\bar{b} = 0$ , inequality constraint matrix to capture the search directions  $\bar{G}$ , and corresponding vector from raytracing  $\bar{h}$ . The optimization problem is

$$\min \bar{x}^T \bar{Q} \bar{x} \quad \text{s.t.} \quad \bar{A} \bar{x} = \bar{b}, \quad \bar{G} \bar{x} \leq \bar{h}.$$

The system is linearized around straight and level flight.  $x_0$  is the difference from those conditions at the current time. Those dynamics are captured by the equality constraints

$$\bar{x} = \begin{bmatrix} \delta x_0 \\ \delta u_0 \\ \vdots \\ \delta x_{n-1} \\ \delta u_{n-1} \\ \delta x_n \end{bmatrix}, \quad \bar{A} = \begin{bmatrix} I & 0 & 0 & \cdots & 0 \\ A & B & -I & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & A & B & -I \end{bmatrix}, \quad \bar{b} = \begin{bmatrix} x_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

where  $A$  and  $B$  are the linearizations of the dynamics with respect to the entire state and control input, respectively. The feasible region found from raytracing is encoded as inequality constraints. The unit length search directions  $g_1, \dots, g_p$  are identical for each point and determine the inequality constraint matrix  $\bar{G}$ :

$$G = [g_1 \quad \cdots \quad g_p]^T, \quad \bar{G} = \begin{bmatrix} I_n \otimes [G \quad 0_{p \times 2}] \\ G \end{bmatrix}.$$

For the experiments in this paper, we used  $p = 8$ .

The raytracing vector is calculated by taking the inner product of the vector to the ‘‘hit’’ point ( $o$ ) from the collocation point and the search direction. The inner product of the new reference coordinate and the search direction must be less than this value of the coordinate to be inside the open region. Therefore, we have

$$h_i = [g_1^T o_{i1} \quad \cdots \quad g_p^T o_{ip}]^T, \quad \bar{h} = [h_0^T \quad \cdots \quad h_n^T]^T.$$

An additional inequality constraint is used to ensure the thrust  $T$  is nonnegative.

The cost function consists of a penalty on the control and the terminal state’s deviations from straight and level flight values. We chose the cost penalty matrix to be

$$\bar{Q} = \begin{bmatrix} I_n \otimes \begin{bmatrix} 0_{6 \times 6} & 0_{6 \times 2} \\ 0_{2 \times 6} & C_{\text{ctrl}} \end{bmatrix} \\ 100 \cdot I_6 \end{bmatrix}, \quad C_{\text{ctrl}} = \begin{bmatrix} 1 & 0 \\ 0 & 3 \end{bmatrix}.$$

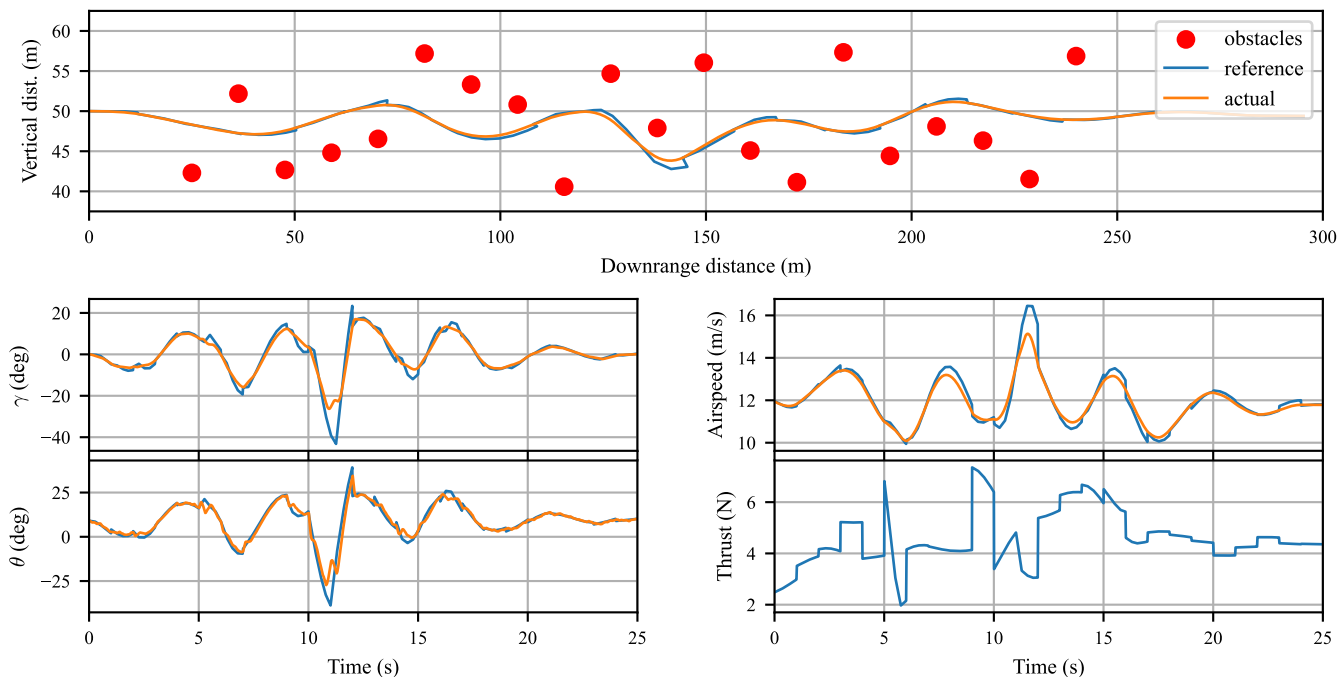


Fig. 7. A sample run of MPP with 20 obstacles and 25 paths per iteration. The reference trajectory is the output of MPP, which is then tracked by the HF controller in closed loop. The reference trajectory (blue) is jagged because it is re-planned every second (with a 4.5 second horizon), at which point it snaps to the current estimate of the system state. The true state is shown in orange.

We chose  $C_{ctrl}$  and the final weight of 100 empirically; using a greater penalty for thrust than elevator produced smoother trajectories. Obstacle avoidance is enforced solely in the constraints and does not appear in the cost function.

The problem is convex by construction since  $\bar{Q}$  is positive semidefinite. Intuitively, the linear constraints provided by the raytracing step yield a convex inner approximation of the original non-convex occupancy grid.

#### D. Dynamic modifications

Trajectory optimization is conducted using a time resolution of 0.25 seconds. This resolution is inadequate for accurately capturing the rotational dynamics exhibited by the system. To address this limitation, we introduced adjustments to the equality constraints associated with the rotational states. The differential equations are transformed into algebraic ones, constraining these states to their steady-state values after oscillations have decayed. In scenarios involving a single variable, this transformation takes on the following structure:  $x_{t+1} = Ax_t + Bu_t \implies (I - A)x_{t+1} = Bu_t$ . In the optimization equations, this takes the form

$$\bar{A}_i = \begin{bmatrix} 0 & \cdots & 0 & B & (I - A) & 0 & \cdots & 0 \end{bmatrix}.$$

In addition to allowing the planner to run at a lower resolution, this procedure prevents negative interactions between the HF controller and the path planner.

#### IV. RESULTS

Fig. 7 shows a typical successful run of MPP using 25 path candidates through a field of 20 1-m radius obstacles. The vehicle’s vision is limited to 45 m ahead and it cruises at

12 m/s. The trajectory is re-planned every 4.5 s (54 m). The reference trajectory and open loop commands are calculated at 1 Hz, while the controller and dynamics resolve at 100 Hz. This simulation demonstrates that MPP can effectively plan feasible trajectories through highly obstructed environments when implemented in a receding horizon fashion.




Fig. 8. MPP success rate based on number of paths and number of obstacles. With a single path, using MPP to refine the path into a feasible trajectory improves RRT\*’s success rate from 19% to 35% in a cluttered environment (20 obstacles). The improvement increases to 90% when using 25 paths. Using 1 path is representative of traditional MPC. Each dot is the average of 100 trials. Shaded region indicates 95% confidence interval.

Fig. 8 shows the aggregate performance of many runs over different randomly generated obstacle fields with different numbers of obstacles. For each run, “failure” occurs if either an extreme state is reached (pitch exceeds  $60^\circ$  or flight path angle exceeds  $45^\circ$ ) or there is a collision with an obstacle. For comparison, we also plot the performance of RRT\* without any refinement. In other words, the feedback controller attempts to directly track the path provided by

RRT\*. This approach often fails because such paths may make abrupt changes that the vehicle cannot track. Refining the path using MPP yields improvement, but there are still many failures due to the fact that the initial candidate path may not lead to a feasible trajectory. As we include more candidate paths, the success rate increases. The tradeoff for this higher success rate is greater computational load, which is proportional to the number of paths.

Code that reproduces all the simulations and figures in this paper is available at <https://github.com/QCGroup/mpp>.

## V. CONCLUSION

MPP is a modification of traditional MPC to plan trajectories around obstacles and is well-suited for low-agility vehicles. The key feature of MPP is to use many candidate paths, thereby improving the chance of finding a feasible trajectory.

We demonstrated MPP on a longitudinal flight model and found that MPP performance improves markedly as we increase the number of candidate paths, particularly when there are many obstacles. Moreover, MPP performs better than using a path-planner alone due to the low agility of the vehicle.

There are several possible avenues for improvement of MPP. First, the RRT\*-AR path planner often produces topologically equivalent paths, which leads to redundant computation. If the path planner could be modified to produce topologically distinct paths, it would lead to a substantial performance improvement. Second, MPP could potentially benefit from multiple rounds of linearization and refinement for more strongly nonlinear systems. Finally, raytracing could be accomplished with a faster parallel implementation, which would solve the primary computational bottleneck of the algorithm.

## REFERENCES

- [1] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [2] A. V. Rao, "Trajectory optimization: a survey," *Optimization and optimal control in automotive systems*, pp. 3–21, 2014.
- [3] P. E. Rutquist and M. M. Edvall, "PROPT-Matlab optimal control software," *Tomlab Optimization Inc*, vol. 260, no. 1, p. 12, 2010.
- [4] 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.
- [5] S. Choudhury, S. Scherer, and S. Singh, "RRT\*-AR: Sampling-based alternate routes planning with applications to autonomous emergency landing of a helicopter," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 3947–3952.
- [6] M. Ryll, J. Ware, J. Carter, and N. Roy, "Efficient trajectory planning for high speed flight in unknown environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 732–738.
- [7] R. A. Shyam, P. Lightbody, G. Das, P. Liu, S. Gomez-Gonzalez, and G. Neumann, "Improving local trajectory optimisation using probabilistic movement primitives," in *International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 2666–2671.
- [8] E. Bulka, *Control and obstacle avoidance for agile fixed-wing aircraft*. McGill University (Canada), 2021.
- [9] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *IEEE International Conference on Robotics and Automation (IRCA)*, 2011, pp. 4569–4574.

- [10] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1714–1721.
- [11] J. Tordesillas, B. T. Lopez, M. Everett, and J. P. How, "Faster: Fast and safe trajectory planner for navigation in unknown environments," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 922–938, 2022.
- [12] B. Zhou, F. Gao, J. Pan, and S. Shen, "Robust real-time UAV replanning using guided gradient-based optimization and topological paths," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 1208–1214.
- [13] D. H. Jacobson, "New second-order and first-order algorithms for determining optimal control: A differential dynamic programming approach," *Journal of Optimization Theory and Applications*, vol. 2, pp. 411–440, 1968.
- [14] O. So, Z. Wang, and E. A. Theodorou, "Maximum entropy differential dynamic programming," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 3422–3428.
- [15] O. de Groot, L. Ferranti, D. Gavrilu, and J. Alonso-Mora, "Topology-driven parallel trajectory optimization in dynamic environments," *arXiv preprint arXiv:2401.06021*, 2024.
- [16] Z. Lin, L. Castano, E. Mortimer, and H. Xu, "Fast 3D collision avoidance algorithm for fixed wing UAS," *Journal of Intelligent & Robotic Systems*, vol. 97, no. 3, pp. 577–604, 2020.
- [17] E. Bulka and M. Nahon, "High-speed obstacle-avoidance with agile fixed-wing aircraft," in *International Conference on Unmanned Aircraft Systems (ICUAS)*, 2019, pp. 971–980.
- [18] R. F. Stengel, "Longitudinal motions," in *Flight Dynamics*. Princeton University Press, 2005, ch. 5.
- [19] S. Fuller, B. Greiner, J. Moore, R. Murray, R. van Paassen, and R. Yorke, "The Python control systems library (python-control)," in *IEEE Conference on Decision and Control*, 2021, pp. 4875–4881.
- [20] C. Meerpohl, M. Rick, and C. Büskens, "Free-space polygon creation based on occupancy grid maps for trajectory optimization methods," *IFAC-PapersOnLine*, vol. 52, no. 8, pp. 368–374, 2019, IFAC Symposium on Intelligent Autonomous Vehicles.
- [21] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Açıkmeşe, "Convex optimization for trajectory generation: A tutorial on generating dynamically feasible trajectories reliably and efficiently," *IEEE Control Systems Magazine*, vol. 42, no. 5, pp. 40–113, 2022.
- [22] M. S. Andersen, J. Dahl, and L. Vandenbergh, "CVXOPT: A python package for convex optimization, version 1.3.2," 2023, available at <https://cvxopt.org>.

## APPENDIX

Lift, drag, and pitching moment are modeled as functions of coefficients ( $C_L$ ,  $C_D$ , and  $C_M$  respectively) that depend on velocity ( $v$ ), elevator angle ( $\delta_e$ ), angle of attack ( $\alpha := \theta - \gamma$ ), and various parameters.

$$L = \frac{1}{2}\rho v^2 S C_L, \quad D = \frac{1}{2}\rho v^2 S C_D, \quad M = \frac{1}{2}\rho v^2 S c C_M,$$

$$C_L = C_{L0} + C_{L\alpha}\alpha, \quad C_D = C_{D0} + K C_L^2,$$

$$C_M = C_{M0} + C_{M\alpha}\alpha + C_{M\dot{\alpha}}\dot{\alpha} + C_{M\delta_e}\delta_e.$$

The parameters used in our simulations are based on estimates for a small hobby remote-controlled aircraft:

$$m = 3.2 \text{ kg}, \quad C_{L\alpha} = 5.73 \text{ rad}^{-1}, \quad C_{L0} = 0.5,$$

$$K = 0.05, \quad C_{M\alpha} = -8.02 \text{ rad}^{-1}, \quad C_{M0} = 0.5,$$

$$S = 0.25 \text{ m}^2, \quad C_{M\dot{\alpha}} = -0.46 \text{ s rad}^{-1}, \quad C_{M\delta_e} = 0.2,$$

$$c = 0.13 \text{ m}, \quad \rho = 1.225 \text{ kg m}^{-3}, \quad C_{D0} = 0.1,$$

$$g = 9.81 \text{ m s}^{-2}, \quad I_{yy} = 0.17 \text{ kg m}^2.$$