Fourteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2021)

Trajectory Planning for Mission Survivability of Autonomous Vehicles in Moderately to Extremely Uncertain Environments

Fanruiqi Zeng
Georgia Institution of Technology
Atlanta, Georgia
fanruiqi.zeng@gatech.edu

Husni R. Idris
NASA Ames Research Center
Moffett Field, California
husni.r.idris@nasa.gov

John-Paul Clarke
The University of Texas at Austin
Austin, Texas
johnpaul@utexas.edu

***Abstract*—Trajectory planning is a particularly challenging task for autonomous vehicles when there are moderate to extreme uncertainties in their operating environment, i.e., where the trajectories of hazards are *partially known* to *completely unknown*. In this paper, we propose a receding horizon control strategy with novel trajectory planning policies that enable dynamic updating of the planned trajectories of autonomous vehicles. The proposed policies utilize two metrics: (1) the number of feasible trajectories; and (2) the robustness of the feasible trajectories. We measure the effectiveness of the suggested policies in terms of mission survivability, which is defined as the probability that the primary mission is accomplished or, if that is not possible, the vehicle lands safely at an alternative site. We show that a linear combination of both metrics is an effective objective function when there is a mix of partially known and unknown uncertainties. When the operating environment is dominated by unknown disturbances, maximizing the number of feasible trajectories results in the highest mission survivability. These findings have significant implications for achieving safe aviation autonomy.**

***Keywords*—non-cooperative traffic; unknown disturbances; non-modeled uncertainty; mission robustness; mission survivability; urban air mobility; trajectory flexibility**

I. Introduction

Highly autonomous vehicles must both operate and make decisions without external support or supervision. This is particularly challenging when there are exogenous uncertainties, i.e., when the locations and future trajectories of weather systems and other vehicles are partially to completely unknown.

Several methods have been proposed to characterize and/or bound the uncertainty associated with dynamic obstacles and thereby account for their occurrence and motion. However, it is difficult to characterize or for that matter bound the behavior of highly unpredictable occurrences, or model unobserved or complex dynamics (such as non-cooperative intruder traffic or a flock of birds crossing the path of a vehicle) [1]. In this study, we focus on the uncertainty caused by the trajectories of uncertain intruder traffic.

To date, researchers who have studied the problem of trajectory planning for highly autonomous vehicles in uncertain operating environments have focused on scenarios where the uncertainty can be characterized statistically. This typically involves two steps: (1) predicting and estimating the potential impacts of the uncertainty on the feasibility of a trajectory; and (2) determining the trajectory that minimizes the predicted risk. For example, Bry and Roy [2] proposed a graph-search based algorithm–the Rapidly-exploring Random Tree algorithm–where, given a nominal trajectory, distributions for future states of the vehicle are first predicted to assess whether the probability of collision, given a state, is bounded below a threshold value. Next, a set of trajectories is incrementally constructed while efficiently searching for the best candidate path. Paths are evaluated based on their probability of being realized by a closed-loop controller. Similar work that focuses on graph-based search algorithm can also be found in [3] and [4]. Separately, roadmap-based approaches that rely on an understanding of the safe states in the environment (or the configuration space in general) have attracted significant research interest. Typically, they utilize one of three techniques: visibility graphs [5], Voronoi diagrams [6], or potential fields [7].

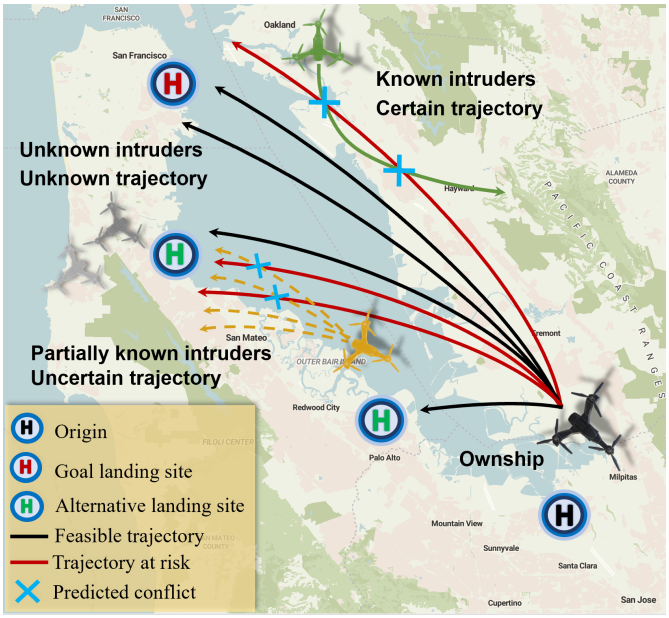

Figure 1: Illustration of the three types of intruders present in the operating environment.

The difficulty imposed on trajectory planning by an inability to characterize and account for the behavior of highly unpredictable obstacles is illustrated in Figure 1. The green intruder has a trajectory that is both deterministic and known to the ownship. In this case, the ownship can maintain the feasibility of its trajectory by following one of many conflict-free trajectories (solid black lines). The yellow intruder is partially known to the ownship because either the time it enters the airspace, its speed, or its path are unknown. In this case, the yellow intruder's trajectory is determined via bounded probability distributions with prescribed parameters. To protect itself, the ownship can then leverage analytical

techniques to assess the chance of collision for each feasible trajectory and follow the one that maximizes a certain objective. The grey intruder represents the extreme case where the intruder is completely unknown to the ownship, i.e., neither the time of entry, the speed, nor the path are known. Thus, from the ownship's point of view, the grey intruder can show up unexpectedly at anytime and anywhere in its operating environment, and the associated statistical distributions are unbounded.

The trajectory flexibility metrics proposed by Idris et. al [8] provide the basis for a promising approach to improving the ability of autonomous vehicles to adapt to unknown disturbances. Specifically, [9] demonstrated that self-separation and self-organizing behaviors may be induced among autonomous agents, and traffic complexity reduced by maximizing trajectory flexibility. Further, building on that research, [10] leveraged adaptability, one of the trajectory flexibility metrics, to estimate airspace capacity under different control schemes.

In this paper, we contribute to the literature by proposing a receding horizon control strategy with a set of novel trajectory planning policies that enable the autonomous vehicle to dynamically update its planned trajectory in environments where potential conflicts are, from a statistical perspective, either *partially known* or completely *unknown*. Most importantly, we demonstrate that maximizing the total number of feasible trajectories is effective in mitigating the consequences of extreme uncertainty.

The remainder of this paper is structured as follows. In Section II we establish the theoretical foundation for measuring the robustness and survivability of a mission. Next, in Section III we propose a backtracking algorithm for counting the number of feasible trajectories as well as a Monte Carlo simulation to estimate the robustness of a trajectory segment. We then, in Section IV, formalize the receding horizon trajectory problem and present a detailed description of four trajectory planning policies. In Section V we outline the computational experiment for investigating and comparing the effectiveness of the proposed trajectory planning policies. The results and analysis of our experiment are presented in Section VI. Finally, in Section VII we provide our conclusions and suggestions for future research directions.

II. Theoretical Framework

We consider the trajectory planning problem for a fixed flight level. To count the number of trajectories, we establish a discrete representation of space and time. Specifically, time is discretized into equal time steps that are Δ_t apart and space is discretized into rectangular cells of dimension $\Delta_x \times \Delta_y$ as shown in Figure 2. The state of a vehicle is denoted by $(x_i, y_i, h_i, v_i, t_i) \in \mathcal{R}^5$ where h_i is the heading and v_i is the speed. A motion primitive is defined by the pairing of heading and speed changes, that is $(\delta_h, \delta_v) \in \mathcal{R}^2$. The set of all possible motion primitives is denoted by $\mathcal{A}$, where the cardinality of $\mathcal{A}$ indicates the maneuverability of a vehicle.

Next, we introduce the mapping $c : \mathcal{R}^3 \mapsto \mathcal{N}^3$ between a way point in continuous space and a cell in discrete space defined by $c(x, y, t) = (\lfloor x \rceil, \lfloor y \rceil, \lfloor t \rceil)$ where $\lfloor \cdot \rceil$ rounds the element $\cdot$ to the nearest integer. We further define the partial trajectory, connecting cell $c(x_i, y_i, t_i)$ to $c(x_{i+1}, y_{i+1}, t_{i+1})$ in discrete space and time, as a segment ls_i. Equivalently, a trajectory l can also be described as a sequence of segments, that is $l = [ls_1, ls_2, \ldots, ls_{T-1}]$, where $[\cdot]$ is an ordered list. Depending on the resolution of the discretized time and space, a segment l_i may degenerate to a cell (i.e., two end cells are identical) or cross multiple cells between two end cells. Cells that are blocked at time t are denoted by $\mathcal{B}(t) \subset \mathcal{N}^3$, where the information on the environment is refreshed at a certain rate. Further, we apply the following Assumptions II.1 and II.2.

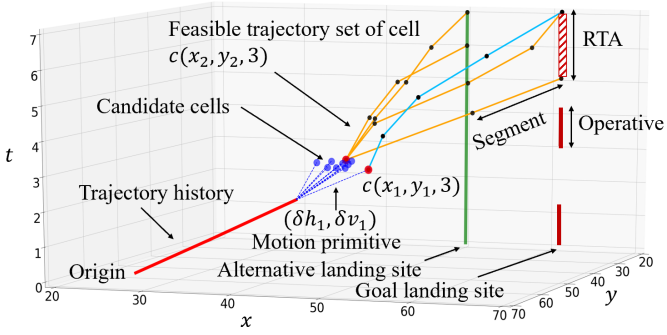

Figure 2: Discrete representation of space and time.

Assumption II.1. *We consider intruders as the only sources of uncertainty in the environment, and that a cell is completely blocked if any part of it is crossed by an intruder's trajectory.*

Assumption II.2. *A straight line passing through multiple cells can be used to approximate a trajectory segment. Let $c(x_i, y_i, t_i)$ and $c(x_j, y_j, t_i + \epsilon_t)$ denote two end cells of a segment. A segment is feasible if all cells that are crossed by the segment are accessible. Specifically,*

$$c(x_i + \beta(x_j - x_i), y_i + \beta(y_j - y_i), t_i + \beta\epsilon_t) \notin \mathcal{B}(t_i),$$

where $\beta \in [0, 1]$.

Cell quantization has an impact on trajectory quality and safety. If the cell dimensions are small relative to the required safe separation, blocking the cells traversed by the trajectory does not provide sufficient separation from other vehicles. In such cases, Assumption II.2 is leveraged to block cells crossed by the trajectory segment as well as those within a defined radius of the trajectory segment. Typically, such a radius corresponds to the required safe separation between two vehicles. If the cell dimensions are large relative to the motion primitives, a trajectory may not be able to transit from one cell to another in a single time step. Further, large cell dimensions and time steps may result in trajectories that are too coarse.

A mission is formally defined as follows: an ownship departs from its origin (x_0, y_0) at time t_0 and arrives at its goal landing site (x_f, y_f) during the time interval $[t_{f_l}, t_{f_u}]$ that defines its Required Time of Arrival (RTA). At the time of arrival, the heading of the vehicle h_f depends on the landing pad's orientation and thus is limited to a finite set of discrete values $\{h_1, \ldots, h_m\}$, where $h_i \in [0, 2\pi]$. In addition, we require that the vehicle maintains a minimum speed $v_f = v_{min}$ at the moment of arrival. We denote by $\Omega(t)$ the set of all valid terminal states.

We focus on the task of computing a feasible trajectory that directs the ownship to the goal landing site subject to the constraint that it must always be (or at least have a very high probability of being) able to reach an alternative landing site should the goal landing site become unreachable. Since the goal/alternative landing sites are built infrastructures, they block cells in space and time that correspond to their physical locations at all time (see the red and green vertical line in Figure 2). While the goal landing site is only open to the

ownership during the RTA period, the alternative landing site is open at all times.

The ownership is capable of adapting its trajectory dynamically to account for uncertainty. Specifically, a mission consists of T decision points set $\epsilon_t \in \mathcal{R}_+$ apart. The choice of ϵ_t depends on the frequency that the information on environment is updated. Typically, we consider $\epsilon_t > \Delta_t$. The mission completion time should not exceed the maximum flight endurance, which is determined by the ownship's fuel capacity. Regarding the vehicle dynamics, we assume that the ownship is allowed to change its heading h and speed v at a decision point by applying a pre-defined control primitive $(\delta_h, \delta_v) \in \mathcal{A}$, where $\delta_h \in [-\delta_{h_l}, \delta_{h_u}]$ and $\delta_v \in [-\delta_{v_l}, \delta_{v_u}]$. The vehicle maintains constant heading and speed until it reaches the next decision point. Therefore, at each decision point, the core decision-making problem involves optimizing the trajectory plan so as to maximize a certain objective function while respecting all constraints imposed on trajectories.

To facilitate exposition, we next discuss the concepts of mission robustness and survivability – the metrics by which the effectiveness of a planning strategy is assessed. Following that, we formally state and prove our theorem II.1.

Definition II.1 (Feasibility of a trajectory plan)**.** *A trajectory is feasible if the following constraints are satisfied:*

(1) Vehicle dynamics: The vehicle is able to follow the trajectory without violating its maneuverability constraints (i.e., turning rate and acceleration ranges, and maximum speed).

(2) Terminal constraint: If the location (x_f, y_f) corresponds to the goal landing site, the arrival time T should be in the range of RTA, that is $T \in [t_{f_l}, t_{f_u}]$. Else, if location (x_f, y_f) corresponds to an alternative landing site, t_f is unbounded. In both cases, the terminal states must adhere to the heading and velocity requirements at the corresponding landing site.

(3) Conflict avoidance constraint: For any segment that belongs to the trajectory l, the cells, crossed by the segment, should be outside of blocked zones $\mathcal{B}(t)$.

Definition II.2 (Robustness of a trajectory plan)**.** *The robustness of a trajectory $\mathcal{P}_l$ is defined as the likelihood that the trajectory l will remain feasible despite the occurrence of disturbances that pose a constraint violation risk. The robustness of a trajectory is expressed as:*

$$\mathcal{P}_l = \Big(\prod_{i=0}^{T-1} p_i \Big) \alpha_T, \tag{1}$$

where p_i is the probability that segment ls_i remains feasible in the presence of disturbances, and α_T is the probability that the corresponding landing site is available at time T.

At a given point in space and time, a vehicle may have a set of feasible trajectories $\mathcal{L} = \mathcal{L}_g \cup \mathcal{L}_a$, where the trajectories in the set $\mathcal{L}_g = \{l_1, l_2, \ldots, l_n\}$ terminate at the goal landing site, while the trajectories in the set $\mathcal{L}_a = \{l_1, l_2, \ldots, l_m\}$ lead to an alternative landing site. A mission is deemed to be **successful** if there is at least one feasible trajectory that reaches the goal landing site (i.e., $\mathcal{L}_g \neq \emptyset$). Similarly, a mission is deemed to be **survivable** if there is at least one feasible trajectory that leads to either the goal landing site or an alternative landing site (i.e., $\mathcal{L} \neq \emptyset$). Therefore, the probability that a mission is successful or survivable is highly dependent on the robustness of each trajectory in the set $\mathcal{L}_g/\mathcal{L}$ and its cardinality.

Definition II.3 (Robustness of a mission)**.** *Consider the set of trajectories $\mathcal{L}_g = \{l_1, l_2, \ldots, l_n\}$ that terminate at the goal landing site. With a slight abuse of notation, we denote the robustness of a trajectory l_i as $\mathcal{P}_{l_i} \in [0, 1]$, and the probability that the mission is successful as:*

$$\mathcal{P}^{succeed}(\mathcal{L}) = 1 - \prod_{i=1, l_i \in \mathcal{L}_g}^{n} \Big(1 - \mathcal{P}_{l_i}\Big) \tag{2}$$

Definition II.4 (Survivability of a mission)**.** *Given the set of feasible trajectories $\mathcal{L} = \{l_1, \ldots, l_m\}$ that terminate at a valid landing site (i.e., at either the goal landing site or an alternative landing site), we denote the robustness of an individual trajectory l_i as $\mathcal{P}_{l_i} \in [0, 1]$, and the probability that the mission is survivable as:*

$$\mathcal{P}^{survive}(\mathcal{L}) = 1 - \prod_{i=1, l_i \in \mathcal{L}}^{m} \Big(1 - \mathcal{P}_{l_i}\Big) \tag{3}$$

Theorem II.1. *Let $\mathcal{L} = \{l_1, l_2, \ldots, l_n\}$ be a set of independent, feasible trajectories, where the robustness of trajectory $l_i \in \mathcal{L}$ is $\mathcal{P}_{l_i} \in [0, 1]$. Then, given the trajectory set $\mathcal{L}$, the survivability of the mission is expressed in Equation (3).*
i. Let $\mathcal{P}'_{l_i} \in [0, 1]$ be a new robustness measurement of trajectory l_i such that $\mathcal{P}'_{l_i} > \mathcal{P}_{l_i}$. Then, the survivability of the mission increases by the amount $\Delta\mathcal{P}_1 \geq 0$.
ii. Let the trajectory set $\mathcal{L}$ increase in size through the addition of a feasible trajectory l_{n+1} that is independent to all existing trajectories, with the robustness $\mathcal{P}_{l_{n+1}} \in [0, 1]$. Then, the survivability of the mission increases by the amount $\Delta\mathcal{P}_2 \geq 0$.

Proof. For simplicity, define $A = \prod_{i=1}^{n}(1 - \mathcal{P}_{l_i})$ and $A \in [0, 1]$. To show $i.$, we express the increment in the survivability of the mission due to the inclusion of an existing trajectory l_i when $\mathcal{P}_{l_i} \in [0, 1)$ as follows:

$$\begin{aligned}\Delta\mathcal{P}_1 = \mathcal{P}^{'survive} - \mathcal{P}^{survive} &= 1 - A\frac{(1 - \mathcal{P}'_{l_i})}{(1 - \mathcal{P}_{l_i})} - (1 - A) \\ &= A\frac{\mathcal{P}'_{l_i} - \mathcal{P}_{l_i}}{1 - \mathcal{P}_{l_i}} \geq 0\end{aligned}$$

In a special case where $\mathcal{P}_{l_i}$, since $\mathcal{P}'_{l_i} \geq \mathcal{P}_{l_i} = 1$ and $\mathcal{P}'_{l_i} \in [0, 1]$, we have $\mathcal{P}'_{l_i} = 1$ which leads to the following conclusion:

$$\Delta\mathcal{P}_1 = \mathcal{P}^{'survive} - \mathcal{P}^{survive} = 0 - 0 = 0$$

To show statement $ii.$, we again express the increment in the survivability of the mission due to the increase in the total number of trajectories by 1 as follows:

$$\begin{aligned}\Delta\mathcal{P}_2 = \mathcal{P}^{'survive} - \mathcal{P}^{survive} &= 1 - A(1 - \mathcal{P}_{l_{n+1}}) - (1 - A) \\ &= A\mathcal{P}_{l_{n+1}} \geq 0\end{aligned}$$

□

Corollary II.1.1. *If the current survivability of the mission is less than 1, improving the robustness of any existing trajectory or adding a new and uncorrelated trajectory to the current trajectory set increases the survivability of the mission.*

III. Methodology

At each decision point, the vehicle should be able to evaluate, for each candidate cell: (1) the robustness of the segment connecting the current cell to the candidate cell, (2) the feasible trajectories that contain the segment to the candidate cell, and (3) the robustness of each trajectory in the feasible trajectory set. Given this information, the vehicle can evaluate

the robustness and survivability of the mission and develop a robust trajectory plan according to its objectives. To this end, in Section III-A, we propose a backtracking algorithm for computing the set of feasible trajectories for each cell in space and time. Next, in Section III-B, we develop a Monte Carlo simulation that propagates the uncertainties from the environment (i.e., intruders' trajectories) to the probability distribution describing the likelihood of "a segment being feasible". The robustness of segments is then integrated into the measures of mission robustness and survivability.

A. Compute Feasible Trajectory Set via Backtracking

A trajectory is feasible if it is free of conflicts, meets operational constraints such as the required RTA, and adheres to the constraints on vehicle dynamics specified by the motion primitives per Definitions II.1. To search for all feasible trajectories starting from the cell of interest to the goal landing site, we use a backtracking algorithm as illustrated in Figure 3.

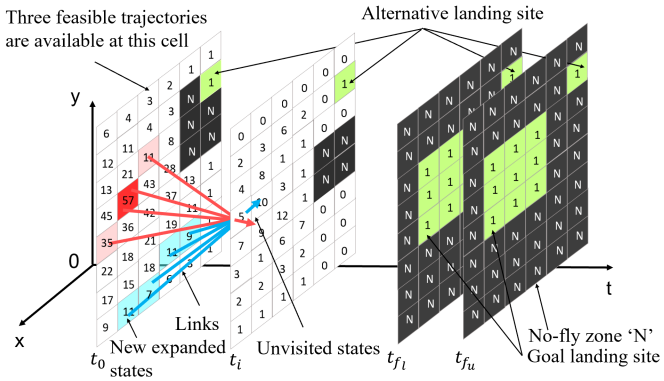

Figure 3: Compute the set of feasible trajectories for a cell in space and time.

We begin the search by marking the set of valid terminal states as unvisited and the corresponding cells as 1. We project intruder trajectories into space and time as blocked cells with the label 'N' to account for the motion of obstacles in the environment. We generate a list of backward reachable states from each unvisited state under consideration and eliminate the set of states that cannot be part of a feasible trajectory. At the end of the iteration, only promising states will be mapped into cells and become unvisited states. It is possible that several states will be assigned to a single cell. Since each promising state corresponds to a feasible trajectory that leads to one of the valid terminal states, the total number of feasible trajectories in a cell equals the number of promising states in that cell. Denote by $c(x, y, t).N_f$ the total number of feasible trajectories and by $c(x, y, t).\bar{S}$ the set of unvisited states in cell $c(x, y, t)$. The main steps are summarized below:

- Initialization: As shown in Figure 3, the green squares represent goal-reaching cells. For any valid terminal state $(x_f, y_f, h_f, v_f, t_f) \in \Omega(t_0)$, we set the number of feasible trajectories for a goal reaching cell $c(x_f, y_f, t_f)$ to one, i.e., $c(x_f, y_f, t_f).N_f = 1$. In addition, each valid terminal state becomes the only unvisited state in the corresponding goal reaching cell, i.e., $c(x_f, y_f, t_f).\bar{S} = (x_f, y_f, h_f, v_f, t_f)$. If a cell is blocked, we set $c(x_t, y_t, t).N_f = c(x_t, y_t, t).\bar{S} =$ 'N'. Otherwise, $c(x_t, y_t, t).N_f = 0$ and $c(x_t, y_t, t).\bar{S} = \emptyset$.
- Expand & Prune: At time frame t, given a non empty cell, any unvisited state (x_t, y_t, h_t, v_t, t) in the cell $c(x_t, y_t, t)$, called ancestor, should be propagated one step backward by applying all possible heading and velocity change $(-\delta_h, -\delta_v)$, where $(\delta_h, \delta_v) \in \mathcal{A}$. All new generated states $(x_{t-1}, y_{t-1}, h_{t-1}, v_{t-1}, t-1)$ should belong to the time frame $t-1$. (see red/blue arrow and its corresponding new generated states). A new generated state $(x_{t-1}, y_{t-1}, h_{t-1}, v_{t-1}, t-1)$ is forward-reachable if the minimum time required for the vehicle to move from the current state to the new generated state does not exceed $t - 1 - t_0$, as stated in Inequality (4).

$$\frac{||(x_t, y_t) - (x_0, y_0)||}{v_{max}} \leq t - 1 - t_0, \tag{4}$$

  where (x_0, y_0, t_0) is the way point corresponding to a departure and v_{max} is the upper limit of vehicle's speed. In the pruning step, a newly generated state is eliminated if the segment, connecting itself to its ancestor, is infeasible; or if the newly generated state is not forward reachable. Otherwise, we assign the newly generated state $(x_{t-1}, y_{t-1}, h_{t-1}, v_{t-1}, t-1)$ to the list of unvisited states in the corresponding cell $c(x_{t-1}, y_{t-1}, t-1)$, increase the number of unvisited states for cell $c(x_{t-1}, y_{t-1}, t-1).N_f$ by 1, and mark the ancestor state as visited.
- We repeat the steps described above for the previous time frame until $t - 1 = t_0$. Once the algorithm terminates, for each cell $c(x, y, t)$ in space and time, we obtain the number of feasible trajectories to the valid terminal cells by looking at its $c(x, y, t).N_f$ parameter.

The implementation details are provided in Algorithm 1, while the details of the $ExpandPrune(\cdot)$ function are provided in Algorithm 2.

B. Estimate the robustness of a segment via Monte Carlo Simulation

The robustness/survivability of a mission at a given point in space and time is dependent on the robustness of all the available feasible trajectories at that point. The robustness of a trajectory is determined by the robustness of its segments as well as the probability that the terminal landing site is operational per Definition II.2. Because the uncertainty in landing site availability is complicated and deserves its own paper, we assume that both the goal landing site and the alternative landing site are operational at all times with probability $\alpha^t = 1, t \in [t_0, t_{f_u}]$. In this paper, we focus on estimating the robustness of trajectory segments.

To this end, we propose a Monte Carlo (MC) simulation based on mathematical models that describe the statistical properties of uncertain events (i.e., the motion of intruders). These events may be known, partially known, or unknown to the ownership. The MC simulation yields the robustness of a segment, which is then used to compute the robustness of a trajectory per Equation (1).

Trajectory model of intruders: We model the intruder's trajectory using a Bezier curve of degree k. Given k random control points $\{(x_i^B, y_i^B, t_i^B) \mid x_i^B \in [X_l, X_u], y_i^B \in [Y_l, Y_u], t_i^B \in (t_0, t_{f_u}), \forall i = 0, \ldots, k\}$, the uniquely defined Bezier curve l^B defines a continuous trajectory that enters the time and space network at cell $c(x_0^B, y_0^B, t_0^B)$ and exists the time and space network at cell $c(x_k^B, y_k^B, t_k^B)$. Adjusting the 0-th and k-th control points, one can specify the entrance and exit location of the intruder as well as its flight duration. By increasing the degree number k, the resulting curve has a greater number of twists and turns, mimicking the motion of an aggressive intruder. We generate a Bezier curve with

Algorithm 1: Compute feasible trajectory set via Backtracking

Input:
Origin:
(x_0, y_0, t_0)
Goal ∪ Alternative landing sites:
$(x_f, y_f, h_f, v_f, [t_{f_l}, t_{f_u}]) \cup (x_f, y_f, -, -, -)$
Flight dynamics:
$\delta_h \in [\delta_{h_l}, \delta_{h_u}], \delta_v \in [\delta_{v_l}, \delta_{h_u}], v \in [v_{min}, v_{max}]$
BlockMap: Space and time with information on no-fly zones
Output: TrajMap: The estimate of total number of trajectories

$t = t_{f_u}$
for $t_f \in [t_{f_l}, t_{f_u}]$ **do**
  $c(x_f, y_f, t_f).N_f = 1$
  $c(x_f, y_f, t_f).\bar{S} \leftarrow (x_f, y_f, h_f, v_f, t_f)$
while $t > t_0$ **do**
  for $c(x_i, y_i, t) \in [X_l, X_u] \times [Y_l, Y_u]$ **do**
    if $BlockMap(x_i, y_j, t) == 0$ **then**
      $c(x_i, y_i, t).N_f = -\infty$
      $c(x_i, y_i, t).\bar{S} \leftarrow$ N
    else if $c(x_i, y_i, t).\bar{S} \neq \emptyset$ **then**
      for $s \in c(x_i, y_i, t).\bar{S}$ **do**
        $\mathcal{I}_{feasible} \leftarrow$ ExpandPrune(s, δ_t, $\mathcal{A}$)
        $c(:, :, t-1).\bar{S} \leftarrow \mathcal{I}_{feasible}$
        $c(:, :, t-1).N_f += |\mathcal{I}_{feasible}|$
  t = t - 1
for $c(x, y, t) \in [X_l, X_u] \times [Y_l, Y_u] \times [t_0, t_{f_u}]$ **do**
  TrajMap$(x, y, t) = c(x, y, t).N_f$
return TrajMap

Algorithm 2: ExpandPrune(·) function

Function `Expand & Prune` $(s, \delta_t, \mathcal{A})$ **:**
  $(x_t, y_t, z_t, h_t, v_t, t) \leftarrow s$
  $\mathcal{I}_{feasible} \leftarrow \emptyset$
  for $(\delta_h, \delta_v) \in \mathcal{A}$ **do**
    $v_{t-1} = v_t - \delta_v$
    $h_{t-1} = h_t - \delta_h$
    $x_{t-1} = x_t - v_{t-1}\delta_t \cos(h_{t-1})$
    $y_{t-1} = y_t - v_{t-1}\delta_t \sin(h_{t-1})$
    $s_{new} \leftarrow (x_{t-1}, y_{t-1}, h_{t-1}, v_{t-1}, t-1)$
    if *segment connecting* $c(s)$ *and* $c(s_{new})$ *is feasible per Assumption II.2 & s_{new} is forward reachable per inequality 4* **then**
      $\mathcal{I}_{feasible} \leftarrow \mathcal{I}_{feasible} \cup s_{new}$
  return $\mathcal{I}_{feasible}$

degree number 1 - a straight line that occupies a fixed location for a time interval - to simulate the trajectory of a hovering intruder.

Uncertainty model of a trajectory Trajectory uncertainty models typically utilize either a probability density function (pdf) or bounded shapes [11]. We utilize a shape-based methodology to facilitate greater utility. With the assumption that the ownship is capable of tracking the prescribed trajectory, the possible trajectories can be bounded by geometric volumes (i.e. sheared cylinders) [12]. Given an intruder trajectory l^B, we further define the trajectory uncertainty volume $V(l^B, r)$ as the union of disks with radius r that are centered at points along the trajectory l^B. A possible trajectory is a continuous curve starting at cell $c(x_0^B, y_0^B, t_0^B)$ and ending at cell $c(x_k^B, y_k^B, t_k^B)$ such that all intermediate points are within the trajectory uncertainty volume $V(l^B, r)$.

With this uncertainty model, we assume that if an intruder is only partially known to the ownship, the nominal trajectory l^B and the corresponding uncertainty level r are certain to the ownship, but the actual trajectory is unknown. Typically, if $r > 0$, the intruder may fly a trajectory that deviates at most r from the planned trajectory. If $r = 0$, the partially known intruder becomes fully known to the ownship. The known intruder will follow l^B exactly through space and time. In the extreme case where an intruder is unknown to the ownship, the ownship is unaware of the intruder's existence until the ownship encounters a conflict with the intruder.

A traffic scenario is defined as a situation in which there are N_k known intruders, N_p partially known intruders, and N_u unknown intruders. The number of partially known intruders, the uncertainty threshold r of each intruder's trajectory, and the number of unknown intruders all influence the severity of the uncertainties in the environment. The proposed MC-based simulator takes into account a segment or set of segments related to the current trajectory planning decision all at once. The expected output for each input segment is the probability of the segment being feasible in the presence of the current traffic scenario. Figure 4 depicts the framework of the MC-based simulator which consists of four major processes:

- Initialization: During the initialization process, we simulate intruder trajectories by randomly generating N_k+N_p Bezier curves of various lengths between $[t_0, t_{f_u}]$. A counter $N_{success}$ is used to record how many times a segment is feasible out of the N_{mc} different traffic realizations.
- Traffic Realization: The MC-based simulator randomly generates realizations of an intruder's trajectory for a given traffic scenario based on each intruder's nominal trajectory l^B and uncertainty threshold r (if applicable). When a new realization is generated, the timer advances by one. If the timer reaches the maximum number of runs N_{mc}, the simulation process terminates and outputs an estimate of a segment's robustness. Otherwise, the simulator proceeds to the next step.
- Feasibility Evaluation: A $BlockMap$ is created by blocking the cells traversed by the realized trajectories l^B. Given a segment of interest, the simulator evaluates the feasibility of the segment via Assumption II.2. The counter for the segment $N_{success}$ is increased by one if the segment is feasible. Otherwise, we do nothing.
- Termination: After N_{mc} runs, the simulation process terminates. The probability of the segment being feasible is measured by the ratio $N_{success}/N_{mc}$.

IV. Trajectory Planning Framework

A. Receding Horizon Control

An autonomous vehicle using a fixed horizon control scheme optimizes a sequence of control actions $a_1, a_2, \ldots, a_T$ over T steps. If, during the T steps, unexpected events occur or the system behaves differently than was expected during the design of the control scheme, the controller will not be able to account for them. This shortcoming can be addressed through Receding Horizon Control (RHC) where

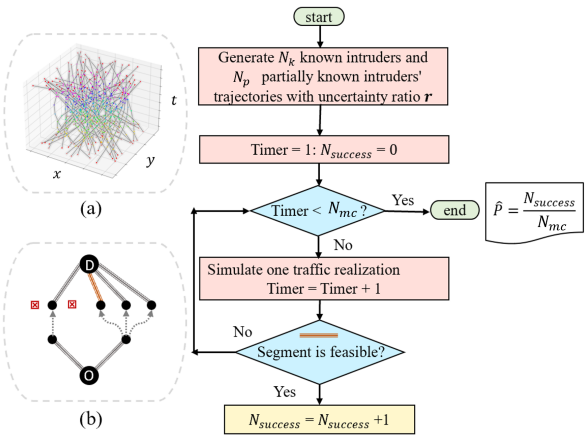

Figure 4: MC simulation framework. (a) Visualization of a traffic scenario in space and time. (b) Input trajectory segment.

control actions are repeatedly optimized over a moving time horizon. Specifically, the controller generates the optimal control inputs over M time steps and executes the first control action. At the next time step, a new control problem with the most recent environmental information will be solved for the remaining M-1 time steps [13]. We summarize the general RHC optimization problem as follows.

$$\min_{z} \quad F(s_t, a_t) + \beta \sum_{k=0}^{T} V(c_{t+k}) \tag{5}$$

$$\text{s.t.} \quad s_{t+k+1} = f(s_{t+k}, a_{t+k}) \tag{6}$$

$$s_t = (x_t, y_t, h_t, v_t, t) \tag{7}$$

$$c_{t+k} = c(x_{t+k}, y_{t+k}, t+k) \tag{8}$$

$$c_{t+k} \notin \mathcal{B}(t) \tag{9}$$

$$s_{t+T} \in \Omega(t) \tag{10}$$

$$a_{t+k} = (dh_{t+k}, dv_{t+k}) \in \mathcal{A} \tag{11}$$

$$s_{t+k} \in \mathcal{R}^3 \times [v_l, v_u] \times [t, t_{f_u}] \tag{12}$$

where we optimize over control commands $z = \{a_t, \ldots, a_{t+T-1}\}$ for the remainder of the mission given the current knowledge of the environment stored in $\mathcal{B}(t)$ at time t. However, only the first command a_t is executed. This process is then repeated until the vehicle reaches the goal, or terminated if no feasible solution was founded. The vehicle dynamics are prescribed in Constraint (6) with the initial state given by constraint (7). The vehicle state is mapped to a cell in discrete space and time in Constraint (8). In addition, we ensure that a vehicle does not cross a blocked region $\mathcal{B}(t)$ via constraint (9). Constraint (10) ensures that the end state is one of the valid terminal states in the set $\Omega(t)$. Note that set $\Omega(t)$ consists of only goal-landing sites at the start of the flight. However, if no feasible trajectory to the goal-landing site is available, the set $\Omega(t)$ will be changed to include alternative landing sites. This ensures that priority is given to the completion of a mission over its survival. Constraint (11) limits actions to the set of motion primitives $\mathcal{A}$. The state-space is specified by constraint (12) where the speed of the vehicle is explicitly bounded by $[v_l, v_u]$ and the mission duration is upper bounded by t_{f_u}.

The objective function (5) is made up of two parts. The first term $F(s_t, a_t)$ represents the robustness of the first trajectory segment, which is the outcome of applying control a_t in state s_t. The second term $V(c_{t+k}), k = \{0, \ldots, T-1\}$ measures the quality of the k-th way point by evaluating the goodness of its corresponding cell. A discount coefficient $\beta \in [0, 1]$ is applied to the second term to adjust the weighting between instant and future response. Generally, the discount can be a function of time, which is not considered in this paper.

B. Trajectory planning policies

A data preparation process includes computing the following parameters:

- $c_t.\mathcal{L}$: Feasible trajectory set available at cell c_t via the Backtracking algorithm introduced in Section III-A.
- $c_t.N_f$: Cardinality of the feasible trajectory set $c_t.\mathcal{L}$.
- $p_{c_t c_{t+1}}$: The robustness of a segment connecting cell c_t to cell c_{t+1} (probability of segment being feasible) using the MC-simulation introduced in Section III-B.
- $\mathcal{P}_{l_i}, \forall\ l_i \in c_t.\mathcal{L}$: Robustness of trajectory l_i.
- $\mathcal{P}^{survive}(c_t.\mathcal{L})$: Survivability of the mission.

The first policy is to maximize the robustness of the resulting trajectory to partially known disturbances.

$$(\pi_R) \quad p_{c_t c_{t+1}} + \beta \sum_{k=1}^{T-1} p_{c_{t+k} c_{t+k+1}}$$

Motivated by a situation where there are only unknown events or estimates of segment robustness are unavailable, the second policy considers the maximization of the number of trajectories only.

$$(\pi_{N_f}) \quad c_{t+1}.N_f + \beta \sum_{k=2}^{T} c_{t+k}.N_f$$

The third policy takes into account both the robustness of a trajectory and the total number of alternative trajectories available at each way point to maximize mission survivability. The resulting trajectory attempts to provide at least one feasible trajectory to guide the vehicle to a safe landing at any point along the trajectory.

$$(\pi_S) \quad p_{c_t c_{t+1}} + \beta \sum_{k=1}^{T} \mathcal{P}^{survive}(c_{t+k}.\mathcal{L})$$

The fourth policy serves as an uninformed baseline against which other policies can be measured, where the vehicle chooses one of the feasible paths at random. The objective function maximizes the chances of surviving the mission under this assumption.

$$(\pi_{Avg(R)}) \quad p_{c_t c_{t+1}} + \beta \sum_{k=1}^{T} \left(\frac{1}{c_{t+k}.N_f} \sum_{l_i \in c_{t+k}.\mathcal{L}} \mathcal{P}_{l_i} \right)$$

V. Experiments Setup

Our computational experiments had two goals: First, to show how the proposed trajectory planning framework can be used to direct a mission. Second, to compare the effectiveness of our four planning policies in mitigating varied levels of uncertainty. The experiment is set up as follows.

a) Mission: We are interested in a mission in the San Francisco Bay Area which is specified as an $85 \times 85\ km^2$ area as shown in Figure 1. A vehicle takes off from the Santa Clara Towers Heliport and heads northwest to land at the UCSF Helipad. We consider two alternative landing sites: San Francisco International Airport and Stanford Hospital Heliport. The flight mission includes a succession of three phases: takeoff, cruising, and landing. All three phases may involve hovering.

The departure time is $t_0 = 00:00$ minutes, and the required arrival time at the goal landing site is $t_f = [00:$

$40,00 : 55]$ minutes. We assume that the vertical take-off and hover, vertical landing and hover phases each take 5 minutes. In addition, the vehicle enters and exits in the cruising phase with a speed of 0 m/s. Due to the topology of the destination helipad, the vehicle must maintain its headings at $h_f \in \{90^o, 135^o, 180^o\}$ upon arrival at the goal landing site before initiating a vertical descent. When the vehicle arrives at an alternative landing site, the heading requirement changes to $h_f \in \{90^o, 180^o, 270^o\}$.

During the cruising phase, we assume that the vehicle's cruise speed is in the range $[0, 40]\ m/s$. At each replanning window, the vehicle can change its heading by $dh \in [-30^o, 30^o]$, in 5 equally spaced discrete increments, and its speed by $dv \in [-10, 10]$ m/s, in 5 equally spaced discrete increments. Therefore, the motion primitive set is composed of 5×5 different combinations of heading-velocity changes. We divide time and space into cells of size $(1\ km, 1\ km, 1\ minute)$. As a result, the mission entails 6, 7, 8, or 9 decision-making windows of 5-minute duration. We assume that the maximum flight endurance is 60 minutes, which is equivalent to 12 decision-making windows.

b) Environment: Intruders are of two types: (1) Observed intruders with uncertainty radius $r = 5\ km$ whose trajectories are partially known to the ownship; and (2) Unobserved intruders whose presence is unknown to the ownship. Our experiment assesses 36 different traffic scenarios. Each scenario is characterized by the total number of partially known intruders $N_p \in \{50, 100, 150, 200, 250, 300\}$, and the total number of unknown intruders $N_u \in \{50, 100, 150, 200, 250, 300\}$.

c) Simulation Setup: For each traffic scenario characterized by (N_p, N_u), we generate $N_{instance}$ different traffic instances. For each instance, N_p trajectories are randomly generated with different entrance and exit points in space and time. In addition, we generate a set of N_u trajectories that are unknown to the ownship during trajectory planning but actually exist in the simulated environment. Given a traffic instance, we simulate the real traffic between t_0 and t_{f_u}. The ownship plans its trajectory by employing each of the policies π_{N_f}, π_S, π_R, and $\pi_{Avg(R)}$. The computation of π_S, π_R, and $\pi_{Avg(R)}$ involves measuring the robustness of feasible trajectory segments via MC simulation introduced in Section III-B. The three possible outcomes of a mission are: success by arriving at the goal landing site (G), survival by arriving at one of the two alternative landing sites (A), or failure, which involves the more risky situation of an emergency landing at an unplanned location (L). The survivability of a mission under traffic scenario (N_p, N_u) when employing a specific policy is estimated by the ratio of the number of outcomes (G) and (A) out of $N_{instance}$. Similarly, the success of the mission is measured by the ratio of the number of outcomes (G) out of $N_{instance}$.

d) Convergence Analysis: The convergence analysis is based on the principle of utilizing sample proportion to estimate an unknown population proportion, as well as quantifying uncertainty in a population proportion estimate. First, we investigate the case of utilizing direct MC simulation to estimate the robustness of a trajectory segment, i.e., the likelihood of a trajectory segment being feasible. Given a segment of interest, let E denote the event that the segment is feasible. The robustness of a segment p is defined as the expected value of the indicator function $\mathbb{I}_E$, that is $p = Exp(\mathbb{I}_E)$. The variance of $\mathbb{I}_E$ is $p(1-p)$. To estimate p, we generate n independent samples $x_1, x_2, \cdots, x_n$ of $\mathbb{I}_E$ using MC simulation and compute the sample mean.

$$\hat{p}_n = \frac{1}{n}\sum_{1}^{n} x_i$$

In this example, we refer to p as the population proportion and $\hat{p}_n$ as the sample proportion. We can further use sample proportion $\hat{p}_n$ to approximate the sample variance of $\mathbb{I}_E$, that is $s_n^2 = \hat{p}_n(1-\hat{p}_n)$. To quantify the uncertainty/error associated with the estimator $\hat{p}_n$, we leverage the central limit theorem to construct a confidence interval associated with $\hat{p}_n$, with the chance of p falling inside this interval equal to $1-\alpha$. The confidence interval is given as follows [14].

$$\hat{p}_n \pm z_c \frac{s_n}{\sqrt{n}} = z_c \frac{\sqrt{\hat{p}_n(1-\hat{p}_n)}}{\sqrt{n}}$$

Where $z_c = 1.96$ for a common choice of confidence level $1-\alpha = 95\%$. Notice that when $\hat{p}_n = 0.5$ the half uncertainty band $h = z_c \frac{\sqrt{\hat{p}_n(1-\hat{p}_n)}}{\sqrt{n}}$ attains its maximal value and therefore is upper bounded by $h_{max} = z_c\sqrt{\frac{0.25}{n}}$. In our experiment, given the computation time for obtaining one data sample and the desired accuracy of estimator $\hat{p}_n$, it is sufficient to limit h_{max} within 2% with confidence level 95%. To achieve this goal, in theory, we need at least $n = 2401$ samples to obtain a point estimate $\hat{p}_n$ of p. Therefore, in the case of using MC simulation to evaluate the robustness of a trajectory segment, we set the number of MC simulation runs to $N_{mc} = 2500$. As an example, we consider measuring the robustness of three[1] independent trajectory segments in an environment with 100 partially known intruders. As shown in Figure 5, for each segment, the 95% confidence interval associated with the trajectory robustness estimate shrinks as N_{mc} increases. Moreover, when $N_{mc} \geq 2500$, the half uncertainty band is smaller than 2% for all three segments.

Given a traffic scenario, we further assess the convergence of a policy's performance measure, such as mission survival rate, with respect to the number of samples $N_{instance}$ used to estimate performance metrics. Consider measuring the mission survival rate when the vehicle employs policy π_S in a traffic scenario parameterized with (N_p, N_u). In this example, let $E^{'}$ indicate the occurrence in which the vehicle survives a mission. The mission survival rate $p^{'}$ is defined by the expected value of the indicator $Exp(\mathbb{I}_{E'})$. With n independent samples of $\mathbb{I}_{E'}$, we can then use the sample average $\hat{p_n^{'}}$ to estimate mission survival rate $p^{'}$. Following the same logic as the MC simulation, we can theoretically obtain a half uncertainty band $\leq 2\%$, centered on the mission survival rate estimate $\hat{p_n^{'}}$, with a 95% confidence level by utilizing at least 2401 samples of $\mathbb{I}_{E'}$. Therefore, to assess the performance of a trajectory planning policy given a traffic scenario, we set the number of traffic instances to $N_{instance} = 2500$. As shown in Figure 6, we consider employing policy π_S in five different traffic scenarios where the ratio of N_p to N_u is kept constant at one while the total number of intruders in the system grows linearly from 100 to 500. When $N_{instance} \geq 2500$, the half uncertainty band associated with the mission survival rate estimate is smaller than 2% across all traffic scenarios.

[1]Three segments are surrounded by relatively light, moderate, and severe traffic, respectively.

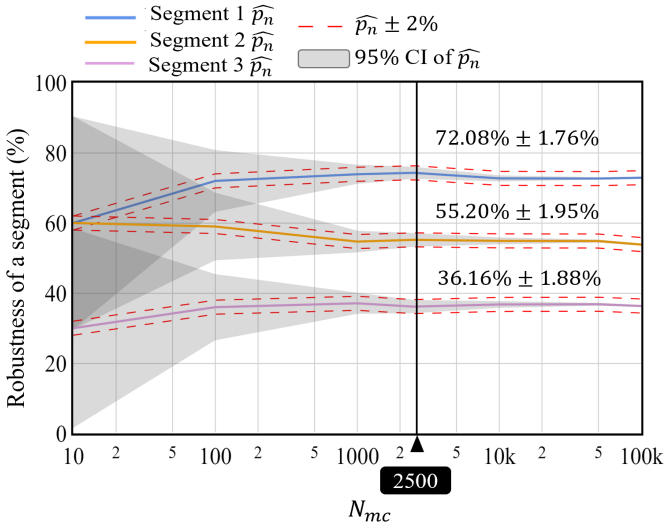

Figure 5: The estimate of the robustness of three independent trajectory segments via MC simulation

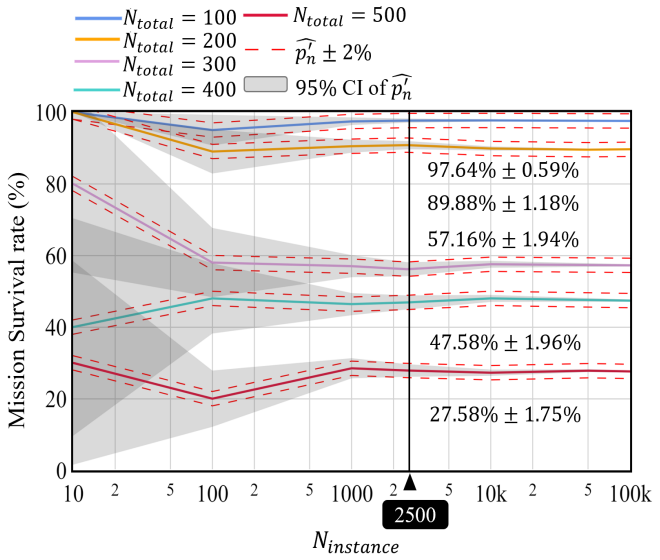

Figure 6: The estimate of mission survivability when employing policy π_S in five different traffic scenarios.

VI. Results and Discussion

A. Comparison between mission success and survival rates

In the first set of analyses, we examine how the traffic volume affects mission robustness and survivability. For this purpose, we focus on two critical performance indicators: mission success rate, which indicates mission robustness, and mission survival rate, which implies mission survivability. The results are summarized in Figure 7 for five traffic scenarios $(N_p, N_u) = \{(50, 50), (100, 100), \ldots, (250, 250)\}$ where the ratio of partially known and unknown intruders is constant at 1. Along the x-axis, the total number of intruders N_{total} in the environment increases from 100 to 500 corresponding to the five traffic scenarios. Each traffic scenario is associated with a group of bars, and each bar in a group reports the vehicle's survival and success rates employing one of the four proposed planning policies including π_{N_f}, π_S, π_R, and $\pi_{Avg(R)}$ as introduced in Section IV-B.

First, as expected, the mission success and survival rates decline as more intruders are introduced into the environment regardless of the policy that has been employed in trajectory planning.

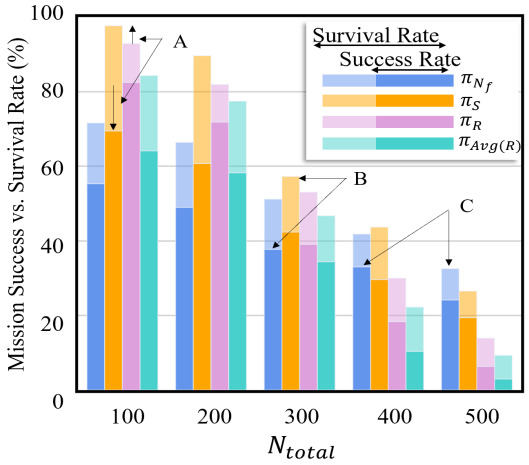

Figure 7: Mission Success Rate vs. Mission Survival Rate.

Second, allowing a vehicle to land at alternative landing sites increases flight safety significantly. This is shown by positive improvements in mission survival rate compared to mission success rate across all columns. In the first set of grouped bars, for example, the orange bar indicates that when the vehicle flies through the airspace that contains 100 actively operating intruders while enforcing policy π_S, the vehicle has a 28.24% higher probability of landing safely with only two additional landing sites provided.

Third, policy π_R may be sufficient to maximize mission success rate at low traffic volume, but this is achieved at the expense of the survival rate, which is improved by also considering the number of trajectories per policy π_S (see A). Considering the number of feasible trajectories becomes more effective as the traffic volume increases, as indicated by the fact that π_S is the best policy for both success and survival at 300 (see B), and the policy π_{N_f} is best for mission success at 400 and for survival at 500 (see C). These observations motivate us to further explore in Section IV-B the policies π_S and π_{N_f} in terms of their ability to mitigate uncertainty that is either partly known or unknown; and to concentrate on mission survival for the remainder of the study.

B. Effectiveness of policies π_S and π_{N_f}

We observed that policy π_{N_f} offers no significant benefit in protecting the ownship from partially known uncertainty. The is borne out by the comparison in Figure 8 of the mission survival rate of the four policies when the number of unknown intruders in the environment is kept constant at 100 and the number of partially known intruders is increased from 50 to 300. It is apparent that employing policy π_S results in the highest mission survival rate across all 5 traffic scenarios. This promising result can be viewed as a by product of the Theorem II.1. In particular, maximizing the total number of feasible trajectories and improving the robustness of a trajectory increase the likelihood that the ownship will maintain at least one trajectory that will lead to a safe landing.

Second, the policy π_{N_f} is shown to be the most effective policy in mitigating extreme uncertainty - the unknown. As shown in Figure 9, where we present the performance of the proposed policies while keeping the number of partially known intruders at 100 and increasing the number of unknown intruders from 50 to 300, we observe a gradual decline in the mission survival rate as more unknown intruders are introduced into the environment no matter which policy is applied.

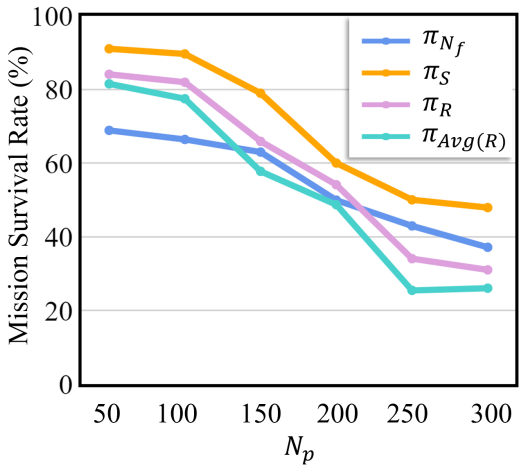

Figure 8: Performance of four trajectory planning policies as a function of the number of partially known intruders. Note that the number of unknown intruders is kept at 100.

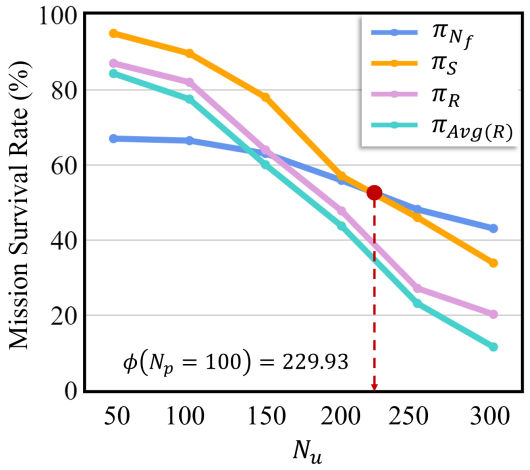

Figure 9: Performance of proposed trajectory planning policies as a function of the number of unknown intruders. Note that the number of partially known intruders is kept at 100.

The phenomenon that occurs when the blue curve (i.e., policy π_{N_f}) and the yellow curve (i.e., policy π_S) intersect at the red dot corresponds to a traffic scenario with around 229.9[1] unknown intruders. This suggests that given a fixed number of partially known intruders in the environment N_p, policy π_{N_f} provides marginal protection against unknown uncertainty over all other policies when the number of unknown intruders exceeds the threshold value $\phi(N_p)$. Therefore, it is beneficial to switch to the policy π_{N_f} at that threshold. As shown in Figure 9, $\phi(N_p = 100) = 229.9$.

The intersection of the blue and yellow curves was identified for values of N_p between 100 and 300, and the corresponding policy switching threshold $\phi(N_p)$ for the different values of N_p are plotted (see red line) in Figure 10. As may be seen, policy π_{N_f} outperforms policy π_S for traffic situations in the blue area where $N_u > \phi(N_p)$. This suggests that policy π_{N_f} is highly effective when the level of uncertainty is moderate to extreme (i.e., when there are more unknown intruders than partially known intruders). We posit that this is due to increased difficulty measuring mission robustness and survivabiity in a highly unpredictable environment (because it only accounts for partially known intruders), and that maintaining the number of trajectories is a more effective strategy when accounting for unknown risks.

[1]For comparison purpose, we keep the number to the nearest tenth. In real application, this number should be an integer.

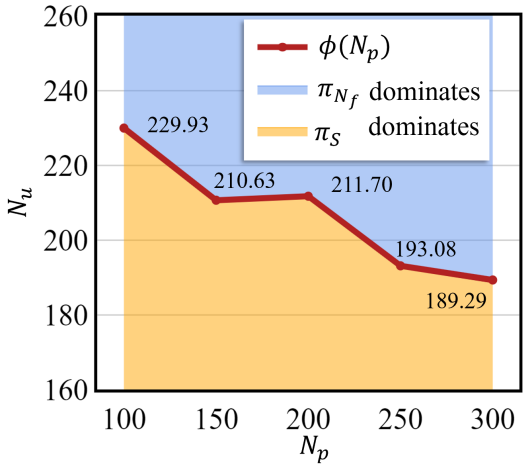

Figure 10: Policy switching boundary as a function of the number of unknown intruders.

Additional evidence for the effectiveness of policy π_{N_f} is provided in Figure 11. In this study, we set the total number of intruders in the environment to $N_{total} = 300$ while adjusting the ratio of unknown intruders in $[\frac{50}{300}, \frac{100}{300}, \cdots, \frac{250}{300}]$, corresponding to traffic scenarios $(N_p, N_u) = \{(50, 250), (100, 200), \cdots, (250, 50)\}$. We ignore the impact of traffic congestion and focus exclusively on the relative severity of the uncertainty. As the ownship's environment becomes more uncertain, policy π_S, π_R, and $\pi_{Avg(R)}$ become less effective, as evidenced by a lower mission survival rate. What stands out in this figure is the increasing trend in mission survival rate associated with policy π_{N_f}. The rising blue curve intersects the orange curve at a point corresponding to when 64.4% intruders in the environment are unknown. Hence, the policy π_{N_f} performs better than policy π_S when more than 64.4% intruders in the environment are unknown. These observations imply that policy π_{N_f} is highly competitive in mitigating extreme uncertainty caused by an increased ratio of unknown events.

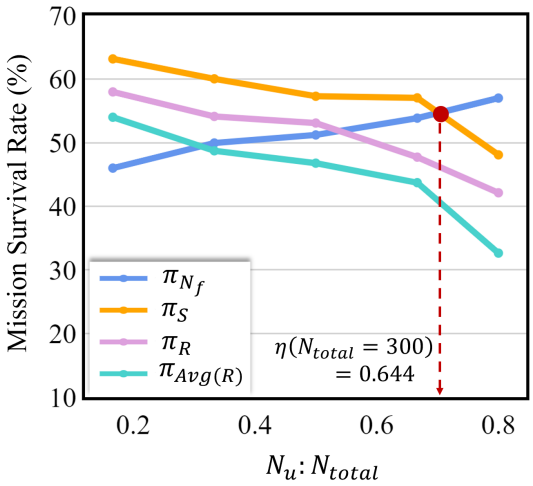

Figure 11: Performance of proposed trajectory planning policies as a function of the ratio of unknown intruders in the environment. Note that the total number of intruders is kept at 300.

The critical percentage $\eta(N_{total})$ of unknown intruders beyond which policy π_{N_f} becomes the most effective is shown in Figure 12 as a function of traffic volume (see red

line). As may be seen, $\eta(N_{total})$ decreases with increasing value of N_{total}, which (as the red line) indicates that as traffic becomes more congested, policy π_S becomes less tolerant of unknown intruders – $\eta(N_{total})$ drops from 64% to 38%. The blue region constitutes traffic scenarios with a large percentage of unknown intruders in the environment. We observe that policy π_{N_f} results in the highest mission survival rate for traffic scenarios falling in the blue region.

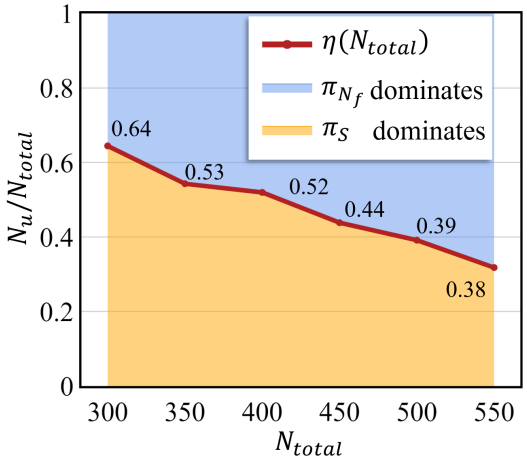

Figure 12: Policy switching boundary as indicated by the ratio of unknown intruders.

VII. Conclusion

In this paper, we studied the trajectory planning problem under uncertainty due to partially known and unknown intruders. We demonstrated that the success and survivability of a mission from a given point in space and time are dependent on two metrics: (1) the number of feasible trajectories available at that point, which can be computed using the proposed backtracking algorithm; and (2) the robustness of each of the trajectories, which can be evaluated using a Monte-Carlo simulation. We performed experiments for 36 traffic scenarios with varying numbers of partially known and unknown intruders. Our findings indicate that a policy π_S that combines the two metrics by maximizing the probability of having at least one feasible trajectory, outperforms all other policies when the vehicle is exposed to partially known uncertainty and moderate levels of unknown uncertainty. When uncertainty is dominated by unknown intruders, however, policy π_{N_f}, which maximizes the first metric, yields the highest success and survival rates.

We believe that the proposed metrics and policies will provide robust trajectory planning for autonomous vehicles operating in environments with moderate to extreme uncertainty, e.g., for on-demand urban air mobility missions in non-cooperative traffic. Finally, further research is required to address the computation time of the proposed solutions, which increases exponentially with the resolution of the discrete space-time network. In addition, the impact of spatial and temporal quantization on metric estimation should be investigated in future work.

Acknowledgment

This work was funded by the NASA Transformative Tools and Technologies Project. The material is based upon work supported by NASA under Contract Number NNA16BD14C, managed by the Universities Space Research Association (USRA).

References

[1] H. A. Blom, S. H. Stroeve, and T. Bosse, "Modelling of potential hazards in agent-based safety risk analysis," Tenth USA/Europe Air Traffic Management Research and Development Seminar, 2013.

[2] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *2011 IEEE international conference on robotics and automation*, IEEE, 2011, pp. 723–730.

[3] J. Miura and Y. Shirai, "Probabilistic uncertainty modeling of obstacle motion for robot motion planning," *Journal of Robotics and Mechatronics*, vol. 14, no. 4, pp. 349–356, 2002.

[4] G. S. Aoude, B. D. Luders, D. S. Levine, and J. P. How, "Threat-aware path planning in uncertain urban environments," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2010, pp. 6058–6063.

[5] L. Lulu and A. Elnagar, "A comparative study between visibility-based roadmap path planning algorithms," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2005, pp. 3263–3268.

[6] O. Takahashi and R. J. Schilling, "Motion planning in a plane using generalized voronoi diagrams," *IEEE Transactions on Robotics and Automation*, vol. 5, no. 2, pp. 143–150, 1989. DOI: 10.1109/70.88035.

[7] T. Khuswendi, H. Hindersah, and W. Adiprawita, "Uav path planning using potential field and modified receding horizon a* 3d algorithm," in *Proceedings of the 2011 International Conference on Electrical Engineering and Informatics*, 2011, pp. 1–6. DOI: 10.1109/ICEEI.2011.6021579.

[8] H. Idris, N. Shen, and D. Wing, "Complexity management using metrics for trajectory flexibility preservation and constraint minimization," in *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, including the AIAA Balloon Systems Conference and 19th AIAA Lighter-Than*, 2011, p. 6809.

[9] H. Idris, D. Delahaye, and D. Wing, "Distributed trajectory flexibility preservation for traffic complexity mitigation," in *Proceedings of the 8th USA/Europe Air Traffic Seminar (ATM'09)*, Citeseer, 2009.

[10] H. Idris and N. Shen, "Estimating airspace capacity based on risk mitigation metrics," in *Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013), Chicago*, 2013.

[11] R. Lange, H. Weinschrott, L. Geiger, A. Blessing, F. Dürr, K. Rothermel, and H. Schütze, "On a generic uncertainty model for position information," in *International Workshop on Quality of Context*, Springer, 2009, pp. 76–87.

[12] G. Trajcevski, O. Wolfson, K. Hinrichs, and S. Chamberlain, "Managing uncertainty in moving objects databases," *ACM Transactions on Database Systems (TODS)*, vol. 29, no. 3, pp. 463–507, 2004.

[13] J. Mattingley, Y. Wang, and S. Boyd, "Code generation for receding horizon control," in *2010 IEEE International Symposium on Computer-Aided Control System Design*, IEEE, 2010, pp. 985–992.

[14] D. C. Montgomery and G. C. Runger, *Applied statistics and probability for engineers*. John Wiley & Sons, 2010.

Biographies

Fanruiqi Zeng is a Ph.D candidate in the Daniel Guggenheim School of Aerospace Engineering at Georgia Tech. Her research interest lies at the intersection of optimization and algorithm design.

Dr. Husni R. Idris is an aerospace research engineer at the NASA Ames Research Center where he leads research in collective autonomous mobility for enabling increasingly autonomous airspace operations. He holds a Ph.D. from MIT and is an associate fellow of AIAA.

Dr. John-Paul Clarke is a professor of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, where he holds the Ernest Cockrell Jr. Memorial Chair in Engineering. Previously, he was a faculty member at Georgia Tech and MIT, and a Vice President of Strategic Technologies at United Technologies Corporation (now Raytheon). He received his Sc.D. from MIT and is a Fellow of the AIAA.