

Feasibility-Guided Safety-Aware Model Predictive Control for Jump Markov Linear Systems

Zakariya Laouar, Qi Heng Ho, Rayan Mazouz, Tyler Becker, Zachary N. Sunberg

Abstract—In this paper, we present a controller framework that synthesizes control policies for Jump Markov Linear Systems subject to stochastic mode switches and imperfect mode estimation. Our approach builds on safe and robust methods for Model Predictive Control (MPC), but in contrast to existing approaches that either optimize without regard to feasibility or utilize soft constraints that increase computational requirements, we employ a safe and robust control approach informed by the feasibility of the optimization problem. We formulate and encode finite horizon safety for multiple model systems in our MPC design using Control Barrier Functions (CBFs). When subject to inaccurate hybrid state estimation, our feasibility-guided MPC generates a control policy that is maximally robust to uncertainty in the system’s modes. We evaluate our approach on an orbital rendezvous problem and a six degree-of-freedom hexacopter under several scenarios and benchmarks to demonstrate the utility of the framework. Results indicate that the proposed technique of maximizing the robustness horizon, and the use of CBFs for safety awareness, improve the overall safety and performance of MPC for Jump Markov Linear Systems.

I. INTRODUCTION

Advances in estimation and control techniques have enabled many applications of cyber-physical systems such as unmanned aerial systems (UAS), including military, industrial, and commercial [1]. As these systems are increasingly deployed alongside humans in the real world, safety and reliability becomes increasingly important. The systems that operate in these environments must contend with the inevitable uncertainties present in the real world, which can be continuous or discrete in nature [2]. In this work, we consider Jump Markov Linear Systems (JMLS), a class of multiple model systems described by a discrete-time linear dynamic model whose parameters depend on discrete modes that the system switches between [3]. Many real-world systems have imperfect models of the uncertain mode switches and imprecise sensors to detect these switches. This makes the problem of controlling these systems challenging due to the need to provide safety guarantees while accounting for the possibility of mode switches. This paper focuses on this challenge and develops a principled framework for safety-aware control of JMLS subject to mode switches.

Example 1. Consider a mineshaft inspection scenario, where a hexacopter is tasked to inspect a narrow mineshaft while avoiding collisions with obstacles. The system is subject to partially observable stochastic rotor faults. At any time, any one of its six rotors may fail with low probability,




Fig. 1: Feasibility-guided MPC for a multiple-model hexacopter subject to partially observable stochastic rotor faults. Our approach computes an adaptive robust control (consensus) horizon at each planning step to improve robustness and feasibility.

causing its dynamics to change and the system to switch mode.

The hexacopter can be modeled as a JMLS subject to stochastic mode switches. Due to imperfect mode switching models and sensors, when a rotor has failed, the system may require some time to detect it. To ensure that the task is completed successfully and safely, a controller needs to ensure that the system is safe even if a mode switch occurs and is undetected. A typical method is to compute a solution robust to the worst case [4]. However, in many scenarios, optimizing for the worst case over the entire control horizon can lead to an infeasible optimization problem; There may not be a single control trajectory that is safe for all possible modes. Therefore, the system should determine the appropriate horizon of controller robustness, known as a *consensus horizon*, by accounting for problem feasibility, as depicted in Fig. 1. In this work, we aim to design a controller that accounts for problem feasibility while being robust to stochastic mode jumps and mode switch detection delays.

A. Related Work

The general problem of planning under stochastic uncertainty can be modeled as a partially observable Markov decision process (POMDP) [5]. Unfortunately, solving a POMDP exactly is intractable due to the curses of dimensionality and history. Approximate solutions are often used in practice [6]–[9], but do not perform well for large or continuous action spaces, which is a characteristic of multiple model systems.

We address safety-critical control for JMLS by proposing a Model Predictive Control (MPC) based approach. Scenario or multi-stage MPC approaches have been used to control systems under uncertainty, through sampling of scenarios [10]. These methods include a robustness horizon up to

Authors are with the department of Aerospace Engineering Sciences at the University of Colorado Boulder, CO, USA {zakariya.laouar@colorado.edu}.

which scenarios are restricted to share the same action at each step in the planning horizon. However, a major limitation of scenario or particle approaches for controlling multiple model systems manifests in ensuring safety, since they have at most asymptotic safety guarantees. In our approach, we design a controller that maximizes robustness over mode uncertainty without sampling scenarios.

Several safety techniques have been proposed for MPC, including barrier approaches [11]–[13], computation of reachable sets [14] and discrete approximations [15]. Particularly, work [16] leverages CBFs to generate policies that consider uncertainty in dynamics. In this paper, we utilize CBFs in MPC and extend it to the case of JMLS mode uncertainty.

Feasibility of MPC problems involving CBF constraints has been studied and enhanced in recent work [17]–[20]. [17], [18] enhance feasibility by reducing the CBF decay rate for single mode systems. [19] increase feasibility by using Generalized CBFs for inconsecutive timesteps. While this increases feasibility, safety is no longer enforced for intermediate timesteps. Another approach is using soft constraints via slack variables, compromising on safety guarantees [20]. In this work, we discuss the impact of the size of the consensus horizon on the feasibility of the JMLS MPC problem with CBF constraints, and propose a method to improve feasibility while maximizing safety for each planning horizon.

B. Contribution

In this work, we present an MPC design that encodes the safety of JMLS with uncertain mode jumps and mode switch detection delays using CBFs. We provide a feasibility and safety analysis of the size of the consensus horizon, and propose a method to obtain the optimal consensus horizon while ensuring the finite-horizon optimization problem is feasible. We evaluate our control design on two case study scenarios. Results indicate that the proposed technique of maximizing the robustness horizon guided by the feasibility of the optimization problem and the use of CBFs are critical for the overall safety and performance of MPC for JMLS.

In summary, the main contributions of this paper are four-fold: (i) A safety-aware control design for controlling JMLS using CBF constraints, (ii) An analysis of the role of the consensus horizon on feasibility and safety of the closed-loop system, (iii) A practical method to obtain the optimal consensus horizon that maximizes robustness while ensuring problem feasibility at each time step, and (iv) Simulation case studies and benchmarks that demonstrate the safety and performance of the proposed approach.

II. SYSTEM SETUP

In this work, we aim to develop safety-critical control policies for multiple-model systems subject to mode switches. Consider a discrete-time Jump Markov Linear System

$$\begin{aligned} \mathbf{x}_{t+1} &= A_{i(t)}\mathbf{x}_t + B_{i(t)}\mathbf{u}_t, \\ \mathbf{y}_t &= H_{i(t)}\mathbf{x}_t + \mathbf{v}_t, \end{aligned} \quad (1)$$

where $t \in \mathbb{N}$ and system state $\mathbf{x}_t \in \mathbb{R}^n$. The control input $\mathbf{u}_t \in \mathcal{U} \subset \mathbb{R}^m$, where \mathcal{U} is a compact set. The dynamics

and control matrices of the $i(t)$ th mode are $A_{i(t)} \in \mathbb{R}^{n \times n}$ and $B_{i(t)} \in \mathbb{R}^{n \times m}$, respectively. Measurement $\mathbf{y}_t \in \mathbb{R}^p$ is governed by $H_{i(t)} \in \mathbb{R}^{p \times n}$, and the additive noise term $\mathbf{v}_t \in \mathbb{R}^p$. The index $i(t) \in \{1, \dots, M\}$ represents the discrete mode at a given time t , of which evolution is governed by a state and input independent finite state Markov chain

$$\mu_{t+1} = \Omega\mu_t, \quad (2)$$

where $\Omega = \{\omega_{ij}\} \in [0, 1]^{M \times M}$ defines the mode transition matrix and $\mu_t \in [0, 1]^M$ defines the categorical mode probability at time t .

Due to noisy measurements, the system's state is estimated by a hybrid state estimator, such as an Interacting Multiple Model (IMM) filter [3] which generates a mode probability distribution $\hat{\mu}_t$ and a mean continuous state $\hat{\mathbf{x}}_t$ at each timestep t from the measurements. Then, starting from an initial state x_0 and mode $i(0) \sim \mu_0$, under an output-feedback controller π , the evolution of the discrete-time closed-loop system is

$$\begin{aligned} \mathbf{x}_{t+1} &= A_{i(t)}\mathbf{x}_t + B_{i(t)}\pi(\mathbf{Y}_t), \\ \mathbf{y}_t &= H_i\mathbf{x}_t + \mathbf{v}_t, \\ i(t) &\sim \mu_t, \\ \mu_{t+1} &= \Omega_{i(t)}, \end{aligned} \quad (3)$$

where $\mathbf{Y}_t = \{\mathbf{y}_0, \dots, \mathbf{y}_t\}$ and $\Omega_{i(t)}$ refers to the $i(t)$ th column of the mode transition matrix.

In many safety-critical systems, the mode transition matrix Ω is an approximation [21]. As such, the state estimator may have inaccurate mode probability distributions, or require additional time to detect mode switches when they occur. In this work, we aim to develop a safety-aware controller π robust to the inaccuracy of the mode estimate.

III. SAFETY-AWARE TRAJECTORY PLANNER DESIGN

The safety-aware trajectory planning problem consists of synthesizing a controller for a robot subject to uncertain mode switches. This section discusses our approach to designing a controller for this problem. The approach uses MPC with control barrier functions and a novel variable consensus horizon that balances between maximizing robustness over modes and feasibility of the optimization problem.

A. Model Predictive Control

To guarantee safety for the entire problem duration, a planner would have to reason about all possible sequences of states and modes across the planning horizon. This presents a significant computational challenge, rendering most problems computationally infeasible. However, open-loop planning can often provide a satisfactory approximation of an optimal closed-loop plan while greatly enhancing computational efficiency [6], [22]. To this effect, we adopt the receding horizon paradigm of MPC.

We provide a base template MPC formulation for JMLS that we build upon in the following sections. At each time step, the Model Predictive Controller plans over a finite horizon H . We represent the stage cost as $q(\mathbf{x}_k^i, \mathbf{u}_k, \hat{\mu}_k)$

and the terminal cost as $p(\mathbf{x}_H^i, \hat{\mu}_k)$, where \mathbf{x}_k^i represents the state for mode $i(k)$ at planning time k . The system is subject to the continuous dynamics in System (1) with mode-switching dynamics (2), initial state x_t , and initial mode probability $\hat{\mu}_t$. The safety criteria, e.g. "avoid obstacle collisions" is represented by constraining the state for each mode \mathbf{x}_k^i to remain within a designer-specified set \mathcal{S}_k . In our approach, we utilize CBFs for \mathcal{S}_k . Finally, under some control constraints, the goal is to generate a control sequence that minimizes the cost function subject to the constraints. This yields the following optimal control problem (OCP):

$$\begin{aligned} \mathbf{U}_t^* := \operatorname{argmin}_{\mathbf{X}_t, \mathbf{U}_t} & p(\mathbf{x}_H^{1:M}, \hat{\mu}_k) + \sum_{k=0}^{H-1} q(\mathbf{x}_k^{1:M}, \mathbf{u}_k^{1:M}, \hat{\mu}_k) \\ \text{s.t.} & \forall k \in \{0, \dots, H-1\}, i \in \{1, \dots, M\}, \\ & \mathbf{x}_{k+1}^i = A_i \mathbf{x}_k^i + B_i \mathbf{u}_k^i, \\ & \mathbf{x}_{k+1}^i \in \mathcal{S}_k, \\ & \mathbf{x}_0^i = x_t, \\ & \mathbf{u}_k^i \in \mathcal{U}, \\ & \hat{\mu}_k = \hat{\mu}_t, \end{aligned} \quad (4a-f)$$

where $\mathbf{X}_t = \{\{\mathbf{x}_k^i\}_{i=1}^M\}_{k=0}^{H-1}$, $\mathbf{U}_t = \{\{\mathbf{u}_k^i\}_{i=1}^M\}_{k=0}^{H-1}$ and the mode probability $\hat{\mu}$ is assumed to remain constant over H . We assume that p and q are quadratic cost functions, making (4) a Quadratic Program.

At each instance of Problem (4), a sequence of open-loop control inputs is generated, and only the first input is applied to System (1). After each control input is applied, the system generates a new hybrid state observation. Then, the state is updated and Problem (4) is solved again. This recursive process allows the closed-loop controller π to overcome the inaccuracies of neglecting the acquisition of future state information due to open-loop planning.

In this current formulation, each mode has a corresponding control input trajectory that may be unique with respect to the other trajectories. Thus, π must somehow choose between these trajectories to execute the first control input. In Section III-C, we discuss how to account for uncertainty in the modes and mode switches to make this decision.

B. Control Barrier Functions

CBFs provide a framework of safety by creating a link between the dynamical system (1) and a designer's expressed safety criteria (4c), guaranteeing the satisfaction of these criteria [23]. For JMLS, we leverage CBFs as state constraints of (4c) to provide robustness against mode uncertainty. We first introduce our notion of safety for closed-loop JMLS using the concept of set invariance [23].

Definition 1 (Forward Invariance & Safety). *A compact set $\mathcal{C} \subset \mathbb{R}^n$ is forward invariant for closed-loop System (3) if $\mathbf{x}_0 \in \mathcal{C}$ implies $\mathbf{x}_t \in \mathcal{C}$ for all $t \in \mathbb{N}$. In this case, we call System (3) safe with respect to set \mathcal{C} .*

Additionally, we say π is *safe* if the closed-loop system under π is safe with respect to \mathcal{C} . Formally, we enforce the

notion of safety using CBFs for discrete systems [24]. To that extent, let set \mathcal{C} be described by the 0-superlevel set of a continuously differentiable function $\beta : \mathbb{R}^n \rightarrow \mathbb{R}$, such that

$$\begin{aligned} \mathcal{C} &= \{\mathbf{x} \in \mathbb{R}^n \mid \beta(\mathbf{x}) \geq 0\}, \\ \partial\mathcal{C} &= \{\mathbf{x} \in \mathbb{R}^n \mid \beta(\mathbf{x}) = 0\}. \end{aligned} \quad (5)$$

Function β is a CBF if $\frac{\partial\beta}{\partial\mathbf{x}} \neq 0, \forall x \in \partial\mathcal{C}$, and there exists an extended class \mathcal{K}_∞ function γ , such that for discrete-time System (1), β satisfies for a given time k

$$\exists \mathbf{u}_k \in \mathcal{U} \text{ s.t. } \Delta\beta(\mathbf{x}_k, \mathbf{u}_k) \geq -\gamma(\beta(\mathbf{x}_k)), \quad (6)$$

where $\Delta\beta(\mathbf{x}_k, \mathbf{u}_k) := \beta(\mathbf{x}_{k+1}) - \beta(\mathbf{x}_k)$. For simplicity, we will use the scalar form of γ throughout the manuscript: $0 < \gamma \leq 1$. From (6) we see that $1 - \gamma$ defines a lower bound on the rate of exponential decrease of $\beta(\mathbf{x})$, i.e., $\beta(\mathbf{x}_{k+1}) \geq (1 - \gamma)\beta(\mathbf{x}_k)$.

For our System described in (1), the safe set is the set of states satisfying (6) at each time k ,

$$\mathcal{S}_k = \{\mathbf{x}_{t+k} \in \mathbb{R}^n : \beta(\mathbf{x}_{t+k}) \geq (1 - \gamma)\beta(\mathbf{x}_{t+k-1})\}. \quad (7)$$

Note, \mathcal{S}_k is implicitly generated from the safe control input in (6). Since our problem is characterized by mode uncertainty due to imperfect modeling and estimation, we first represent safety with respect to each mode. When planning over M modes, the mode-specific set is, $\forall i \in \{1, \dots, M\}$,

$$\mathcal{S}_k^i = \{\mathbf{x}_{t+k}^i \in \mathbb{R}^n : \beta(\mathbf{x}_{t+k}^i) \geq (1 - \gamma)\beta(\mathbf{x}_{t+k-1}^i)\}. \quad (8)$$

Eq. (4c) requires the controller to be safe for each mode, separately. In the next section, we show how to enforce finite horizon safety across all modes using control consensus.

C. Robustness via Control Consensus across Modes

The optimal solution to (4) contains the optimal trajectory for each mode i . When there is high uncertainty over modes, or when detection of mode switching is inaccurate, executing the control input for one mode optimistically, based on maximum likelihood, may lead to unsafe behavior if the mode is estimated incorrectly. A safer and more conservative approach to tackle mode uncertainty is by enforcing the control inputs of the modes to be the same across all modes for some number of time steps during planning [25]. This number of steps is called the *consensus horizon*.

Definition 2 (Consensus Horizon). *A consensus horizon $h \in \{0, \dots, H-1\}$ is a horizon up to which the control input is constrained to be the same for all modes during planning, i.e., $\mathbf{u}_k^i = \mathbf{u}_k^j \quad \forall i, j \in \{1, \dots, M\}, \forall k \in \{0, \dots, h\}$.*

Thus, Problem (4) requires a constraint specifying that all modes share the same control input for h planning time steps.

This yields the updated formulation:

$$\mathbf{U}_t^* := \underset{\mathbf{x}_t, \mathbf{U}_t}{\operatorname{argmin}} p(\mathbf{x}_H^{1:M}, \hat{\mu}_k) + \sum_{k=0}^{H-1} q(\mathbf{x}_k^{1:M}, \mathbf{u}_k^{1:M}, \hat{\mu}_k) \quad (9a)$$

$$\text{s.t. } \forall k \in \{0, \dots, H-1\}, i, j \in \{1, \dots, M\},$$

$$\mathbf{x}_{k+1}^i = A_i \mathbf{x}_k^i + B_i \mathbf{u}_k^i, \quad (9b)$$

$$\mathbf{x}_{k+1}^i \in \mathcal{S}_k^i, \quad (9c)$$

$$\mathbf{x}_0^i = x_t, \quad (9d)$$

$$\mathbf{u}_k^i \in \mathcal{U}, \quad (9e)$$

$$\hat{\mu}_k = \hat{\mu}_t, \quad (9f)$$

$$\forall l \in \{0, \dots, h-1\},$$

$$\mathbf{u}_l^i = \mathbf{u}_l^j. \quad (9g)$$

The interaction between mode-dependent CBF constraints (9c) and consensus horizon control constraints (9g) leads to robustness over all modes over the consensus horizon.

With (9) now parametrized with a consensus horizon h , we are faced with the following question: *what should h be?* The typical approach for multi-modal MPC is to use a user-defined fixed consensus horizon h [25]. One common approach is to set $h = 1$, for one-step consensus, resembling hindsight optimization techniques. This approach reasons through a certainty equivalent lens that assumes the robot will resolve its uncertainty after the first timestep. Another common approach is to set $h = H$, for full consensus, corresponding to robust control methods. Conversely, such an approach plans under the assumption that the uncertainty will not be resolved throughout the planning horizon. Nonetheless, selecting h is challenging due to the difficulty in estimating how quickly the mode can be determined during execution. Substantial trial and error has to be conducted to choose a sufficient consensus horizon a priori. Additionally, as we discuss in the following, a fixed a priori chosen consensus horizon may lead to safety and feasibility issues during closed-loop execution.

D. Feasibility Analysis

In this section, we discuss the potential safety and feasibility drawbacks of planning with a pre-defined consensus horizon. Specifically, we show that too large h may lead to infeasible OCP while too small h may have safety issues. First, note that a larger consensus horizon leads to a smaller feasible set. This is formalized in the following proposition.

Proposition 1. *Let F_h be the feasible set for (9) with a consensus horizon h . Given two consensus horizons $h_1 < h_2 \leq H$, $F_{h_1} \supseteq F_{h_2}$.*

Proof. The feasible set of the optimization problem is the intersection of the feasible set for each independent constraint. A larger consensus horizon strictly adds constraints to the optimization problem. Therefore, the feasible set for h_2 cannot be a superset of the feasible set for h_1 . \square

Increasing the consensus horizon may shrink the feasible set and may lead to an infeasible optimization problem (if

the feasible set becomes empty). Thus, we need to ensure that (9) is parameterized with a *feasible consensus horizon*.

Definition 3 (Feasible Consensus Horizon). *A feasible consensus horizon h is a consensus horizon for which (9) is feasible.*

On the other hand, a lower consensus horizon may be unsafe during closed-loop execution.

Proposition 2. *Let π_A be a safe controller that generates a feasible sequence of control inputs that repeatedly solves (9) with a consensus horizon $h_A = H$ for $t \in \{1, \dots, \infty\}$. Let controller π_B repeatedly solve (9) with a consensus horizon $h_B < H$. π_B may be an unsafe controller.*

Proof Sketch. Let $F_{h_{A,1}}$ and $F_{h_{B,1}}$ be the feasible set at $t = 1$ for π_A and π_B , respectively. From Proposition 1, $F_{h_{B,1}} \supseteq F_{h_{A,1}}$.

Let $(\mathbf{X}^*, \mathbf{U}^*)$ be the optimal state and control inputs output by π_B at $t = 1$. Since $F_{h_{B,1}} \supseteq F_{h_{A,1}}$, it is possible that $(\mathbf{X}^*, \mathbf{U}^*) \notin F_{h_{A,1}}$. That is, $(\mathbf{X}^*, \mathbf{U}^*)$ is not a feasible decision set for the horizon $h_{A,1}$. Therefore, when the first control input \mathbf{u}_1 is executed for π_B , the system may reach a state $\mathbf{x}_{B,2}$ that is not contained within the reachable set of states for π_A at $t = 2$, i.e.

$$\begin{aligned} \mathbf{x}_{B,2} &\in \\ &\{\mathbf{x} : \mathbf{x}_2 = A_i \mathbf{x}_1 + B_i \mathbf{u}_1, \mathbf{u}_1 \in F_{h_{B,1}}, i \in \{1, \dots, M\}\} \\ &\supseteq \{\mathbf{x} : \mathbf{x}_2 = A_i \mathbf{x}_1 + B_i \mathbf{u}_1, \mathbf{u}_1 \in F_{h_{A,1}}, i \in \{1, \dots, M\}\}. \end{aligned}$$

Then, at $t = 2$, if $\mathbf{x}_{B,2}$ is not in the reachable set of states for π_A , there exists scenarios in which the system starting from $\mathbf{x}_{B,2}$ may be unsafe. In the worst case scenario, the optimization problem (9) with $\mathbf{x}_{B,2}$ may immediately be infeasible since the previous optimization at $t = 1$ did not account for the control consensus constraints for planning time steps $k = h_B$ to $k = h_B + 1$, leading to an empty feasible set for the optimization problem (9) starting from $\mathbf{x}_{B,2}$, whereas $\mathbf{x}_{A,2} \in \{\mathbf{x}_2 : \mathbf{x}_2 = A_i \mathbf{x}_1 + B_i \mathbf{u}_1, \mathbf{u}_1 \in F_{h_{A,1}}, i \in \{1, \dots, M\}\}$ with $h_A = H$ is safe. \square

Therefore, if a controller that maximizes the consensus horizon is safe, a controller operating with a lower consensus horizon is not necessarily safe. Propositions 1 and 2 exemplify the dilemma in pre-determining a consensus horizon.

Consider Example 1 again. If the hexacopter adopts a longer consensus horizon and is close to an obstacle, there may be no single sequence of control inputs that allows it to avoid a collision with the mineshaft walls across all possible mode dynamics. In this case, the OCP becomes infeasible. However, reducing the consensus horizon may expand the feasible set and render the OCP feasible. On the other hand, when the problem is already feasible, we generally want the consensus horizon to be as large as possible to be more robust to mode uncertainty. A pre-defined fixed consensus horizon is unable to achieve such a behavior.

Algorithm 1: Maximally Feasible Consensus Horizon Binary Search

Input : x, H
Output: h^*

```

1  $h_{max} \leftarrow H - 1$ 
2  $h_{min}, h_{best} \leftarrow 0$ 
3 while  $h_{min} < h_{max}$  do
4    $h \leftarrow \lfloor (h_{min} + h_{max})/2 \rfloor$ 
5   if  $\text{feasible}(h)$  then
6      $h_{min} \leftarrow h + 1$ 
7      $h_{best} \leftarrow h$ 
8   else
9      $h_{max} \leftarrow h - 1$ 
10 if  $\text{feasible}(h_{max})$  then
11    $h \leftarrow h_{max}$ 
12 else
13    $h \leftarrow h_{best}$ 
14 return  $h^* = h$ 

```

where

- $\text{feasible}(h)$: returns a boolean expressing the feasibility of (10) with consensus horizon h .
-

E. Adaptive Consensus Horizon Planning

Planning with a pre-defined fixed consensus horizon has drawbacks in feasibility (if consensus horizon is too large) and safety (if consensus horizon is too small). This implies that there exists a possibly varying optimal sequence of consensus horizons that is best able to balance between robustness and feasibility to guarantee the safety of the robot. However, since the controller uses a finite planning horizon in a receding horizon MPC, computing the optimal consensus horizon sequence beforehand is not generally possible; Designing a recursively feasible MPC controller is an open area of research [26]. Therefore, using the results of Proposition 1 and 2 as a heuristic, we propose a *greedy* controller design. At each time step, we select a consensus horizon that is as large as possible while maintaining feasibility.

In the first stage of each planning step, we compute the consensus horizon that maximizes h such that the OCP is still feasible. We define such a consensus horizon as a *maximally feasible consensus horizon*.

Definition 4 (Maximally Feasible Consensus Horizon). *A maximally feasible consensus horizon h^* is a feasible consensus horizon in which all horizons below h^* are feasible and all horizons above h^* are infeasible, i.e.,*

$$h^* := \underset{\mathbf{x}_k^i, \mathbf{u}_k^i, h}{\text{argmax}} \quad h \quad (10)$$

s.t. (9) is feasible.

To compute h^* efficiently, we first see that the feasibility of consensus horizons is an ordered monotonic sequence. That is, if h_i is a feasible consensus horizon, $h_j < h_i$ is also a feasible consensus horizon for all $j \in \{0, \dots, i-1\}$. Using

this insight, we propose to solve (10) through a search over an ordered value¹. This can be efficiently done using, e.g., a binary search over the possible consensus horizon values $\{1, \dots, H-1\}$ as shown in Algorithm 1. For a given h , we construct an optimal control problem identical to (9) with the objective function (4b) replaced with an arbitrary constant. During each search iteration, we can tractably determine if a given h is a feasible consensus horizon by solving the newly constructed problem parametrized with h and obtaining a feasibility certificate. This reduces the optimization program from a Quadratic Program to a Linear Program. After computing h^* , we solve (9) by setting the consensus horizon at that time step to $h = h^*$.

F. Computational Complexity

In comparison to single-mode MPC OCP [28], each control input call from our approach is more computationally intensive in the following ways. First, a JMLS inherently admits a larger optimization problem, as there are more decision variables and constraints. The number of decision variables increases linearly in the number of modes, and the number of constraints increases linearly in the number of modes and consensus horizon. That is, a JMLS MPC OCP has $O(M \cdot T \cdot H \cdot m \cdot n)$ decision variables, compared to $O(T \cdot m \cdot n)$ of single-mode MPC OCP. Second, as compared to using a fixed consensus horizon, computing h^* in Alg. 1 computes the feasibility of varying consensus horizon $\log_2 H$ number of times. Let T be the maximum time taken to solve an OCP with a fixed h . Then, an upper bound on the time taken by our approach to solve (9) is $(\lceil \log_2 H \rceil + 1)T$. However, each feasibility check is a Linear Program with a constant objective function, which can generally be solved faster than finding an optimal solution to the full OCP.

This added computation may lead to a slower solution output, which may lead to performance issues for on-board computation, depending on the size of the state space, number of system modes, and the solution time requirements of the system. If the ability to run in high frequency, e.g. real-time, is an important consideration, a practical solution is to start from $h = 1$ and increment until h^* is found or the planning time limit is reached. Another solution is to solve (9) with $h = 1$ as a backup controller in case a solution is not found in the planning time limit. Further, the search for h^* can be parallelized, mitigating the added computation from maximizing the feasible consensus horizon.

IV. SIMULATION EVALUATION

In this section, we evaluate our approach on two robotic systems in simulation: a 6-dimensional spacecraft conducting orbital rendezvous and a 12-dimensional hexacopter system subject to rotor faults. We demonstrate the performance of our approach against several baselines. Simulations are implemented in Julia, and the quadratic programs are solved using the interior-point solver Clarabel.jl [29]. The optimizations are all performed single-threaded on a computer with

¹This search procedure bears similarities to the bisection method for quasiconvex optimization [27].

a nominally 3.7 GHz CPU. The computational performance of the algorithms is discussed at the end of this section.

For both domains, we represent the cost as a quadratic stage-wise function that includes a state reference error term $\|(\mathbf{x}_k^i - \mathbf{x}_r)\|_Q^2$ and a penalty on the control effort $\|\mathbf{u}_k\|_R^2$, where $Q, R \succ 0$:

$$\sum_{i=0}^M \mu_k(i) \left[\sum_{k=0}^{H-1} (\|(\mathbf{x}_k^i - \mathbf{x}_r)\|_Q^2 + \|\mathbf{u}_k\|_R^2) \right]. \quad (11)$$

For simplicity, the terminal cost is set to zero. We weight the stage-wise cost per mode by its corresponding probability.

The purpose of our evaluation is to assess the ability of our approach to plan robustly with respect to mode uncertainty and estimator inaccuracy. To this end, at every time step, each system is given a categorical mode probability estimate and mean state estimate by an oracle filter. In each simulation, a mode switch is induced at a time unknown to the controller. Additionally, the oracle filter undergoes a specified delay, also unknown to the controller, in detecting when a mode switch occurs. We compare our approach vs. baselines across a range of mode-switch and estimation-delay times in addition to problem-specific parameters discussed below.

We compare our approach to the following baselines:

- **First-Step Consensus:** Use consensus horizon $h = 1$.
- **Full-Step Consensus:** Use consensus horizon $h = H$.
- **Non-Robust:** No consensus constraints. Instead, plan for the maximum likelihood mode according to the categorical mode probability $\hat{\mu}$ given by the oracle filter,

A. Case Studies

1) *Spacecraft Orbital Rendezvous:* We first demonstrate our approach on a spacecraft rendezvous involving a chaser and target satellite. Following [30], the linearized system dynamics are characterized by the following Clohessy-Wiltshire-Hill equations:

$$\dot{\mathbf{x}} = \begin{bmatrix} \mathbf{0}_3 & & I_3 \\ 3n^2 & 0 & 0 & 0 & 2n & 0 \\ 0 & 0 & 0 & -2n & 0 & 0 \\ 0 & 0 & -n^2 & 0 & 0 & 0 \end{bmatrix} \mathbf{x} + \begin{bmatrix} \mathbf{0}_3 \\ I_3 \end{bmatrix} \mathbf{u} \quad (12)$$

where the state $\mathbf{x} = [x_r, y_r, z_r, \dot{x}_r, \dot{y}_r, \dot{z}_r]$ represents the relative positions and velocities between the chaser and target satellites, the control input $\mathbf{u} = [F_x, F_y, F_z]$ represents the thrust of the chaser satellite in the x, y and z directions. The parameter $n = \sqrt{\mu/a^3}$ is the mean motion of the target satellite, where μ is Earth's gravitational constant and a is length of the target satellite orbit's semi-major axis.

The chaser satellite must maintain a safe relative distance to the target satellite defined the safe set \mathcal{C} : the intersection of the following halfspaces (in kilometers),

$$\begin{aligned} \beta_1(x) &= \{x : 6 \leq x_1 \leq -6\} \\ \beta_2(x) &= \{x : 0 \leq x_2 \leq 4\} \\ \beta_3(x) &= \{x : -10 \leq x_3 \leq 10\} \end{aligned} \quad (13)$$

We constructed the barriers with various values of γ and found $\gamma = 0.05$ balances efficiency in reaching the objective with conservativeness with respect to the obstacles.

We model (12) as a JMLS by constructing a 2-mode system parameterized by the target's nominal mean motion n and off-nominal mean motion n_p . This models uncertainty in the chaser satellite's uncertainty over the mean motion of the target satellite due to, e.g., the target satellite operating under two possible orbit maneuver strategies. We compare our approach against the baselines across a range of off-nominal mean motion value, mode-switch times and delay times. We set the nominal mean motion $n = 0.061$ and range the off-nominal mean motion n_p from 0.041 to 0.101 with a step size of 0.01. Additionally, we range the mode-switch times from 1 to 40 with a step size of 5 and the delay times from 0 to 5 with a step size of 1. Simulation parameters are shown in Tab. I.

TABLE I: Simulation Parameters

Parameters	Orbital Rendezvous	Mineshaft Inspection
Planning Horizon, H	300 s	0.5 s
Sampling time-step, Δt	10 s	0.05 s
Minimum control, u_{\min}	-0.1 N	0.1 N
Maximum control, u_{\max}	0.1 N	20 N
State weight matrix, Q	$\text{diag}[50_{3 \times 3}, 10_{3 \times 1}^{-1}]$	$\text{diag}[50_{3 \times 3}, 10_{9 \times 1}^{-1}]$
Control weight matrix, R	$0.01 \cdot \mathbf{I}_{3 \times 3}$	$0.01 \cdot \mathbf{I}_{6 \times 6}$
Starting position, \mathbf{x}_0	$[0.01, 3.8, 0]^T$ km	$[0, 0, 0]^T$ m
Target location, \mathbf{x}_{ref}	$[1.0, 1.0, 0]^T$ km	$[-0.7, 0.7, -5]^T$ m

TABLE II: Orbital Rendezvous simulation results of our approach against baselines. The best entries are bold.

	First-Step	Full-Step	Non-Robust	Ours
Total Trials	336	336	336	336
Successes	188	230	65	336
Success (%)	56.0	68.4	19.3	100.0
Average Cost	2.05	2.29	2.54	1.95

2) *Hazardous Mineshaft Inspection:* Next, we consider a safety-critical trajectory planning example whereby an autonomous hexacopter is tasked to inspect a hazardous mineshaft. The hexacopter must descend the mineshaft and reach a defined inspection point. At any time, the system can experience a complete rotor failure. Throughout its descent to the goal, the hexacopter must be safe with respect to the safe set \mathcal{C} , a three-dimensional rectangular polyhedron with no ceiling. Here \mathcal{C} is defined as the intersection of the following halfspaces,

$$\begin{aligned} \beta_1(x) &= \{x : -1 \geq x_1 \geq 1\} \\ \beta_2(x) &= \{x : -1 \geq x_2 \geq 1\} \\ \beta_3(x) &= \{x : x_3 \geq -6\} \end{aligned} \quad (14)$$

with $\gamma = 0.05$ for all barriers.

We consider a simplified representation of the hexacopter system, where two rotors are prone to failure at any time. Thus, the system has three modes, one for its nominal dynamics and two for the modes corresponding to the rotor




Fig. 2: Spacecraft Orbital Rendezvous: example state trajectories for our approach and the baselines: First-Step Consensus, Full-Step Consensus, and Non-Robust. Only the in-track (y -direction) relative position is displayed to highlight the chaser spacecraft’s (in)ability to stay behind the target satellite. The mode switch is induced at 35s and the estimation delay is 7s after. The simulation duration is 280s. Lower opacity illustrates problem infeasibility. When the problem is infeasible, the first control input from the mode with the highest likelihood is chosen and executed.

failure. We compare our approach against the baselines across 9 equidistant initial state conditions in the xy -plane representing the entrance of the mineshaft, over a range of mode-switch times (from 1 to 40) and delay times (from 0 to 5). Simulation parameters are shown in Tab. I.

TABLE III: Mineshaft Inspection simulation results of our approach against baselines. The best entries are bold.

	First-Step	Full-Step	Non-Robust	Ours
Total Trials	108	108	108	108
Successes	40	21	58	108
Success (%)	37.0	19.4	53.7	100.0
Average Cost	5.81	4.94	4.74	3.92

B. Results Discussion

Figs. 2 and 3 depict illustrative example simulations of the compared approaches for the Spacecraft Orbital Rendezvous problem and Mineshaft Inspection problem, respectively. In Fig. 2, the Full-Step consensus planner fails immediately after the mode switch is induced due to its inability to recourse after encountering infeasibility. Shortly after, the Non-Robust planner fails. The First-Step consensus planner also encounters infeasibility yet in this example, it manages to control the chaser satellite to the reference point, albeit at a higher average cost with respect to our planner. Meanwhile, our planner safely guides the spacecraft to the in-track (y -direction) reference point with the lowest average cost. In Fig. 3, the Full-Step consensus planner fails first as well, this time followed by the First-Step consensus planner (1.1s) and finally the Non-Robust planner (1.25s). Meanwhile, our planner safely recovers the hexacopter after rotor failure and reaches the desired inspection reference point. Here, the discontinued x and y trajectories result from the colliding with the barrier in the z -direction.

Tab. II and Tab. III show the success rate and average cost of the compared planners for the Spacecraft Orbital

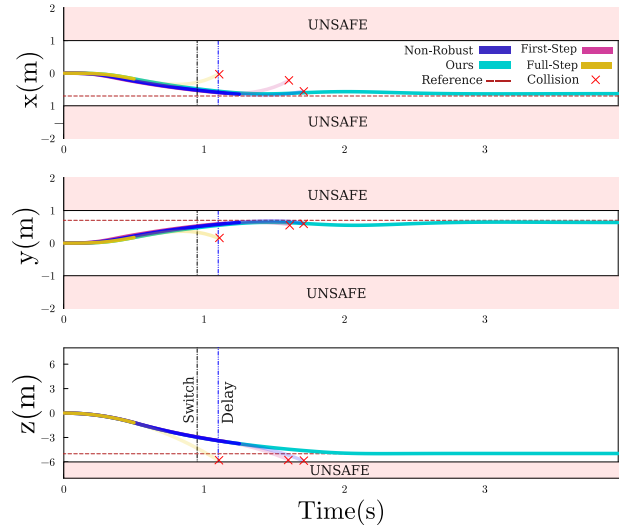


Fig. 3: Hazardous Mineshaft Inspection: example state trajectories for our approach and the baselines: First-Step Consensus, Full-Step Consensus, and Non-Robust. The mode switch is induced at 0.9s and the estimation delay is 0.15s after. The simulation duration is 4 seconds. Lower opacity illustrates problem infeasibility. When the problem is infeasible, the first control input from the mode with the highest likelihood is chosen and executed.

Rendezvous problem and Mineshaft Inspection problem, respectively. The results show that our approach of adaptive consensus horizon planning is the safest and has the lowest cost across both problem environments. In the Spacecraft Orbital Rendezvous problem, the over-optimistic Non-Robust planner performs the worst with the highest average cost across all trials. The First-Step planner with $h = 1$ performs relatively less poorly, with lower cost and higher success rate. Nonetheless, its poor performance is attributed to its short consensus horizon causing feasibility issues when mode switch detections are delayed. The Full-Step planner with $h = H$ performs slightly better than both the First-Step and Non-Robust planners since the problem generally allows feasibility with a maximum consensus horizon. Our planner reaches the desired reference in 100% of the trials, validating our approach of maximizing the consensus horizon subject to feasibility. In the Mineshaft Inspection problem, the Full-Step planner performs the worst out of all the planners in terms of success rate, while the First-Step planner obtains the highest cost. The Non-Robust planner performs better than the other baselines, but worse than our approach. In this problem, since the dynamics are significantly different between each mode and the safe sets are generally smaller, planning too optimistically (First-Step and Non-Robust) or too conservatively (Full-Step) leads to unsafe closed-loop behavior. Our approach is able to balance between robustness over mode uncertainty and problem feasibility, achieving a 100% success rate with the lowest cost.

The results indicate that our approach of computing the maximally feasible consensus horizon allows more flexibility over different problem environments without needing to tune the consensus horizon for each new environment. The high success rate and low cost of our approach over both environ-

ments and varying parameters imply that the combination of the CBF constraints and the consensus horizon is critical for the overall safety and performance of MPC for JMLS.

Finally, Table IV shows the control input rate. Our controller optimizes for control inputs at an average rate of $9.89Hz$ for Orbital Rendezvous and $11.21Hz$ for Mineshaft Inspection. Our approach, although slower than the baselines, is able to maintain an operational control input computation rate while delivering a lower cost and higher success rate.

TABLE IV: Control Input Computation Rate (Hz)

Algorithm	Orbital Rendezvous		Mineshaft Inspection	
	Mean	Std. Dev.	Mean	Std. Dev.
One-Step	48.74	9.42	63.61	11.68
Full-Step	44.37	10.78	67.08	15.29
Non-Robust	53.97	6.51	66.16	10.88
Our Approach	9.89	1.14	11.21	0.81

V. CONCLUSION & FUTURE WORK

Our work presents an MPC framework that synthesizes safety-aware control policies for Jump Markov Linear Systems subject to stochastic mode switches. The MPC framework uses CBFs and tackles the robustness-feasibility trade-off when asserting control consensus constraints over the planning horizon. Simulation experiments illustrate the utility of our proposed approach. For future work, we plan to extend the CBF constraints to include nonlinear safe sets and to account for noise in the dynamics. Additionally, we plan to explore techniques to address the added computation of computing the maximally feasible consensus horizon at each time step. Finally, we plan to validate our control design on hardware experiments.

VI. ACKNOWLEDGMENTS

We thank Michael H. Lim for his insightful discussions. This work was supported in part by NSF grant 203906.

REFERENCES

- [1] B. Canis, "Unmanned Aircraft Systems (UAS): Commercial Outlook for a New Industry," 2015.
- [2] L. Laurenti, M. Lahijanian, A. Abate, L. Cardelli, and M. Kwiatkowska, "Formal and Efficient Synthesis for Continuous-time Linear Stochastic Hybrid Processes," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 17–32, 2020.
- [3] H. A. P. Blom and Y. Bar-Shalom, "The Interacting Multiple Model Algorithm for Systems with Markovian Switching Coefficients," *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, Aug. 1988.
- [4] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*. Springer, 2007, pp. 207–226.
- [5] A. R. Cassandra, "A Survey of POMDP Applications," in *Working notes of AAAI 1998 fall symposium on planning with partially observable Markov decision processes*, vol. 1724, 1998.
- [6] M. J. Kochenderfer, *Decision Making under Uncertainty: Theory and Application*. MIT press, 2015.
- [7] Z. Sunberg and M. Kochenderfer, "Online Algorithms for POMDPs with Continuous State, Action, and Observation Spaces," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 28, 2018, pp. 259–263.

- [8] M. H. Lim, T. J. Becker, M. J. Kochenderfer, C. J. Tomlin, and Z. N. Sunberg, "Optimality guarantees for particle belief approximation of pomdps," *Journal of Artificial Intelligence Research*, vol. 77, pp. 1591–1636, 2023.
- [9] H. Gupta, B. Hayes, and Z. Sunberg, "Intention-aware navigation in crowds with extended-space pomdp planning," in *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022, p. 562–570.
- [10] F. Micheli and J. Lygeros, "Scenario-Based Stochastic MPC for Systems with Uncertain Dynamics," in *2022 European Control Conference (ECC)*. IEEE, 2022, pp. 833–838.
- [11] S. Prajna and A. Jadbabaie, "Safety Verification of Hybrid Systems using Barrier Certificates," in *HSCC*. Springer, 2004, pp. 477–492.
- [12] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control Barrier Function based Quadratic Programs with Application to Adaptive Cruise Control," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 6271–6278.
- [13] R. Mazouz, K. Muvvala, A. Ratheesh, L. Laurenti, and M. Lahijanian, "Safety Guarantees for Neural Network Dynamic Systems via Stochastic Barrier Functions," *Advances in Neural Information Processing Systems*, 2022.
- [14] M. Althoff, C. Le Guernic, and B. H. Krogh, "Reachable Set Computation for Uncertain Time-Varying Linear Systems," in *Proceedings of the 14th international conference on Hybrid systems: computation and control*, 2011, pp. 93–102.
- [15] A. Girard, C. L. Guernic, and O. Maler, "Efficient Computation of Reachable Sets of Linear Time-Invariant Systems with Inputs," in *International workshop on hybrid systems: Computation and control*. Springer, 2006, pp. 257–271.
- [16] R. Cosner, P. Culbertson, A. Taylor, and A. Ames, "Robust safety under stochastic uncertainty with discrete-time control barrier functions," in *Robotics: Science and Systems*, 2023.
- [17] J. Zeng, B. Zhang, and K. Sreenath, "Safety-Critical Model Predictive Control with Discrete-Time Control Barrier Function," in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 3882–3889.
- [18] J. Zeng, Z. Li, and K. Sreenath, "Enhancing feasibility and safety of nonlinear model predictive control with discrete-time control barrier functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 6137–6144.
- [19] H. Ma, X. Zhang, S. E. Li, Z. Lin, Y. Lyu, and S. Zheng, "Feasibility enhancement of constrained receding horizon control using generalized control barrier function," in *IEEE International Conference on Industrial Cyber-Physical Systems*, 2021, pp. 551–557.
- [20] K. P. Wabersich and M. N. Zeilinger, "Predictive Control Barrier Functions: Enhanced Safety Mechanisms for Learning-Based Control," *IEEE Transactions on Automatic Control*, 2022.
- [21] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, 2002, ch. 11, pp. 453–457.
- [22] T. Erez, K. Lowrey, Y. Tassa, V. Kumar, S. Kolev, and E. Todorov, "An integrated system for real-time model predictive control of humanoid robots," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2013, pp. 292–299.
- [23] W. Xiao, C. G. Cassandras, and C. Belta, *Safe autonomy with control barrier functions: Theory and applications*. Springer, 2023.
- [24] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Robotics: Science and Systems*, vol. 13. Cambridge, MA, USA, 2017, pp. 1–10.
- [25] R. Dyro, J. Harrison, A. Sharma, and M. Pavone, "Particle MPC for Uncertain and Learning-Based Control," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 7127–7134.
- [26] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017.
- [27] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex Optimization*. Cambridge university press, 2004.
- [28] F. Nan, S. Sun, P. Foehn, and D. Scaramuzza, "Nonlinear mpc for quadrotor fault-tolerant control," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5047–5054, 2022.
- [29] P. Goulart and Y. Chen, "Clarabel.jl: Interior point conic optimization for julia," <https://github.com/oxfordcontrol/Clarabel.jl>, 2022.
- [30] C. Jewison and R. S. Erwin, "A spacecraft benchmark problem for hybrid control and estimation," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, 2016, pp. 3300–3305.